



Lightbits vs Ceph Storage: A Practical Performance and Resiliency Comparison

February 2026

Author: Tom Dau, Principal Solutions Engineer

Table of Contents

Executive Summary	2
1. Introduction	3
2. Test Methodology	3
2.1. Cluster Topology and Hardware Configuration	3
2.2. Volume Replication Configuration	4
2.3. Storage Access Methods	4
2.4. Workload Profiles	4
3. Baseline Performance: Canonical 4K Workloads	4
3.1. Observed Behavior	4
3.2. Implications	5
4. Scaling Characteristics Under Concurrency	6
4.1 Multi-Client 4K 100% Random Read Workloads	6
4.2. Observations	6
5. Write Performance Analysis	7
5.1 Multi-Client 4K 100% Random Write Workloads	7
6. Resiliency Testing: Node Failure and Recovery	7
6.1. Lightbits Failure Behavior	8
6.2. Ceph Failure Behavior	9
6.3. Observations	9
7. Rebuild Behavior and Operational Impact	9
8. Augmenting Ceph with Lightbits	10
9. Extended Workload Validation	10
10. Conclusion	12
About Lightbits Labs™	13



Executive Summary

Ceph is widely recognized as one of the most flexible and broadly adopted distributed storage platforms in the industry. Its ability to provide block, object, and file storage within a unified architecture has made it a foundational component of many modern infrastructure deployments. This flexibility, however, can introduce trade-offs, particularly for latency-sensitive and performance-critical block storage workloads.

This white paper presents a structured, lab-driven comparison of Lightbits software-defined storage and Ceph, with a specific focus on the block storage behavior under realistic operating conditions. Using identical hardware classes, an equal number of storage nodes, and industry-standard benchmarking tools, the analysis evaluates steady-state performance, scaling characteristics, and system behavior during node failure and recovery. Replication factors were selected based on test objectives: a replication factor of two (RF2) was used for performance characterization, while a replication factor of three (RF3) was applied for resiliency and failure testing. This methodology was intentionally chosen to enable an apples-to-apples comparison and to isolate architectural behavior from scale-driven effects.

While Ceph performance can improve in very large-scale deployments, these gains are typically achieved by increasing the cluster size and node count. However, adding more nodes to improve performance would also affect failure behavior, operational simplicity, and hardware efficiency, all of which can directly influence the total cost of ownership. The focus of this study is on performance and resilience comparisons of commonly deployed small-to-medium clusters under latency-sensitive block workloads; TCO analysis is outside the scope of this study.

The results show that Lightbits block storage consistently delivers better performance, linear scaling, and faster recovery time for latency-sensitive workloads, while Ceph continues to excel as a general-purpose, unified storage platform. Rather than positioning these solutions as direct competitors, this paper illustrates how Ceph Storage environments can be augmented by offloading performance-critical block workloads to Lightbits software, allowing Ceph to retain its role as the primary platform for file and object storage.

Intended audience: Infrastructure architects, storage engineers, platform engineers, SREs, and technical decision-makers evaluating block storage performance and resiliency.

Note

A technical white paper based on Lightbits Labs' own testing in a controlled lab environment.

1. Introduction

Distributed storage platforms are often tested using synthetic benchmarks that focus on peak throughput or specific performance metrics. While these tests are useful, they often fail to capture system behavior under real-world operating conditions, such as sustained load, increasing concurrency, and infrastructure failures.

This paper aims to provide a practical comparison of Lightbits and Ceph by examining how the systems behave under common operational tools and typical workloads. The goal isn't to declare a single "best" solution, but to highlight the architectural differences and operational trade-offs that affect how well each platform fits different use cases.

2. Test Methodology

All testing was conducted by Lightbits Labs in a controlled laboratory environment designed to minimize external variables and avoid benchmark-specific optimizations.

2.1. Cluster Topology and Hardware Configuration

Both platforms were deployed using a three-node storage cluster configuration

For the Ceph deployment:

- Three storage nodes
- Seven NVMe SSD-backed OSDs per node
- Twenty-one total OSDs across the cluster
- Equal OSD weights and balanced distribution across hosts
- Ceph version : 20.2.0 (Tentacle w/Bluestore)

Ceph monitor (MON) services were distributed across the storage nodes, consistent with commonly deployed small-to-medium cluster architectures. No dedicated monitor-only nodes were used in this configuration.

For the Lightbits deployment:

- Three storage nodes
- Seven NVMe/SSD devices per node
- Twenty-one total storage devices across the cluster
- Lightbits version: 3.17.1



An identical disk count and hardware class were used for both platforms to ensure an apples-to-apples comparison and eliminate media or node-count bias.

All storage media consisted of NVMe devices connected via 100Gb Ethernet networking.

2.2. Volume Replication Configuration

Replication factors were selected based on test objectives and were applied identically across both platforms:

- **Performance testing:** Replication factor of 2 (RF2)
- **Resiliency testing:** Replication factor of 3 (RF3)

2.3. Storage Access Methods

- **Lightbits:** NVMe/TCP
- **Ceph:** RBD over TCP/IP

2.4. Workload Profiles

The following workload categories were evaluated:

- Canonical 4K block size (queue depth = 1, single job)
- Multi-volume scaling (one versus two volumes)
- Multi-client, high queue depth workloads
- 100% read, 100% write, and mixed read/write profiles
- Sustained I/O during node failure and recovery events

No artificial caching, throttling, or platform-specific tuning was applied. The results reflect default, production-oriented configurations intended to represent typical deployment behavior.

3. Baseline Performance: Canonical 4K Workloads

The canonical 4K workload (queue depth of one, single job) establishes a baseline for latency by minimizing queueing effects. This workload highlights architectural efficiency rather than concurrency management.

3.1. Observed Behavior

Lightbits demonstrated consistently lower, more stable latency across tested volumes, and near-linear IOPS scaling as additional volumes were introduced. Latency remained largely unchanged as throughput increased.

Ceph showed a higher baseline latency. While IOPS increased with additional volumes, latency improvements were marginal, and variability was more noticeable. (Figures 1 and 2)

Figure 1: Canonical 4K Random Write (QD=1) IOPS and average latency for Lightbits and Ceph.

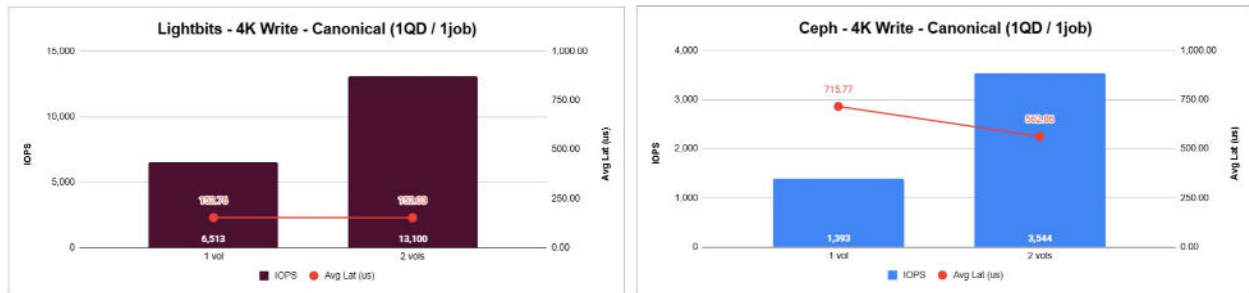


Figure 2: Canonical 4K Random Read (QD=1) IOPS and average latency for Lightbits and Ceph.



3.2. Implications

This workload closely represents latency-sensitive applications such as databases, VM boot volumes, and infrastructure control-plane services, where predictable response times are often more critical than peak throughput. In practice, this includes common operations such as metadata access, index lookups, transaction log writes and flushes, small synchronous reads, and state updates performed by control-plane components.

4. Scaling Characteristics Under Concurrency


Increasing concurrency exposes differences in how distributed storage systems manage internal coordination, replication, and resource contention.

4.1 Multi-Client 4K 100% Random Read Workloads

Lightbits scaled IOPS proportionally with increasing job counts and client concurrency. Latency increased gradually and predictably, with no early saturation point observed. (Figure 3)

Ceph showed an early plateau in achievable IOPS as concurrency increased. Latency rose sharply beyond this point, indicating increased queuing and internal contention rather than additional usable throughput.

Figure 3: Multi-client 4K 100% random read scaling (IOPS and average latency) under increasing concurrency.



4.2. Observations

Lightbits demonstrated outward scaling under concurrency, while Ceph reached a throughput ceiling beyond which additional load was absorbed primarily through increased latency.

5. Write Performance Analysis

Write workloads provide greater visibility into replication behavior and internal data movement.


5.1 Multi-Client 4K 100% Random Write Workloads

Lightbits exhibited near-linear IOPS scaling as job counts increased, while maintaining sub-millisecond average latency even at high throughput levels. Replication overhead was largely transparent to the client.

Ceph showed modest write IOPS scaling accompanied by a steep increase in latency as concurrency increased. Back-pressure effects became visible to client applications under sustained load. (Figure 4)

These results are consistent with Ceph’s object-centric replication model compared to Lightbits’ block-native architecture.

Figure 4: Multi-client 4K 100% random write scaling comparison highlighting throughput and latency behavior.



6. Resiliency Testing: Node Failure and Recovery

To evaluate the resiliency characteristics of both storage options, a mixed 50/50 read/write workload was implemented while both platforms were subjected to sustained I/O and a primary storage node was intentionally taken offline.

Figure 5: Sustained I/O during primary node failure and recovery, showing IOPS and latency over time (Lightbits Cluster).





Figure 6: Sustained I/O during primary node failure and recovery, showing IOPS and latency over time (Ceph Cluster).



6.1. Lightbits Failure Behavior

(Figure 5)

- I/O interruption measured in seconds
- Automatic NVMe/TCP ANA path failover
- Stable performance during degraded operation
- Rapid rebuild following node recovery
- Minimal latency impact during rebuild activities

6.2. Ceph Failure Behavior

(Figure 6)

- I/O interruption measured in tens of seconds
- Noticeable throughput degradation during degraded mode
- Significant latency spikes during recovery
- Longer rebuild duration
- Direct contention between client I/O and recovery traffic

6.3. Observations

Lightbits prioritizes path resiliency and fast re-routing before initiating rebuild operations, allowing I/O to continue with minimal disruption. Ceph, by design, begins restoring data durability shortly after a failure by updating the OSDMap and re-evaluating data placement through the CRUSH map. Affected placement groups are marked as degraded, and replica recovery and data redistribution are initiated across the cluster. This recovery activity occurs concurrently with client I/O and can impact latency during failure and rebuild events.

While Ceph's monitoring processes in a production environment are typically distributed across multiple failure domains, this comparison maintains equal node counts to preserve test symmetry. The observed recovery behavior reflects Ceph's replica restoration model rather than monitor placement alone.

7. Rebuild Behavior and Operational Impact

Rebuilding operations directly affects application performance and operational stability.

Lightbits rebuilds are completed within predictable timeframes and affect a limited portion of system resources, resulting in minimal impact on active workloads.

Ceph rebuilds consume cluster-wide resources, extending recovery windows and introducing measurable performance variability for active applications.

These effects become increasingly significant as cluster utilization and scale increase.

8. Augmenting Ceph with Lightbits

The results of this analysis support a complementary deployment model rather than a replacement strategy.

Ceph remains well-suited for:

- Capacity-oriented block workloads
- Object storage
- File storage

Lightbits is particularly well-suited for:

- Latency-sensitive block workloads
- Virtual machine boot and root disks
- Databases and transactional systems
- Environments requiring rapid and predictable failure recovery, including OpenStack, Kubernetes, OpenShift, and KVM-based platforms

By offloading performance-critical block workloads to Lightbits, organizations can preserve Ceph's flexibility while mitigating its performance trade-offs for demanding applications.

9. Extended Workload Validation

While this paper emphasizes 4K workloads for clarity, additional testing produced consistent trends and reinforced the conclusions presented in this analysis. Additional tests were performed using:

- 8K block size
- 100% random read/write
- Mixed read/write ratios
- Varying queue depths

These results may be included as supplemental material without altering the overall findings.

Figure 7: Multi-client 8K 100% random read




Figure 8: Multi-client 8K 100% random write

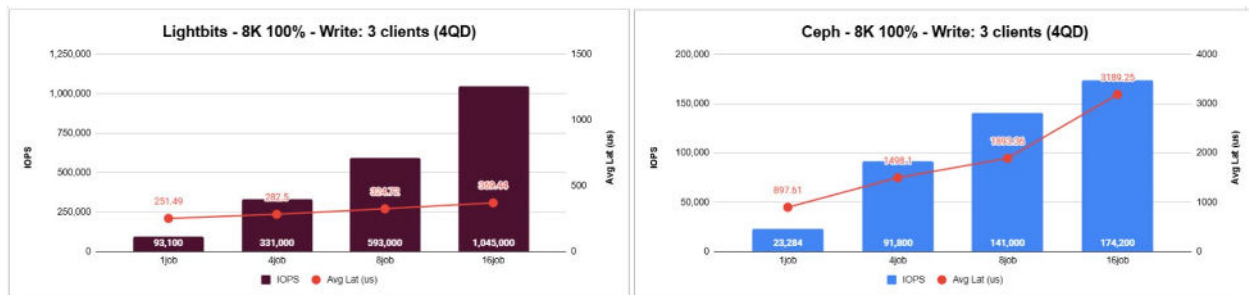


Figure 9: Multi-client 8K random mixed 70/30 (r/w)

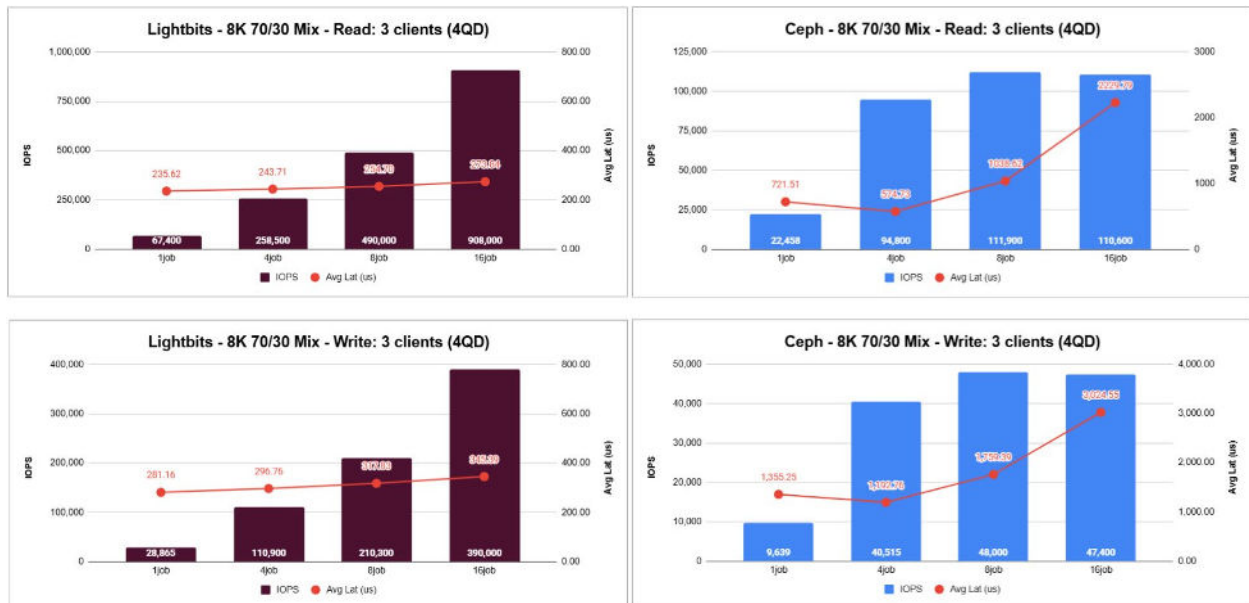
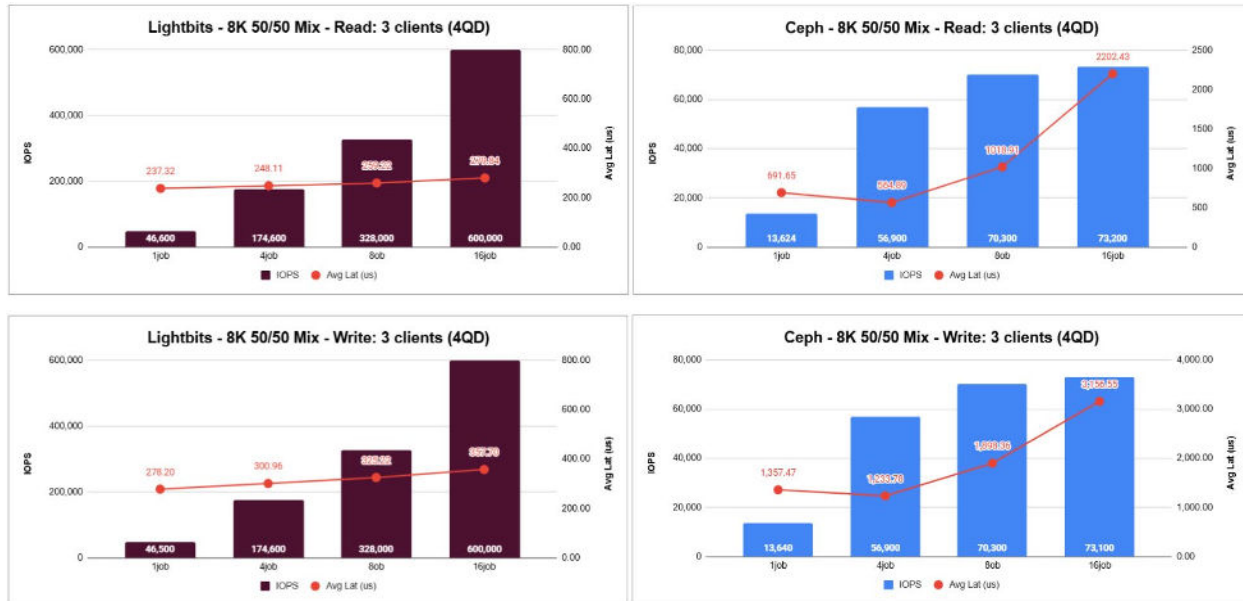


Figure 10: Multi-client 8K random mixed 50/50 (r/w)



10. Conclusion

While isolated benchmarks can highlight peak performance, they do not reflect how storage platforms behave under sustained load or during failure and recovery events.

These tests illustrate that Lightbits delivers consistently low latency, linear scaling, and fast, predictable recovery for block storage workloads, even under sustained load and failure scenarios. Ceph remains a capable and flexible unified storage platform; however, when deployed alongside Lightbits, organizations can better support performance-critical block workloads while leveraging Ceph to support capacity-oriented and multi-protocol use cases.



About Lightbits Labs™

Lightbits Labs (Lightbits) is leading the digital data center transformation by making high-performance elastic block storage available to any cloud. Creators of the NVMe® over TCP (NVMe/TCP) protocol, Lightbits software-defined storage is easy to deploy at scale and delivers performance equivalent to local flash, accelerating cloud-native applications in bare-metal, virtual, or containerized environments. Backed by leading enterprise investors including Cisco Investments, Dell Technologies Capital, Intel Capital, JP Morgan Chase, Lenovo, and Micron, Lightbits is on a mission to make high-performance elastic block storage simple, scalable and cost-efficient for any cloud.

www.lightbitslabs.com

info@lightbitslabs.com

US Offices
1830 The Alameda,
San Jose, CA 95126, USA

Israel Office
17 Atir Yeda Street,
Kfar Saba 4464313, Israel

The information in this document and any document referenced herein is provided for informational purposes only, is provided as is and with all faults and cannot be understood as substituting for customized service and information that might be developed by Lightbits Labs Ltd for a particular user based upon that user's particular environment. Reliance upon this document and any document referenced herein is at the user's own risk.

The software is provided "As is", without warranty of any kind, express or implied, including but not limited to the warranties of merchantability, fitness for a particular purpose and non-infringement. In no event shall the contributors or copyright holders be liable for any claim, damages or other liability, whether in an action of contract, tort or otherwise, arising from, out of or in connection with the software or the use or other dealings with the software.

Unauthorized copying or distributing of included software files, via any medium is strictly prohibited.

LBWP21/2026/2

COPYRIGHT© 2026 LIGHTBITS LABS LTD. - ALL RIGHTS RESERVED