

Cost Efficient, High-Performance and Resilient Storage for the Azure VMware Solution





# **Table of Contents**

Azure VMware Solution	3
Enhancing AVS with Cost-Efficient, High-Performance, and Resilient External Block Storage	3
Why Lightbits for Azure VMware Solution?	3
Azure Managed Application for AVS	4
Integration with AVS	5
Zone-Redundant External Storage for Highest Availability	6
AVS-Certified	7
Storage Options for Azure VMware Solution	8
vSAN	8
Azure NetApp Files	8
High Performance NVMe/TCP VMFS Datastores for AVS	8
Storage Performance - FIO	8
Benchmark Methodology	9
Performance Results	9
Application Performance - SQL Server with SQLIOSim	10
Benchmark Methodology	10
Performance Results	10
Support Models	12
The Managed Application Shared Responsibility Model	12
Additional Support Information	13
Get Started	13
About Lightbits Labs	13



# **Azure VMware Solution**

Azure VMware Solution (AVS) provides a VMware Software-Defined Data Center (SDDC) in the Azure cloud. It allows you to focus on developing and running VMware workloads in Azure, while Microsoft manages and maintains the infrastructure and software. AVS is a VMware validated solution with ongoing enhancements and upgrades Azure VMware Solution

For enterprises looking to migrate VMware workloads to the public cloud, AVS provides a fully consistent VMware platform that allows them to 'lift-and-shift' to Azure without time-consuming and costly application refactoring.

# Enhancing AVS with Cost-Efficient, High-Performance, and Resilient External Block Storage

AVS clusters are based on hyper-converged infrastructure built on top of a set of host types with specific amounts of CPU, memory, disk and network. See table (1) for details. While these nodes are generally effective for hosting many workloads, the real challenge arises when migrating storage-intensive workloads to AVS, such as SQL or NoSQL databases. These workloads demand high throughput, low latency, and various storage-to-compute ratios. If your only tactic for scaling performance or capacity is to



add AVS nodes that include compute or storage that you don't actually require, you might encounter budgetary constraints, especially as the storage demand of your applications increases.

Lightbits enhances the current storage offerings within AVS, providing the most cost-efficient external storage solution for workloads that demand low latency and zone-redundancy.

Lightbits is a VMware-certified external storage option for AVS. Using Lightbits, you can create VMFS datastores backed by NVMe/TCP volumes. This combination gives you the high performance and low latency of NVMe-oF external storage with enterprise data services, and all of the flexibility and scalability of VMFS. VMFS has been a popular choice because it's specifically designed to handle the requirements of virtualization and is optimized for VM disk operations. It provides features such as file locking, thin provisioning, snapshots, and more - all of which are essential for efficiently managing VMs in VMware environments.

### Why Lightbits for Azure VMware Solution?

Lightbits for AVS offers many advantages and benefits, including:

- The ability to efficiently run performance-sensitive virtualized applications on AVS
- Ideal for large-scale deployments of virtualized databases: SQL Server, Oracle, MySQL, PostgreSQL, MongoDB, and more.
- Scaling storage independently from compute resources.
- Controlling storage costs with simple, predictable pricing.
- Improving availability and enabling multi-tenancy.
- Certified by VMware and fully integrated with AVS.



Host Type	CPU (Cores/GHz)	RAM (GB)	vSAN Cache Tier (raw)	vSAN Capacity Tier (raw)	Network Cards
AV36	Dual Intel Xeon Platinum 6140 CPUs, 36 physical cores total	576	3.2ТВ	15.20TB	4x 25 Gb/s NICs
AV36P	Dual Intel Xeon Gold 6240 CPUs, 36 physical cores total	768	1.5TB	19.20TB	4x 25 Gb/s NICs
AV52	Dual Intel Xeon Platinum 8270 CPUs, 52 physical cores total	1,536	1.5TB	38.40TB	4x 25 Gb/s NICs

 Table 1: Host types available for Azure VMware Solution

# **Azure Managed Application for AVS**

Lightbits for Azure VMware Solution is available through the Azure Marketplace as a fully-managed service. All of the data resides in your subscription, while Lightbits has access to maintain, operate, and monitor the storage for you.

As shown below in figure (1), when you deploy the Lightbits Managed Application, all of the required resources are created within a managed resource group. These resources include a set of Lv3 series storage-optimized virtual machines, equipped with local NVMe devices connected to AVS using Express Route Ultra Gateway with Fast Path enabled.



Figure -1- Lightbits Deployment for AVS



Lightbits software runs within these virtual machines and creates a storage cluster. A minimum of three VMs is required. The Lightbits storage cluster aggregates all of the NVMe devices available in the Lv3 VMs, and exposes them as a pool of storage that you can use to provision high performance NVMe/TCP VMFS datastores for AVS.

The capacity and performance of the Lightbits cluster depend on your selected virtual machine type and the number of virtual machines you deploy for the managed application. See table (2) for available options and their respective capacity and performance per VM.

For instance, a cluster with 16 L64sv4 VMs provides 139TB capacity (16 x 8.7TB), 13.7 million read IOPS (16 x 860K), and 3.4 million write I/Os (16 x 215K) - all while maintaining data protection through three replicas. Lightbits can easily increase capacity and performance by simply adding more VMs to the existing cluster as demand grows.

Azure VM type	Usable capacity per Lv3 VM (3X replication, 2:1 compression)	Max Read IOPS per Lv3 VM (4KB block size)	Max Write IOPS per Lv3 VM (4KB block size, 3X replication)
L32asv3 7.68 TB raw, 16Gbe	4.4 TB	430,000	107,500
L32sv3 7.68 TB raw, 16Gbe	4.4 TB	430,000	107,500
L48asv3 11.52 TB raw, 24Gbe	6.5 TB	645,000	161,250
L48sv3 11.52 TB raw, 24Gbe	6.5 TB	645,000	161,250
L64asv3 15.36 TB raw, 32Gbe	8.7 TB	860,000	215,000
L64sv3 15.36 TB raw, 32Gbe	8.7 TB	860,000	215,000
L80asv3 19.20 TB raw, 32Gbe	10.9 TB	860,000	215,000
L80sv3 19.20 TB raw, 32Gbe	10.9 TB	860,000	215,000

Table -2- Virtual Machine Types Available for Lightbits with Their Respective Capacity and Performance

# Integration with AVS

Lightbits has partnered with Microsoft to enable NVMe/TCP and VMFS datastores within AVS, through the Run Command infrastructure. With Run Command, you can execute operations using a set of PowerShell cmdlets. This functionality is now integrated into the Microsoft.AVS.Management run command package.

We also introduced an orchestration tool to streamline the setup and administration of Lightbits datastores. This tool offers a straightforward set of commands that automate tasks such as enabling NVMe/TCP for your AVS hosts, connecting AVS to a Lightbits cluster, and creating, deleting, mounting, unmounting, rescanning, and listing VMFS datastores. Behind the scenes, the tool coordinates these operations by invoking AVS run commands and the Lightbits API, as shown in figure (1) above.



For example, to connect your SDDC to a Lightbits Cluster you can run:

Ibavs network connect MY\_SDDC\_CLUSTER\_NAME

To provision a new 1 TiB datastore with 3 replicas and compression enabled, simply run:

Ibavs datastore create MY\_DATASTORE\_NAME MY\_SDDC\_CLUSTER\_NAME 1TiB 3 true

You can find a comprehensive list of available commands and detailed documentation in our <u>Lightbits on</u> <u>Azure VMware solution quick start guide</u>.

# Zone-Redundant External Storage for Highest Availability

The Lightbits cluster can be deployed across multiple availability zones (AZs) to protect your AVS datastores from zone failures. As shown in figure (2), Lightbits will synchronously replicate writes across three AZs. Reads can be served locally for highest throughput and lowest latency.

Your AVS deployment can reside in a single AZ. Even if this AZ experiences an issue, your datastores remain protected by Lightbits, with replicas stored in other AZs. You can then mount these datastores in either an existing or a new SDDC located in a healthy AZ.



Figure -2- Data is Synchronously Replicated to Different AZs by Lightbits' Storage Cluster



AVS also offers support for deploying VMware-stretched clusters on a subset of the regions and on specific host types. When utilizing stretch clusters, you will need to deploy AVS hosts in pairs (minimum 6, maximum 16) and then they will be evenly placed across two AZs. A VMware witness appliance is then deployed in a third AZ. See figure (3) for details.

If you opt to deploy VMware stretch clusters, your Lightbits datastores will be accessible by the AVS nodes running in any of the two AZs.



Figure -3- AVS stretched clusters

# **AVS-Certified**

Lightbits is a VMware Technology Alliance Partner and Day 0 design partner for NVMe/TCP external storage. Lightbits has been collaborating extensively with VMware on development and testing of the new vSphere NVMe/TCP feature since 2021. This joint effort has enabled Lightbits to optimize its solution for vSphere integration. Lightbits was also the first NVMe/TCP storage solution certified for on-prem vSphere 7.0U3.

In addition to being certified as an NVMe/TCP external storage solution since 2021, Lightbits recently worked in coordination with the VMware certification team to certify our solution specifically for AVS. We completed and passed this certification using the VIVa test suite, which included:

- Connectivity tests to validate NVMe/TCP integration between Lightbits and AVS.
- Storage provisioning and VMFS datastore operations.
- Multipathing and high availability.
- Resiliency testing via simulated failover and fault injection.



We also performed rigorous internal validations, spanning:

- Control plane functionality and management at scale.
- Performance benchmarking across real-world workloads and access patterns.
- Cluster scale testing under contention.

Completing this intensive VMware certification process underscores Lightbits' ability to deliver the enterprise-grade performance, availability, and scale required for production workloads on AVS.

# Storage Options for Azure VMware Solution

In AVS, you have two storage choices with general availability: VMware vSAN and Azure NetApp Files (ANF). These options are suitable for a wide range of workloads and use-cases:

### vSAN

Provides cluster-wide storage built on top of the direct attached block devices available in each AVS host. The total capacity depends on the number of hosts and the host type used by your AVS cluster. A small portion of the capacity is utilized as a cache tier while the majority is utilized as a capacity tier. Table (1) shows the raw capacity you can get for each AVS host type. You can scale capacity by adding more AVS hosts to the cluster. By default, vSAN protects you from up to one failure via mirroring (RAID-1 with 2 replicas).

### **Azure NetApp Files**

ANF can be used to attach Network File System (NFS) datastores as persistent storage for AVS connected over ExpressRoute. ANF supports three service levels, ranging from 16MiB/s up to 128MiB/s per TiB. According to benchmarks published on the Azure website, traffic latency from AVS to Azure NetApp Files datastores varies from a sub-millisecond (for environments under minimal load), to up to 2-3 milliseconds (for environments under medium to heavy load). The total cost depends on the service level and the total provisioned capacity (minimum 4TiB with increments of 1TiB).

Lightbits enhances these existing storage options in AVS by introducing a new choice specifically designed for demanding workloads that necessitate resilient, high-performance external block storage - such as SQL and NoSQL databases. Data is protected with three replicas and is not affected by any failures on the AVS side. You can deallocate AVS clusters without losing your data and mount your datastores later in a new or existing AVS cluster.

# High Performance NVMe/TCP VMFS Datastores for AVS

We ran Flexible I/O Tester (or FIO) to measure the overall storage performance of AVS with Lightbits and SQLIOSim. We did so in order to measure the storage performance of Microsoft SQL Server using Lightbits datastores.

### **Storage Performance - FIO**

FIO is an open-source benchmarking and stress testing tool used to evaluate the performance of storage devices and systems. It allows you to simulate various I/O workloads, measure disk performance, and assess



how different storage configurations handle different types of I/O operations. We used FIO to stress the storage and simulate different loads and read/write patterns to the datastores.

#### **Benchmark Methodology**

We ran performance benchmarks on the Azure VMware solution using Lightbits as the storage option. We deployed an AVS cluster of 16 x AV36P hosts and a Lightbits storage cluster of 16 x L64asv3 VMs. The Lightbits cluster was connected to the AVS SDDC via an express route ultra gateway with fast path enabled. We used a total of 16 VMFS datastores. Each ESXi server was hosting a virtual machine running FIO.

#### **Performance Results**

Figure (4) shows the aggregated throughput of the AVS cluster while using Lightbits for different block sizes, and while running a mix of 70% reads and 30% writes (all random). The benchmarks were done while the Lightbits cluster was maintaining three replicas to keep the datastores protected from up to two failures on the L64asv3 VMs. Lightbits was able to achieve 56GB/s for block sizes of 32KB and higher. Smaller block sizes also showed impressive results: 13GB/s for 4KB, 25GB/s for 8KB, and 45GB/s for 16KB.



Figure -4- Bandwidth for 70% Read/30 % Write for Different Block Sizes

Figure (5) shows the aggregated IOPS for increasing loads and the read and write latency measured at each load while running a mix of 50% reads and 50% writes, 8KB random. Lightbits was able to achieve more than one million IOPS while keeping sub-millisecond latency for both read and writes. Also, in this case the benchmarks were run while the Lightbits cluster was maintaining three replicas to keep the datastores protected from up to two failures on the L64asv3 VMs.





Figure -5- IOPS and Latency for 50% Read/50% Write for 8KB Block Size

### **Application Performance - SQL Server with SQLIOSim**

SQLIOSim (SQL I/O Simulator) is a tool developed by Microsoft to simulate and test the I/O (Input/ Output) subsystem performance of a SQL Server database. It helps database administrators and developers evaluate the performance of their storage systems by generating a workload that mimics the I/O patterns of SQL Server. SQLIOSim allows you to perform stress testing and benchmarking to ensure that your storage infrastructure can effectively handle the demands of SQL Server workloads.

#### **Benchmark Methodology**

We ran SQLIOSim using a small cluster of 3 x AV36 hosts and a small Lightbits storage cluster of 3 x L64asv3 VMs, using three datastores. The Lightbits cluster was connected to the AVS SDDC via an express route ultra gateway with fast path enabled. Each ESXi server was hosting two virtual machines, with a total of six in the AVS cluster. Each VM was configured with 24 vCPUS and three virtual disks, with each virtual disk in a different datastore.

We ran the SQLIOSim configured for 1ms target I/O duration and 10ms target I/O duration. We did so using the default profiles (Random, Audit, LogAudit, ReadAhead, BulkUpdate, Shrink), and aggregated the results. In SQLIOSim, the "target I/O Duration" is a setting that specifies the desired time for completing a single I/O operation. This setting helps control the speed and intensity of the I/O workload generated by SQLIOSim. By adjusting the Target I/O Duration, you can simulate different levels of I/O activity and measure how well your storage system can handle SQL Server workloads with varying I/O response times.

#### **Performance Results**

In Figure (6), we can see the running average I/O durations measured by SQLIOSim - using Lightbits datastores and the default storage option available in AVS, our reference. The "running average I/O duration" in SQLIOSim is a metric that represents the average time it takes to complete I/O operations over a specific period of time during the execution of SQLIOSim tests. It provides insights into the I/O responsiveness and efficiency of the underlying storage system. The lower the metric, the better the performance.



We observed that Lightbits achieved a running average I/O duration 3.8 times lower for a 10ms target I/O duration and 2.7 times lower for a 1ms target I/O duration for the data files. In Figure (7), we can also see that Lightbits datastores experienced significantly fewer IO throttles for the data files, ranging from 3.7 to 4.2 times less compared to other configurations.



#### SQLIOSIim - Running Average I/O Duration





### SQLIOSIim - Number of Times I/O was Throttled

Figure -7- Data files: Number of Times IO Throttled for 1ms and 10ms Target IO Duration (Lower is Better)



# Support Models

### The Managed Application Shared Responsibility Model

The Lightbits managed application on Azure has two support models; customer-managed and fully managed. In a fully managed Azure application, Lightbits delivers high-performance storage software-as-a-service, while taking complete ownership of deployment, management, updates, and support. Customers simply deploy Lightbits from the Azure Marketplace without having to maintain the software.

Lightbits handles end-to-end infrastructure monitoring, scaling, security patching, troubleshooting, and upgrades for the application. We provide 24/7 support and guaranteed SLAs on availability. Customers retain control and security of their data, subscription, identity management, and compliance adherence - while benefiting from our operational excellence. The customer's only responsibility is to pay Azure infrastructure usage charges based on metered consumption.

In the customer-managed model, clients have full control of the application resources and infrastructure, while Lightbits provides reactive break-fix support upon request. Lightbits and customers together perform monitoring, maintenance, and updates.

The two models provide flexibility for customers to choose the level of operational control they need, along with Lightbits managed services based on the customer's IT preferences.



Figure -8- Lightbits Managed Application Shared Responsibility Model



### **Additional Support Information**

To find out more about how our support models are implemented, or to view Service Level Agreements (SLAs) and contact Lightbits Support, visit the <u>Lightbits Support Documentation</u>.

### **Get Started**

To learn more about how Lightbits for AVS can set your cloud migrations up for success, <u>request a demo</u> today. The Lightbits demo features a review of your specific cloud challenges, goals, and use cases - as well as a demonstration of Lightbits in action, with a discussion on deployment options and pricing as well.

# About Lightbits Labs

Lightbits Labs (Lightbits) is leading the digital data center transformation by making high-performance elastic block storage available to any cloud. Creators of the NVMe® over TCP (NVMe/TCP) protocol, Lightbits software-defined storage is easy to deploy at scale and delivers performance equivalent to local flash to accelerate cloud-native applications in bare metal, virtual, or containerized environments. Backed by leading enterprise investors including Cisco Investments, Dell Technologies Capital, Intel Capital, JP Morgan Chase, Lenovo, and Micron, Lightbits is on a mission to make high-performance elastic block storage simple, scalable, and cost-efficient for any cloud.

Lightbits and Lightbits Labs are registered trademarks of Lightbits Labs, Ltd.

All trademarks and copyrights are the property of their respective owners.

E: info@lightbitslabs.com www.lightbitslabs.com

### **US Offices**

1830 The Alameda, San Jose, CA 95126 | 1412 Broadway 21st Floor, New York, NY 10018

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 Itd 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 noninfringement. 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© 2023 LIGHTBITS LABS LTD. - ALL RIGHTS RESERVED

### Israel Office

17 Atir Yeda Street, Kfar Saba 4464313

13