



OpenStack with Block Storage Solutions for High-Performance Workloads

Technical White Paper

JUNE 2025

This white paper presents a performance evaluation of OpenStack integrated with Lightbits software-defined block storage, natively designed with NVMe® over TCP, focusing on I/O patterns commonly found in real-world-like applications, including random read/write and 70/30 mixed workloads. Using 32 virtual machines (VMs) across a four-node OpenStack cluster, the results demonstrate strong IOPS scalability and consistent throughput as concurrency increases. The findings highlight Lightbits as a high-performance, software-defined storage solution well-suited for OpenStack environments requiring reliable, scalable, and efficient block storage.

Table of Contents

1. Executive Summary	2
2. Introduction	2
3. OpenStack Fundamentals	3
3.1. Core Components	3
3.1.1. Compute (Nova)	3
3.1.2. Block Storage (Cinder)	3
3.1.3. Image Service (Glance)	3
3.1.4. Networking (Neutron)	4
3.1.5. Control Plane vs. Data Plane	4
3.1.6. Supporting Control Plane Services	4
3.2. OpenStack Cluster Configuration Best Practices	5
3.2.1. Controller Nodes	5
3.2.2. Compute Nodes	6
3.2.3. Storage Backend (Cinder)	6
4. Optimized Storage Solutions	6
4.1. Lightbits Storage Solution	6
4.2. Lightbits Storage for OpenStack	7
5. Testbed Overview	8
5.1 High-Level Architecture	8
5.2 Testbed Objectives	9
5.2.1 Testbed Objectives	9
5.2.2 OpenStack Services and Configuration	9
5.2.3 Lightbits Configuration	10
6. Performance Evaluation	10
7. Conclusion	16
8. Appendix	17
About Lightbits Labs	18

1. Executive Summary

This white paper presents a performance-focused evaluation of an OpenStack environment integrated with Lightbits software-defined block storage. The goal is to demonstrate how Lightbits can serve as a high-performance storage backend for OpenStack Cinder, enabling scalable and efficient infrastructure for a wide range of virtualized workloads.

Key highlights of this deployment include:

- **Optimized Block Storage:** Lightbits provides a scalable, high-throughput Cinder backend that supports demanding application workloads in OpenStack, including transactional systems, databases, and general-purpose Virtual Machines (VMs).
- **Seamless Integration:** Lightbits integrates natively with OpenStack via the Cinder driver, allowing for straightforward adoption without requiring architectural changes to existing deployments.
- **Strong Performance and Scalability:** Benchmark results using 32 VMs show strong IOPS scaling and consistent throughput under common I/O patterns such as random read and 70/30 mixed workloads, validating Lightbits' superior performance capabilities in OpenStack environments.
- **Simple, Standards-Based Architecture:** The solution is built using standard networking (NVMe/TCP) and commodity AMD-based hardware, reducing complexity and cost while maximizing performance.
- **Operational Efficiency:** With its software-defined design and native OpenStack support, Lightbits simplifies storage management while delivering predictable performance across a wide range of workloads.

By adopting this approach, organizations can modernize their OpenStack infrastructure with high-performance block storage that is easy to deploy, scale, and operate—while meeting the needs of both current and future applications.

2. Introduction

OpenStack remains a popular choice for building private cloud infrastructure due to its flexibility, open-source ecosystem, and support for a wide range of workloads. As organizations modernize their infrastructure, selecting the right storage backend becomes critical to ensuring consistent performance, scalability, and operational simplicity.

This paper evaluates the performance of OpenStack when paired with Lightbits NVMe/TCP storage. Lightbits integrates with OpenStack Cinder as a block storage backend, providing high throughput and low-latency access using standard Ethernet networking. The goal is to assess how well this combination performs under typical workload patterns, including random read and mixed read/write I/O.

The paper outlines the test environment, describes the workload methodology using FIO, and presents results based on 32 VMs running across a four-node OpenStack cluster. The findings help illustrate how Lightbits can be a strong fit for OpenStack deployments that require high-performance block storage for scaling VMs.

3. OpenStack Fundamentals

OpenStack is a modular cloud platform that allows organizations to build and operate private or public cloud infrastructure using commodity hardware. Its core components provide services for compute, networking, and storage—all of which are critical in delivering infrastructure-as-a-service (IaaS). Understanding how these components interact is essential to the performance results highlighted in this paper.

3.1. Core Components

While DAS offers high performance, it lacks dynamic scaling and simplicity. The tight coupling to individual servers leads to inefficient resource utilization, limited scalability, and increased management overhead. Applications locked to servers and storage result in low utilization of CPU and storage, while application-based replication and recovery over the network lead to service degradation.

3.1.1. Compute (Nova)

- Nova is the compute service responsible for provisioning and managing VMs across a pool of physical servers, referred to as compute nodes. Nova interacts with a scheduler that determines where new instances should run based on available CPU, memory, storage, and other placement criteria. In this test, one node functioned as both the controller and a compute node, while three additional nodes were dedicated compute hosts, allowing Nova to balance VM creation across all four.
- Each VM runs on a KVM hypervisor and is managed through libvirt. When a VM is launched, Nova coordinates with the image service (Glance), the network service (Neutron), and the block storage service (Cinder) to provide the instance with the necessary boot image, networking, and persistent volume.

3.1.2. Block Storage (Cinder)

- Cinder provides persistent block storage volumes to VMs. In this evaluation, Cinder is configured with the Lightbits Cinder driver, enabling volume provisioning from a Lightbits storage cluster using the NVMe/TCP protocol. When a volume is created or attached to a VM, Cinder communicates with the Lightbits storage backend to allocate and map the volume appropriately.
- This modularity enables OpenStack users to decouple compute and storage, allowing them to scale these resources independently. More importantly, it enables the use of high-performance storage platforms, such as Lightbits, while still maintaining standard OpenStack workflows (e.g., volume snapshots, resizing, detaching, and reattaching).

3.1.3. Image Service (Glance)

- Glance is responsible for storing and serving disk images used to boot VMs. These images can be backed by various storage systems and can be reused across many VMs. During VM boot, Nova retrieves the image from Glance and writes it to the target root disk.

3.1.4. Networking (Neutron)

- Neutron provides software-defined networking (SDN) for OpenStack environments. It handles IP address management, DHCP, routing, and virtual network isolation. Neutron ensures that VMs across different compute nodes are connected consistently and securely, regardless of where they are scheduled.
- For this test, basic L2 and L3 connectivity was configured, allowing all 32 VMs to participate in FIO workloads without networking bottlenecks. While Neutron is not directly involved in block storage performance, its role is crucial in enabling the execution of scalable and distributed workloads across the cluster.

3.1.5. Control Plane vs. Data Plane

- The OpenStack control plane consists of services like Nova scheduler, Cinder scheduler, Glance API, Neutron server, Keystone (identity), and the Horizon dashboard. These services orchestrate the provisioning and management of resources but are not involved in the actual data path once VMs are running.
- The data plane, on the other hand, encompasses the real-time I/O path between VMs and their storage. In this setup, once a Lightbits volume is attached to a VM, all subsequent reads and writes happen directly between the VM and the Lightbits storage over NVMe/TCP. This separation allows control plane operations to remain lightweight while enabling high-speed data transfers directly between compute and storage nodes.

3.1.6. Supporting Control Plane Services

In addition to Nova, Cinder, Glance, and Neutron, several essential services and databases form the backbone of OpenStack's control plane. These services coordinate authentication, messaging, caching, and persistent state—ensuring that OpenStack functions reliably at scale.

- **Keystone (Identity Service)**
 - Keystone is the central authentication and authorization service for OpenStack. It manages user credentials, project (tenant) scopes, roles, and service tokens. All API calls to OpenStack services are authenticated through the Keystone service. In this setup, Keystone ensured secure access to Nova, Cinder, and other APIs during orchestration of the VM lifecycle and volume attachments.
- **MySQL (MariaDB)**
 - MySQL (or MariaDB, depending on the distribution) serves as the relational database backend for nearly every OpenStack service. It stores metadata, including instance states, volume information, network assignments, and service configurations. A single MySQL database cluster or server can support multiple OpenStack services, each using a separate schema. In this environment, MySQL stored persistent state for services like Nova, Cinder, Glance, and Keystone.
- **RabbitMQ (Message Queue)**
 - RabbitMQ is the message bus used for communication between OpenStack services in a decoupled and asynchronous fashion. For example, when Nova schedules a VM or Cinder provisions a volume,

the actual actions are executed through messages sent over RabbitMQ to various service agents. RabbitMQ ensures high throughput and resilience in service-to-service coordination.

- **Memcached (Caching Layer)**
 - Memcached is used to improve the performance of Keystone token validation and API request processing across the control plane. It reduces the load on the database and speeds up repeated authentication checks and metadata retrieval. In this environment, Memcached helped accelerate interactions between Keystone and other services.

3.2. OpenStack Cluster Configuration Best Practices

Designing a production-ready OpenStack deployment requires careful consideration of control plane services, compute node resources, and storage backends. The following best practices focus specifically on controller nodes, compute nodes, and the Cinder block storage service. While there are many other components and configurations in OpenStack, this section highlights only these core areas—especially in the context of integrating high-performance storage such as Lightbits with NVMe/TCP.

3.2.1. Controller Nodes

Controller nodes host the core OpenStack services that manage authentication, orchestration, image management, scheduling, and inter-service communication.

- **High Availability:** For production environments, it is recommended to have at least three controller nodes to ensure high availability. This configuration enables quorum-based services like MySQL Galera and RabbitMQ to continue functioning during node failures.
- **Load Balancing:** Utilize a load balancer in front of the API services on the controller nodes to evenly distribute user and service requests. This approach ensures consistent API performance and availability.

- **Service Placement:**

Controller nodes typically run:

- Keystone (identity)
- Nova scheduler and API
- Cinder scheduler and API
- Glance API
- Neutron server and DHCP/L3 agents
- RabbitMQ and Memcached
- MySQL/MariaDB or Galera cluster

Ensure these nodes are provisioned with adequate CPU, memory, and I/O capacity, especially if services are colocated.

3.2.2. Compute Nodes

Compute nodes run the KVM hypervisor and host the VM workloads managed by Nova.

- **Scaling by Workload:** Start with a baseline of two to three compute nodes and scale horizontally according to workload demands. Horizontal scaling provides better fault isolation and resource elasticity.
- **Storage Considerations:** Avoid placing VM data on local disks if you plan to use external block storage (e.g., Lightbits via Cinder). Lightbits offers disaggregated, high-performance storage, enabling improved resilience and scalability.

3.2.3. Storage Backend (Cinder)

- **External Block Storage Integration:** Using Lightbits with NVMe/TCP for Cinder volumes enables high IOPS and low latency without requiring specialized fabrics. It decouples compute from storage, improving data durability and enabling dynamic provisioning.
- **Volume Types and Quality of Service (QoS):** Define Cinder volume types to represent different storage backends or policies (e.g., replication factor, compression). Lightbits supports various volume features, and defining them in Cinder allows tenants to choose the right backend for their workloads.
- **Snapshot and Backup Strategy:** Ensure your block storage strategy includes snapshot support for fast recovery and cloning, and configure backups for long-term retention where needed.

4. Optimized Storage Solutions

4.1. Lightbits Storage Solution

Lightbits represents a paradigm shift in how storage resources are managed and utilized in modern data centers. Lightbits' unique, disaggregated, software-defined [NVMe over TCP block storage](#) architecture, combined with Intelligent Flash Management and enterprise-rich data services, makes it the only solution for high-performance block workloads in OpenStack that's easy to provision and simple to manage, while also being highly scalable, performant, and cost-efficient.

Lightbits Labs invented NVMe/TCP storage protocol technology, natively designed into the software, to deliver low-latency, [high-performance block storage](#) over a standard Ethernet network. This not only reduces the cost and complexity compared to other technologies, such as Fibre Channel or InfiniBand, but also democratizes access to high-speed [NVMe storage](#).

The software-defined nature of Lightbits' solution means it can be deployed on commodity servers and scaled across a distributed architecture without the need for specialized networking equipment or proprietary hardware. This flexibility, combined with advanced features such as Intelligent Flash Management, ElasticRAID, in-line compression, software encryption, thin provisioning, snapshots, and clones, produces a solution that enhances data durability, optimizes storage efficiency, and reduces overall operating costs.

4.2. Lightbits Storage for OpenStack

Integrating Lightbits with OpenStack provides a compelling advantage for environments that require high performance and efficiency in managing storage operations. Even if a single application doesn't initially demand significant performance, at scale, with hundreds or thousands of applications running concurrently, storage performance requirements can become extremely high. Lightbits can sustain such loads without needing to scale the storage infrastructure simply to meet performance demands.

The following are the key benefits of integrating Lightbits with OpenStack:

1. **High Performance:** Lightbits leverages NVMe over TCP to deliver high-speed, low-latency storage, which is ideal for performance-critical applications such as databases and real-time analytics running in OpenStack. This ensures that applications can maximize their performance without the traditional bottlenecks associated with networked storage.
2. **Simplified Scalability:** With Lightbits, OpenStack environments can scale storage independently of compute resources with ease. This decoupling enables more flexible resource management, allowing storage to be added or adjusted without impacting the compute nodes. Such scalability is crucial for dynamic containerized environments, where workloads can change rapidly.
3. **Enhanced Data Durability:** Lightbits provides built-in data protection and resilience mechanisms that enhance the durability of volumes provisioned through OpenStack Cinder. This is especially valuable in cloud environments where workload availability and data integrity are critical. Features such as thin provisioning, snapshots, encryption, and synchronous replication ensure that data remains protected against hardware failures and is easily recoverable when needed.
4. **Cost Efficiency:** By optimizing data storage with compression, Lightbits helps reduce the overall storage footprint and costs. This efficiency enables the management of more data with less storage, which is economically beneficial for growing organizations seeking to optimize their IT investments.
5. **Ease of Integration:** Lightbits integrates with OpenStack through the Cinder service, making it straightforward to add Lightbits storage solutions to existing OpenStack clusters. This seamless integration simplifies the deployment and management of storage resources, allowing IT teams to focus on deploying and scaling applications rather than managing underlying storage infrastructure. Furthermore, Lightbits provides multi-tenancy, enabling OpenStack administrators to efficiently manage storage access, thereby further enhancing the operational efficiency and security of the cluster.
6. **Cloud Infrastructure Compatibility:** Lightbits is designed to support cloud infrastructure with its software-defined storage architecture, making it a strong fit for OpenStack environments. Its seamless integration with OpenStack Cinder enables agile, scalable, and high-performance block storage that aligns with the dynamic provisioning and elasticity principles central to OpenStack.
7. **Multi-Tenancy Support:** Lightbits supports multi-tenancy by enabling secure, isolated access to storage resources across different OpenStack projects or tenants. Through native integration with OpenStack Cinder, volumes can be provisioned, managed, and scoped per tenant, ensuring data isolation and operational separation in shared infrastructure environments.

By integrating Lightbits with OpenStack, organizations can build a high-performance, scalable, and cost-effective storage environment that meets the demands of modern, data-driven workloads. This integration enhances the

operational efficiency of storage management through OpenStack Cinder while improving the overall performance and resilience of applications running on OpenStack VMs.

5. Testbed Overview

5.1 High-Level Architecture

To evaluate the performance and scalability of OpenStack with Lightbits storage, a dedicated testbed environment was configured using AMD-based servers and standard networking infrastructure. This environment was designed to simulate a realistic deployment model while enabling focused performance testing across multiple VMs. Although this setup does not represent a formal production reference architecture, it provides a solid baseline for assessing Lightbits as a backend block storage option for OpenStack.




Figure 1: High-Level Architecture

5.2 Testbed Objectives

The goal of this testbed was to:

- Evaluate IOPS scalability across multiple VMs, including how performance scales with increasing FIO job concurrency within each VM.
- Analyze storage performance under synthetic workloads using FIO.
- Observe the behavior of the OpenStack control and compute services under increasing load.

5.2.1 Testbed Objectives

Table 1: Hardware Layout

Role	Server Count	CPU	Memory	Networking
OpenStack Controller & Compute Node	1	AMD EPYC 9655 (multi-core)	1.5 TB	100GbE for data + 1GbE mgmt
OpenStack Compute Only Nodes	3	AMD EPYC 9655 (multi-core)	1.5 TB each	100GbE for data + 1GbE mgmt
Lightbits Storage Nodes	3	AMD EPYC 9575F (multi-core)	768 GB each	100GbE for NVMe/TCP

- **Storage Devices:** Each Lightbits node was equipped with 8 x 3.84TB Samsung NVMe SSDs.
- **Network Layout:** A dedicated 100GbE network was used for NVMe/TCP traffic between compute and storage nodes. Management and control plane services are operated on a separate 1GbE network.

5.2.2 OpenStack Services and Configuration

- **Deployment Type:** Using OpenStack 2024.2 (Dalmation) on Ubuntu 22.04 operating system
- **OpenStack Services:** Nova, Cinder, Keystone, Neutron, Glance, RabbitMQ, Memcached, and MySQL
- **Cinder Backend:** Lightbits Cinder driver using NVMe/TCP
- **Glance Image Backend:** Local filesystem on the controller node
- **Networking:** Flat network for simplicity, sufficient for internal performance testing

5.2.3 Lightbits Configuration

- **Volumes:** Each VM was attached to a 20GB Lightbits-provisioned volume via the OpenStack Cinder service.

- **Data Placement:** Volumes were distributed across the three Lightbits nodes to ensure balanced I/O.
- **Replication:** Configured with a replication factor of 3 (RF3) to ensure data durability and resilience across the Lightbits storage cluster while still achieving high performance under load.
- **Access Protocol:** NVMe/TCP over dedicated 100GbE interfaces.

6. Performance Evaluation

With Lightbits fully integrated as a backend for OpenStack Cinder, the next step was to evaluate the storage performance it can deliver in a virtualized OpenStack environment. To do this, a series of FIO benchmark tests were executed across 32 VMs deployed on an OpenStack cluster backed by Lightbits with NVMe/TCP storage.

Each VM was configured with 12 vCPUs and attached to a 20GB Lightbits volume via a predefined Cinder volume type. Before testing, volumes were preconditioned to ensure consistent and repeatable results.

The evaluation focused on two I/O patterns that closely resemble real-world-like application behavior:

- **Random Read:** Common in workloads such as web applications, content delivery platforms, and search engines, where many users or services access small, dispersed chunks of data.
- **70/30 Mixed Read/Write:** Representative of transactional systems like OLTP databases, NoSQL stores (e.g., Oracle DB, MongoDB, Cassandra), and general-purpose enterprise applications where read operations dominate but frequent writes still occur.

Each workload was tested with **4K and 8K block sizes**, and FIO's numjobs parameter was varied across **1, 4, 8, and 12** to assess how performance scales with increasing concurrency per VM.

Rather than emphasizing raw performance numbers alone, the analysis highlights scalability across the VM fleet and the impact of increasing I/O parallelism within each instance. These insights demonstrate how the combination of OpenStack and Lightbits enables consistent, high-throughput block storage for performance-sensitive applications.

The following charts illustrate:

- IOPS and throughput trends for random read/write and 70/30 workloads at varying block sizes
- Latency behavior as concurrency increases
- Performance scaling as numjobs increases from 1 to 12 across 32 VMs
- Performance scaling as VM increases from 1 to 32.




Figure 2: 4K Random Read Random Write Charts




Figure 3: 8K Random Read Random Write Charts




Figure 4: 4K 70/30 Mixed Charts




Figure 5: 8K 70/30 Mixed Charts




Figure 6: 4K Random Read VMs Scaling

Across all test scenarios, Lightbits demonstrated strong and consistent performance. In 4K random read benchmarks, performance scaled linearly as the number of FIO jobs per VM increased, ultimately reaching over 5.3 million aggregate IOPS across 32 VMs—highlighting the ability of Lightbits to support highly parallel read-heavy workloads with minimal latency impact.

For 70/30 mixed read/write workloads—common in transactional and general-purpose application environments—performance scaled effectively up to 8 jobs per VM before leveling off. Even as the I/O mix introduced write operations, Lightbits maintained solid throughput and responsiveness, showing its capability to support real-world-like workloads with a balanced access pattern.

Overall, the data reinforces Lightbits' suitability for performance-sensitive OpenStack environments. Whether supporting transactional workloads, virtualized databases, or distributed applications, Lightbits provides predictable scaling, strong throughput, and low-latency access across a wide range of I/O profiles.

In addition to 4K and 8K block sizes, 32K random read performance was also measured to illustrate how throughput behaves with larger I/O sizes. While IOPS naturally decrease as block size increases, Lightbits still achieved over 1 million IOPS at 32K under optimal test conditions, demonstrating its ability to sustain high throughput in read-heavy scenarios that are less sensitive to latency.




Figure 7: IOPS achieved with 4K, 8K, 32K block sizes under optimal test conditions

7. Conclusion

This paper explored the performance characteristics of OpenStack when integrated with Lightbits storage. The goal was to understand how well Lightbits performs as a Cinder backend in a virtualized environment running a mix of synthetic I/O workloads. The results show that Lightbits delivers strong read performance and scales well across multiple VMs, particularly in read-intensive scenarios.

Even under write and mixed workloads, performance remained stable, highlighting Lightbits' ability to handle a variety of application patterns commonly found in OpenStack deployments. The test setup also demonstrated that integrating Lightbits into OpenStack is straightforward, using native Cinder support and standard networking.

For organization building or expanding OpenStack infrastructure, Lightbits offers a solid storage option that provides good performance and doesn't require re-architecting the environment. This paper also provided practical steps for installing the Lightbits Cluster and configuring the Lightbits Cinder driver, helping make the case for its use in real-world OpenStack environments.

To learn more about Lightbits Labs, visit <https://www.lightbitslabs.com>.

8. Appendix

The following appendix provides the FIO job file template used in the performance testing described in this paper. This file outlines the parameters used to evaluate various I/O patterns and concurrency levels across the OpenStack environment with Lightbits storage.

A. FIO Job File Template

The following FIO configuration was used by each VM for benchmarking. Adjust the **bs** (block size), **rw** (read/write mode), **numjobs** (parallel jobs), and **iodepth** (queue depth) parameters to reflect the desired workload characteristics.

```
[global]
ioengine=libaio
direct=1
prio=0
rw=randread
bs=4k
numjobs=1
iodepth=1
time_based
ramp_time=5
norandommap
randrepeat=0
runtime=120
group_reporting=1
[job]
filename=/dev/xvdb
```



About Lightbits Labs

Lightbits Labs® (Lightbits) invented the NVMe over TCP protocol and offers best-in-class software-defined block storage that enables data center infrastructure modernization for organizations building a private or public cloud. Built from the ground up for low consistent latency, scalability, resiliency, and cost-efficiency, Lightbits software delivers the best price/performance for real-time analytics, transactional, and AI/ML workloads. Lightbits Labs is backed by enterprise technology leaders [Cisco Investments, Dell Technologies Capital, Intel Capital, Lenovo, and Micron] and is on a mission to deliver the fastest and most cost-efficient data storage for performance-sensitive workloads at scale.

 www.lightbitslabs.com

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

 info@lightbitslabs.com

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.

COPYRIGHT© 2025 LIGHTBITS LABS LTD. - ALL RIGHTS RESERVED

LBWP13/2025/05