Custom Ceph Storage

Software-Defined Storage at Petabyte Scale

ONE PLATFORM, EVERY INTERFACE

A single cluster delivers all three storage types at once: S3- and Swift-compatible **object** storage for data lakes and apps, **block** storage (RBD) for VMs and databases, and a POSIX **file** system (CephFS) for shared HPC scratch and home directories — one pool of capacity, no silos.

ENGINEERED TO SCALE

Capacity and performance grow together. Add nodes and the cluster expands from hundreds of TB to many petabytes with no central bottleneck. Ceph's CRUSH algorithm places and finds data without a metadata server to choke under load. Rebalancing runs automatically in the background.

RESILIENT BY DESIGN

Data is protected through replication or erasure coding and distributed across drives, nodes, racks, and sites to prevent a single point of failure. When a drive or server fails, Ceph detects it and self-heals, re-creating lost copies with no admin intervention and no downtime. Durability is configurable per pool.

ADS CUSTOM CEPH AT A GLANCE

Storage types	Object, block, and file; unified
Protocols	S3, Swift, RBD, iSCSI / NVMe-oF, NFS, CephFS
Data protection	Replication or erasure coding; self-healing
Scale	Hundreds of TB to multi-PB; linear scale-out
Media tiers	NVMe flash, SAS/SATA HDD, hybrid
Networking	Up to 1.6Tb Ethernet; up to 800Gb InfiniBand
Management	Linux Ceph dashboard; Prometheus Grafana telemetry

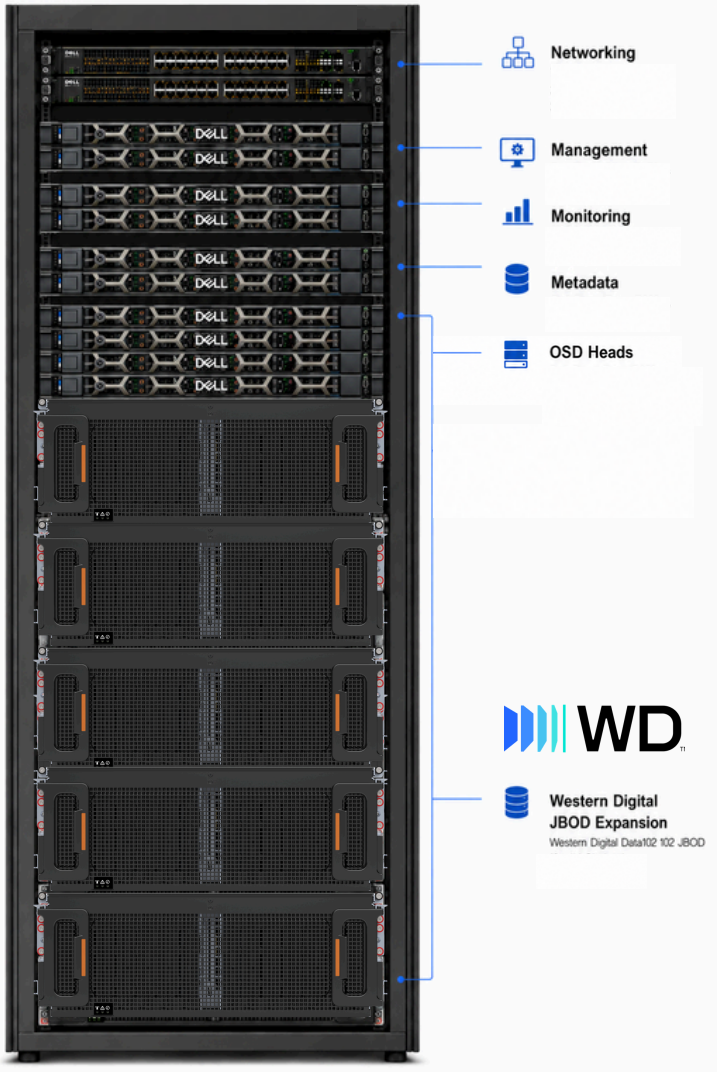

OPEN SOURCE APPLIANCE. ANY SCALE. NO LOCK-IN.

With training data, sensor and ISR feeds, simulation output, and model checkpoints measured in petabytes, proprietary storage appliances can only keep pace with that growth through per-terabyte licensing, rigid scaling, and forklift upgrades. ADS engineers custom Ceph clusters into a single open, software-defined appliance that serves object, block, and file data from a single namespace, scales linearly as you add nodes, and heals itself when hardware fails, with no per-terabyte license and no vendor lock-in.

WHERE IT FITS

- AI / ML training data lakes feeding GPU clusters
- HPC shared scratch, home directories & checkpointing
- Sensor, video & ISR repositories for autonomy and drone programs
- On-prem, S3-compatible object storage to replace public cloud

ADS CEPH ARCHITECTURE

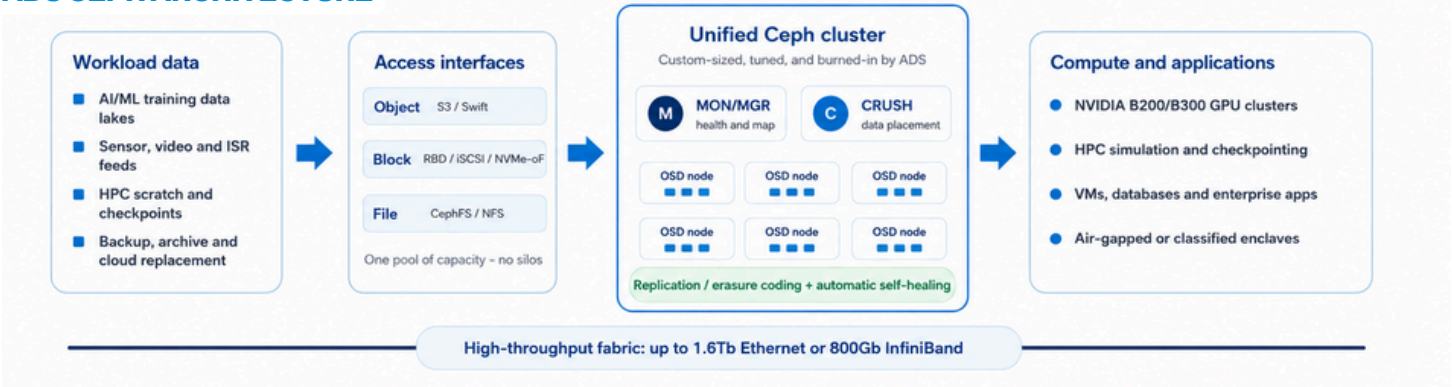

THE ADS CUSTOM CEPH ADVANTAGE

✓ **Custom-architected, not a black box**
Every cluster is sized, tuned, and burned-in for your workload at the ADS integration center in Poway, CA, then delivered ready to run.

✓ **Open source, zero license lock-in**
Ceph is open and hardware-agnostic. Scale on standard servers and NVMe/HDD media with no per-terabyte fees or proprietary controllers.

✓ **One integrated stack**
Storage, compute, networking, security (CrowdStrike, Palo Alto, Fortinet), and OS (Rocky Linux, Red Hat) designed and supported together.

✓ **TAA-compliant & federal-ready**
Built on TAA-compliant hardware for ITAR, CUI, and CMMC environments — including air-gapped and classified enclaves.

✓ **GPU-adjacent performance**
Engineered to feed NVIDIA B200/B300 training and inference clusters over up to 1.6Tb Ethernet or 800Gb InfiniBand, so GPUs stay fed, not idle.

✓ **One partner**
Architecture, integration, and lifecycle support from one vendor: no finger-pointing across vendors.

LIFECYCLE SUPPORT UP TO 7 YEARS

DS backs every Custom Ceph cluster with hardware support terms extending up to seven years, sized to the multi-year horizon of federal programs and long-lived AI/HPC infrastructure. Rather than forcing a refresh on a vendor's timetable, you keep the cluster supported and serviceable across its full operational life, protecting the capital investment and avoiding a disruptive forklift upgrade mid-program.

MISSION-CRITICAL SLAS UP TO 24X7X4

Support tiers scale to the criticality of the workload, up to 24x7x4 coverage — around-the-clock monitoring and response with four-hour parts and engineering turnaround. For always-on inference, production data lakes, and classified operations where downtime isn't an option, the cluster stays covered every hour of every day.

OPTIONAL PROFESSIONAL SERVICES & SUPPORT

Beyond the hardware, ADS offers optional professional services to stand the cluster up and keep your team self-sufficient — onsite installation and integration at your facility, full deployment and configuration, and administrator training and knowledge transfer. Optional Ceph software support contracts are also available for organizations that want a vendor-backed subscription with a defined escalation path, pairing ADS hardware support with ongoing software-level coverage for the Ceph platform itself.

FLEXIBLE MEDIA: ALL-NVME OR HYBRID SSD + HDD

Clusters are configured to the workload. All-NVMe flash builds deliver maximum throughput and the lowest latency for the most demanding jobs — feeding GPU training pipelines, high-IOPS databases, and metadata-intensive operations. Where capacity economics matter more than peak speed, hybrid SSD + HDD configurations pair high-density drives for bulk petabyte storage with flash for caching and metadata acceleration, keeping cost per terabyte low without sacrificing responsiveness. Both approaches can be combined into tiered pools within a single cluster.

ROLE-DEDICATED NODES FOR PERFORMANCE

ADS architects each cluster with nodes dedicated to specific Ceph functions rather than collapsing every role onto shared hardware. Monitor, manager, metadata (MDS), and object gateway (RGW) services run on their own nodes, separate from the OSD nodes that hold data. Isolating roles eliminates resource contention, delivers predictable low-latency performance under load, and lets each tier scale independently as the cluster grows.

LET'S ARCHITECT YOUR CLUSTER

ADS starts with your data: capacity, throughput, protocols, and growth curve, and designs a Ceph cluster around it.

Solution Brief

For more information from Applied Data Systems:
844.371.4949 | info@applieddatasystems.com | www.applieddatasystems.com