



PRODUCT OVERVIEW

LightOS™ 2.0

Software Defined Composable Storage

Lightbits Labs LightOS™ is software-defined block storage bringing hyperscale efficiency and flexibility to all. It delivers composable, high-performance, scale-out and redundant NVMe/TCP storage that performs like local flash.

LightOS delivers an unmatched combination of performance, resiliency and scalability. It simplifies infrastructure management and operations while simultaneously lowering cost.

HIGH PERFORMANCE

LightOS storage is disaggregated, redundant and offers rich data services yet performs like local NVMe flash. A single LightOS target server with drive and target server redundancy enabled can serve up to 3 million random 4K read IOPs and up to 800,000 random 4K write IOPs. The same server can provide up to 16GB/s of read bandwidth and up to 8GB/s of write bandwidth.

High IOPs and bandwidth are nice, but NVMe based storage excels at low latency. This is what makes it the media of choice for databases, analytics and many more modern, scalable cloud-native applications.

LightOS performs at local NVMe latencies even under load. The same target server as described above delivers 4K IO response time averages of 200µs at over 2 million read IOPs and 300µs at nearly 600,000 write IOPs.

LOWER TOTAL COST OF OWNERSHIP

LightOS lowers your total cost of ownership both for the initial purchase, as well as over time with greater operational efficiency.

Application server environments with local NVMe are often only 15-25% utilized. When moving to a LightOS clustered, centralized storage service there are vast improvements in capacity and performance utilization.

This means less money is spent on NVMe flash while providing a more operationally efficient environment. High availability features mean servers and applications remain available in light of drive failures and applications are not bound to specific server hardware.

WHY LIGHTOS 2.0?



Accelerate Databases
Build Cloud Native Storage
Manage Big Data Analytics
Power Your Private Cloud



kafka



cassandra

RICH DATA SERVICES


Lower TCO is not only achieved by improving capacity and performance utilization. LightOS offers rich data services that are not generally associated with NVMe storage, at NVMe performance latencies. All LightOS logical volumes are thin provisioned and when combined with compression support (that can be enabled/disabled on a per-volume basis), LightOS can achieve total data reduction levels of as much as 10:1 in service provider and private cloud environments.



10:1 Total Data Reduction Levels

ENABLING QLC FLASH

QLC flash is inexpensive, but often not suitable for use locally in application servers where the write pattern is small, or unpredictable. Write performance of QLC flash is poor when compared to more expensive TLC and MLC devices. Lastly, unless write patterns are sequential and in large chunks, it's possible to wear out QLC media quickly. Thus, it's difficult to take advantage of the lower cost of QLC flash directly in application servers, especially in service environments.



LightOS lowers cost by enabling the use of QLC flash in LightOS targets. LightOS ensures all writes are staged to the QLC media sequentially and in large chunks, extending the endurance of QLC by up to 5 times. By aggregating writes over multiple devices, LightOS ensures high write performance with consistent response times. Thus, LightOS enables the use of low-cost QLC media without compromising performance or flash endurance.

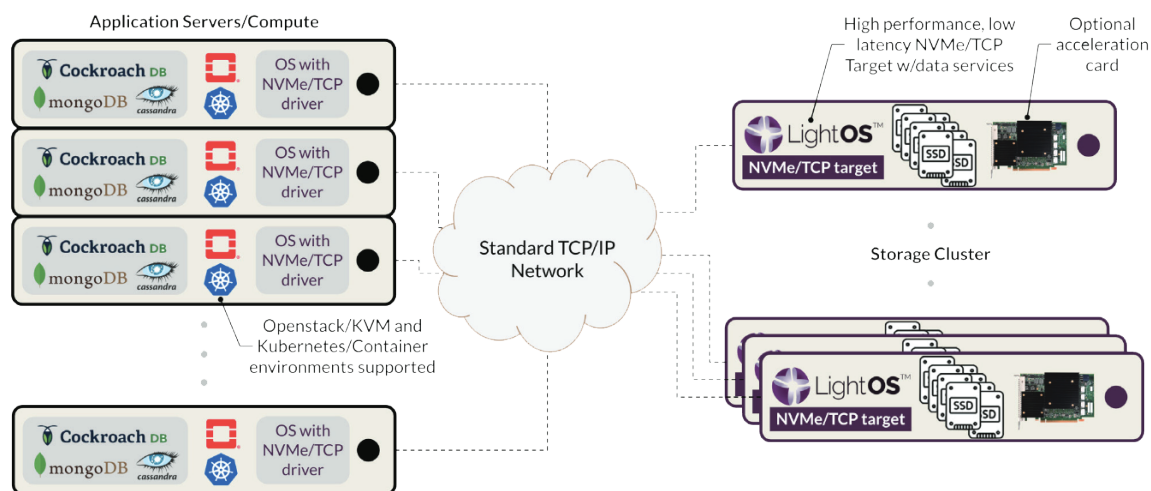
FLEXIBILITY AND EASE OF DEPLOYMENT

The LightOS solution is extremely flexible and well suited to integrate with your private cloud, service provider or modern enterprise data center environment. It can be deployed on your existing data center infrastructure without replacing any component of the data center infrastructure.

STANDARD X86 SERVERS AND NVME DRIVES

LightOS provides the freedom to tailor storage server configurations to the unique needs of any environment. Servers can be configured with standard NVMe drives, Network Interface Cards (NICs) and in various form factors. Different LightOS servers in the same cluster can have varying numbers and sizes of drives, and drives can be added on-the-fly when desired.

You simply populate LightOS target servers with the drives of your choice (instead local to the application servers) and then any application server on the network can enjoy NVMe performance without the inefficiencies of direct attached flash.



LightOS on standard servers showing various applications, software components and topology

STANDARDS BASED, “NO TOUCH NETWORKING” ON CLIENTS AND SWITCHES

LightOS implements NVMe over Fabrics (NVMe-oF) utilizing the TCP protocol on Ethernet (NVMe/TCP). Other NVMe-oF solutions that require RDMA protocols (RoCE or iWarp) require special Ethernet NICs from a limited set of vendors and are more expensive than their non-RDMA counterparts. RoCE protocol also requires special network switch settings.

LightOS NVMe/TCP not only utilizes the same network cards and infrastructure that might be used by protocols like iSCSI, but the application server block driver is also 100% standard and included in all major recent Linux distributions. Thus, there is no proprietary software to install on application servers. Upgrades are easy as the required drivers are part of Linux distributions and in the upstream kernel. Network switches don't require any special settings for the NVMe/TCP storage protocol. LightOS is a target side only solution that works with the networking hardware and practices already in place in the data center.

BARE METAL, CONTAINER AND VIRTUAL MACHINE ENVIRONMENTS SUPPORT

LightOS meets Linux application service needs regardless of the type of deployment . LightOS is a set of packages that is applied to popular Linux distributions including Red Hat Enterprise Linux, CentOS, Ubuntu, SUSE SLES, OpenSuse, Debian and Fedora. For bare metal and image maintenance, LightOS is provided with Ansible playbooks.

Virtualization environments such as Openstack and KVM/QEMU are supported via a Cinder driver. Kubernetes and other container orchestration environments are natively supported with LightOS persistent storage via CSI.




For cloud (public or private), Openstack and Kubernetes environments scalability is key and LightOS meets those scalability challenges by supporting up to 32K client connections per target server and up to 64K logical volumes per cluster. Clients are free to connect to multiple clustered simultaneously, so scalability of capacity and performance is practically limitless.

For questions on Microsoft Windows environment support, please contact us.

SCALABILITY

LightOS offers unprecedented scalability with options to both scale up and/or scale out. LightOS target servers need not be deployed with 100% of their drive bays filled. LightOS Elastic RAID supports hot-plug addition of drives at any time without disruption and immediately adds capacity to that target server's NVMe pool.

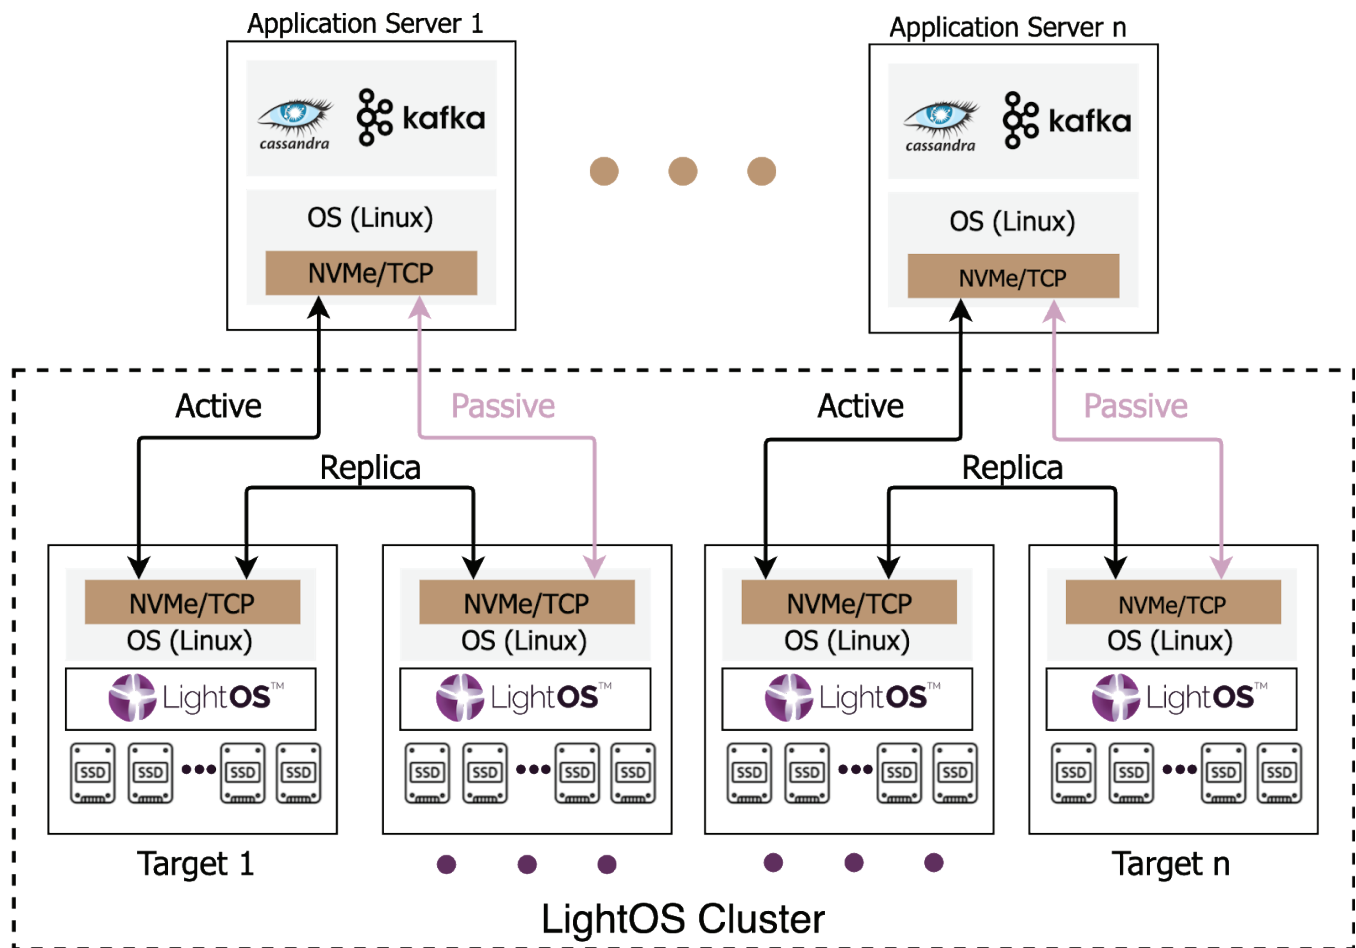


Scale out is achieved by adding LightOS target servers to an existing cluster, adding additional performance and capacity for new volumes.

LightOS 2.0 clusters can scale to an aggregate 48 million IOPS and as large as 15 petabytes of usable capacity. These numbers are expected to increase over time. As software, LightOS takes advantage of improvements in CPUs, NVMe drives and networking to constantly evolve to higher levels of performance and reduced latency.

SELECTABLE HIGH AVAILABILITY ON STANDARD SERVERS

LightOS 2.0 is the world's first software defined NVMe/TCP to offer full redundancy against both drive failures as well as storage target server failures. This is achieved via Asymmetric Namespace Access (ANA), more commonly multi-pathing, or multi-path I/O. ANA is part of the NVMe specification (1.4) for local, PCIe based NVMe and also across fabrics with the 1.1 NVMe-oF specification. With LightOS volumes set to 2x or 3x replication, an application server utilizing the NVMe/TCP driver is aware of multiple network paths to a logical volume (namespace) identifier. One path will be the active (optimized) path and the additional 1 or 2 other paths will be passive (non-optimized).



Network link/path failover can also be facilitated (if desired) with port bonding methods (such as LACP) with LightOS 2.0.

For drives, LightOS utilizes Elastic RAID within each target server to protect against drive failures. This adaptable erasure coding algorithm adjusts and rebalances in response to drives failing (or being removed) and drives being added. These features working in unison result in a fully redundant storage solution at extremely high availability levels all built on standard components and servers. No special “dual path” drives, nor special chassis are required to build a highly available, high performance and scalable block storage solution with LightOS.

LIGHTOS 2.0 KEY FEATURES

UNMODIFIED SOFTWARE ON CLIENTS

- Leveraging standard NVMeoF 1.1 multipathing (ANA)
- Clients can be connected to multiple storage clusters enabling effectively unlimited capacity

CLUSTERED/FAILOVER STORAGE SOLUTION

- Distributed and durable cluster management with fast failover handling (2-3 seconds average)
- No single point of failure in Data and Control paths
- Cluster size: 3-16 servers
- Automatic volume placement to balance capacity and IOPs
 - Replicas are stored indifferent failure domains

HIGH PERFORMANCE AND LOW LATENCY

- Max IOPS per Target Server:
- 4K Random Read: 3M IOPs
- 4K Random Write: 800K IOPs
- Latency Per Target Server, 2x replication:
 - Average 4K Random Read $\leq 200\mu s$ @ 2.1M IOPs
 - Average 4K Random Write $\leq 300\mu s$ @ 600K IOPs
- Max Bandwidth per Target Server:
 - 16GB/s reads
 - 8GB/s writes

MAXIMUM USABLE CAPACITY (16 TARGETS)

- 15PB (1x replication, no Elastic RAID, 2:1 compression)
- 7.3PB (2x replication, Elastic RAID, 2:1 compression)
- 2.4PB (3x replication, dual Elastic RAID groups, no compression)

STORAGE SERVICES

- Thin provisioning
- Compression/decompression configurable per volume
- Elastic RAID (per target) for protection against drive failures
 - No interruption of services on transient or permanent SSD failures
 - Automatic rebuild
- Volume replication (1x, 2x or 3x) configurable per volume
- Non disruptive drive additions for scale-up

NODE MANAGEMENT

- Replication and failover handling
 - Rebuild after permanent failure
 - Rapid partial rebuild after transient failures for unlimited time
- Network failures handling using NVMe/TCP multipath
- Non-disruptive (for 2x and 3x replicated volumes) cluster upgrades

APPLICATION ENVIRONMENT SUPPORT

- Kubernetes v1.13, v1.15 and above via CSI
- Openstack Queens, Rocky & Stein via Cinder
- Bare Metal
- Ansible playbooks for various Linux distributions

MANAGEMENT

- RESTful API provides a standard HTTPS-based interface
- CLI support for scripts and monitoring
- Metrics and Alerts based on Prometheus
- Pre-configured Grafana dashboards


DEPLOYING LIGHTOS

LightOS is available as software only, software with an optional hardware accelerator (LightField), or as a preconfigured storage target (SuperSSD).

THE POWER OF CHOICE



Select the x86 server platform hardware and NVMe drives from a vendor of choice



The optional LightField™ accelerator card offloads data services, reducing CPU needs and lowering latency for compressible data workloads



SuperSSD™ makes deployment easy with pre configured hardware and software with a choice of capacities and full hardware and software support



SOFTWARE-DEFINED STORAGE

LightOS 2.0 as software is licensed per storage server on an annual subscription basis. It runs on x86 servers and utilizes standard Ethernet cards and NVMe drives. In general, the minimum CPU requirement is 10 cores. Target servers, in general, should use 100Gbps Ethernet interfaces with each interface capable of supporting up to 8-10 GB/s of storage bandwidth, largely depending on the number of CPU cores. Lightbits Labs is happy to provide reference platform guides and/or consulting on the right server configuration tailored to workload requirements.

HARDWARE ACCELERATION OPTION

Lightbit Labs can provide a LightField card for offload of data services from the server CPU. The card is 100% optional and can accelerate compression/decompression and Elastic RAID functionality. It has the potential to save money as well by allowing the selection of lower core-count CPUs in target servers.

DEPLOYMENT READY APPLIANCES

For those that value convenience in a heavily tested and optimized platform, LightOS is available as a SuperSSD appliance. This 2U, 24 drive platform comes in various pre-configured capacities with software and hardware support. This is the fastest and easiest way to get LightOS deployed and is backed by world-wide warranty and support services.

WHY LIGHTBITS 2.0

LightOS 2.0 represents a revolution in high performance, scale-out block storage. The need for special fabrics and/or network protocols to achieve low latency, high IOPs and high bandwidth is a thing of the past. With LightOS, you can use ubiquitous TCP/IP on Ethernet and achieve performance levels 5-10 times higher than all-flash-arrays while paying a fraction of the cost of those same arrays. Moreover, data services don't have to be sacrificed as LightOS provides rich data services commonly associated with legacy storage solutions.

Performance. Lower TCO. Flexibility. Let LightOS show you the way.

FIND OUT MORE

To learn more, please visit our website, www.lightbitslabs.com

To contact our team, email us at info@lightbitslabs.com

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.

All third party product and company names and/or logos are trademarks™ or registered® trademarks of their respective holders. Use of them does not imply any affiliation with or endorsement by them.