# Lightbits on Microsoft Azure

## A scalable, high-performance cloud data platform for storage-intensive workloads in the cloud

September 2023
Rev 1.0

# Table of Contents

# Introduction

It's no secret that storage-intensive workloads have been among the last to migrate to the public cloud. While cloud storage is convenient to deploy and easy to manage, many enterprises have found that high-performance block storage options are unpredictably expensive, and often do not maintain the required low latency as application demand increases.

At the same time, cloud storage options often lack the simplified provisioning and enterprise data services of their on-premises SAN. For these reasons, many business-critical databases, analytics, and transactional workloads are stuck firmly as on-premises, despite corporate cloud policies.

## Enterprise SAN in the Cloud

To meet the needs of demanding enterprise workloads, Lightbits creates a high-performance enterprise SAN in the cloud using clusters of Azure Lv3- series storage optimized virtual machines and the power of NVMe® over TCP (NVMe/TCP). With Lightbits, you can provision persistent volumes that deliver low latency and high IOPS and scale seamlessly, so you can migrate VMware workloads and SQL and NoSQL databases to Azure knowing that they'll run successfully at scale without incurring excessive cloud charges.

Lightbits enables you to provision, manage, and pay for storage resources as a centralized pool instead of volume by volume. For large-scale migrations, you won't need to spend weeks sizing each volume individually, and you won't have to over-provision resources to meet the maximum expected throughput or capacity.

The Lightbits cluster can be dynamically scaled up or down based on your organization's storage needs. Furthermore, Lightbits delivers enterprise data services — such as replication, snapshots, and data encryption — are available at no additional charge to maintain data integrity and security in the cloud.

## An Azure Managed Application

Lightbits is available through the Azure Marketplace as an Azure Managed Application, making it as easy to deploy and manage as one of Azure's native storage offerings. As a fully managed service, Lightbits controls operations within your Azure subscription. You choose the type of pre-configured virtual machine (VM) and the number of nodes in your cluster, and Lightbits manages all storage cluster operations with self-healing, auto-scaling, replication, rolling upgrades, and more. Critical functions are performed behind the scenes to ensure availability, performance, and data protection as well.

# Azure Block Storage Options

This whitepaper centers on the highest-performing block storage alternatives suitable for storage-intensive applications like SQL and NoSQL databases. Additional block storage device options also exist that offer lower performance; however, we will not delve into these less-performant options in this context.

## Managed Disks

1. **Premium SSD v2**: provides high-performance, low-latency disk support for storage-intensive workloads, but is limited to a maximum disk size of 64 TiB and a maximum IOPS of 80,000. We see Application teams often "stack" these disks together and use tools such as LVM to stripe across the devices, giving more performance to a single Azure Virtual Machine (VM), attempting to save costs compared to Ultra Disk. Premium SSDv2 disks cannot be attached to zonal VMs, or VMs managed by availability sets — and they also don't support backups.
2. **Ultra Disk**: offers maximum performance and low latency for data-intensive applications but is limited to a maximum disk size of 64 TiB and a maximum IOPS of 160,000. The size of an Ultra Disk cannot be expanded without either deallocating the VM or detaching the disk. Snapshots for Ultra Disks are available as well but have additional limitations and the disks cannot be backed up.

## Other Options

1. **Elastic SAN:** currently in preview, Elastic SAN offers a SAN experience in Azure that provides up to 2 million IOPS (region-dependent), across the cluster — with single digit millisecond latency. Each volume can provide up to 64,000 IOPS and uses iSCSI across the standard VM network interface, which overcomes some of the performance limitations of managed disks when aggregating volumes. Elastic SAN only supports a subset of the features of managed disks and does not support snapshots. The performance of Elastic SAN is directly proportional to the size of the cluster, which means that you must build a cluster of substantial capacity to meet performance requirements in the hundreds of thousands of IOPS range. Elastic SAN is tailored for applications capable of accommodating latency in the range of single-digit milliseconds.
2. **Azure VMs with Local NVMe:** Provide a Direct Attached Storage (DAS) experience for workloads with high performance requirements in the millions of IOPS that replicate at the application level or can afford to potentially lose data. Local NVMe devices are non-persistent, which means that data can be lost during maintenance events, power-downs, and VM size changes. Azure VMs that support local NVMe are provided with a set amount of storage, which can result in overprovisioning for capacity and performance. Only a subset of Azure VMs supports local NVMe devices, and only the Ls and Las families can be used. This could result in unoptimized storage configurations. Local NVMe devices in Azure do not support backups, snapshots, or replication.

For additional insights on cloud block storage offerings, consult the Azure documentation.

# The Lightbits Cloud Data Platform

Lightbits eliminates the limitations found in Azure native cloud storage offerings. Compared to purchasing performance disks, Lightbits is a fraction of the cost on a per-gigabyte basis, and scaling performance doesn't require purchasing additional, underutilized capacity; You can scale compute and storage independently. Unlike equipping your VMs with local NVMe devices — and unlike Elastic SAN — Lightbits provides storage-level redundancy for data persistence, snapshots, and other essential enterprise data services.

Lightbits can be deployed as an Azure Managed Application. The solution is implemented within a managed resource group, enabling Lightbits to support, operate, and assist in managing the solution, enhanced by built-in automation features.

Azure Managed Applications are fully managed IT solutions deployed in customer environments. They offer hassle-free maintenance, updates, and support, while allowing users to focus on strategic tasks rather than operational overhead — thus increasing return on investment (ROI). They allow vendors such as Lightbits to provide customized solutions to organizations leveraging the Azure cloud while feeling like a first-party service.

Lightbits on Azure is installed on Azure storage-optimized VMs inside your subscription, so you can view all the resources and get the benefit of using any Azure discounts that you already have. The Lightbits managed application leverages Lsv3 or Lasv3 instances, providing the extreme high performance of local NVMe with the benefits of a SAN — including replication, thin provisioning, data reduction, backups, and scalability.

## Why Lightbits on Azure?

Implementing the Lightbits high-performance, low-latency, software-defined storage software in the Azure cloud addresses several challenges:

1. **Performance Bottlenecks:** Native storage often becomes a bottleneck in data-intensive applications in the public cloud. A high-performance, low-latency solution can provide faster data access and processing, improving application performance.
2. **Scalability:** As data grows, it can be challenging to scale cloud native storage infrastructure. Lightbits can scale out easily to accommodate increasing capacity and performance, providing flexibility and preventing the need for over-provisioning.
3. **Cost Efficiency and ROI:** High-performance Azure storage solutions can be expensive to provision, especially for performance-intensive workloads. Lightbits on Azure can be more cost-effective, as it leverages thin provisioning, thin clones, and thin snapshots — reducing capacity while maintaining high performance and low latency. The managed application deployment enables a more “hands-off” maintenance approach, which means minimal management overhead and a greater ROI.
4. **Data Availability and Durability:** With replication across the cluster, Lightbits volumes provide highly available storage.
5. **Data Management:** Managing data across multiple storage types and segregated environments can be complex. Lightbits simplifies data management by providing a single platform for all storage resources, thus reducing storage sprawl and overprovisioning while improving total cost.
6. **Simple Architecture:** Lightbits uses a scalable single-tier architecture, which means that all instances in the cluster support management and data operations. The cluster can be easily scaled out by adding one instance at a time. This simplifies the architecture and reduces the total number of instances required, thus lowering costs and operational overhead compared to multi-tier architectures.

# Lightbits as an Azure Managed Application

## Architecture

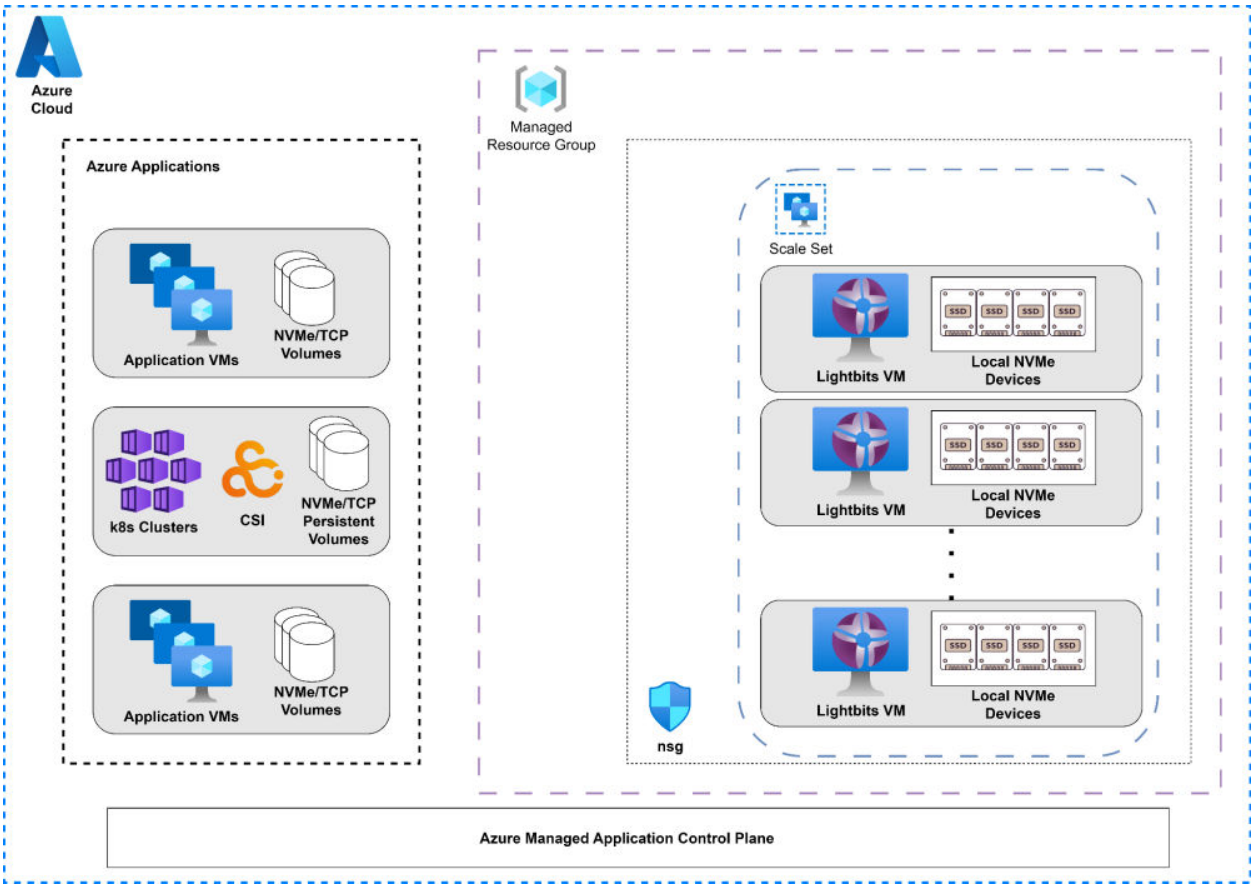


Lightbits Managed Application in Azure

## Features

The Lightbits Managed Application has a vast feature set to help reduce operational overhead for organizations, as well as provide essential enterprise features that cloud native storage may be lacking. Some of these features are outlined below.

### Deployment Models

- Bring Your Own License (BYOL)
- Azure Marketplace Managed Application
- Runs in your cloud environment
- Autonomous operation with auto scaling, healing, and maintenance
- Self-managed with Lightbits support
- Automated by Lightbits cluster functions

### Data Services

- Logical volumes with online resize
- Thin-provisioning
- Inline compression
- Redirect-on-write snapshots and clones
- High performance, consistent low latency
- Quality of Service (QoS)
- Encryption (PMK)

### Cloud Integration

- Provide storage to any Azure instance, including k8s
- Runs on Azure Lsv3 and Lasv3 VMs
- Hashicorp
- Terraform
- Azure CLI
- Azure Console

### Cluster Size and Configuration

- 3-16 storage servers per cluster
- Up to 24K volumes (64K coming soon)
- Up to 64K clients per cluster
- Online automatic node add/remove
- Per-volume ACLs and IPACLs
- Dynamic data rebalancing

### High Availability and Data Protection

- Per-volume replication policies
- User-defined failure domains
- DELTA log recovery (partial rebuild)
- Data integrity
- Automatic cluster node failure recovery
- Data replication within a single availability zone for highest performance
- Data replication across multiple availability zones for highest durability
- Proactive auto scaling
- Resiliency through auto healing
- Predictive auto maintenance
- Backups to Azure Blob

### Management and Monitoring

- gRPC, REST APIs and CLI interfaces
- Multitenancy with RBAC
- Ecosystem:
  - Kubernetes (CSI)
  - Linux OS
  - VMware
- Monitoring stack: Prometheus and Grafana
- Cluster, node, capacity, performance metrics
- Rolling cluster upgrades

### Client Connectivity

Clients connect to the Lightbits cluster through NVMe/TCP via the IP address of the Azure VM NIC — utilizing the VNet facing network interfaces. NVMe/TCP has been built into the Linux kernel natively since version 5 and backported to version 4.15, so almost all modern Linux distributions will connect to Lightbits out of the box with no proprietary software on any clients.

NVMe/TCP multipathing provides automatic, non-disruptive path failover if a Lightbits target node loses connectivity.

If you are deploying your application using K8s, each worker node contains the NVMe/TCP kernel modules. The Lightbits CSI plugin provides orchestration for dynamic provisioning of persistent volumes (PVs) — as well as snapshots and clones.

NVMe/TCP operates on ports 4420 for data and 8009 for discovery. Thus, each client must establish connectivity across these ports to every target node within the Lightbits cluster to ensure proper functionality. Note that you should permit these ports within your Azure Network Security Groups (NSGs) to facilitate uninterrupted communication.

# Performance

Lightbits on Azure can provide up to a million IOPS per volume and tens of millions of IOPS per cluster while maintaining a sub-millisecond average and even tail latency at the 99th percentile. Overall performance is a factor of the VM series and size — as well as the number of VMs used for the Lightbits cluster.

## Performance Benchmarks

Below is an example of how Lightbits can provide consistently low tail latency as the cluster scales to provide higher overall IOPS.

The test involved deploying a Lightbits three-node cluster in Azure using L64sv3 VMs as the target nodes and connecting to three Azure L64sv3 client VMs. Each client VM was given its own volume on each Lightbits target node, and FIO was run with increasing queue depth to push more load in each test configuration.

The test was run with a three-node Lightbits cluster, and then three nodes were added to create a six-node cluster. Three nodes were added again to create a nine-node cluster.

This is an example of the performance that Lightbits can provide. Because Lightbits can scale up to 16 nodes per cluster, there is plenty of room to increase cluster IOPS even further if necessary.

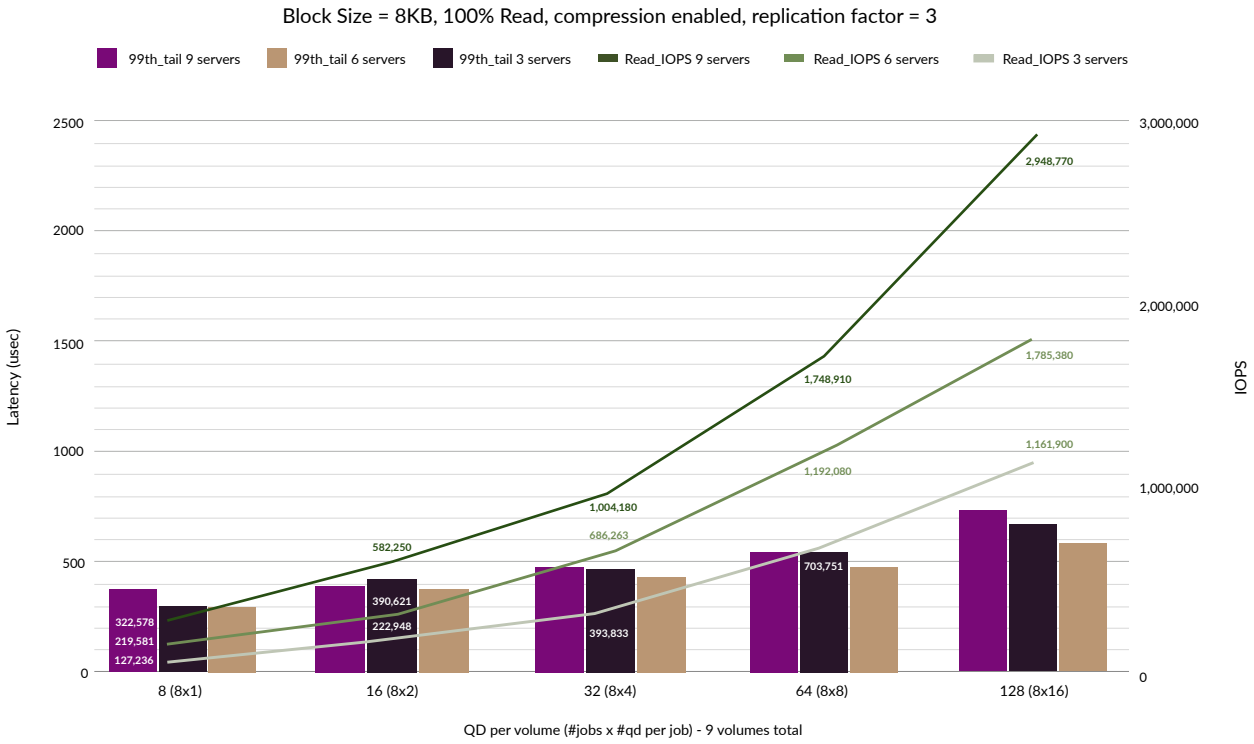


Read IOPS and 99th Tail Latency of Lightbits as Load Is Increased

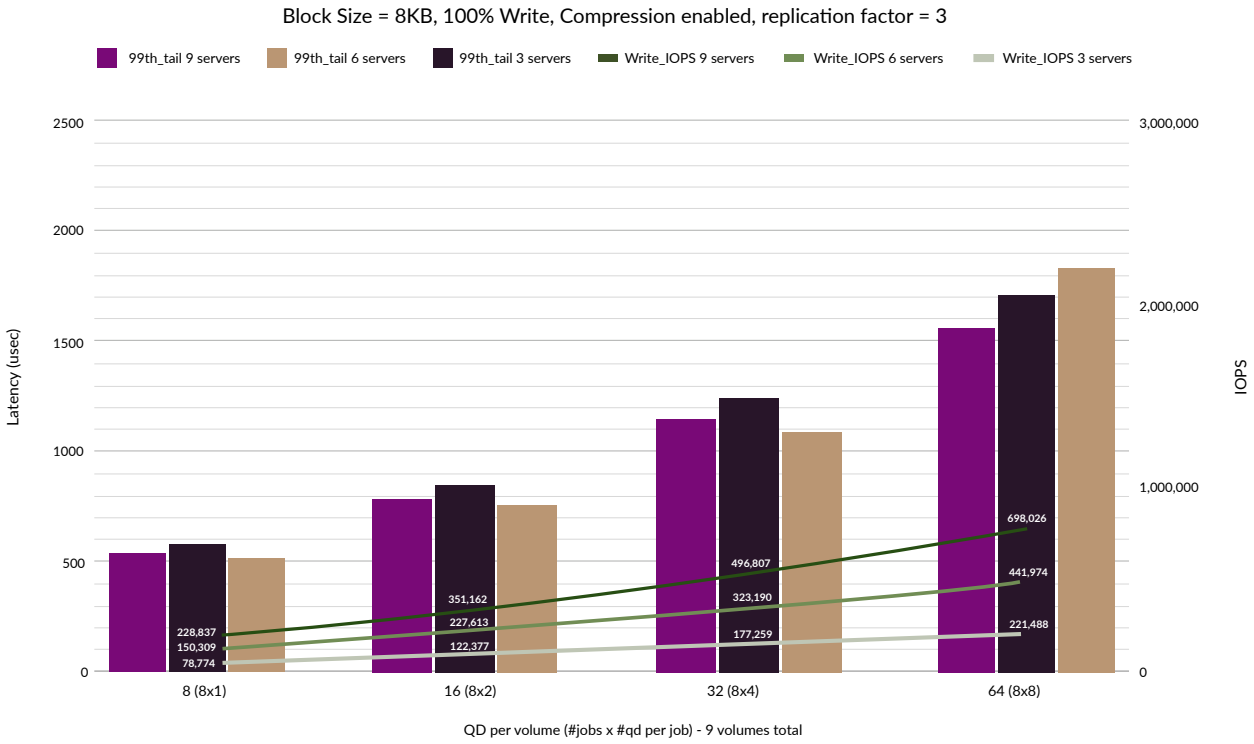


Write IOPS and 99th Tail Latency of Lightbits as Load Is Increased

## Performance Considerations

Consider these factors to achieve the optimal Lightbits performance for your applications:

- **Client:** The Azure client VMs consuming Lightbits should have sufficient network bandwidth to reach the desired performance level. Azure provides the maximum network bandwidth (Mbps) value per VM size and type in their documentation. This will be the upper limit of write performance that the VM can achieve to the Lightbits storage. Read performance is unaffected by this value, but it will be affected by the Lightbits target node type. NVMe/TCP uses multiple TCP connections to parallelize storage processing when running on multicore VMs equipped with multi-queue NICs. This is how the protocol gains much of its performance improvement over iSCSI.
- **Lightbits Targets:** The performance of a single volume will be set by the choice of target VMs used for the Lightbits cluster and read/write mix. With Lightbits on Azure, storage clusters can be scaled out to meet workload demands, which means that performance-sensitive workloads can scale based on increased requirements.
- **Block Size:** Block size is a significant factor in performance. Lightbits has proven that we can saturate the Mbps available with a block size of 8KB given the appropriate Mbps and CPU cores on the client. This means that any workload with <8KB block size will likely not hit the maximum potential throughput of a single Lightbits volume. However, this value has been proven to be >1 million IOPS when using a Lightbits target VM size of at least L64sv3.
- **Quality of Service (QoS):** Lightbits is a multi-tenancy cloud storage platform that can support thousands of applications and clients. As a result of this, QoS policies can be assigned per volume to ensure that all application performance needs can be met without the risk of introducing “noisy neighbors,” or applications that will consume more than their allowance of read or write throughput from the pool of performance.

## Client Optimization

Compared to Ultra Disk and other Azure native cloud block storage devices, Lightbits allows for applications to use optimized client VMs, which require less managed disk IOPS in order to reach application performance requirements. This is because Azure provides a greater throughput for VNet-facing NICs than for the managed disk network. This can greatly reduce TCO for applications that require a high level of performance and therefore previously had to use large VMs with a high level of managed disk throughput.

Additionally, it’s important to note that only specific VMs can access the provisioned IOPS disks such as ultra, which further reduces the application owner’s choice of client instances. With Lightbits, this is not the case, since every VM type can access NVMe/TCP volumes over the VNet.

## Sizing

A Lightbits cluster size is based on the local NVMe capacity of the VMs that were chosen as targets and the number of VMs in the cluster. The cluster capacity can be extended by adding more nodes up to the maximum cluster size of 16. For more information, consider the example below or contact a Lightbits representative for bespoke sizing for any environment.

### Sizing Example: L64sv3

This example shows a cluster of L64sv3 Azure VMs. Note that this does not take into account compression and considers a replication factor of three (three copies of each volume spread across the cluster).

| Number of VMs | Max 8KB Read IOPS | Max 8KB Write IOPS Replication factor=3 | Max read throughput - 128KB (GB/s) | Max write throughput - 128KB (GB/s) | Effective capacity (TB) |
|---|---|---|---|---|---|
| 3 | 1,161,900 | 221,488 | 10.00 | 3.23 | 12.288 |
| 4 | 1,549,200 | 295,317 | 13.33 | 4.31 | 16.384 |
| 5 | 1,936,500 | 369,146 | 16.67 | 5.39 | 20.480 |
| 6 | 2,323,800 | 442,976 | 20.00 | 6.47 | 24.576 |
| 7 | 2,711,100 | 516,805 | 23.33 | 7.54 | 28.672 |
| 8 | 3,098,400 | 590,634 | 26.67 | 8.62 | 32.768 |
| 9 | 3,485,700 | 664,464 | 30.00 | 9.70 | 36.864 |
| 10 | 3,873,000 | 738,293 | 33.33 | 10.78 | 40.960 |
| 11 | 4,260,300 | 812,122 | 36.67 | 11.86 | 45.056 |
| 12 | 4,647,600 | 885,952 | 40.00 | 12.93 | 49.152 |
| 13 | 5,034,900 | 959,781 | 43.33 | 14.01 | 53.248 |
| 14 | 5,422,200 | 1,033,610 | 46.67 | 15.09 | 57.344 |
| 15 | 5,809,500 | 1,107,440 | 50.00 | 16.17 | 61.440 |
| 16 | 6,196,800 | 1,181,269 | 53.33 | 17.24 | 65.536 |

Table showing the performance and capacity as a Lightbits cluster of L64sv3 instances scales.

# Cost

## Purchasing Models

Lightbits on Azure can be purchased in multiple ways, depending on how your organization prefers to pay for and consume resources. These models can be categorized as “OpEx” and “CapEx” models.

It’s worth noting that for both payment models, this only includes the cost of the Lightbits license and not the cost of the Azure infrastructure utilized by the cluster (VMs, storage accounts, functions, etc.). Those will be charged in your standard methods just like any other Azure resource. This also means that you will be eligible for any current Azure discounts on that infrastructure based on your current Azure agreements.

### PayGo (OpEx)

The PayGo (Pay-as-you-Go) model enables customers to pay for Lightbits per hour based on utilized capacity. The model is available through the Azure Marketplace and will be billed by Azure via the marketplace. The model is best for organizations that wish to pay for the product through their Azure subscription, and for customers where consumption may be elastic and change often over the period of a year. It’s also great for customers that would like to test the product temporarily. For customers with long-term commitments, there are methods to provide built-in discounts for larger capacities.

### BYOL (CapEx)

This BYOL (Bring-Your-Own-License) model allows for customers to transact their licensing directly with Lightbits, allowing for 1-, 2-, or 3-year pricing models based on the VM sizes. This option is great for customers that are migrating workloads into Azure, or that already understand their storage capacity and performance requirements. It allows for further bulk discounts for sustained, long-term cluster usage.

# Use Case

Due to the high performance and low latency of Lightbits, databases such as Oracle are a perfect use case to run on the platform in Azure. Below is an example of how Lightbits can help migrate large numbers of Oracle databases while optimizing the cost of running those applications in Azure.

## Performance

First, measuring the performance of Ultra Disk compared to Lightbits, both running Oracle and measured by SLOB:

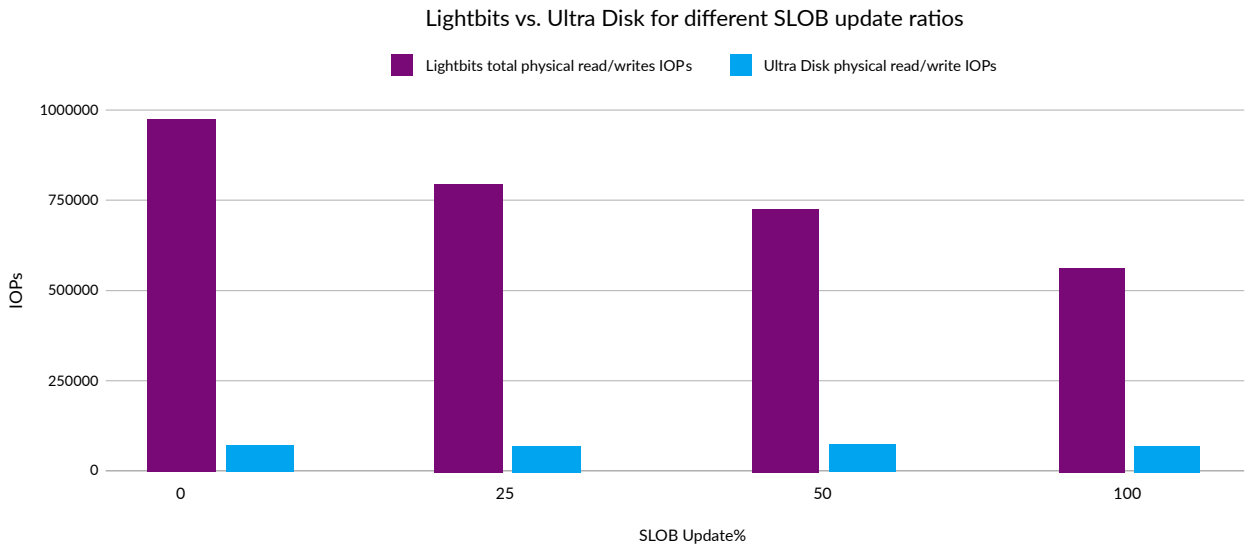


Lightbits Performance Running SLOB for Oracle vs. Ultra Disk

As shown in the chart above, Lightbits running Oracle databases can far outperform the maximum configuration of Azure Ultra Disks. This solution allows for high performance, dense databases to run more efficiently on the Lightbits platform. It also enables previously immovable databases to migrate onto the Azure cloud.

## Cost Reduction

From the performance information above, the cost vs. Ultra Disk for running Oracle on Azure can now be calculated. The calculation is based on a setup that can reach 900K physical reads from Oracle.

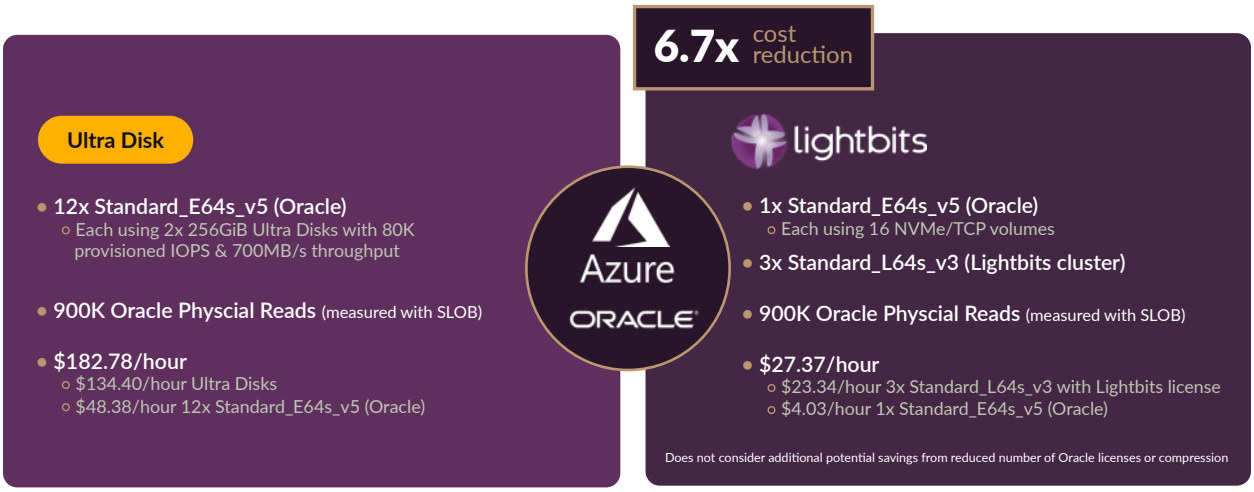


As shown above, due to the extreme performance of the Lightbits data platform on Azure, workloads can be better optimized to save costs — up to 6.7x lower costs than Azure Ultra Disks. This is not just applicable to Oracle, but for optimizing most applications that require high performance and consistent low latency. For more information about how to run Lightbits on Azure VMs, please consult the Azure documentation.

# Support

## 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 Lightbits' operational excellence. The only responsibility of the customer 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 perform their own monitoring, maintenance, and updates. The two models give customers the flexibility to choose the level of operational control versus Lightbits managed services based on their IT preferences.

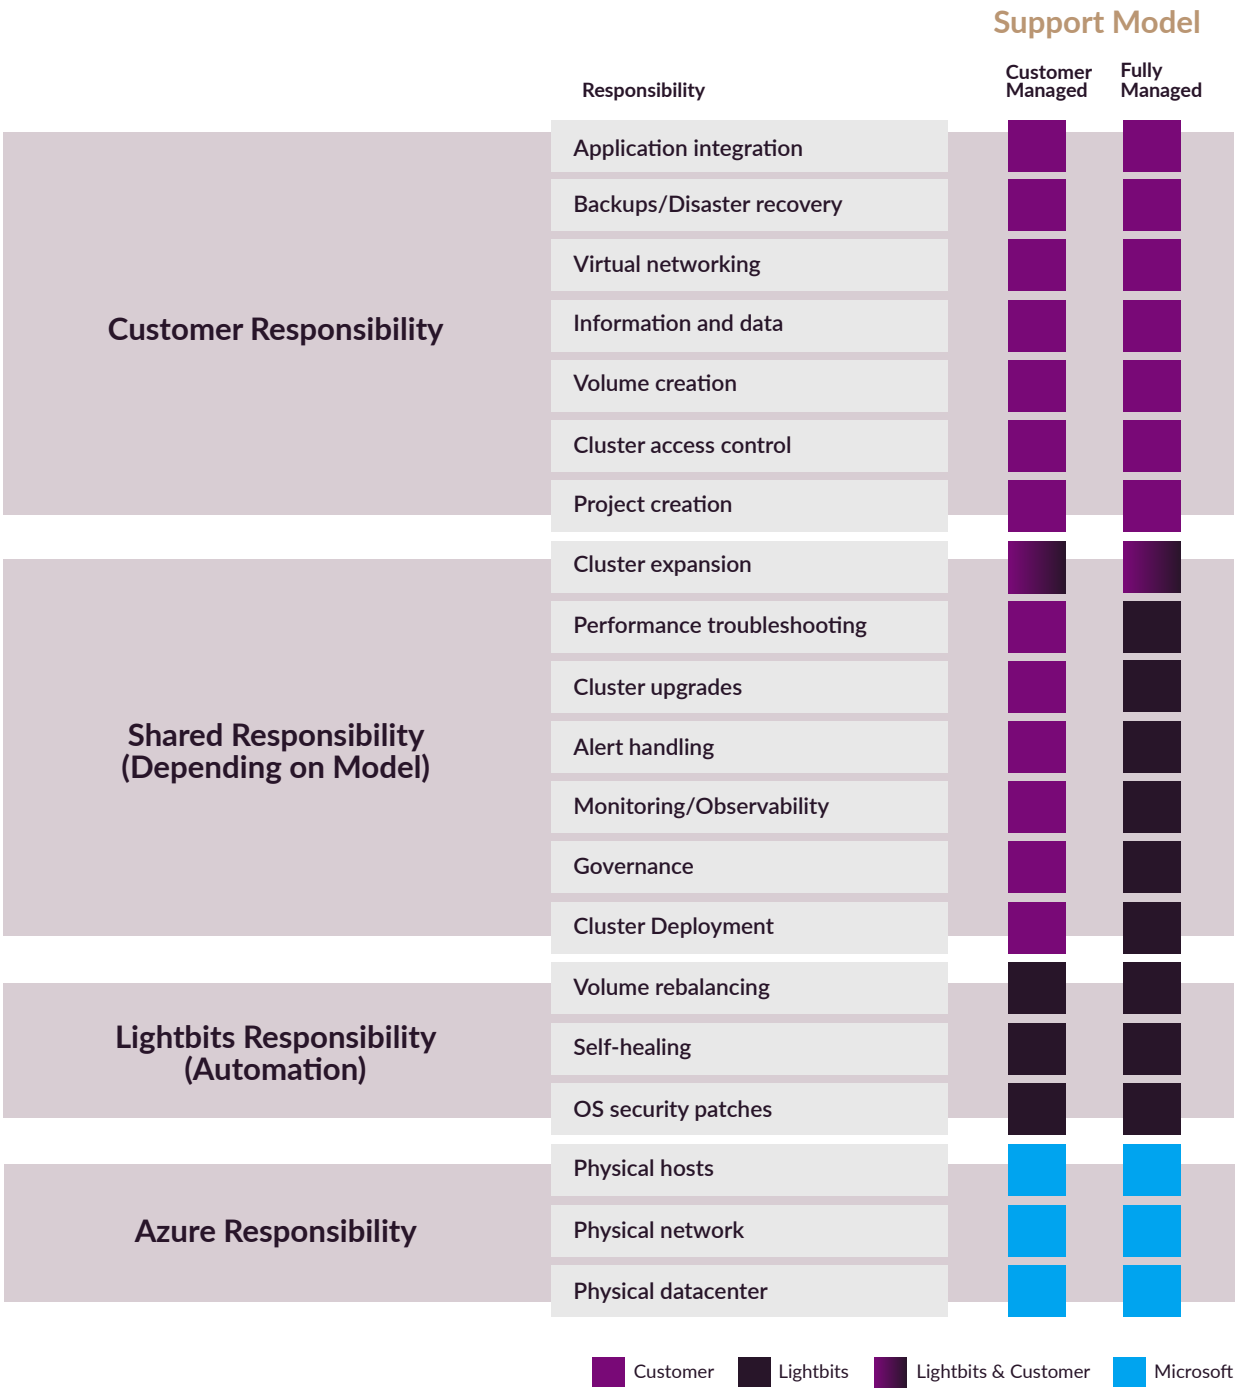


The Lightbits managed application shared responsibility model.

## More Support Information

To find out more about support models implementation, service level agreements (SLAs), and support, please consult the Lightbits support documentation.

# Get Started

To explore how Lightbits on Azure can set your cloud migrations up for success, request a demo today. The demo features a review of your specific cloud challenges, goals, and use cases, and a demonstration of Lightbits in action, with a discussion on deployment options and pricing.

## About Lightbits Labs™

Lightbits Labs® (Lightbits), offers a complete data platform that enables organizations with a cloud-first strategy to migrate storage-intensive workloads to the cloud or to build a high-performance cloud. It eliminates the performance bottlenecks commonly associated with legacy, traditional, or cloud-native storage and delivers high performance, consistent low latency, and predictable low costs. The native NMVe/TCP and clustered architecture, coupled with essential data services, solve the common cloud storage challenges of performance, efficiency, and cost. Lightbits is backed by enterprise technology leaders [Cisco Investments, Dell Technologies Capital, Intel Capital, Lenovo, and Micron] and is on a mission to deliver a robust cloud storage platform with unmatched simplicity, efficiency, agility, and flexibility.


www.lightbitslabs.com

info@lightbitslabs.com

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

1412 Broadway, 21st Floor
New York, NY 10018, USA

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