Solution brief

HPE CRAY SUPERCOMPUTING STORAGE SYSTEMS K3000

Supercomputing storage for the AI for science era

Unleash the most data-intensive AI applications running on your supercomputer with fast, yet cost-effective all flash storage systems.

The storage spending growth challenge in converged HPC and AI supercomputing

Just a few ago, storage accounted for only 10% to 15% of the budget[1] for the next generation supercomputer of organizations that depend on supercomputing-scale high performance computing (HPC) to achieve their mission or business objectives.

As the #1 vendor in the converged HPC/AI market,[2] every year HPE has thousands of meetings with organizations like this and recently, there is one consistent theme in many of those meetings when it comes to supercomputing storage that gets included into the procurement process for next-generation supercomputers.

With the convergence of classic HPC workloads like physics-driven modeling and simulation and data-driven artificial intelligence (AI) models on modern supercomputers, many bids for next-generation supercomputers from other vendors contain proprietary AI storage from storage-only companies that now consume 25% to 30% of the budget for the next-generation supercomputer. This diverts resources from compute nodes, hindering application performance, and ultimately slowing scientific progress.

If this supercomputing storage spending growth trend continues unchecked, more and more of available budget for the next-generation supercomputer of those organizations will be consumed by the storage portion of the request for proposal (RFP). This diverts resources from maximizing investment into compute nodes that ultimately drives application performance and scientific discovery.

HPE Cray SC Storage Systems to the rescue

The founder of Cray—Seymour Cray—once famously said "Anyone can build a fast CPU. The trick is to build a fast system."[3]

Adjusted to the supercomputing spending growth challenge of our time, that quote can be modified as follows:

In times of solid-state storage media anyone can build a fast supercomputing storage system. The trick is to build a fast and cost-effective storage system that does not divert more resources away from compute nodes than absolutely necessary.

This exactly is the design mandate for HPE Cray SC Storage Systems.

Before we dive into the portfolio of HPE Cray SC Storage Systems, let explore some of the fundamental design mandates that all HPE Cray SC Storage Systems share.

[1] HPE internal analysis of supercomputing RFP responses, 2025.

[2] "Welcome to the ISC25 Hyperion Research market update," Hyperion Research, June 2025.

[3] "About Seymour Cray," IEEE Computer Society.

Figure 1 lists those mandates in a table.

Design mandates	Customer benefit
Integrate open-source supercomputing storage software owned by a community and not by a single company	Less vendor lock-in compared to proprietary file systems
Supercomputing storage systems are products of the supercomputing system vendor and not third-party products	Prevents multiple layers of markups by streamlining pricing through a single vendor, resulting in greater cost savings for our customers
Absence of capacity-based software licensing per terabyte storage capacity or storage drive	Lower storage cost compared to proprietary file systems
24x7 technical support through in-house R&D team for hardware and the embedded open-source storage software	Less downtime compared to open-source software with community support only
Strong read and write performance for both AI and classic HPC workloads	Higher compute node utilization compared to file systems that are read-optimized only
Protocol support for SMB, NFS, and S3	Seamless integration into the data pipeline of the organization
Optimized storage servers that extract the maximum performance from the storage media	Lower storage cost compared to nonoptimized storage servers
Factory-tested storage systems with turnkey implementation services and 24x7 support	Faster time to results and lower maintenance efforts compared to reference architectures

Figure 1. Design mandates for HPE Cray SC Storage Systems

With the convergence of classic HPC workloads like physics-driven modeling and simulation and data-driven artificial intelligence (AI) models on modern supercomputers, all HPE Cray SC Storage Systems are purpose-engineered to support both converged workloads—that could not be more different from each other (see Figure 2)—exceptionally well.

Classic HPC storage systems	HPE Cray SC Storage Systems	AI storage
Optimized for **physics-based modeling & simulation**	**Both**	Optimized for **data-driven AI**
Optimized for serving **flagship x86 CPU**	**Both**	Optimized for serving **flagship GPU**
Optimized for **write-intensive I/O**	**Both**	Optimized for **read-intensive I/O**
Optimized for **large file sizes**	**Both**	Optimized for **all file sizes**
Optimized for **sequential data access**	**Both**	Optimized for **random data access**
Optimized for **data throughput**	**Both**	Optimized for **IOPS**

Figure 2. Characteristics of classic HPC storage systems and AI storage

The portfolio of HPE Cray SC Storage Systems

The portfolio consists of HPE Cray SC Storage Systems E2000 with embedded open-source Lustre[4] file system and HPE Cray SC Storage Systems K3000 with embedded open-source distributed asynchronous object storage (DAOS)[5] software.

Figure 3 provides an overview of the portfolio.

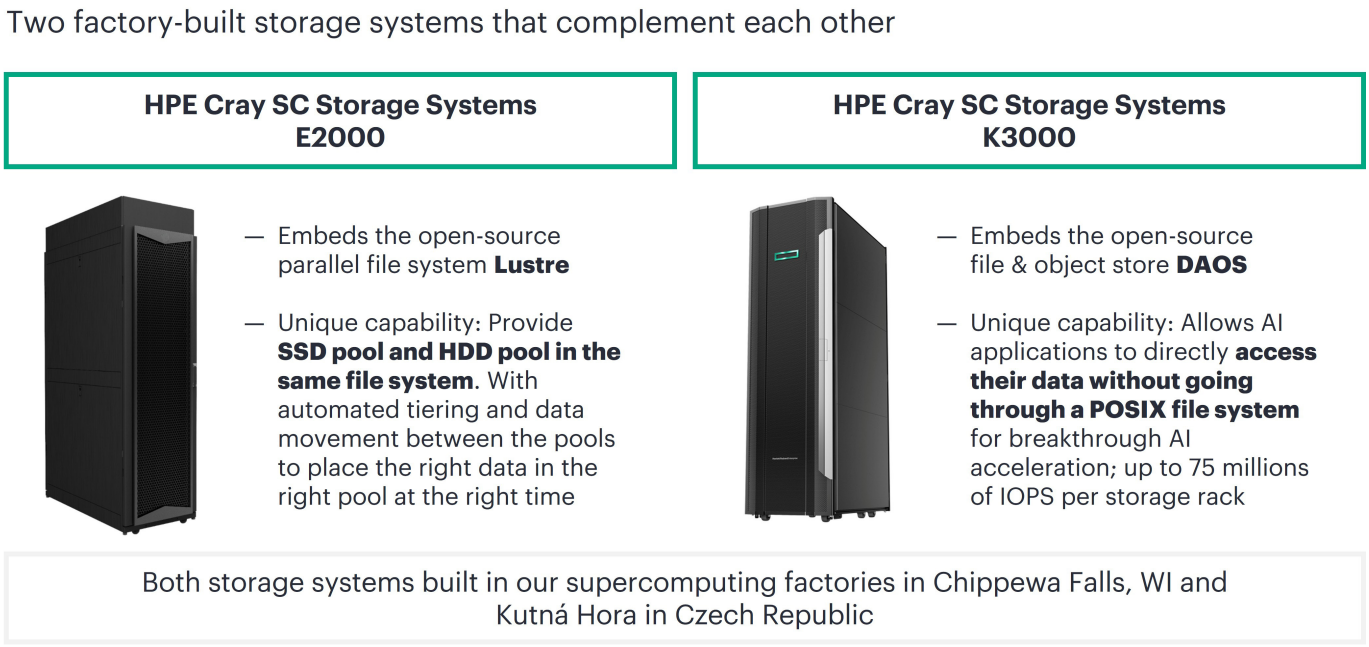

Figure 3. Portfolio of HPE Cray SC Storage Systems

The deployment scenarios for both systems are very straightforward:

— HPE Cray SC Storage Systems E2000 gets deployed with every HPE Cray SC EX4000 or GX5000 solution to provide the bulk of fast and cost-effective shared file storage for the next-generation supercomputer of converged HPC/AI cluster of the customers.

— If the customer also has emerging AI applications that have I/O performance requirements beyond about 50 million IOPS per storage rack, then HPE Cray SC Storage Systems K3000 gets deployed alongside HPE Cray SC Storage Systems E2000.

HPE Cray SC Storage Systems K3000

HPE Cray SC Storage Systems K3000 embed the innovative open-source DAOS software running on HPE ProLiant Compute DL360 Gen12 servers into a factory-tested and fully integrated storage system.

DAOS is a persistent data software platform for converged HPC/AI workloads that

— **Fully leverages performance of modern flash storage media**

- Ultra-low-latency networked data transfers
- High-concurrency transactional operations without locks
- Direct interaction of applications with networked flash

[4] Lustre.

[5] DAOS.

— **Scales linearly—from on rack to exascale-class systems**

- Hundred percent user address space with kernel bypass
- Object-oriented namespace, no centralized metadata
- Racks of x86 NVMe servers as storage nodes

— **Protects data against failures**

- Intelligent data sharding across small or large clusters
- Erasure coding across servers, flexible object by object
- Fault domains with rebuild across servers

— **Supports multiple application interfaces**

- Native key-value, byte-addressable object, and arrays
- Parallel POSIX through layered DAOS File System (DFS)
- Many application-specific middleware packages (see Figure 4)

— **Is fully open source**

- BSD 2-Clause licensed
- DAOS Foundation (under Linux® Foundation)

This unique combination of features along with a differentiated software ecosystem for AI, analytics, and modeling and simulation workflows (show in Figure 4) makes HPE Cray SC Storage Systems a perfect companion to be deployed alongside HPE Cray SC Storage Systems E2000.

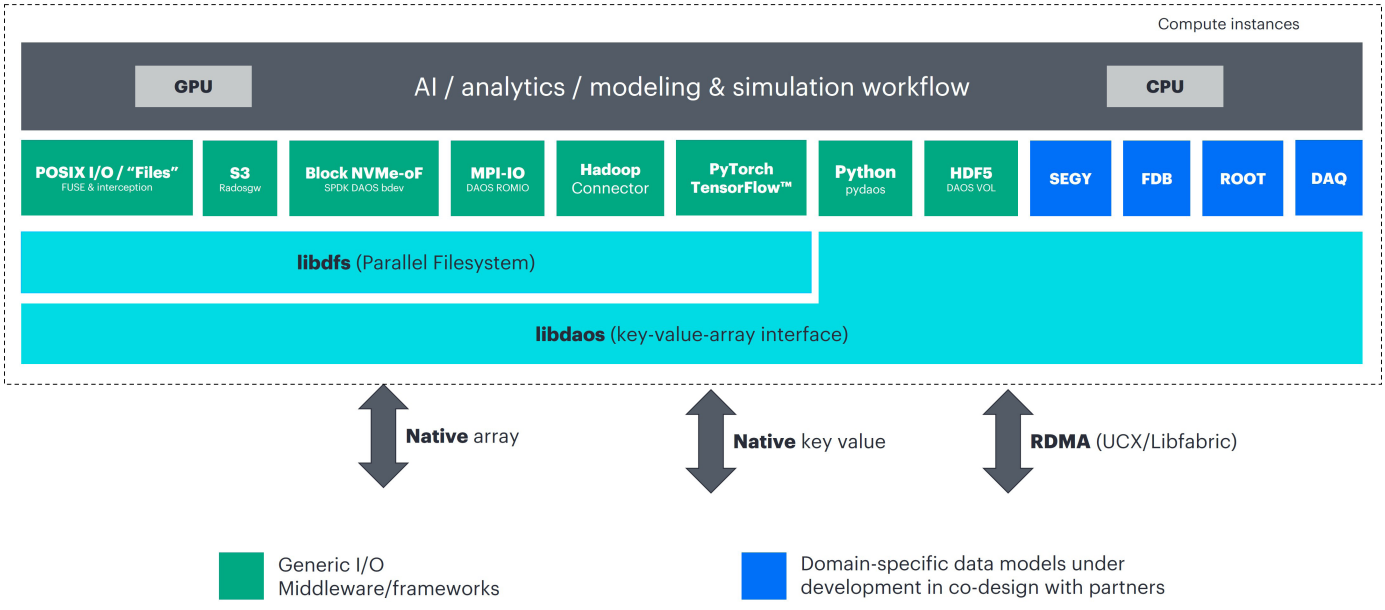

Figure 4. The software ecosystem of HPE Cray SC Storage Systems K3000

The HPE ProLiant Compute DL360 Gen12 is a perfect DAOS storage server—if it would not have already existed, we would have to built it for HPE Cray SC Storage Systems K3000.

A minimum configuration of HPE Cray SC Storage Systems K3000, for example, for a proof of concept (POC) is:

— One rack

— One DAOS management node (HPE ProLiant Compute DL325 Gen11)

— At least four DAOS server nodes (HPE ProLiant Compute DL360 Gen12) in identical configuration with

- Up to 20 E3.S 1 DWPD NVMe drives per server node
  - 2 Intel® CPUs with 48 cores each (96 total cores per server node)
  - 20 NVMe drives for the capacity-optimized configuration
  - 8, 12, or 16 NVMe drives for performance-optimized configurations
- Memory: 512 GB, 1,024 GB, or 2,048 GB DRAM configurations based on selected drive size
- Available drive sizes: 3.84 TB, 7.68 TB, and 15.36 TB
- Available interconnect interfaces: HPE Slingshot 200/400, InfiniBand NDR, 400GbE

With HPE Cray SC Storage Systems K3000 you can start small, but if needed scale to 1000s of DAOS server nodes in one storage system that supports 10000 of compute nodes.

Unprecedented market penetration and unbelievable storage benchmark results

Today, our Lustre-based supercomputing storage systems are powering all of the world's exascale supercomputers, 6 of the world's top 10 most powerful supercomputers, 55% of the aggregate compute power of the world's top 100 most powerful supercomputers, and 70% of the aggregate compute power or the world's most energy-efficient supercomputers.[6]

The recent addition of DAOS-based HPE Cray Supercomputing Storage Systems K3000 further strengthens our supercomputing storage leadership position.

HPE Cray SC Storage Systems K3000 already have been selected alongside our Lustre-based supercomputing storage systems as the storage for Oak Ridge National Laboratory's next-generation exascale supercomputer "Discovery".[7]

The performance leadership position of DAOS is evidenced by the fact that the #1 and #2 systems in the global supercomputing storage benchmark IO500 are DAOS-based and together have four times the storage benchmark performance than the next 30 storage systems on the list combined.[8]

[6] HPE internal analysis of Top500 and Green500 lists, June 2025.

[7] HPE to build two systems for Oak Ridge National Laboratory: Next-generation exascale supercomputer "Discovery" and AI cluster "Lux".

[8] "Production ISC25 list," IO500.

Spend less on storage and get more compute nodes with HPE

If you are concerned that more and more of the available budget for your next-generation supercomputer is consumed by the storage portion—diverting resources from boosting investment into compute nodes that ultimately are driving application performance and scientific discovery—then HPE Cray SC Storage Systems can help you through a unique combination of

— Embedded open-source file systems that are owned by vibrant communities and not by a single company

— Absence of capacity-based licensing per terabyte or per storage drive

— Optimized storage servers that extract the breakthrough performance from the underlying storage media

All of that delivered to you with a turnkey implementation and a 24x7 operational support experience.

Get started

Every HPE Cray SC Storage Systems configuration is uniquely tailored to meet your specific needs. Each design is customized to address your individual workloads, with application performance optimization driving the configuration to help ensure peak efficiency. The process begins with a complimentary discovery workshop, where your team collaborates with our experts to explore the HPC and AI applications critical to achieving your mission or business objectives. Based on these insights, our specialists design a supercomputing solution optimized for the highest application performance within your budget—so you can spend less on storage and get more compute nodes.

If you are ready to spend less on storage and more on compute, click the **Chat now** button to ask for a meeting with our HPE Cray SC Storage Systems experts and schedule your personalized workshop today.

Learn more at

HPE.com/Cray

Visit HPE.com

Chat now

a00154232ENW, Rev. 1

HEWLETT PACKARD ENTERPRISE

hpe.com