Reach for the Stars

Many IT organizations that rely on Kubernetes are compromising application portability for performance. **Are you one of them?**

More and more businesses are turning to containers and Kubernetes for:

Scalability, manageability, and portability of apps

Portability of cloud-native applications that support web-based and service-oriented architectures

29%

expected compound annual growth rate (CAGR) for the container market over the next five years[1]

48%

Kubernetes adoption rate over the next five years (up from up to 27% today)[1]

When SQL, NoSQL, and in-memory database applications are deployed on bare metal, most software vendors recommend using direct-attached local NVMe flash for the best application performance.

CAUTION:
With direct-attached local flash for persistent storage, you lose application portability, which can result in poor flash utilization.

YOUR MISSION

Hit the stratosphere with persistent storage Kubernetes—without compromise. Shoot for the moon with rich data services and a lower TCO.

Choose your route:

One option is local persistent volume functionality:
This can lead to compromised application portability, and it introduces poor flash utilization.

Another option is an open source or proprietary block-based storage solution with a Container Storage Interface (CSI) driver:
This can lead to compromised simplicity and performance, while introducing hardware vendor lock-in.

There's a better option to get persistent storage for Kubernetes while navigating your way out of the performance-versus-portability compromise.

Many have already plotted a successful Kubernetes journey by using high-performance, software-defined block storage and a wide range of Intel® technologies to power their modern cloud-native applications.

Lightbits LightOS + Kubernetes + Intel technologies align with Kubernetes values

Performance
High IOPS and low latency enabled by LightOS and a wide breadth of Intel technologies

Flexibility
Persistent storage over standard TCP/IP networks

Lower TCO[2]
Coupling LightOS, optimized for Intel's high-performance hardware platform, with low-cost SSDs

Land with no-compromise performance, light years ahead of other software-defined or storage array solutions.

LightOS

Up to 4x higher throughput for 4K reads[3]

Up to 17x more throughput for 4K writes[3]

Up to 14x more throughput for 8K read/write workloads[3]

Up to 10x more throughput for 16K read/write workloads[3]

Up to 5x more throughput for 32K read/write workloads[3]

LightOS, optimized for low-cost SSDs and Intel high-performance technologies, delivers a powerful cost-effective platform for Kubernetes environments.

Why wait? Start navigating your way to better storage architecture by visiting

lightbitslabs.com/kubernetes-persistent-storage/

[1] Research and Markets with Mordor Intelligence. "Application Container Market - Growth, Trends, COVID-19 Impact, and Forecasts (2021-2026)." January 2021. **researchandmarkets.com/reports/4845968/application-container-market-growth-trends.**

[2] LightOS can help lower your total cost of ownership (TCO) both for the initial purchase and over time, with greater operational efficiency. Source: Lightbits Labs. "Kubernetes and LightOS: Performance, Persistence, Simplicity." June 2021. **lightbitslabs.com/ty-solutions-brief-kubernetes/**

[3] Performance of Lightbits LightOS and the Intel technology platform compared to Ceph storage on Red Hat OpenShift container platform. Source: Internal testing by Lightbits Labs. **Configurations: Lightbits LightOS platform:** Intel® Xeon® Gold 6338 processor (2.00 GHz), 256 GB RAM, 8 x 14 TB Intel QLC NAND SSDs, and Intel Ethernet Network Adapter E810 (single port used in 25 gigabit Ethernet [GbE] mode). **Ceph on Red Hat OpenShift container storage platform worker nodes:** Intel Xeon Platinum 8173M processor (2.00 GHz), 96 GB RAM, Intel Ethernet Network Adapter XXV710 (single port used in 25 GbE mode). **Workload description:** Software: vdbench50407. 12 worker nodes; each worker node has 12 pods running vdbench, and each pod has 1 TB persistent volume claim (PVC). All PVCs are first fully written, and then the different types of workloads (block size and read/write ratio) run for 45 minutes (nine times each), where the first 15 minutes out of the 45 are considered warm up, so the performance measurement is from minute 16 to 45.

Performance varies by use, configuration and other factors. Learn more at **www.Intel.com/PerformanceIndex.**

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Printed in USA 1121/SM/PRW/PDF Please Recycle 348263-001US

intel + lightbits