OPEN COMPOSABLE COMPATIBILITY LAB

Solution Brief

Next-Generation Tiered Object Storage with Ceph Tentacle, OpenFlex® Data24, and Ultrastar® Data60 3000

Challenges

Organizations dealing with AI, analytics, backup, and cloud-native workloads face multiple challenges:

- Rapid data growth requiring scalable and cost-efficient storage.
- Balancing performance and cost across diverse workloads.
- Managing hot (frequently accessed) vs. cold (archival) data efficiently.
- Eliminating data silos across structured and unstructured storage.
- Ensuring high availability, durability, and data protection.
- Simplifying storage operations across large-scale distributed environments.
- Optimizing storage footprint while maintaining performance SLAs.

Highlights

- Tiered storage ensures the right data is stored on the right media.
- OpenFlex Data24 accelerates performance-critical workloads.
- Ultrastar Data60 3000 reduces cost for long-term data storage.
- Ceph Tentacle provides the intelligence and scalability backbone
- Massive scalability with object-based storage architecture.
- Strong data protection via replication or erasure coding.
- Optimized TCO through dynamic data placement.
- Distributed architecture with MON (cluster state), OSD (data storage), MGR (automation & monitoring), Intelligent Data Placement.
- Automatic data distribution, Dynamic rebalancing using the CRUSH algorithm without central bottlenecks.
- Cloud-Native Object Storage: S3/Swift compatible via RADOS Gateway (RGW).

Modern storage environments require intelligent data placement to balance performance and cost. This solution leverages Ceph Tentacle's tiering capabilities to distribute data across high-performance and capacity-optimized storage platforms.

The architecture introduces a tiered storage model that improves performance for critical workloads while reducing storage costs through automated data movement between tiers. Resources utilization is optimized across the infrastructure, and organizations gain enhanced scalability without a linear increase in cost.

Hot Tier: Manages latency-sensitive workloads such as analytics, AI/ML, and transactional operations.

Cold Tier: Stores infrequently accessed data, archives, and backups at lower cost.

This design leverages:

- Ceph's distributed, software defined architecture.
- OpenFlex's NVMe™ based, disaggregated storage fabric.
- Ultrastar Data60 3000 is a high-density, hybrid HDD-based storage platform.

Key Components

OpenFlex Data24 (Hot Tier)

- Provides ultra-low latency NVMe storage
- Used for active datasets requiring high throughput

Ultrastar Data60 (Cold Tier)

- Provides dense, cost-efficient HDD storage
- Ideal for archival and infrequently accessed data

Benefits of Tiering

- Improved performance for critical workloads
- Reduced storage cost through intelligent data placement
- Automated movement of data between tiers
- Optimized resource utilization across the infrastructure
- Enhanced scalability without linear cost increase

OpenFlex Data24 4000 Series Storage Platforms

The OpenFlex Data24 4000 series NVMe-oF™ storage platform extends the high-performance of NVMe flash to shared storage. The 4000 series provide low-latency sharing of NVMe SSDs over a high-performance Ethernet fabric to deliver similar performance to locally attached NVMe SSDs. WD RapidFlex™ NVMe-oF controllers allow up to six dual pathed hosts to be attached without a switch. The OpenFlex Data24 4200 series uses WD's RapidFlex C2000 Fabric Bridge Adapters to provide up to 12 ports of 100GbE which can connect to RDMA and/or TCP configured host ports.

Next-Generation Tiered Object Storage with Ceph Tentacle, OpenFlex Data24, and Ultrastar Data60 3000

Features

- NVMe-over-Fabrics (NVMe-oF) architecture
- Bandwidth match between PCIe Gen4 SSDs (Storage) and I/O (Network) – No oversubscription
- RESTful API support for simplified management
- N+2 fan redundancy eliminates the need for hot- swappable fan modules Data24 4100
- Connect directly to an Ethernet NIC in your host or through an Ethernet switch
- Device Sharing optional configuration utilizes NVMe Non-Transparent Bridging to allow all Ethernet ports to access any SSD
- High IOPS and throughput performance
- Software-defined composability
- Supports Ethernet-based disaggregated storage
- High availability with redundant components
- Ideal for AI/ML, real-time analytics, and transactional workload

Ultrastar Data60 3000 Overview

The Ultrastar Data60 3000 meets the evolving demands where data growth is constant and scale changes everything. By delivering up to 1.92PB of raw storage in a compact footprint, the Data60 3000 allows organizations to expand storage efficiently without increasing operational complexity. Twelve ports of high-bandwidth 24Gb/s SAS-4 connectivity ensures that data continues to move efficiently across workloads, supporting sustained performance rather than isolated peak metrics. ArcticFlow™ cooling improves airflow and thermal efficiency, while IsoVibe™ vibration isolation protects drive integrity and stability. Together, these innovations reduce failures and improve long-term reliability by up to 62%. The result is a storage platform that enhances efficiency, lowers total cost of ownership, and supports modern disaggregated architectures. Built for cloud, neo-cloud, and HPC systems, Ultrastar Data60 3000 helps organizations scale data infrastructure confidently while maintaining performance, reliability, and economics.

Features

- Up to 60 Ultrastar HDDs (SAS or SATA)
- Up to 1.92PB of raw storage in 4U
- Up to four units may be daisy-chained for a total raw capacity of 7.68PB
- Patented IsoVibe technology ensures maximum performance even in heavy workloads
- Improved cooling from innovative ArcticFlow technology
- Choose dual-port SAS for high availability or single-port SATA for low cost
- Up to 12 x 24Gb/s SAS-4 host connections
- Enterprise-grade redundant and hot swappable PSUs, IO Modules, and fans
- Rack-mounted top cover for quick and easy service

Ceph Tentacle Object Storage

Ceph offers object storage tiering capabilities to optimize cost and performance by seamlessly moving data between storage classes. These tiers can be configured locally within an on-premises infrastructure or extended to include cloud-based storage classes, providing a flexible and scalable solution for diverse workloads. With policy-based automation, administrators can define lifecycle policies to migrate data between high-performance storage and cost-effective archival tiers, ensuring the right balance of speed, durability, and cost-efficiency. Ceph Tentacle introduces performance and efficiency improvements, particularly for erasure coded workloads.

Next-Generation Tiered Object Storage with Ceph Tentacle, OpenFlex Data24, and Ultrastar Data60 3000

Horizontal scalability – Offers horizontal scalability with no single point of failure

Self-healing and self-managing – Automatically detect failures and rebalances data with minimal administrative effort.

Native S3 and Swift APIs – Provides industry standard object interfaces for seamless application integration.

Strong consistency model – Ensures immediate data consistency, unlike many eventually consistent object stores.

Hardware and vendor agnostic – Runs on commodity hardware and supports diverse storage media (HDD, SSD, NVMe).

Multi-tenancy with quotas and isolation – Enables secure sharing of object storage across teams or workloads.

Unified storage platform – Object storage coexists with block and file storage under the same Ceph cluster.

Open-source and cloud-native – Backed by a strong community, avoiding vendor lock in and supporting hybrid or private clouds. Ceph object storage provides S3 compatible access with strong consistency, self healing, and horizontal scalability.

Tentacle Innovations – Offers Faster BlueStore WAL and improved compression, Optimized RGW bucket listing via faster OMAP iteration, Enhanced dashboard and Mult-cluster management, Performance improvements for EC (FastEC).

Multi-site replication and disaster recovery support – Enables automatic data replication across geographically distributed sites to ensure business continuity and rapid recovery in case of failures or disasters.

Advanced lifecycle management for automated data tiering – Automates data movement between storage tiers based on access patterns and policies, optimizing performance and cost without manual intervention

Reference Architecture

The solution architecture integrates Ceph Tentacle object storage on industry-standard servers with OpenFlex Data24 and Ultrastar Data60 3000 platforms, enabling a high-performance, cost-efficient tiered storage environment. In this design, OpenFlex Data24 acts as the hot tier, storing and serving frequently accessed, latency-sensitive data, while Ultrastar Data60 3000 functions as the cold tier, hosting older or infrequently accessed data in a dense and cost-efficient manner. The Ceph platform intelligently manages data placement and movement across these tiers, ensuring transparency and operational simplicity.

This architecture leverages Ceph Tentacle's advanced tiering capabilities to seamlessly distribute data across high-performance NVMe storage and high-capacity HDD-based storage. Ceph OSDs are deployed following the Tentacle architecture, enabling independent scaling of compute and storage resources. OpenFlex Data24 provides shared NVMe storage, accessed over low-latency fabrics such as NVMe-over-Fabrics (RoCEv2/TCP), delivering exceptional performance for active workloads. In parallel, Ultrastar Data60 offers high-density storage enclosures, designed to deliver scalable, resilient, and cost-optimized capacity for colder datasets.

At its core, Ceph operates as a distributed, software-defined storage platform built on RADOS, where Object Storage Daemons (OSDs) manage data across both tiers. Within the Ultrastar Data60 environment, OSDs primarily reside on high-capacity HDDs, optionally enhanced with SSDs for BlueStore WAL/DB acceleration, enabling a hybrid storage approach that combines flash performance for metadata and hot data with economical disk storage for bulk data.

Ceph MON and MGR services run on the storage nodes to maintain cluster state, orchestration, and health monitoring, while RADOS Gateway (RGW) provides S3- and Swift-compatible APIs for seamless cloud-native application access. The architecture separates client (public) and cluster (replication) networks, ensuring optimal performance, fault isolation, and efficient data movement. Data placement is governed by the CRUSH algorithm, enabling intelligent, decentralized distribution without bottlenecks, while ensuring balanced utilization across the cluster and tiers.

OpenFlex Data24 provides shared NVMe drives accessed over a low latency fabric such as NVMe-oF (RoCEv2/TCP). This separation enables granular resource scaling, improves utilization, and removes local disk constraints from OSD nodes. The result is a highly resilient, cloud native object storage solution that delivers consistent performance, simplified operations, and enterprise grade durability, well suited for demanding workloads such as analytics. Each Ultrastar Data60 3000 enclosure contributes dense, high-capacity storage nodes, supporting petabyte-to-exabyte scale growth with linear scalability. The result is a highly resilient, cloud-native object storage solution that delivers consistent performance, enterprise-grade durability, and simplified operations with backup and tired storage. This tiered approach is ideally suited for data-intensive workloads such as analytics platforms, AI/ML data lakes, backup repositories, and large-scale private cloud environments, where both performance and cost optimization are critical.

Next-Generation Tiered Object Storage with Ceph Tentacle, OpenFlex Data24, and Ultrastar Data60 3000

Ceph: Object Storage Deployment Diagram with OpenFlex Data24 and Ultrastar Data60 3000

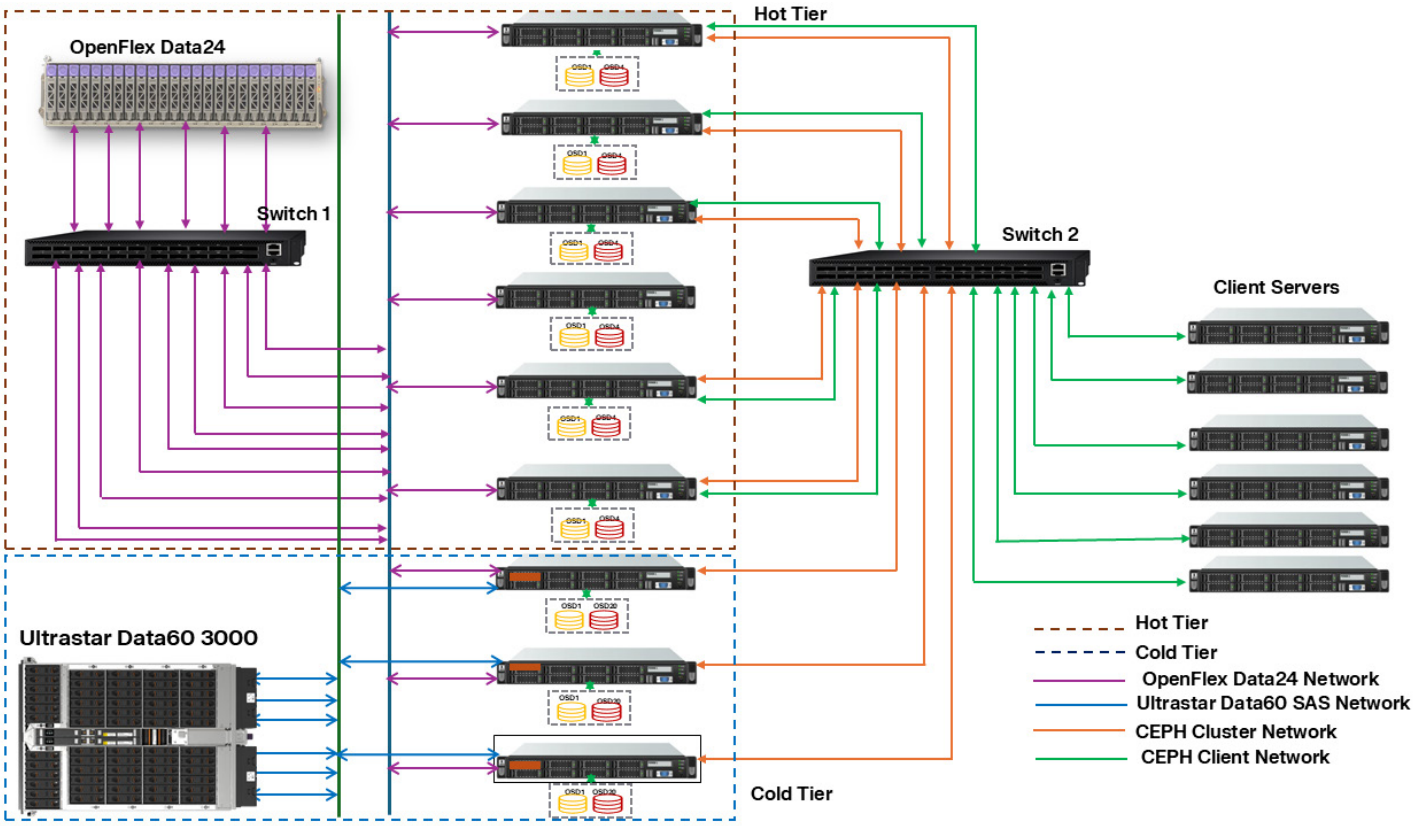

Ceph Storage class Tiering Details

Ceph storage classes define different data tiers (like HOT, COLD) mapped to specific storage pools, and lifecycle rules automatically move objects over time based on policies, enabling cost and performance optimization without manual intervention. Here OpenFlex Data24 drives are configured as part of the hot tier pool and Ultrastar Data60 3000 HDDs are configured as part of the cold tier pool.

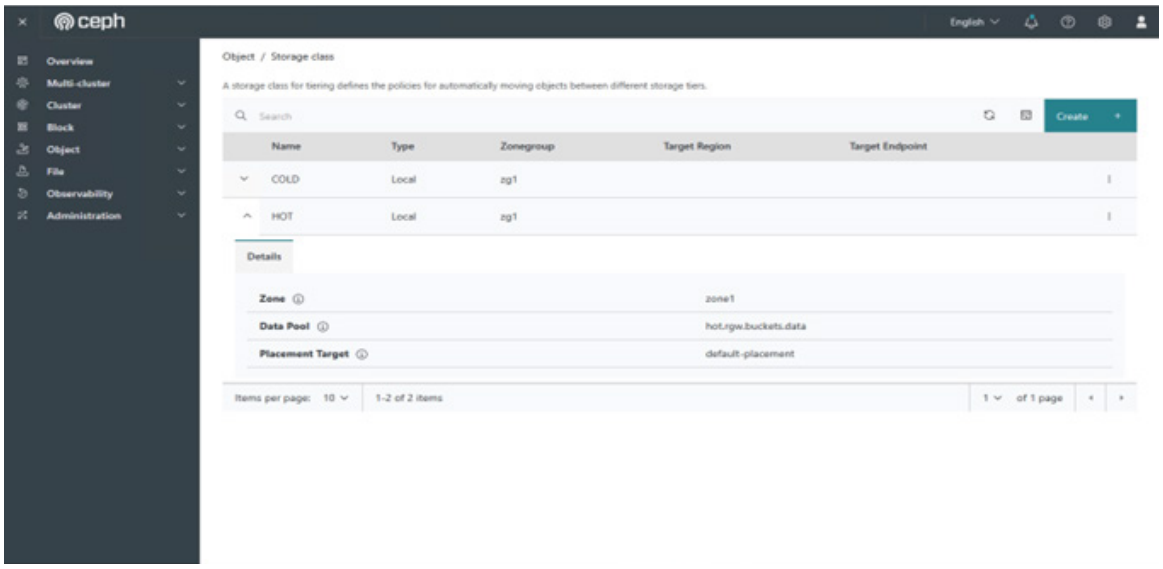

Next-Generation Tiered Object Storage with Ceph Tentacle, OpenFlex Data24, and Ultrastar Data60 3000

Ceph Bucket Life Cycle Rule for Storage Tiring and Life Cycle Progress Status

Ceph A lifecycle transition in Ceph RGW is triggered by lifecycle policies and occurs when:

- **Object age condition is met**: The object reaches the number of days defined in the lifecycle rule (e.g., Days: 1).
- **Lifecycle (LC) engine execution**: The RGW lifecycle engine, a background process, periodically scans buckets and evaluates lifecycle rules.
- **LC worker scheduling**: The transition happens when the lifecycle worker runs and processes eligible objects during its scheduled cycle

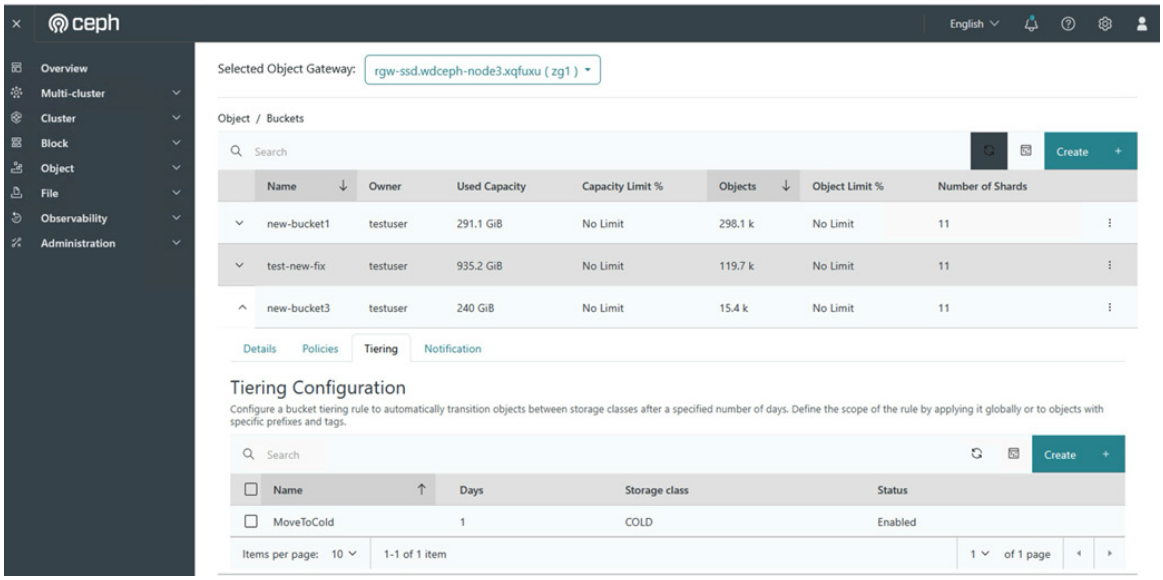

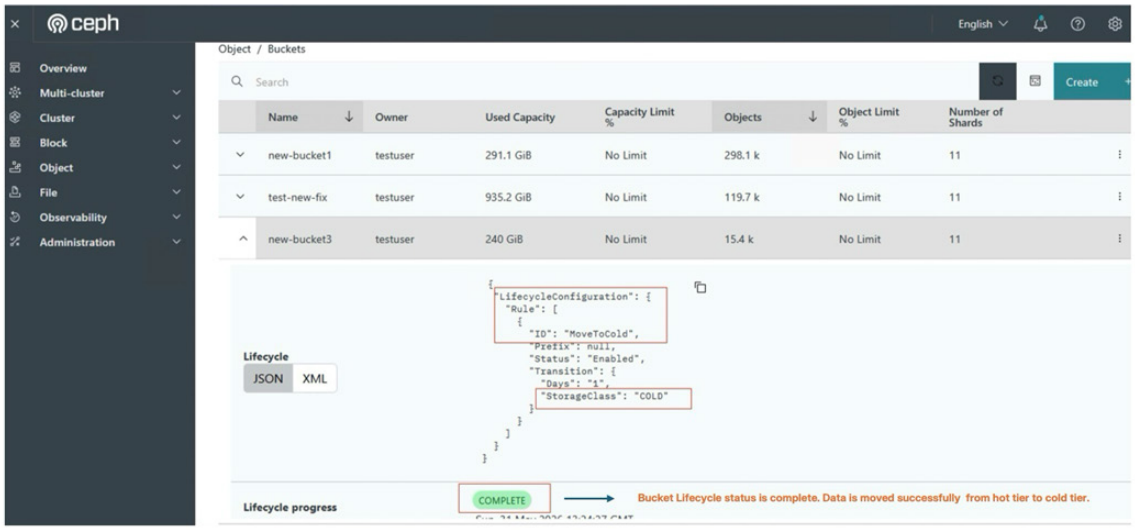

Next-Generation Tiered Object Storage with Ceph Tentacle, OpenFlex Data24, and Ultrastar Data60 3000

Performance Details

The measured performance numbers are specific for the current configuration of Ceph Tentacle with OpenFlex Data24 topology serving as the hot tier, in conjunction with the uses of 24 NVMe SSDs/osds with multipath path connection via six storage nodes. Warp tool is used to run the performance benchmarking on the system. The performance test is run using warp, simultaneously running tests from all the 6 Ubuntu clients.

Ceph Object Storage Replication Performance with OpenFlex Data24 4000

Policy Details	Profile	Throughput
Replication (3-Copies)	GET	29.03 GBps
	PUT	6.37 GBps

Ceph Object Storage Erasure Coding Performance with OpenFlex Data24 4000

Policy Details	Profile	Throughput
Erasure Coding (4+2)	GET	14.40 GBps
	PUT	6.38 GBps

Note: The performance results are captured after running multiple iterations with different data size (object size, objects, concurrency, policies).
Disclaimer: All the performance data are subject to change depending on the Drive model, capacity, workloads, servers, CPU, tunings and HA topology used.

Configurations Details

Metrics	Details
Storage Product	OpenFlex Data24 4000 and Ultrastar Data60 3000
Ceph Version	Ceph Tentacle version 20.2.1
Storage Interface	**OpenFlex**: 12 x QSFP28 Interface (12 x 100Gbps per port) **Ultrastar**: SATA
HBA Details	eHBA 9600-16e Tri-Mode Storage Adapter, FW: 8.17.1.0-00000-00001, Driver: 8.17.1.0.0
Server Details	PowerEdge® R550
Host OS	Ubuntu 22.04.05
Host NIC	CX6
Firmware Version	24.10-4.1.4 (22.32.2004 (MT_0000000437))
CPU	Intel® Xeon® Gold 5318Y CPU @ 2.10GHz or Higher CPU Configuration
CPU Core Details	Dual socket server with 24 core CPU each. 96 logical cores in total with HT enabled
Memory	256 GiB
Switch Details	NVIDIA® Onyx™ 3700
Drive Details	Ultrastar DC HC560 20TB CMR SATA Drives, WUH722020BLE604, PQGSWGC8 Sandisk® SN655 (7.68TB), WUS4BA176DSP3X4, R2210801 – one NVMe drive connected to each Storage server
Tool	Warp v1.4.0

Next-Generation Tiered Object Storage with Ceph Tentacle, OpenFlex Data24, and Ultrastar Data60 3000

Advantages

Optimized Cost-Performance Ratio

Achieves the right balance between performance and cost across storage tiers:

- High-performance NVMe tier (OpenFlex Data24) for latency-sensitive workloads.
- Cost-efficient HDD tier (Ultrastar Data60 3000) for bulk and archival data.

Seamless Scalability Across Tiers

Enables linear and flexible growth as data expands:

- Scale hot and cold tiers independently based on workload demand.
- Add storage nodes without disruption or reconfiguration.
- Supports petabyte-to-exabyte scale environments.

Automated Data Lifecycle Management

Delivers intelligent and policy-driven data placement:

- Automatically moves data between hot and cold tiers.
- Reduces manual intervention and operational overhead.
- Ensures optimal resource utilization across storage layers.

High Availability with No Single Point of Failure

Built for enterprise resilience and continuous availability:

- Distributed Ceph architecture eliminates single points of failure.
- Self-healing and automatic data rebalancing capabilities.
- Supports replication and erasure coding for data protection.

Reduced Latency for Critical Applications

Optimized performance for demanding workloads:

- NVMe-based OpenFlex Data24 ensures ultra-low latency access.
- Accelerates analytics, AI/ML, and transactional operations.
- Minimizes response time for frequently accessed datasets.

Efficient Storage Utilization

Maximizes value from storage infrastructure:

- Stores hot data on flash and cold data on cost-effective disks.
- Optimizes capacity usage through intelligent data placement.
- Reduces unused or overprovisioned storage resources.

Simplified Operations with a Unified Platform

Streamlines storage management across environments:

- Single Ceph platform for object storage across tiers.
- Unified management with cephadm orchestration and automation.
- Reduced complexity compared to siloed storage systems.

Future-Ready Architecture for AI/ML and Big Data Workloads

Designed to support next-generation data-driven use cases:

- Supports AI/ML pipelines and large-scale analytics.
- Ready for cloud-native and containerized environments.
- Enables multi-site deployments with geo-replication support.

Conclusion

This solution combines Ceph Tentacle's intelligent object storage capabilities with the performance of OpenFlex Data24 and the efficiency of Ultrastar Data60, delivering a powerful and flexible tiered storage architecture. It enables organizations to manage large-scale data environments efficiently while optimizing both performance and cost.

Key Takeaways

- Tiered storage ensures the right data is stored on the right media.
- OpenFlex Data24 accelerates performance-critical workloads.
- Ultrastar Data60 3000 reduces cost for long-term data storage.
- Ceph Tentacle provides the intelligence and scalability backbone.
- The architecture is highly adaptable for modern data-driven enterprises.

5601 Great Oaks Parkway
San Jose, CA 95119, USA
www.westerndigital.com

June 2026 / Revision 02