A Guide to Infrastructure Modernization for Cloud Providers and Service Platforms

Unlock efficiency, scalability, flexibility, and performance for your service cloud infrastructure.

MAY 2025

Rev 2.0

This whitepaper introduces Lightbits disaggregated, software-defined storage (SDS) natively designed with NVMe® over TCP (NVMe/TCP), as a solution designed to modernize and future-proof data infrastructure for cloud and platform service providers. It addresses the limitations of traditional storage solutions like DAS, SAN, and HCI – which do not deliver scalability, performance, and cost-efficiency for demanding block workloads at scale and in dynamic cloud environments.

The paper illustrates how Lightbits empowers service providers to grow their business by building a highly performant, scalable, and resilient data platform capable of achieving hyperscale-like efficiency. And the flexibility to effectively support mixed workloads at any scale in virtualized and containerized environments, lowering CAPEX and OPEX costs and reducing the data center footprint.

Table of Contents

1. The Complexities of Building a Cloud Service or Platform Service

There are subtle differences between cloud service providers (CSPs) that build infrastructure as a service (IaaS) and platform service providers. CSPs build a platform to provide on-demand, scalable computing resources. For example, Microsoft Azure is a hyperscale CSP. On the other hand, platform service providers offer a technology-based platform that others can use to conduct business or access specific functionalities. For example, an online retailer builds an e-commerce platform to conduct customer transactions. While there is a major difference in infrastructure objectives, both face similar complexities when architecting their infrastructure. This paper will address the data infrastructure for both types of businesses..

When building a publicly-facing cloud service or platform service, a few things are undeniable: fluctuating and unpredictable customer demands require rapid or dynamic scalability; end user satisfaction requires consistent high performance; and cost-efficient systems are needed to maintain desired revenue margins. Legacy block storage solutions – DAS, SAN, and Ceph Storage – were, and still are, balancing acts of scalability, performance, operational overhead, and cost efficiency - often requiring compromises and tradeoffs.

Legacy systems have constrained limitations and are unsuitable to support today's large-scale cloud-based workloads. In the case of DAS, it's an architecture inflexibility that can lead to resource inefficiencies as applications become wedded to servers and storage. In the case of SANs, it's a matter of cost, as proprietary hardware and specialized networking come at a premium. In the case of Ceph Storage, there is inadequate block storage performance at scale and inherent operational complexities. For CSPs and platform services, balancing business growth with efficiency and customer satisfaction demands a flexible, yet high-performing, reliable, and cost-effective storage system.

It is well known that hyperscale CSPs such as Facebook, Google, Microsoft Azure, and Amazon Web Services (AWS) cracked the code of infrastructure efficiency because it was critical to their business. What's less well known is how they achieved that level of hyper-efficiency without compromising on any of the following:

- Performance
- Scalability
- Reliability

Next-gen block storage from Lightbits Labs has empowered many successful service providers and platform service builders to emulate the infrastructure model of the hyperscalers and achieve the best price/performance value, while supporting mixed workloads at scale without an army of engineers. This whitepaper will illustrate how.

2. Hyperscale Efficiency for All

Specialized cloud and platform services businesses seek to emulate the flexibility, efficiency, scalability, and reliability benefits of public cloud hyperscalers. The challenge for most service provider startups is that building such an infrastructure is a complex and daunting task due to the engineering talent required.

These challenges can be overcome with software-defined storage (SDS) in general, and Lightbits software-defined block storage in particular. To achieve the benefits that hyperscalers can provide (on-premises), technologies like disaggregated, software-defined storage natively designed with NVMe over TCP storage protocol combined with open-source orchestration systems like Kubernetes, OpenShift, OpenStack, and KVM have been key infrastructure enablers for service clouds and platforms.

Whether building data infrastructure from the ground up using KVM, Kubernetes, or Red Hat OpenStack or OpenShift, Lightbits SDS powers mixed workloads in virtualized and containerized environments with optimal price/performance, enabling service provider and platform business growth, while simplifying operations.

3. The Limitations of Legacy Storage Models

In various ways, DAS, FC SAN, and Ceph Storage do not deliver on the fundamental benefits of a hyperscale-like infrastructure and fall short of fully supporting service clouds. They often involve more manual intervention, have tighter coupling between compute and storage (especially DAS), and can be more complex and costly to scale and manage in a dynamic cloud service environment. Service clouds demand a more abstracted, automated, and scalable block storage infrastructure that can adapt to changing application needs with minimal manual effort.

3.1. DAS Limitations

While DAS offers high performance, it fails to offer 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 lead to low utilization of CPU and storage, and application-based replication and recovery over the network cause service degradation.

3.2. HCI Limitations

Hyperconverged Infrastructure (HCI) offers benefits like operational efficiency and scalability, but fails in the price/performance scenario. Initial investment can be high, and vendor lock-in should be a significant concern. And for the budget-conscious, HCI fails in cost-efficiency for high-performance workloads. Because storage and compute must be scaled together, this results in overprovisioning and wasted resources. Additionally, the hypervisor design yields inconsistent latency, and it does not allow for mixing virtualized and non-virtualized workloads.

3.3. SAN Limitations

Storage Area Networks (SANs), especially Fibre Channel-based (FC) – while delivering efficiency and performance, fail the price/performance scenario by introducing complexity with a dedicated network fabric, costly infrastructure, and challenges in achieving elastic scalability. Proprietary expensive appliances are involved, more complex networking is required to deliver high performance, a specialized storage network with expensive equipment and trained staff is required, and it introduces vendor lock-in with less flexibility.

4. Block Storage for Open-Source Clouds

Block storage is a critical component of service clouds and platforms. It is offered natively by the leading hyperscalers and, undoubtedly, will continue to serve as the backbone of primary storage tiers. It is valued for its ability to provide high-speed access to structured data. However, as workloads become more diverse, demanding, and cloud-native, the role of block storage has expanded and is now considered a foundational element of service cloud and platform architecture.

4.1. Block Storage for Red Hat OpenStack, OpenShift, and Kubernetes

Utilizing block storage in OpenStack implemented through a Cinder service, or in OpenShift or Kubernetes through Container Storage Interface (CSI) drivers, offers several key benefits for service clouds and platforms. It enables the creation and management of persistent storage volumes (PVs) attached to virtual machines (VMs); provides reliable and scalable storage independent of the instance lifecycle; and offers the flexibility to choose a storage solution optimized for price/performance.

Service providers can support a broader range of workloads, including high-performance databases and other applications requiring durable, persistent storage. Tenants can easily provision and manage their own storage needs through APIs, reducing operational overhead for providers while offering greater flexibility to their customers.

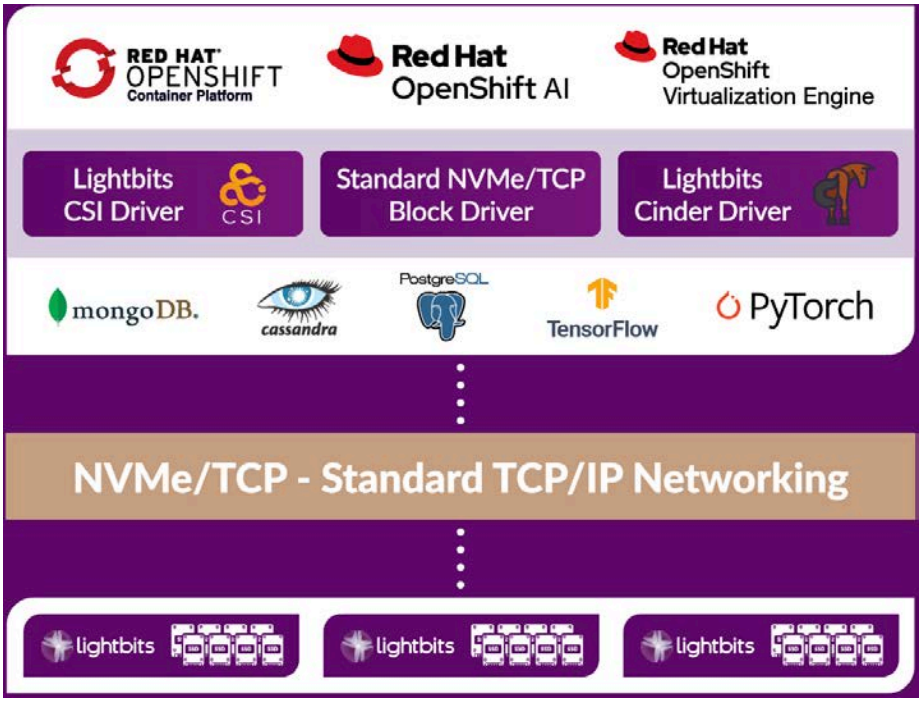

Figure 1. Lightbits seamlessly integrates with open-source environments, like Red Hat OpenShift.

5. Build Your Cloud and Service Platform with Next-Gen SDS for Performance and Efficiency

---

What customers say before they discover Lightbits software-defined storage:

“I don't get enough performance from Ceph Storage."

"We can’t keep up with our scaling requirements with DAS or our SANs.”

“My storage systems are too costly, and I don’t have control over which hardware I use.”

Lightbits’ disaggregated SDS for platform and service cloud builders brings hyperscale efficiency and flexibility to your infrastructure. Designed for high performance and mixed workloads, Lightbits delivers unmatched price/performance, availability, and scalability. With Lightbits in your data center, you can simplify infrastructure management and operations while lowering costs.

5. 1 Disaggregated Architecture

Disaggregation is not a new concept. Hard drives have been disaggregated with SANs, and CPUs with virtualization. It’s simply about sharing server resources across applications to increase utilization and lower infrastructure costs. However, with Lightbits SDS, only recently has high-speed, high-throughput, low-latency, easy-to-manage disaggregated storage been made available for cloud services.

Lightbits allows you to disaggregate storage from compute, making it easier to tailor the number of commodity servers in use, which improves resource utilization. Disaggregating ensures the storage can be shared across multiple applications and scaled independently.

With Lightbits, you can build your own disaggregated storage system and operate like the hyperscalers, buying fewer types of server configurations, while planning compute/performance capacity around commodities like DRAM and CPUs, rather than being constrained by local flash capacity for applications. Servers are designed for specific hardware commodities and not necessarily for applications running on them.

5.2 Software-Defined Storage

The primary difference between SDS and traditional storage lies in how the storage infrastructure is managed and implemented. SDS offers a more flexible, scalable, and cost-efficient approach than traditional storage. It's inherently flexible; its hardware independence eliminates reliance on proprietary infrastructure, allowing you to choose any hardware vendor that meets your needs and avoid vendor lock-in. It offers an attractive architecture model because it maximizes storage capabilities while minimizing administrative overhead by automating and orchestrating storage management. With this model, storage can be scaled up or scaled out with minimal disruption to operational procedures and reduced capital investment.

SDS is designed to abstract storage resources from physical hardware – creating a flexible, scalable, and centrally managed storage environment. Each core architecture component is present in the implementation across various software-defined storage vendors.

The following is a list of key components of software-defined storage architecture (see Figure 2 below).

- **Abstraction Layer:** The core component that decouples the storage services from the physical hardware and makes virtual storage pools appear as a single storage pool to applications and users.
- **Control Plane:** The centralized management layer for storage operations, ensuring efficient resource utilization by orchestrating and automating storage tasks: provisioning, monitoring, policy enforcement, and metadata.
- **Data Plane:** The layer that manages actual data storage and retrieval by executing read/write operations and essential data services (replication, deduplication, and compression). It ensures data integrity and high performance during data access.
- **Management Interface:** The tool used to interact with the SDS system for visibility and control over the storage infrastructure, which in most cases can provide a dashboard for configuring, monitoring, managing storage resources, and simplifying administrative tasks.

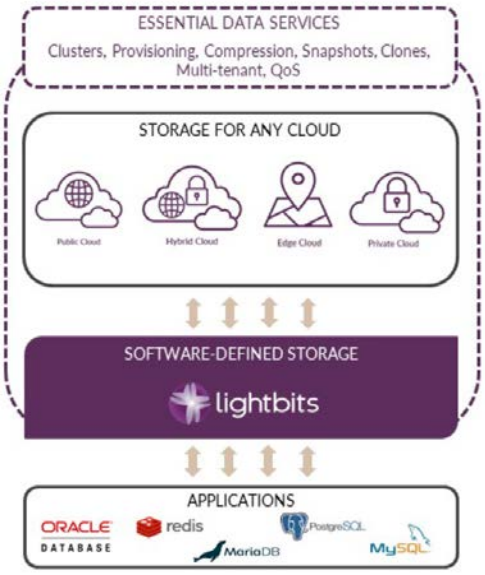

Figure 2. Key components of software-defined storage architecture

5.3 Clustering Support

To prevent data loss, software-defined storage solutions typically replicate data across different storage servers, guaranteeing service and data availability whenever storage servers experience transient or permanent failures. The Lightbits software clusters the resources for redundancy, compression, security, and provisioning.

Lightbits clusters unify storage resources like a SAN and provide high availability by design. A cluster of Lightbits servers replicates data internally. It keeps it fully consistent and available – guaranteeing service reliability and data availability even if one or several storage servers fail sporadically (for a limited time), or permanently (forever).

From the perspective of clients accessing the data, Lightbits' clustering architecture is transparent, and server failover in the event of failure is seamless. Data remains consistent and available. With Lightbits clustering support, client CPU cycles are used for applications and not for storage processing (see Figure 3 below).

Lightbits' clustering maximizes resource utilization and flexibility, and provides high availability and durability with a configurable number of replicas per volume. A Lightbits cluster consists of servers residing in different failure domains and located anywhere within the data center. The architecture maximizes resource utilization and flexibility, and provides high availability and durability with a configurable number of replicas per volume.

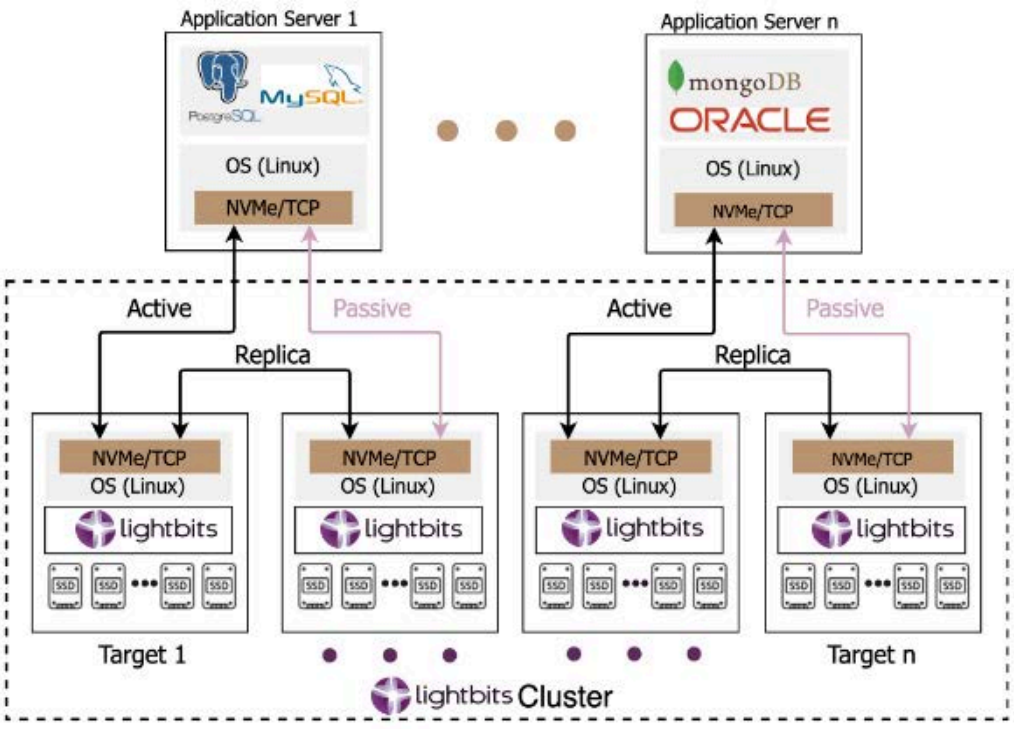

Figure 3: Lightbits clustering multipath replication design

5.4 NVMe over TCP Storage Protocol

NVMe has transformed the storage industry since it emerged as the state-of-the-art protocol for high-performance SSDs. Initially designed for high-performance direct-attached PCIe SSDs, NVMe was later expanded with NVMe over Fabrics (NVMe-oF™) to support a rack-scale remote pool of SSDs. The industry has widely accepted that this new NVMe-oF model will replace the iSCSI protocol as the communication standard between compute servers and storage servers, and become the default protocol for disaggregated storage.

The initial deployment options with NVMe-oF were limited to FC and Remote Direct Memory Access (RDMA) fabrics. NVMe/TCP offers a more modern, powerful technology that provides the speed and performance of NVMe-oF, but without the prohibitive deployment costs and complexity of other protocols. It extends NVMe across the entire data center using a simple and efficient TCP/IP fabric. It is a better technology for data center modernization initiatives, given the business and technical benefits it offers.

Lightbits invented the NVMe/TCP storage protocol, which is natively designed into the software. This enables high-performance, low-latency block storage over standard Ethernet networks, eliminating the need for expensive and specialized FC infrastructure required by legacy SANs. Clients do not require any special drivers or agents beyond the standard NVMe/TCP drivers that are available in all major Linux distributions. This simplifies deployment, reduces costs, and leverages existing network investments. Built on top of the TCP/IP software stack, NVMe/TCP enables efficient and streamlined block storage, optimized for today's multi-core application servers (see Figure 4 below).

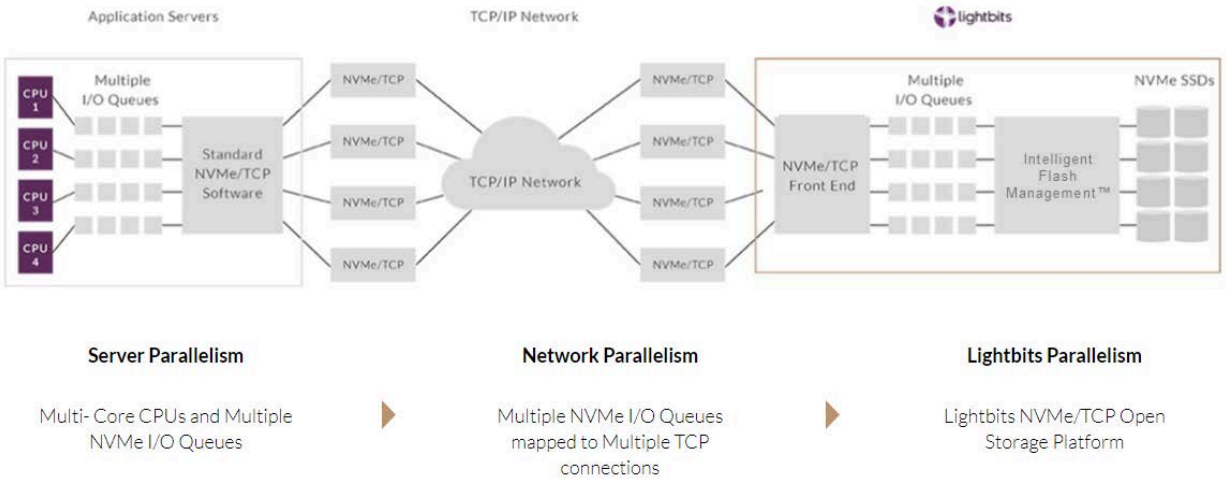

Figure 4. Lightbits with NVMe over TCP are designed for parallelism to fully exploit the benefits of flash.

The industry as a whole embraces the role of NVMe over TCP in data center modernization:

"By 2027, 25 percent of enterprise organizations will deploy NVMe-oF as a storage network protocol, up from less than 10 percent in mid-2023." Storage Analyst, Gartner

“NVMe-over-TCP makes networked NVMe storage easy to deploy at scale, while delivering performance similar to local flash.” Principal, Enterprise Strategy Group

“NVMe/TCP allows organizations to provision scalable storage without fundamentally having to change their network architecture, and provides latencies akin to that provided from conventional direct-attached storage.” Analyst, IDC

“As network bandwidth advances and Ethernet technology innovations continue, organizations will move away from the constrained limitations of traditional fibre channel SAN to more flexible and scalable network-based storage, with NVMe/TCP a clear frontrunner.” Field CTO, Red Hat

---

5.5 Intelligent Flash Management

QLC flash is performant and inexpensive, but not suitable for use locally in application servers where the write pattern is unpredictable. The write performance of QLC flash is poor compared to that of more expensive TLC and MLC devices. Unless write patterns are sequential and in large chunks, QLC media can be worn out quickly. It’s therefore difficult to take advantage of the lower cost of QLC flash directly in application servers, especially in high-performance environments.

Lightbits software has advanced Intelligent Flash Management, specifically optimized for cost-effective QLC SSDs. Lightbits ensures that all writes are staged to the QLC media sequentially and in large chunks, extending the endurance of QLC by up to 20x. By aggregating writes over multiple devices, Lightbits ensures high write performance with consistent response times. Lightbits thus enables the use of low-cost QLC without compromising performance and improves overall utilization of the flash media, leading to better price performance compared to traditional SANs that often rely on more expensive types of flash.

5.6 Enterprise-Ready Data Services

Lightbits offers the essential data services of legacy SAN storage. To maximize your storage resources, Lightbits optimizes SSD media through smart data placement, thin provisioning, and compression for up to 4:1 total data reduction. Space-efficient snapshots provide near-instant restores for data protection, and fast clones enable DevOps to innovate at the speed of NVMe. Multi-zone synchronous replication across racks or data centers provides resiliency for business continuity. To support mixed workload environments, Lightbits provides multi-tenancy with Quality of Service capabilities to prevent resource hogging.

6. Benefits of Modernized Infrastructure for Service Providers

To effectively compete and thrive, a modern infrastructure that consists of disaggregated, software-defined storage natively designed with NVMe/TCP in open-source environments enables cloud and platform service providers to offer differentiated, high-performance services tailored to niche markets and rapidly innovate with new services – attracting customers whom the hyperscalers might underserve. This ultimately fosters customer loyalty and sustainable growth in a market increasingly demanding technological sophistication.

6.1 High Performance and Low Latency at Scale

SDS with NVMe/TCP in KVM, OpenStack, OpenShift, or Kubernetes empowers service providers with significantly enhanced storage performance by disaggregating storage resources and utilizing a high-speed, low-latency protocol over standard Ethernet. This architecture allows for greater parallelism and higher IOPS than DAS and SAN, leading to faster application response times and improved user experiences for hosted workloads. The ability to offer premium, performance-tiered storage services at potentially lower costs than FC SAN can attract performance-sensitive customers, improve overall infrastructure efficiency and resource utilization, accelerate data-intensive workloads like AI/Machine Learning, and provide a more agile and scalable storage foundation for future service innovation and growth.

Lightbits SDS scales to hundreds of petabytes and delivers performance of up to 75 million IOPS and consistent sub-millisecond tail latency under a heavy load. Lightbits' cluster performance testing, consisting of three servers, is shown below (see Figure 5 below) (see Table 1).

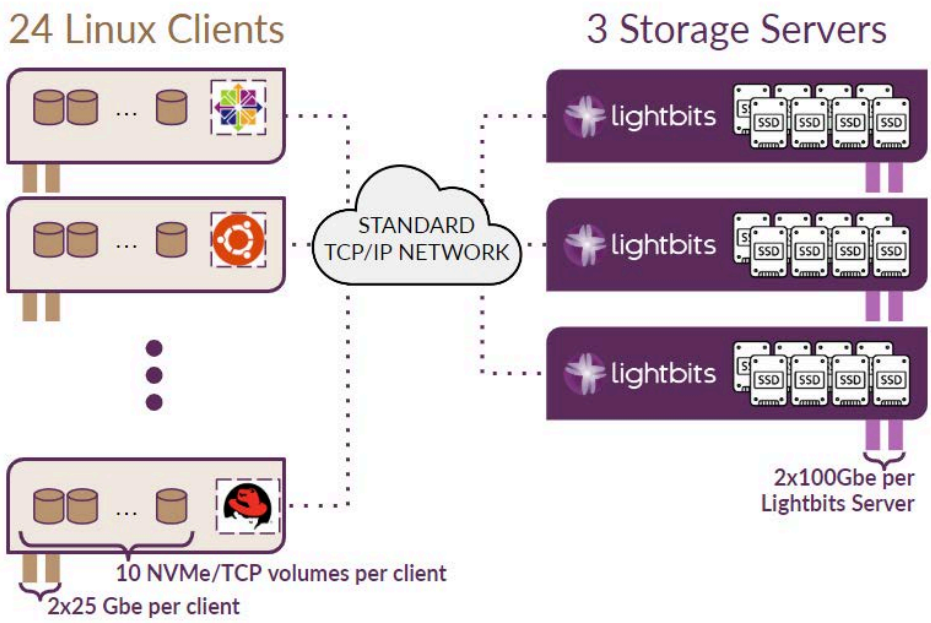

Figure 5. Lightbits cluster

Table 1

Random I/O, Replication Factor 2			
Workload	Block Size	IOPs	BW
Read Only	4KB	14M	53GB/s
	128KB	535K	65GB/s
70% Read/30% Write	4KB	5.3M	20GB/s
	128KB	325K	40GB/s
50% Read/50% Write	4KB	3.8M	14.6GB/s
	128KB	205K	25GB/s
Write Only	4KB	2.3M	8.9GB/s
	128KB	107K	13GB/s

6.2 Flexibility and Agility

DAS, SAN, and HCI storage deployments often suffer from low utilization and a lack of flexibility. Low usage means that purchased infrastructure is sitting idle. Lack of flexibility means that the operational expenses are higher than they should be because it's hard to adapt to changing workload conditions and increase either storage or compute capacities as needed.

With SAN, you must always scale both compute and storage together. Low efficiency and rigid infrastructure contribute to higher costs, and proprietary hardware requirements can introduce supply chain risk.

Eliminate vendor lock-in and reduce CapEx through the use of commodity hardware, improve time-to-market for new services and features by reducing supply chain bottlenecks, improve resource utilization and cost efficiency for lower OpEx, and improve responsiveness to evolving customer needs and workload requirements.

6.3 Lower Storage Costs, Better Cost-Efficiency

For service and platform builders, unpredictable customer and workload demands often result in overprovisioned resources in all legacy models – DAS, SAN, and HCI. Vendor lock-in introduces risk for service and platform providers with fixed pricing structures, upgrade cycles, and supply chain shortages that lead to escalating infrastructure expenses and eroded margins.

Where application servers with local NVMe are often only 15-25% utilized, Lightbits' clustered, disaggregated SDS allows you to scale performance and capacity independently and dynamically to maximize flash utilization.

By enabling the use of cost-effective commodity hardware instead of proprietary storage arrays and by maximizing the utilization of high-performance NVMe flash through efficient data placement and management, Lightbits can lower your TCO by up to 80%.

SDS storage in KVM, Kubernetes, OpenStack, or OpenShift environments eliminates vendor lock-in and allows service providers to right-size their infrastructure based on actual demand, avoiding over-provisioning. The benefits of this model include reduced CapEx on storage hardware, lower OpEx related to power, cooling, and maintenance, improved profitability, increased return on investment for infrastructure assets, and the ability to offer more cost-effective services to attract a broader range of customers while maintaining high performance and reliability.

6.4 Scalability

Monolithic storage arrays come with inherent scalability limitations. Hyperscalers achieve their incredible scaling not through building one large system that grows indefinitely, but by constructing larger systems out of many loosely aggregated smaller systems.

Lightbits' clustering support builds storage clusters from many smaller storage servers – loosely aggregated and managed as one. Ultimately, the optimized software and hardware infrastructure serve remote block devices to millions of servers. This scale-out architecture allows for non-disruptive growth and the ability to elastically adapt to fluctuating customer demands and increasing data volumes.

Service providers and platform builders can:

- Accommodate rapid growth without significant upfront investments or complex upgrades.
- Improve cost-efficiency by scaling resources incrementally as needed.
- Enhance service agility in responding to market opportunities.
- Increase customer satisfaction by meeting changing workload requirements.
- Handle unpredictable spikes in customer demand while maintaining performance and availability.

6.5 Reliability

The challenge for cloud and platform service providers is that systems fail. Reliability is critical for data infrastructure, which must be reliable despite being built out of unreliable components such as commodity servers, NICs, and NVMe SSDs. Single points of failure are inherent in traditional storage arrays. From the perspective of clients accessing the data, data replication has to be transparent, and server failover must be seamless.

SDS storage solutions ensure data availability and service continuity even during hardware failures by leveraging commodity hardware with built-in redundancy and software-based data protection mechanisms like replication and erasure coding. Lightbits' clustered architecture eliminates service disruptions if nodes or drives fail or become inaccessible, heals itself when nodes or drives are replaced, and supports non-disruptive rolling software upgrades.

A modern software-defined infrastructure will result in increased customer trust and satisfaction due to consistent uptime and data protection, reduced downtime and associated revenue loss, lower operational overhead related to reactive maintenance and data recovery, enhanced Service Level Agreement (SLA) adherence, and a stronger reputation for providing dependable and resilient cloud services. This ultimately leads to greater customer retention and acquisition.

7. Conclusion

For cloud services and platform builders, data infrastructure modernization is not a choice - it's imperative for business success. The constrained limitations of legacy systems like DAS, SAN, and HCI supporting mixed and high-performance workloads at scale can lead to higher CapEx and OpEx, reduced margins, poor customer satisfaction, and less revenue due to system downtime. Foundational to data infrastructure modernization is disaggregated, and software-defined storage is natively designed with NVMe over TCP in open-source environments like Kubernetes, KVM, OpenStack, and OpenShift.

Compared to DAS, Lightbits can deliver higher utilization for lower TCO and better utilization of flash for longer endurance of QLC. Compared to SANs, Lightbits NVMe/TCP delivers high performance without the proprietary hardware stack. Compared to HCI, Lightbits can manage mixed (virtualized and non-virtualized) workloads in a single cluster, manage resources more efficiently, and yield low, consistent latency – putting it in a performance leadership position. Lightbits can be seamlessly integrated into OpenShift Virtualization, OpenStack, KVM, Kubernetes, and more – and managed through Cinder, Nova, and Glance through the Cinder API.

Lightbits reduces business risk, running on commodity hardware to eliminate supply chain disruptions and eliminate single points of failure on data and control paths to keep services up at all times. The software only needs to use the standard NVMe/TCP drivers, leading to simple and easy deployment on existing application servers and with existing TCP/IP networks. Multiple Lightbits clusters can also exist in the same cloud data center, which can be easily aggregated and managed as one large cloud-scale block storage solution.

Lightbits' disaggregated, SDS brings hyperscale efficiency, scalability, flexibility, and reliability to your infrastructure operations. Cloud and platform service providers can power any application in virtualized and containerized environments, at any scale, while benefiting from the industry's best price/performance scenario.

"Legacy storage infrastructure can and will impact application performance of data-driven environments. Thus, storage is fundamental and must be the first consideration of any modernization effort….the importance of flexible, disaggregated storage solutions becomes critical. Scale matters, and performance scale is even more important. Companies like Lightbits Labs deliver performance, scale, and cost savings realized by some of the largest organizations." Vice President & Principal Analyst at Moor Insights & Strategy

---

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

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

info@lightbitslabs.com

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

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

LBWP13/2025/05