



Hewlett Packard
Enterprise

HPE'S DAOS SOLUTION PLANS

Lance Evans
Chief Storage Architect
November 19, 2020

RESEARCH PHASE (WRAPPED UP C2020Q3)

Prepare for emerging storage hardware, with HPC-capable data path software

Evaluate DAOS architecture, code quality, maturity, development team, processes

Consider related IP that competes/complements/contributes

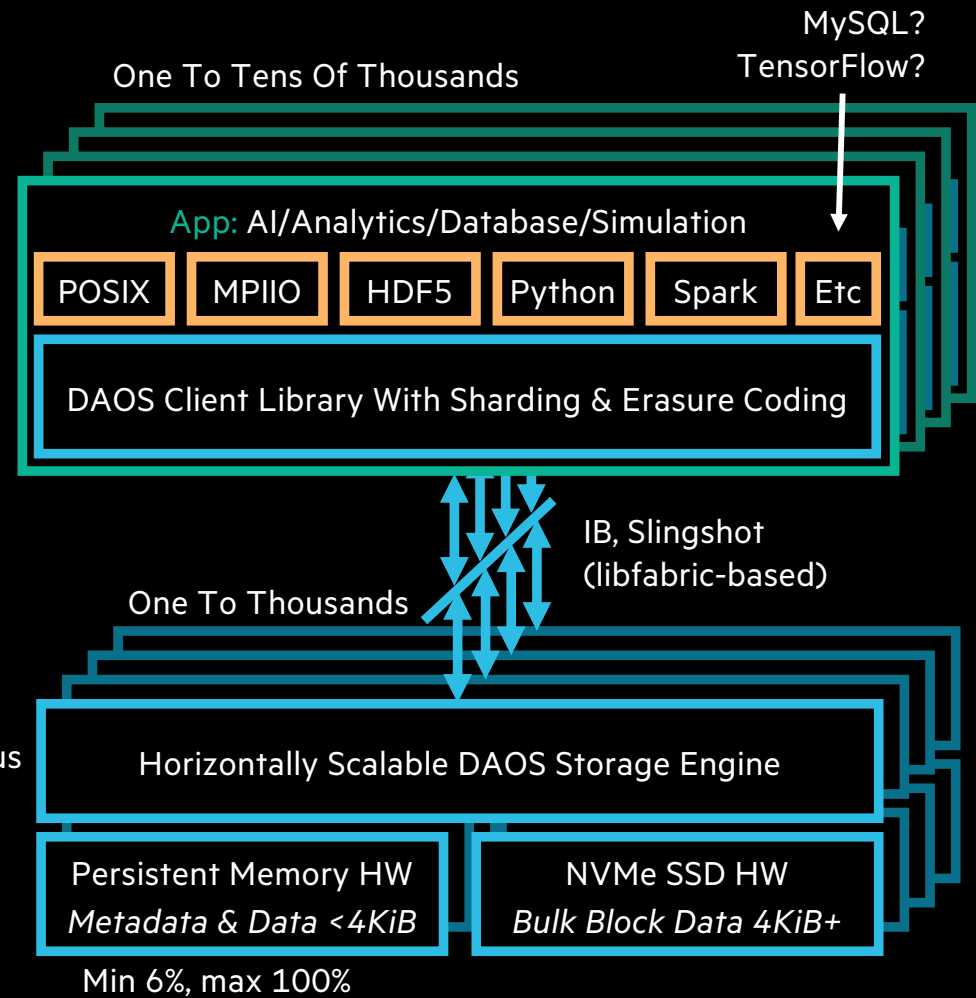
Recommend a path to productization for next-gen HPC data infrastructure

Outcome: Proceed with solution prototyping and technology transfer to Storage R&D



DAOS SOLUTION USE CASES

- **Middleware**
 - Adapts core services to specific use cases
 - No common baggage that all apps must endure
 - All user address space, all open-source
- **Traditional HPC apps**
 - Large shared file with no interference
 - File-per-process on shared devices at speed
 - Simulation checkpoints – burst buffer
- **Non-traditional HPC apps**
 - Machine Learning training and inference
 - Scaled high-level languages
 - Multi-tenant workloads
- **Potential apps**
 - Graph analytics
 - Database
 - Streams processing
 - Computational storage



REFERENCE IMPLEMENTATION GOALS

Harden DAOS data path software, targeting version 2.0

Bundle with released HPE hardware and customized Cray management software

Enable initial customer deployments with repeatability and low risk

Prepare HPE for productization path within the Cray ClusterStor product line




REFERENCE IMPLEMENTATION COMPONENTS

HPE PROLIANT DAOS SERVERS



**INTEL
ICE LAKE,
DCPMM,
NVME**



HPE PROLIANT MANAGEMENT SERVERS



HPE ARUBA MANAGEMENT SWITCHES



MELLANOX 200GB INFINIBAND




HPE SLINGSHOT 200GB ETHERNET



REFERENCE IMPLEMENTATION CHARACTERISTICS

- Single-Rack Solution with Maximums:
 - 16x Servers; 128TB DCPMM, 2PB flash
 - 4x 200Gb Switches (Mellanox or HPE Slingshot)
 - ~1,400GBps/700GBps est raw read/write throughput
 - ~100M/50M est raw read/write OPS
- Unbundled Repeatable Solution Delivery Method
 - Qualified hardware and software BOM
 - Installation / configuration framework and scripting
 - Reference doc set for field or factory integration
 - Customer system administration skills needed
- Availability
 - Seeking early adopters for package trials C2021Q3
 - Individual elements sold and supported individually
 - DAOS, solution scripts, docs via engineering collaboration
 - Full HPE DAOS product pacing is TBD



2x 1/10GbE Aruba Mgmt Switches
3x HPE Proliant Management Servers

4x 200Gb Switches (optional):

- HPE Slingshot (80 uplinks max)
- Mellanox QM8700 (72 uplinks max)

16x HPE DAOS Servers:

- 2-Socket Ice Lake
- Gen4 NVMe SSD up to 128TB
- Optane Memory 8TiB

WHAT'S NEXT?

- Our results so far lead us to believe DAOS is a viable data path for future HPC Storage designs.
- HPE's HPC division (formerly Cray) will participate with Intel in delivering the Aurora supercomputer to Argonne National Labs, one of the largest computers in the world, and containing over 200PB of DAOS.
- HPE's HPC CTO lab will host a small number of potential customers who want to run POCs in our testbed.
- We will collaborate with early adopters, integrating working DAOS-based storage solutions in their data centers, collecting product requirements and feedback.
- Assuming this trial phase is successful, we intend to proceed toward full productization of DAOS in the Sapphire Rapids timeframe.
- Combined with our scalable HPC system management, HPE server designs, and Slingshot networking, HPE will continue to offer the most scalable and performant storage for the world's largest computers.



Hewlett Packard
Enterprise

CONTACT:

lance.evans@hpe.com