OPEN COMPOSABLE COMPATIBILITY LAB

Solution Brief

OpenFlex® Data24 4200 with Ceph Object Storage: The Foundation for Modern Data Platforms

Challenges

Modern data platforms must balance performance, scalability, resiliency, and cost efficiency while supporting unstructured data growth.

Key challenges addressed in this solution include:

- Scaling object storage performance beyond traditional SAN/NAS limits.
- Managing cost and capacity efficiency at multi petabyte scale.
- Achieving predictable performance across replication and erasure coded (EC) policies.
- Efficiently utilizing NVMe SSDs without being bottlenecked by CPU, network, or software layers.
- Simplifying deployment and lifecycle management of distributed storage infrastructure.

Highlights

- High performance Ceph Tentacle object storage deployed on OpenFlex Data24.
- NVMe optimized, disaggregated storage architecture.
- Validated performance across replication and EC (4+1, 4+2) policies.
- Demonstrated balance between throughput, resilience, and storage efficiency.

Solution

This solution integrates Ceph Tentacle Object Storage with OpenFlex Data24 to deliver a scalable, high throughput object storage platform.

The design leverages:

- Ceph's distributed, software defined architecture.
- OpenFlex's NVMe™ based, disaggregated storage fabric.
- Modern erasure coding and replication policies optimized for object workloads.

This results in a flexible object storage platform suitable for analytics, AI/ML data lakes, backup, archive, and cloud native applications.

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.

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.

CRUSH-based data placement – Eliminates centralized metadata and enables deterministic, scalable, and highly resilient data distribution.

Massive horizontal scalability – Scales from a few nodes to exabytes of object data without architectural changes.

Self-healing and self-managing – Automatically detects 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).

Integrated erasure coding and replication – Offers flexible data protection schemes to balance durability, performance, and cost.

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.

OpenFlex Data24 4200 with Ceph Object Storage: The Foundation for Modern Data Platforms

Reference Architecture

The Ceph Tentacle object storage reference architecture with OpenFlex Data24 combines Ceph's massively scalable, software-defined object storage with high performance NVMe disaggregation to deliver a modern, efficient storage platform. In this design, Ceph OSDs are deployed using the Tentacle architecture to independently scale compute and storage, while 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, AI/ML data lakes, and large scale private cloud environments.

Ceph: Object Storage Deployment Diagram with Open Flex Data24

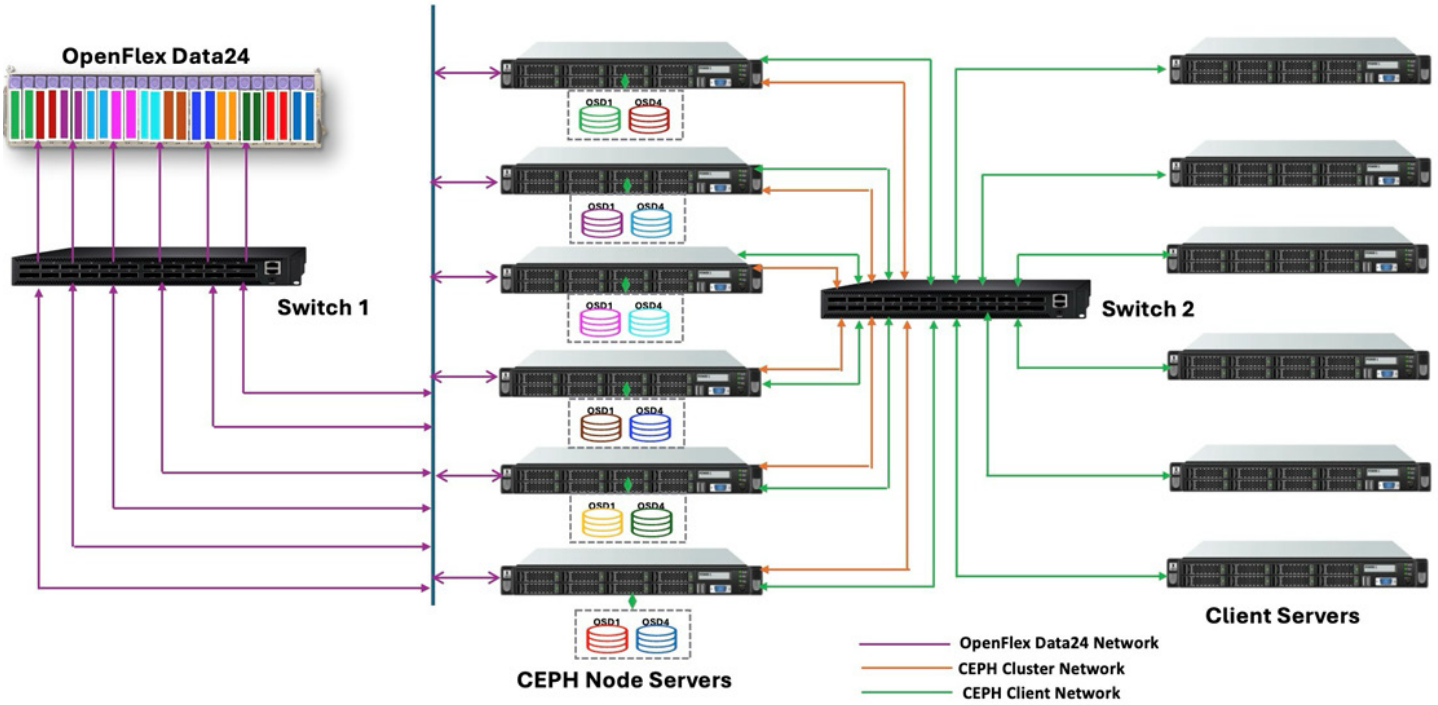

Performance Details

The measured performance numbers are specific for the current configuration of Ceph Tentacle with OpenFlex Data24 topology in conjunction with the uses of 24 NVMe Drives/Volumes with multipath connection via 6 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

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

Ceph Object Storage Erasure Coding Performance with OpenFlex Data24

Policy Details	Profile	Throughput EC (4+2) Policy	Throughput EC (4+1) Policy
Erasure Coding	GET	14.40 GBps	16.25 GBps
	PUT	6.38 GBps	8.15 GBps

OpenFlex Data24 4200 with Ceph Object Storage: The Foundation for Modern Data Platforms

Ceph Replication(3-Copies) 6Node Performance with OpenFlex Data24

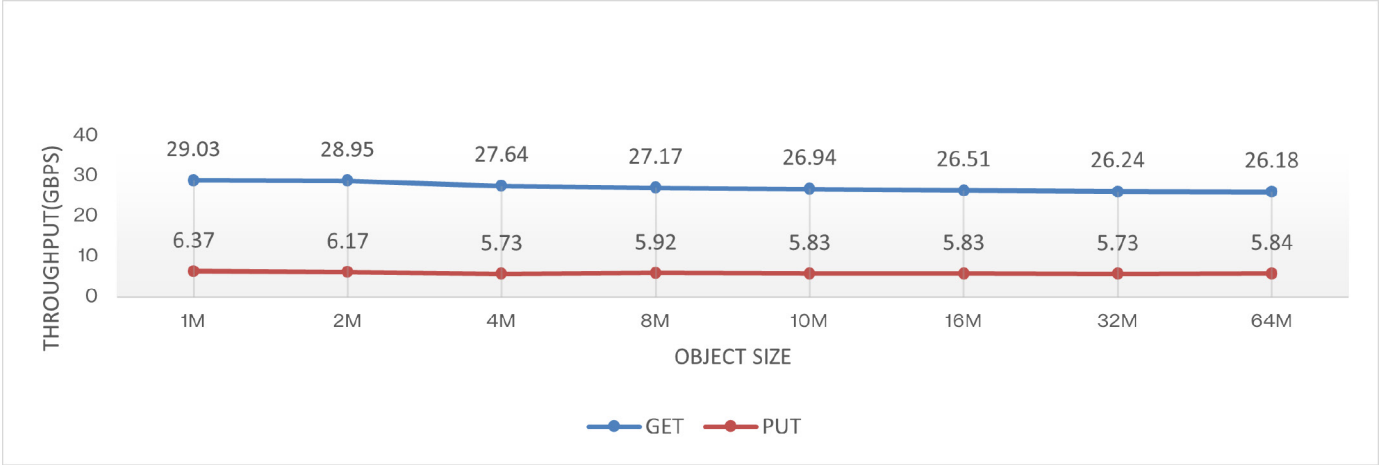

Ceph EC(4+2) 6Node Performance with OpenFlex Data24

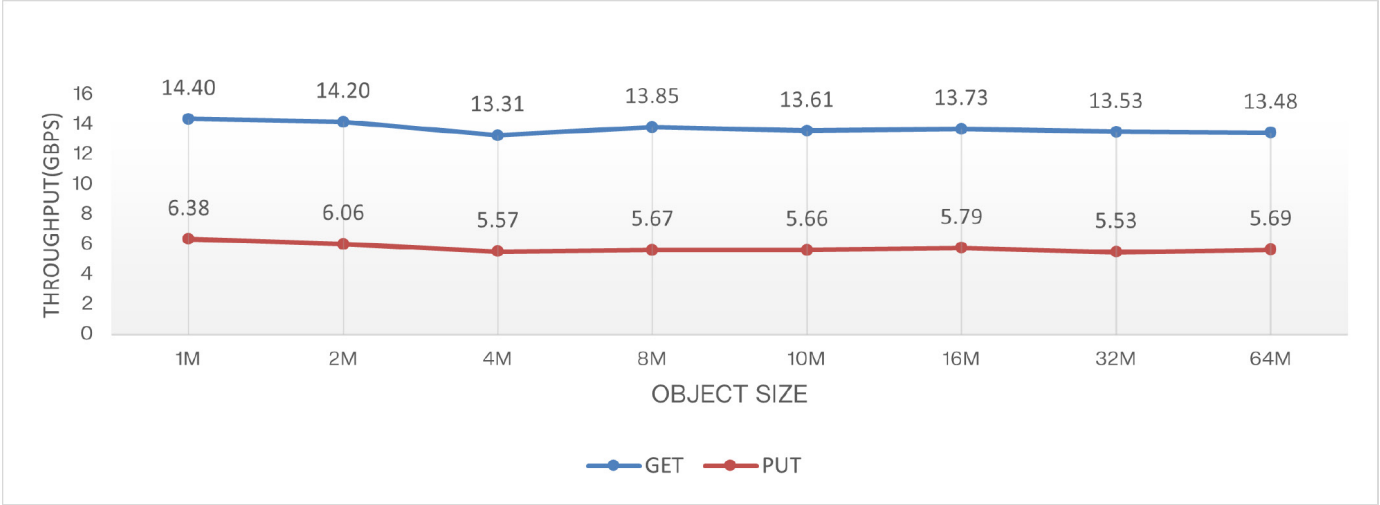

Ceph EC(4+1) 6Node Performance with OpenFlex Data24

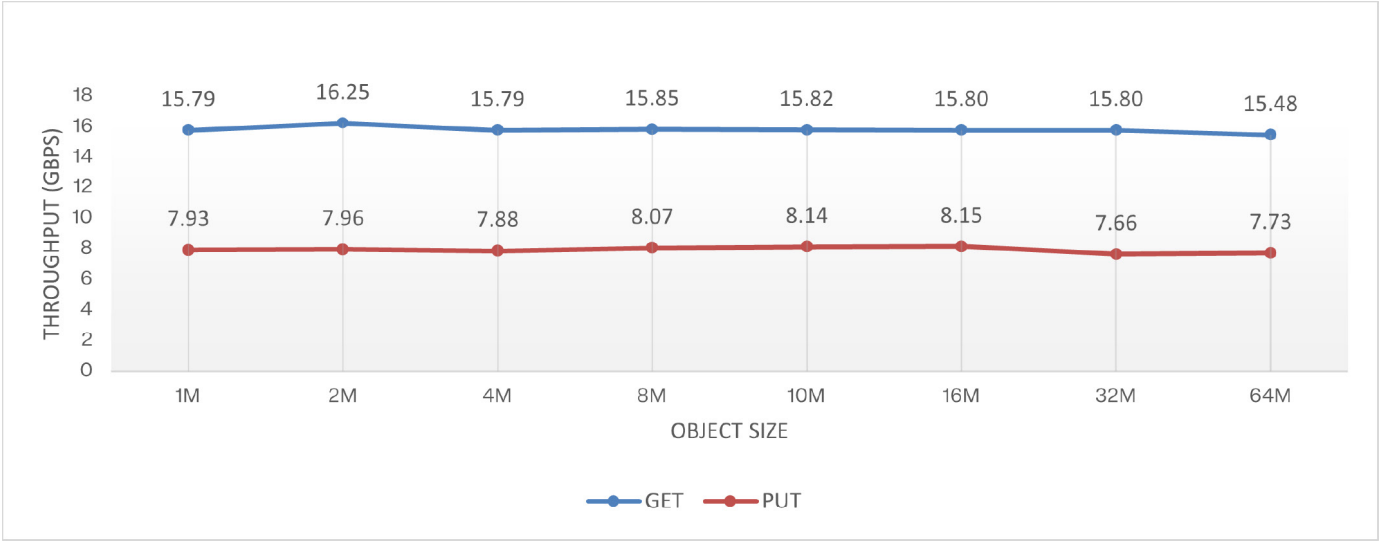

Note: The performance results are captured after running multiple iterations with different data size (object size, objects, concurrency, policies).
Disclaimer: All the performance data are subject to change depending on the drive model, capacity, workloads, servers, CPU, tunings and HA topology used.

OpenFlex Data24 4200 with Ceph Object Storage: The Foundation for Modern Data Platforms

Configurations Details

Metrics	Details
Storage Product	OpenFlex Data24 4200
Ceph Version	Ceph Tentacle version 20.2.1
Storage Interface	12 x QSFP28 Interface (12 x 100Gbps per port)
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	Sandisk® SN655 (3.84TB), WUS5EA138ESP7E4, RC610008
Tool	Warp v1.4.0

Advantages

- **High throughput:** Multi tens of GB/s object performance on NVMe
- **Scalability:** Linear performance scaling with additional storage nodes
- **Efficiency:** EC reduces raw storage overhead compared to replication
- **Resilience:** Host level fault tolerance and self healing
- **Flexibility:** Support for mixed replication and EC tiers
- **Operational simplicity:** Cephadm based lifecycle management

Conclusion

The combination of Ceph object storage with OpenFlex Data24 represents a forward looking architecture that unites software defined scalability with high performance, disaggregated NVMe infrastructure. Ceph's unique strengths—such as CRUSH based deterministic data placement, strong consistency, and unified object storage - are further enhanced by OpenFlex Data24's composable NVMe over Fabrics design, which delivers predictable low latency, high throughput, and exceptional drive density. By decoupling storage media from compute, this architecture enables independent scaling, improves resource utilization, and simplifies lifecycle management while maintaining resilience and durability. Together, Ceph and OpenFlex Data24 provide a flexible, future ready object storage platform ideally suited for modern workloads like AI/ML pipelines, cloud native applications, and large scale data lakes, all while reducing operational complexity and total cost of ownership.

Ceph Tentacle Object Storage on OpenFlex Data24 delivers a robust, high performance object storage platform optimized for modern data intensive workloads. The solution demonstrates that:

- Disaggregated NVMe architectures can unlock significant object storage throughput
- Ceph Tentacle efficiently balances performance and capacity through replication and erasure coding
- Proper CPU, core, and concurrency tuning is critical to maximizing EC performance

This solution is well suited for enterprises seeking a scalable, cost efficient, and future ready object storage foundation.

[1] Actual result depend on CPU memory bandwidth, RNIC count/link speed, PCIe layout, and fabric configuration.
[2] One terabyte (TB) is equal to one trillion bytes and one petabyte (PB) is equal to 1,000 TB. Actual user capacity may be less due to operating environment.

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

May 2026 / Revision 01