Technical Review

Solving Hyperconvergence (HCI) at the Edge with Sunlight

Date: May 2021 **Author:** Alex Arcilla, Senior Validation Analyst

Abstract

This ESG Technical Review documents the evaluation of the Sunlight NexVisor hyperconverged infrastructure (HCI) platform. We reviewed and analyzed how the platform achieves high levels of manageability, efficiency, and performance especially in edge computing environments.

The Challenges

Organizations continue to make significant investments into hyperconverged infrastructure, as ESG research uncovered it to be one of the top five data center modernization initiatives in the upcoming year (see Figure 1).[1]

Figure 1. Top Five Data Center Modernization Investments

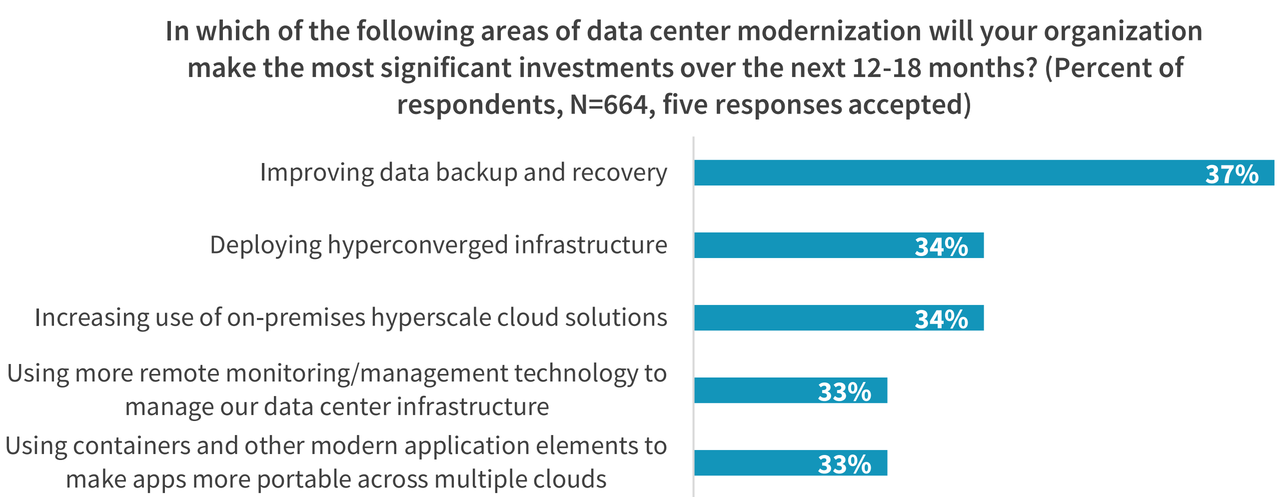

Source: Enterprise Strategy Group

At the same time, ESG research reveals that 37% of respondents consider improving their data analytics for real-time business intelligence and customer insight as one of the top five business initiatives that will drive the most technology spending in their organizations over the next 12 months.[2] While organizations run more data intensive workloads, such as analytics, big data, and artificial intelligence (AI), the volume of data to be processed is increasingly occurring at the IT network's edge. However, extracting value from the data is delayed as they are processed in a centralized location, either within a corporate data center or in the public cloud.

While the obvious solution would be to distribute processing power to the edge, it is often not possible using traditional data center hardware or hyperconverged infrastructure. Deploying the large amount of hardware required for such a task is not feasible, as data is generated and collected within constrained environments, with little room to deploy and manage the required hardware. And while hardware utilization can be maximized via virtualization, current hypervisors still require

[1] Source: ESG Master Survey Results, *2021 Technology Spending Intentions Survey*, December 2020.
[2] Ibid.

This ESG Technical Review was commissioned by Sunlight.io and is distributed under license from ESG.
© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

additional CPU and memory capacity, leading to overprovisioning servers at the edge. Even deploying high performance hardware technologies—NVMe storage, 100G networking infrastructure—cannot compensate for the "virtualization tax," which can cost up to 70% of bare-metal performance.

The Solution: Sunlight Virtualization

Sunlight is a hyperconverged infrastructure stack that cost-effectively maximizes the efficiency and performance of a consolidated mix of virtualized and containerized applications. As shown in Figure 2, Sunlight is designed to remove the performance bottlenecks associated with legacy hypervisors and virtualization software stacks running on the latest high-speed hardware (e.g., NVMe, Intel Optane, NVMe-oF, ARM, 100 GbE). The small footprint and performance efficiency of Sunlight virtualization helps organizations achieve near bare-metal performance for a consolidated mix of data-hungry applications deployed on-premises, in the cloud, and, importantly, in constrained compute environments at the edge.

Figure 2. Reducing IT Costs and Maximizing Workload Performance with Sunlight

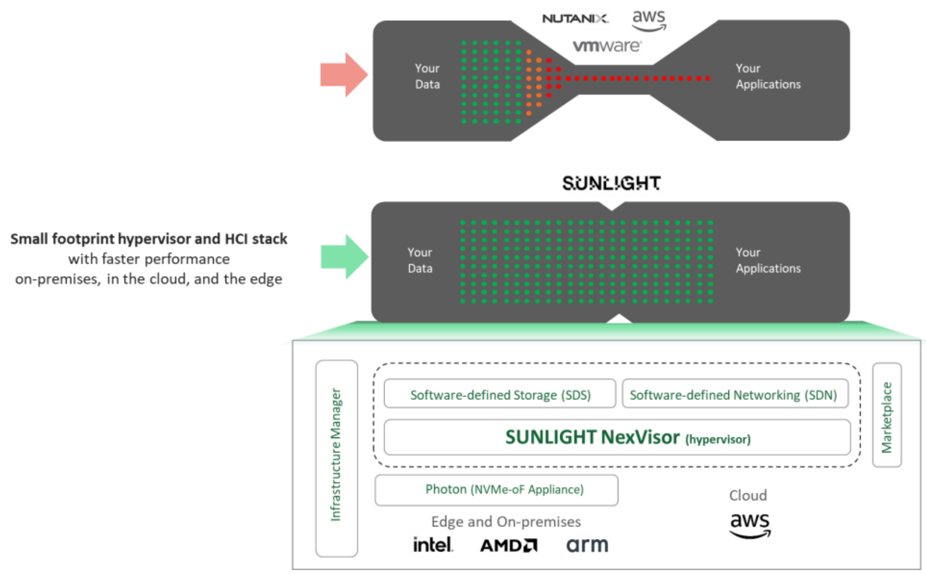

Source: Enterprise Strategy Group

At the heart of the platform is the Sunlight NexVisor, a multi-tenant hypervisor that enables organizations to extract bare-metal performance of high-performance hardware technologies. The Sunlight NexVisor supports virtualization of both storage and networking technologies via Sunlight software-defined storage (SDS) and Sunlight software-defined networking (SDN). With Sunlight SDS, organizations can aggregate local storage drives into a virtual storage pool for deploying any amount of storage to individual VMs. Sunlight SDN enables users to create virtual pools of bandwidth supplied by multiple 10/100G NICs, enabling fully virtualized network deployments. NexVisor can also run AWS Amazon Machine Images (AMIs) natively to simplify migration.

Sunlight Photon is a stackable hardware appliance with up to 128TB of storage capacity per unit and built-in NVMe over fabric adapters that have been pre-installed and configured with the Sunlight virtualization software.

Organizations can manage individual Sunlight-enabled machines via a local element manager, the Sunlight Dashboard, which enables fine-grained control of CPU, storage, and network configuration and performance. It is complemented by the Sunlight Infrastructure Manager (SIM), which enables the remote deployment, management, and monitoring of all Sunlight clusters, both on-premises and in the public cloud, across edge and core networks via a single pane of glass. The SIM can be deployed on demand from the AWS Marketplace as a free product or on-premises.

Sunlight offers the Sunlight Marketplace, integrated with the SIM, to enable "single-click" deployment of third-party applications, such as Splunk or Tensorflow, onto Sunlight clusters. It also supports container-based applications.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Additionally, users can develop "recipes" or playbooks based on Red Hat Ansible for deploying their own custom applications and build custom application marketplaces via the Sunlight open-source recipe hub.

Organizations that wish to automate the creation and deployment of Sunlight clusters can leverage sunlight APIs to integrate with products such as Terraform.

ESG Tested

ESG reviewed the Sunlight platform to examine the ease of manageability with the SIM and the efficiency and performance delivered by the Sunlight NexVisor.

Ease of Deployment and Management

ESG first examined the ease of deploying a two-node Sunlight cluster navigating to the SIM interface (see Figure 3). After choosing the region in which the Sunlight cluster would reside, we inputted the cluster name, the number of nodes (choices range from one to six nodes), AWS instance type for each node, and the Availability Zone for providing redundancy for each node. The process from creating and deploying this cluster lasted between 20 and 25 minutes.

Figure 3. Using SIM to Deploy a Sunlight Cluster in AWS

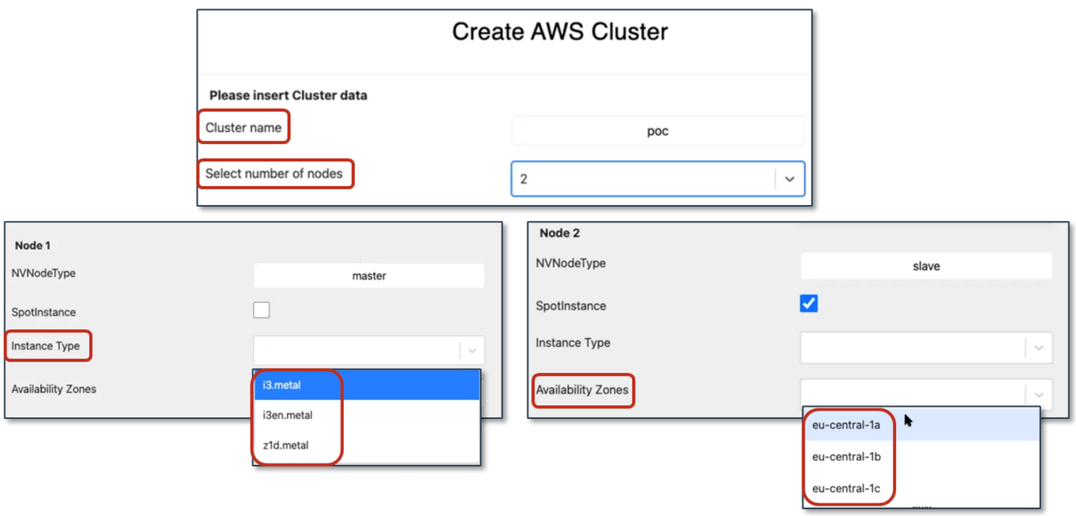

We then navigated to the ***Infrastructure*** tab of the SIM to view more detail about the existing clusters (see Figure 4). We viewed cluster type counts, VM count, and cluster names. Clicking on individual cluster names revealed consumption of CPU, memory, and storage capacity allocated to the cluster, as well as the number of VMs and nodes created.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Figure 4. Viewing Multiple Sunlight Clusters with SIM

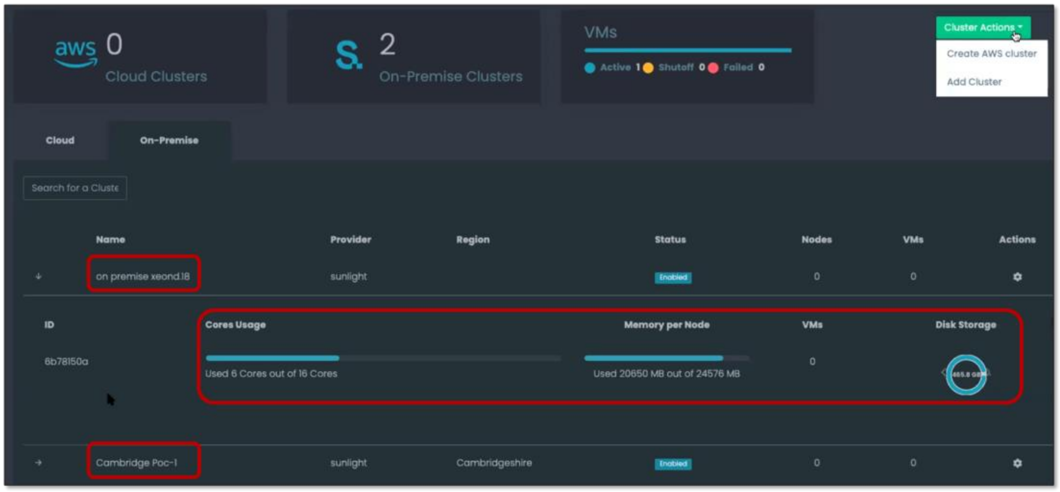

Finally, ESG viewed the Resource Group tab in the SIM (see Figure 5). A resource group contained an allocation of CPU cores, memory, storage tiers, and IP addresses (via NIC ports) in which a user could assign different VMs, depending on a workload's performance requirements. In this case, we saw one VM that existed within the listed resource group called "instance_1." We also noted the capability to delete the VM, move the VM to another resource group, or create a template that could be used to spin up multiple VMs with this specific resource group.

Figure 5. Managing Resource Groups with SIM

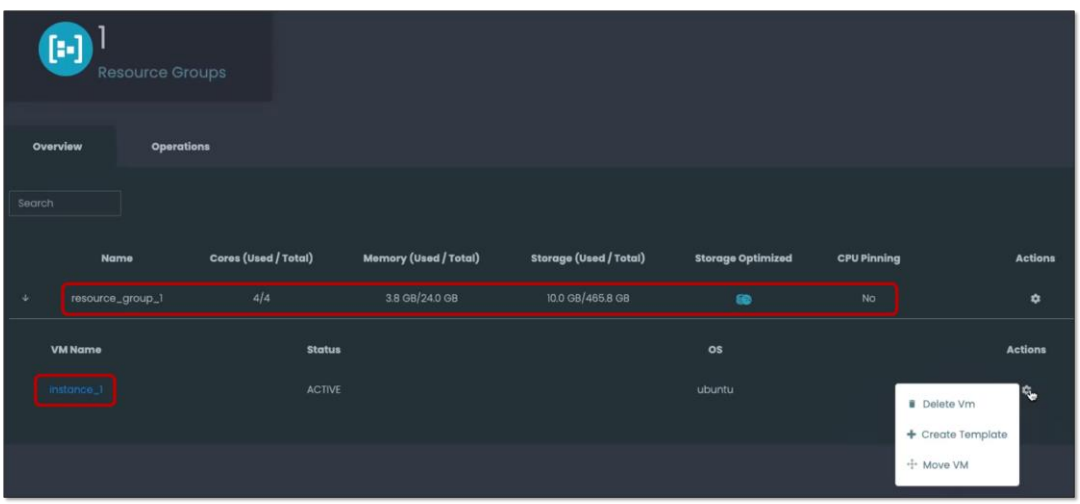

Throughout our review of the SIM, ESG verified that organizations can "deploy and forget" and manage all aspects of Sunlight clusters with a single interface. Less time is spent on management and administration, resulting in accelerated time to value and lowered operational expenses.

Why This Matters

Deploying and managing VMs, either on-premises at the core or edge or within the public cloud, would ideally be quick and easy to save on IT operational costs and meet business demands expediently.

ESG found that deploying and managing Sunlight virtualized workloads is indeed quick and easy when using the SIM. We observed the simplicity of setting up a multi-node Sunlight cluster in AWS and deploying it in a matter of minutes. We also saw the comprehensive view of Sunlight clusters that organizations can gain with SIM, down to the device level.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Efficiency

ESG then examined Sunlight's efficiency, specifically how much CPU and RAM the Sunlight virtualization stack consumes compared with two of the industry-leading HCI vendors. We compared the measured overhead of the Sunlight stack with the publicly available sizing and capacity planning guidelines of competitive HCI solutions. The CPU and RAM utilization of the Sunlight software stack was measured on a four-node cluster with a total of 16 physical processor cores with 128 hyper-threads and 784GB of RAM.[3] The results are summarized in Figure 6 and Table 1.

Figure 6. Comparing Sunlight CPU and RAM Footprint Against Competitive Offerings

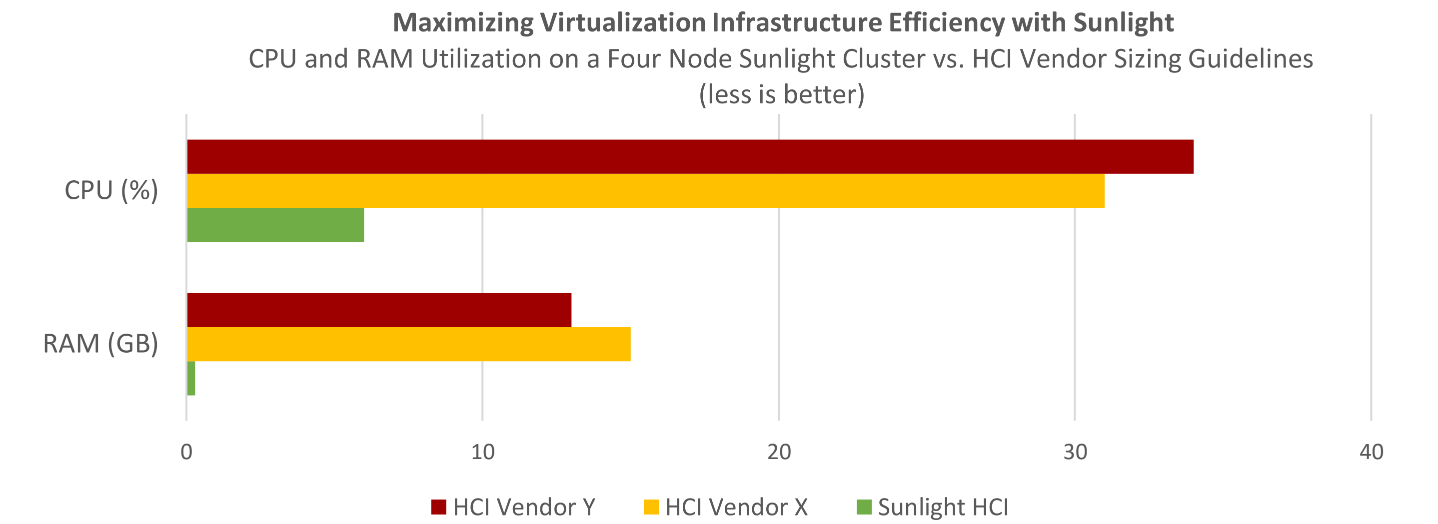

Source: Enterprise Strategy Group

Table 1. Comparing Sunlight CPU and RAM Footprint Against Competitive Offerings

	HCI Vendor X	HCI Vendor Y	Sunlight HCI
CPU (%)	31.0	34.0	6.0
RAM (GB)	15.0	13.0	0.3

Source: Enterprise Strategy Group

What the Numbers Mean

- ESG confirmed that the Sunlight HCI stack consumes 98% less RAM and 81% to 82% less CPU than industry-leading HCI vendors.
- The lower resource utilization and higher efficiency of the Sunlight HCI stack makes more resource available for virtualized applications.
- Consolidating more workloads on less infrastructure lowers the cost of the hardware that needs to be deployed to meet the needs of the business.
- The increased efficiency of the Sunlight HCI stack not only reduces hardware costs, it also reduces power, space, and cooling requirements. This is especially valuable for data-intensive workloads that are being deployed at the edge of the network (e.g., IoT sensors) where space and power are often in short supply.

ESG then examined how the low overhead of the Sunlight stack delivers near bare-metal performance. With a goal of stress testing the overhead of the Sunlight stack, these tests were executed with one of the fastest nonvolatile storage media devices (Intel Optane P4800x) and fastest storage networks (NVMe-oF) that were available at the time this report was published. The performance of a 16-core Sunlight VM with a single Intel Optane P4800 drive was compared to the drive

[3] See the Appendix for more test bed configuration details.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

manufacturer's performance specs.[4] Random 4KB read and write workloads were tested with the de facto standard vdbench utility. The results are shown in Figure 7 and Table 2.

Figure 7. Comparing Sunlight Performance with Intel Optane Drive Performance Specifications

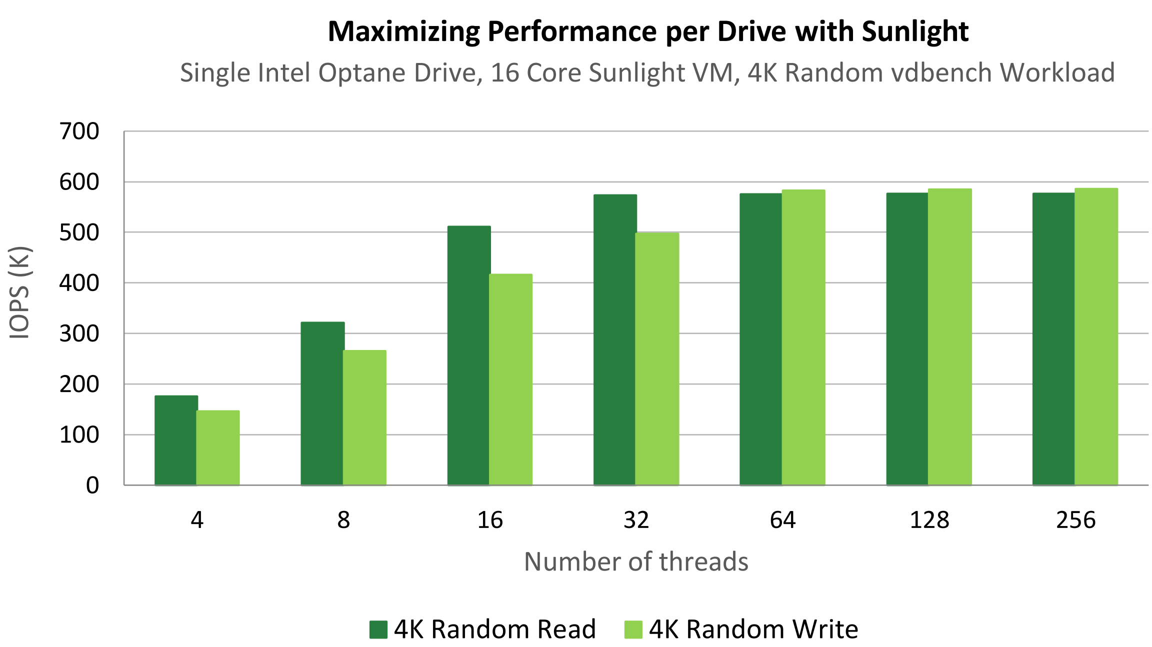

Source: Enterprise Strategy Group

Table 2. Comparing Sunlight Performance with Intel Optane Drive Performance Specifications

	Drive Manufacturer Performance Specifications		Sunlight Performance		Sunlight versus Drive Specifications	
I/O Profile	IOPS	Response Time (us)	IOPS	Response Time (us)	IOPS	Response Time (us)
4KB Random Read	550,000	10	575,000	20.6	4.5%	10.6
4KB Random Write	550,000	10	585,000	22.8	6.4%	12.8

Source: Enterprise Strategy Group

What the Numbers Mean

- ESG confirmed that the low overhead of the Sunlight virtualization stack sustained 4.5% to 6.4% more IOPS than the manufacturer's ratings for an Intel Optane drive while adding a very low level of latency per I/O (10-13 microseconds).
- In other words, the Sunlight stack delivered better than bare-metal IOPS with near bare-metal I/O response times.

Application-level Performance

ESG then evaluated how Sunlight maximizes application performance. We evaluated tests simulating the use of MariaDB, an open-source MySQL relational database management system, running natively on AWS instances—an r5 instance type[5] and r5 instance with provisioned IOPS storage[6]—and on a similarly configured bare-metal instance with Sunlight. Using

[4] See the Appendix for more test bed configuration details.
[5] An AWS EC2 r5 instance is designed for memory-intensive applications such as high-performance databases, distributed web scale in-memory caches, mid-size in-memory databases, and real-time big data analytics.
[6] Provisioned IOPS storage is designed by AWS to meet the needs of I/O-intensive workloads, particularly database workloads that require low I/O latency and consistent I/O throughput.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

sysbench,[7] we generated eight clients running 70/20/10 read/write/other workloads against a 500GB data set.[8] Queries per second (QPS) and latency were measured. Results are shown in Figure 8 and Table 3.

Figure 8. Comparing Performance of MariaDB on AWS with and without Sunlight

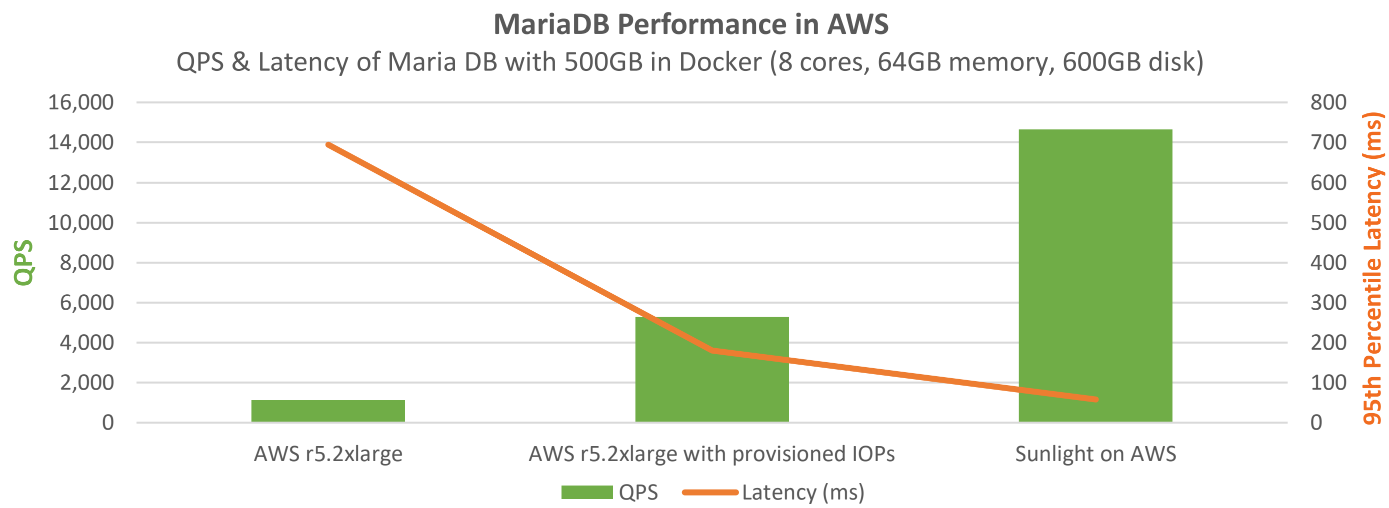

Source: Enterprise Strategy Group

Table 3. Comparing Performance of MariaDB on AWS with and without Sunlight

	AWS r5.2xlarge	AWS r5.2xlarge with Provisioned IOPS	Sunlight on AWS
QPS	1,133.26	5,281.56	14,654.96
Latency (ms)	694.45	179.94	57.87

Source: Enterprise Strategy Group

What the Numbers Mean

- A MariaDB application deployed on a Sunlight virtualization stack on an AWS bare metal instance performed significantly better than a similarly configured AWS instance that leverages the native AWS virtualization stack (up to 12x more QPS with latency improvement of up to 92%).
- Testing on a more expensive AWS instance with provisioned IOPS narrowed the Sunlight performance advantage (2.4x more QPS).
- While these tests were performed in the AWS public cloud, in ESG's opinion, they prove that the efficiency of the Sunlight virtualization stack can be used to improve the performance of data-intensive application workloads running on-premises, in the cloud, or at the edge.

[7] Sysbench is a benchmark tool used to evaluate systems supporting databases under intensive load.
[8] See the Appendix for more test bed configuration details.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Why This Matters

Organizations recognize the need to collect and process the growing amounts of data, especially at the edge, to extract the most value and gain competitive advantage. However, the “virtualization tax” that organizations are paying with legacy hypervisors on the latest high-speed hardware is hindering the performance of data-intensive applications. While deploying more hardware can overcome such performance issues, it adds cost and complexity, particularly in constrained environments at the edge.

ESG verified the lower overhead of the Sunlight HCI stack can cost-effectively maximize the performance of a consolidated mix of data-intensive workloads with less hardware infrastructure. We also confirmed that the Sunlight HCI stack achieves near bare-metal performance due to less CPU and RAM consumption. ESG also verified that Sunlight’s efficient hypervisor maximizes application performance when running Sunlight on AWS.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

The Bigger Truth

Organizations are bolstering their capabilities in extracting value from the growing amount of data they collect to improve overall business agility. As a matter of fact, ESG research has uncovered that 63% of organizations will be increasing 2021 IT spending in AI over 2020 levels, while 47% are increasing investments into other data-intensive workloads centered around business intelligence, analytics, and/or big data.[9] As more organizations are gathering and processing data in constrained edge IT environments, traditional virtualization technologies are not equipped to handle such data-intensive workloads and provide near real-time insights. Organizations thus face over-provisioning of hardware to overcome the "virtualization tax." Even equipping such environments with high-performance hardware technologies, such as NVMe and 100G networking, will not resolve issues brought up by this virtualization overhead.

Julian Chesterfield, an architect of the Citrix Xen hypervisor,[10] founded Sunlight to solve the "virtualization tax" problem for resource-constrained edge hardware and built a lightweight hypervisor to run workloads as if they were running on bare-metal servers. The result is the Sunlight HCI stack. Centered on the Sunlight NexVisor, the platform was designed to help organizations maximize application performance of data-intensive workloads by virtually eliminating the virtualization tax and minimizing the memory and CPU footprint.

ESG validated that Sunlight provides an efficient hypervisor that enables organizations to achieve near bare-metal performance while simplifying deployment and manageability of all Sunlight VM clusters. We first verified that deploying and administering a Sunlight cluster, on-premises and in AWS, can be done in mere minutes via the SIM using a single interface. ESG also verified that the Sunlight stack delivered better than bare-metal IOPS with near bare-metal I/O response times. Finally, we validated that the efficiency of the Sunlight virtualization stack can improve the performance of data-intensive application workloads running at the edge, on-premises, and in the cloud.

Hypervisor technology revolutionized how organizations deployed and managed workloads. It paved the way for today's public cloud infrastructure services. However, this technology has not seen any innovation until now. ESG believes that Sunlight offers the innovation that organizations need to ultimately maximize the value extracted by the volume of data they are collecting. If you seek to maximize the performance of data-intensive workloads, particularly in edge networks, without the burden of the typical virtualization overhead, we strongly urge you to take a closer look at Sunlight.

[9] Source: ESG Master Survey Results, *2021 Technology Spending Intentions Survey*, December 2020.
[10] The Xen hypervisor served as the original foundation of many public cloud infrastructure services, including AWS.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Appendix

Table 4. Efficiency Test Bed

Configuration (Total = 64 Physical Cores, 128 Logical Cores and 784 GB RAM)	
Servers	4
Physical Cores	16
Logical Cores (number of hyperthreads)	32
RAM per Server (Gb)	196
Cluster Nodes	4
HCI Vendor X – Sizing Support	Sizing information was obtained from Vendor X’s publicly available sizing tools.
HCI Vendor Y – Sizing Support	Sizing information was obtained from Vendor X’s publicly available sizing tools.

Table 5. Single Drive Overhead Test Bed

Configuration	
NVMe drive	Intel Optane P4510 (3.8 TB)
100% Random Read (4K Block Size) Manufacturer Published IOPs & Latency	Manufacturer specs were obtained via publicly available data sheets.
100% Random Write (4K Block Size) Manufacturer Published IOPs & Latency	Manufacturer specs were obtained via publicly available data sheets.
Environment Tested	1 x 16-Core VM

Table 6. MariaDB Test Bed

Configuration	
MariaDB	Version: 10.5.4
Host OS	Ubuntu 18.04
Kernel Version	4.15.0-45-generic
Environments Tested	AWS r5.2xlarge & Sunlight 2xlarge

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Validation reports is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Validation reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team’s expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.