SQL Server with Lightbits on Azure VMware Solution (AVS)

A disaggregated, high-performance storage option for running storage-intensive SQL Server workloads on Azure VMware Solution.

May 2024 | Rev 1.0

Authors:

Dèjì Akọmọláfẹ

Staff Solutions Architect - Microsoft Applications Lead

VMware by Broadcom

Felix Melligan

Principal Solutions Engineer

Lightbits

Table of Contents

Introduction

Storage capacity and performance are common considerations when virtualizing enterprise-class Microsoft SQL Server (SQL Server) workloads in Azure VMware Solution (AVS) infrastructure. This document presents a technical discussion of how these storage capacity and performance challenges can be efficiently solved through the cost-beneficial external storage option presented by the Lightbits software-defined storage solution in Azure.

The document is targeted at AVS and SQL Server Administrators looking for technical guidance on how to easily, reliably, and optimally augment storage capacity needs for their SQL Server instances in their AVS environments - while still providing the required level of performance. We provide a list of configuration options, architectural guidelines, and operational considerations - as well as observations of the solution's performance based on the standard practices documented in VMware's published best practices for virtualizing SQL Server workloads on VMware vSphere-based platforms such as AVS.

The contents of this document do not include comparative performance analysis nor are the test results included intended to be considered or interpreted as implying a benchmarking exercise. Our objective and observations are intentionally limited in scope – to evaluate the feasibility of using Lightbits disaggregated software-defined block storage to optimally satisfy storage capacity requirements in an AVS infrastructure, without the cost implications of additional ESXi hosts and without introducing any performance bottlenecks.

We hope that Administrators will be able to extract enough information from this document to help them, in turn, perform the required due diligence in their evaluation of the presented solution, and come to the same conclusion that we reached at the end of our exercise – with Lightbits for AVS.

Customers can now easily and reliably expand AVS' storage capacity for their SQL Server workloads incrementally, while adding both capacity and performance with a consistently low latency, without sacrificing workload performance in the process.

Azure VMware Solution

AVS is one of the most well-known VMware vSphere-based hyperscaler solutions. It integrates multiple core software-defined, engineered components from the vSphere-based solution stack into a cost-beneficial, highly-available, and performant public cloud service offering. AVS extends the boundaries of customers; investments in vSphere-based solutions into the popular Microsoft Azure public cloud infrastructure. This enables customers to simultaneously benefit from operational efficiencies, flexibility, and interoperability - without complex application transformation or onerous re-tooling that would otherwise be required in a typical private-to-public cloud exercise.

AVS - like every other vSphere-based cloud infrastructure - primarily uses VMware vSAN as the storage option for all virtualized guests in the infrastructure. At a high level, vSAN is a storage aggregation solution that uses the local storage capacity available on individual ESXi hosts to create a flexible, scalable, and highly-available storage platform consumable by all VMs in the environment.

Each ESXi host in the environment contributes storage capacity according to its capacity. If additional storage capacity is required, additional hosts will be required to contribute their set of storage to expand vSAN capacity.

Storage capacity is a common challenge in an AVS infrastructure, particularly for workloads such as Microsoft SQL Server, where storage requirements for application data often exceed what can be contributed by the local storage of ESXi hosts. While adding additional ESXi hosts to an AVS infrastructure is a relatively simple operation, the cost

© 2024 Lightbits Labs™

implications of such additions might be undesirable, especially when the additional compute resources (CPU, RAM) contributed by the additional hosts are not particularly required for the virtualized workloads.

We are, therefore, presented with a dilemma: How do we increase storage capacity in an AVS environment at a desirable cost? And how do we do so to conform to the reduced TCO proposition that makes AVS so attractive to enterprises in the first place?

The answer that Azure has provided is with disaggregated storage solutions that can present datastores to the ESXi host to expand the storage footprint. The Lightbits Azure Managed Application is a part of this unique set of options that enables users to expand their AVS storage without sacrificing performance, by using Virtual Machines hosted within their own Azure Subscription. This allows cloud administrators to provide higher capacities incrementally, while maintaining flexibility to provide workloads to both AVS and other Azure environments from a single Lightbits storage cluster.

Microsoft SQL Server

Thirty-five years after its initial release, Microsoft SQL Server remains one of the top 3 Relational Database Management Systems (RDBMS) for enterprises globally. SQL Server's footprints can be seen in almost every Windows-based application, ranging from order-processing systems all the way through billion-dollar CRM/ERP solutions. With the release of SQL Server Linux in 2016, Microsoft has further increased the reach and utility of SQL Server beyond its traditional target market and audience. SQL Server on Linux is Microsoft's database platform for containers.

Traditionally, after initial resistance against the feasibility, viability, and reliability of x86 hypervisors as a proposition for enterprise-class mission-critical applications, virtualization has since displaced physical instances as the default option for running and operating these important applications. It is inarguable that VMware vSphere-based virtualization and cloud computing solutions have been at the forefront of these impressive paradigm shifts.

Likewise, Microsoft SQL Server has historically been the most virtualization-friendly business-critical application in the enterprise to date. SQL Server has supported almost every conceivable virtualization feature since the evolution of virtualization on the x86 platform. For example, while many enterprise applications continue to be unable to leverage the benefits of NUMA, SQL Server is not just NUMA-aware, it is natively so. "Soft NUMA" has been a SQL Server feature for over a decade. SQL Server natively and seamlessly supports hyper-threading.

Although "Dynamic Memory" or "Dynamic Processor" have varying interpretations on different flavors of hypervisor, SQL Server has long supported hot-plugging of compute resources (memory and CPUs) for a long time. SQL Server is even now able to dynamically detect and consume hot-added compute resources without a service restart or interruption.

These factors have combined to make vSphere-based virtualization solutions the primary and major platform for virtualized instances of Microsoft SQL Server. Correspondingly, the Microsoft SQL Server application is the single most virtualized application on vSphere-based platforms today.

The continued growth and ubiquity of virtualized Microsoft SQL Server workloads in the enterprise have continued to be the most singular validating testament to the trust and confidence these solutions have engendered in the minds of IT practitioners, administrators and architects over the past few decades. The fact that these solutions continue to be superior to competing offerings has made those competitions less compelling to IT decision-makers worldwide, even when enticed with what appears (on the surface) to be less financially costly options. The VMware vSphere infrastructure continues to remain the desired target platform for virtualizing most major mission-critical applications.

© 2024 Lightbits Labs™

Why Run SQL Server on AVS?

In recent years, Microsoft has made major strides in aggressively positioning Azure, its public cloud enterprise solution. The investments in these efforts have continued to yield positive results for Microsoft, allowing Azure to gain an increasingly larger foothold in the global public cloud space. This success cannot be attributed solely (or even primarily) to a sudden, organic embrace by enterprises and practitioners. The intentionality of Microsoft in overtly (and consciously) steering its customers away from the on-premises versions of its flagship products (e.g., SQL Server and Exchange Server) into Azure-based cloud infrastructure is a contributor to Azure's growth. By making Azure instances of these products comparatively attractive (for example, by releasing new product features and capabilities first into the Azure-based and hosted versions, and through various recent changes in its support and pricing policies), Microsoft openly broadcasts its desired end-state for where it expects customers to gravitate to – the cloud.

Although most enterprises have made significant investments in operating their own private (on-premises) virtualization and cloud infrastructure, the flexibility afforded by the public cloud (among other attributes) makes the weighty issue of (at least) a partial embrace of that infrastructure unavoidable in the long term.

Many enterprises have moved a significant portion of their infrastructure from their physical data centers into one or more of the various public cloud environments, and experts are predicting an upsurge in this trend in the coming years. While VMware has its own public cloud offerings, the aforementioned ubiquity and prevalence informed VMware's strategic alliances and partnerships with almost all of the major public cloud solution providers - through which VMware Cloud Foundation (VCF) has become a staple solution offering in all but a handful of these public cloud infrastructures.

Microsoft's AVS is based on the same VMware technology in VCF, and is - like almost all other such offerings - a recognition by both VMware and Microsoft that enables their mutual enterprise customers to embrace the promises of the public cloud while retaining the superior advantages of the VCF platform. This has served (and continues to serve) them well.

AVS enables vendors' customers to more easily realize their goals of adopting and using the Azure public cloud infrastructure for the mission-critical applications they have already standardized and virtualized on VCF. This can be achieved with the least possible disruption and service interruptions, while avoiding the associated costs of refactoring such workloads from a familiar construct to another – an undertaking more easily imagined than attempted or accomplished.

AVS is built on top of bare metal servers running specially-engineered VMware vSphere-based hypervisor and a management suite and orchestration capabilities - all inside the Microsoft Azure Cloud infrastructure. Because the underlying technologies are the same as what customers are already familiar and comfortable with in its physical, on-premises incarnation, AVS minimizes the learning curves, costs, and administrative efforts required for an enterprise to transition all or a significant part of its existing virtualized workloads into the public cloud space.

AVS is a solution validated, offered, sold, maintained, and supported by Microsoft. This not only provides a much-needed level of assurance and peace of mind to customers; it also helps ease the concerns about the fate of sunk cost and investments in customers' current infrastructure - while at the same time clarifying and streamlining the cloud adoption and infrastructure modernization journey for them as well.

Regardless of the scale of the enterprise, minimizing costs consistently remains a significant concern. Even when distinctly superior and compelling, the total cost of ownership of a given solution must be sufficiently affordable to continue to command acceptance and adoption by customers. On the other end of the spectrum, smart vendors understand that enterprises look for much more than retail costs when choosing between or among competing solutions. When a customer is given the option to combine the best of both options (affordability and functionalities), everyone wins.

© 2024 Lightbits Labs™

Application support, security, and servicing are other compelling reasons for customers choosing to complement their on-premises VCF-based virtual infrastructure and investment with AVS. Since a significant number of customers running old versions of SQL Server have virtualized them on VMware vSphere and are not necessarily Microsoft's Azure customers nor ready to move their workloads into native Azure (given the attendant complexities of application, infrastructural, and architectural re-plumbing), this presents a challenge for everyone concerned. For customers who aren't ready to upgrade their Microsoft products to newer, mainstream supported versions, or move their vSphere-hosted mission-critical applications to another platform, the added cost of Microsoft's extended security updates have become a significant concern.

Because AVS-hosted workloads are covered under Microsoft Azure Hybrid Benefit, VMware vSphere/VCF customers looking to take advantage of the free Extended Security Updates offer are able to do so by moving some or all of these workloads to AVS without incurring the additional cost that would have been otherwise required, without the added complexities of refactoring, re-provisioning, or re-learning new administrative and infrastructure management tasks. Azure Hybrid Benefit allows customers with SQL Server and Windows Server licenses with Software Assurance to redeploy those licenses to their AVS VMs and avoid the need to purchase new licenses.

Because the Management plane of AVS is the familiar and ubiquitous VMware vCenter UI, AVS affords customers the opportunity and ability to embrace the public cloud they desire, at a price point much more affordable than the alternative, with the least possible disruption and level of effort. With AVS and Azure Hybrid Benefits, enterprises are able to meet the cost-minimization objectives in their overall posture and continue to enjoy the benefits of the superior virtualization and cloud computing solutions they have long become accustomed to over the years.

Microsoft provides primary and direct support for AVS. This means that ESXi host lifecycle management (patching, upgrade, etc.) is handled by Microsoft in an automated, unobtrusive fashion - completely transparent to customers. It also means that SQL Server administrators have direct access to Microsoft Support if they encounter any technical challenges related to the AVS-hosted SQL Server instances. There is no longer a need for them to open Support request tickets simultaneously with both Microsoft and VMware by Broadcom.

VCF is the unified suite of vSphere virtualization platform and related technologies from VMware by Broadcom. It is a tightly-coupled, integrated offering of most of the products and components a customer previously has to purchase, license, deploy and manage separately. This tight integration ensures an even easier management capability, which helps substantially reduce administrative and maintenance efforts. For example - thanks to tight integration of VCF components like HCX, Aria Automation, and Aria Operations - moving, deploying, or monitoring SQL Server workloads at a large scale from other vSphere-based platforms is now much easier to do. VMware Live Site Recovery Manager for Hyperscalers(VLSR, formerly Site Recovery Manager (SRM)) is the component that provides a comprehensive, holistic, and integrated solution to protect critical SQL Server instances against disaster events in AVS.

© 2024 Lightbits Labs™

Clustered SQL Server Instances Out of Scope

The solution discussed in this document does not apply to clustered SQL Server workloads. This use case will be addressed in either a future update to this document, or in a different document dedicated specifically to the use case. The following diagram shows the behavior of attempting to migrate shared VMDKs to an NVMe/TCP datastore.

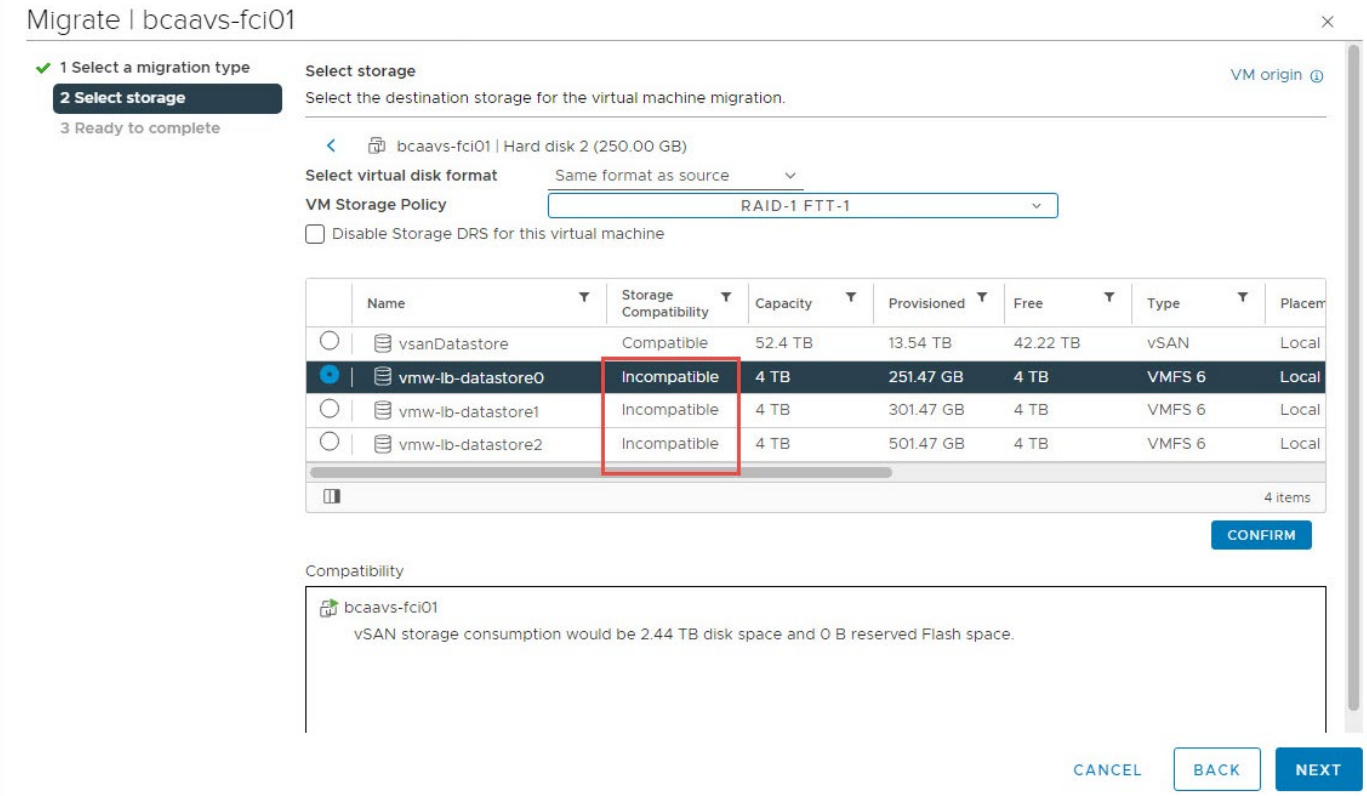

Figure 1 - Incompatible NVMe/TCP Datastores

Lightbits Software-Defined Storage

The Lightbits Cloud Data Platform is a software-defined, block storage solution that can run anywhere - on private, public, or edge clouds. In Azure, Lightbits is deployed on Lasv3 or Lsv3 instances backed by Azure's local NVMe storage. By clustering these virtual machines (VMs) and synchronously replicating data between storage nodes, Lightbits brings the high throughput and low latency of locally-attached NVMe and the durability, resiliency, and scalability of a SAN in the cloud.

With the capability to reach almost a million IOPS per volume, the Lightbits solution in Azure is perfect when workloads need to be fast and resilient to infrastructure failures.

© 2024 Lightbits Labs™

Why Lightbits with AVS?

AVS clusters are based on hyper-converged infrastructure built on top of a set of host types with specific amounts of CPU, memory, storage, and network. 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 Virtual Machine File System (VMFS) datastores backed by NVMe®/TCP volumes. This combination gives you the high performance and low latency of NVMe-oF external storage with essential enterprise data services, and all of the flexibility and scalability of VMFS.

VMFS has become popular because of its ability to handle the requirements of virtualization and to optimize 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.

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.

© 2024 Lightbits Labs™

Architecture

Lightbits for AVS 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 2, 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.

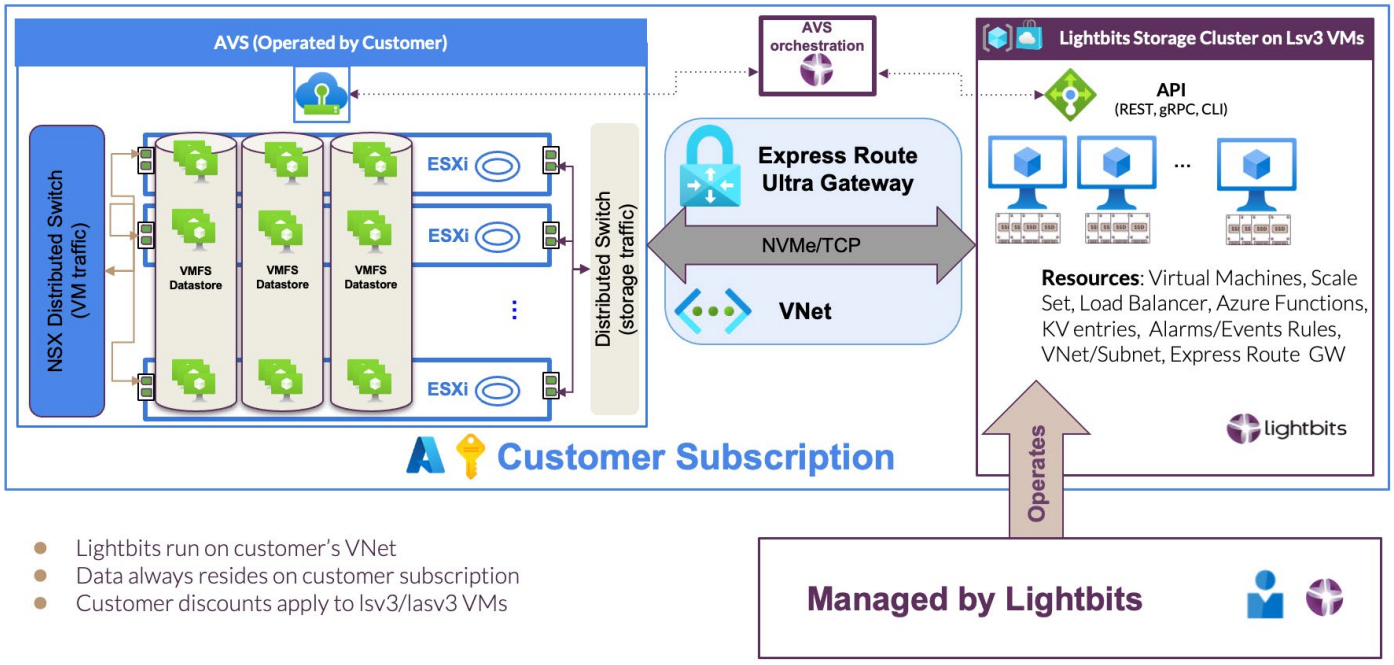

Figure 2 - Lightbits Deployment for AVS

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 1 for the 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 1 - Virtual Machine Types Available for Lightbits with Their Respective Capacity and Performance

Connectivity

Lightbits is connected via an Azure Express Route connection with an Ultra gateway and FastPath enabled. This provides the highest performance with less than half a millisecond round trip latency. In the Deployment guide below, the connectivity can be created during deployment of the Lightbits cluster, but if an existing vNet is already being used, then the Ultra gateway and Express Route connection must be created manually either before or after the Lightbits cluster has been deployed.

Because Lightbits is fully certified with AVS - for the AV36, AV36P, AV52 and AV64 AVS SKUs, Lightbits connects through the Distributed Virtual Switch, leaving the VM traffic network free for applications to use. This is a huge benefit over using the VM network to attach volumes “in-guest”, directly inside the VM’s operating system.

© 2024 Lightbits Labs™

Deployment

Deploying Lightbits

This section will explain how to deploy Lightbits inside Azure and connect the cluster through Express Route to the AVS SDDC.

NOTE: these instructions are relevant for Lightbits release version 3.7.1. For the most recent documentation, please refer to the Lightbits Labs Azure documentation for AVS.

Overview

The LBAVS tool is a binary installed on any Linux machine that has network access to the AVS SDDC, as well as the Lightbits cluster being used to provide datastores. LBAVS can manage many SDDCs and Lightbits clusters from a single machine, as well as manage multi-tenancy through Lightbits projects. By default, any Lightbits cluster v3.6.1 and above deployed in Azure through the Managed Application has LBAVS pre-installed.

The LBAVS tool authenticates with Azure using a service principal or managed identity to run Azure AVS run command cmdlets, which enable users to perform actions on the SDDC through vCenter that would usually require elevated privileges not available in Azure.

Prerequisites

To use Lightbits with AVS, you will need the following prerequisites.

Run Commands for VMFS

To perform the actions required on the SDDC, the appropriate run commands will have to be enabled on your subscription for your SDDC. During the private preview, to activate these run commands, please reach out to your Lightbits representative.

Azure Subscription Permissions

To perform the full install as described, you will need the following permissions on the subscription that is being used for the Lightbits cluster and the AVS SDDC:

- Managed Identity Contributor
- Role-Based Access Control Administrator
- Virtual Machine Contributor

If you do not have these permissions, you may have to create the Managed Identity and Azure Virtual Machine another way.

AVS SDDC

An AVS SDDC needs to have already been deployed prior to starting this process. For more information about deploying an AVS SDDC, please consult the Microsoft Azure documentation.

© 2024 Lightbits Labs™

Enable External Storage for SDDC

This feature can be activated in your Azure subscription by running the following commands in Azure CLI:

```
Unset
az feature register --namespace Microsoft.AVS --name earlyAccess
az feature register --namespace Microsoft.AVS --name iSCSIMultipath
az provider register -n Microsoft.AVS
```

This can take a few minutes, check the status is “Registered” by running:

```
Unset
az feature list --query
“[?name=='Microsoft.AVS/earlyAccess'||name=='Microsoft.AVS/iSCSIMultipath'].{Name:name,
State:properties.state}”
```

Once the features are enabled, this API builds two new subnets and VMKernel adapters within your SDDC, which we will use to connect to the Lightbits cluster.

Start by providing an IP block for deploying external storage. Navigate to the Storage tab in your Azure VMware Solution private cloud in the Azure portal. The address block should be a /24 network.

© 2024 Lightbits Labs™

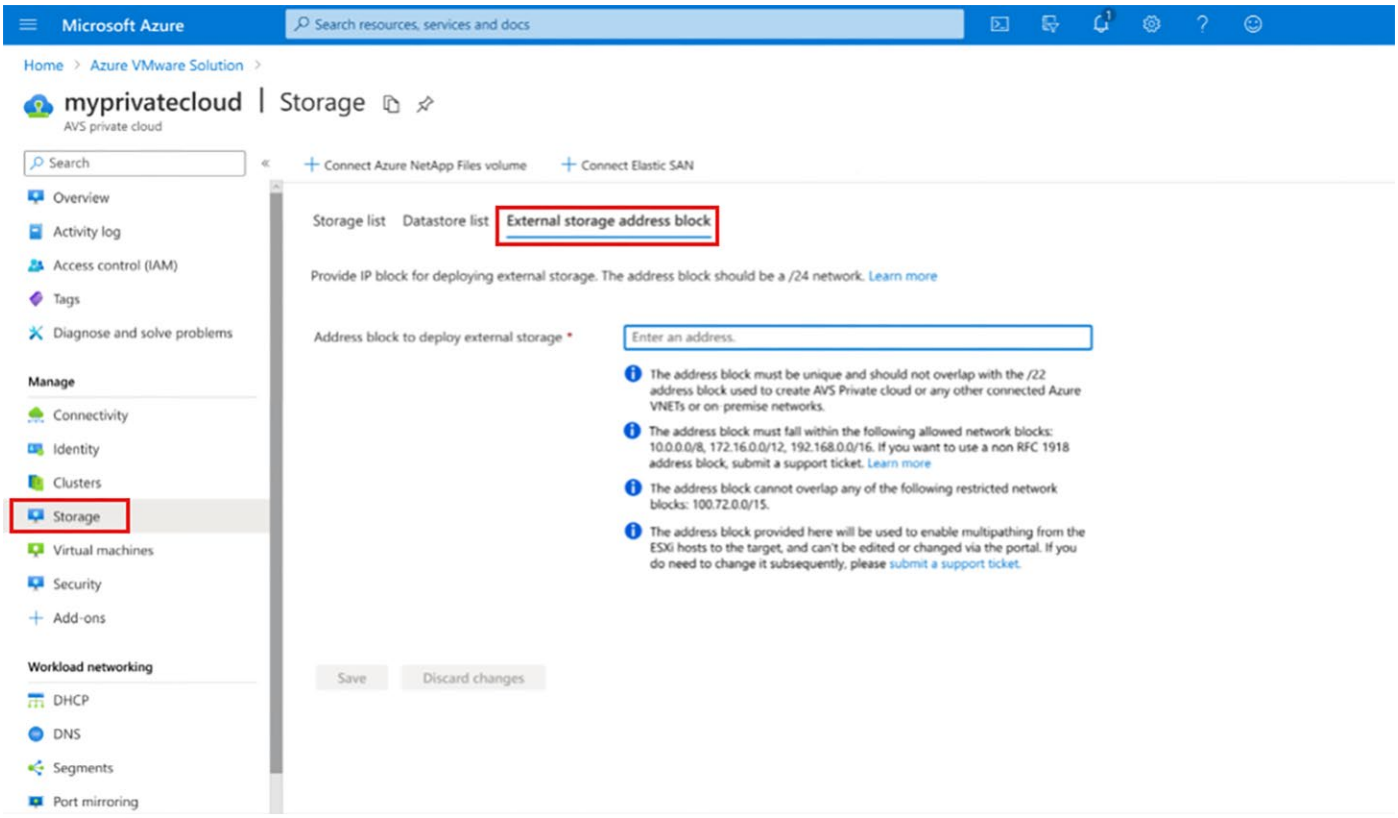

Figure 3 - Enable External Storage

- The address block must be unique and not overlap with the /22 used to create your Azure VMware Solution private cloud or any other connected Azure virtual networks or on-premises network.
- The address block must fall within the following allowed network blocks: 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16. If you want to use a non-RFC 1918 address block, submit a support request.
- The address block cannot overlap any of the following restricted network blocks: 100.72.0.0/15.
- The address block provided is used to enable multipathing from the ESXi hosts to the target. It cannot be edited or changed. If you need to change it, submit a support request.

Terminology

- AVS: Azure VMWare Solution
- SDDC: Software Defined Datacenter - AVS private cloud environment.
- lbavs: Lightbits command-line interface to interact with AVS and Lightbits storage.
- lbcli: Lightbits command-line interface to interact with Lightbits storage on an admin level.
- SDS: Lightbits Scale-Out Disaggregated Software-Defined Storage
- VM: Virtual Machine

© 2024 Lightbits Labs™

Architecture

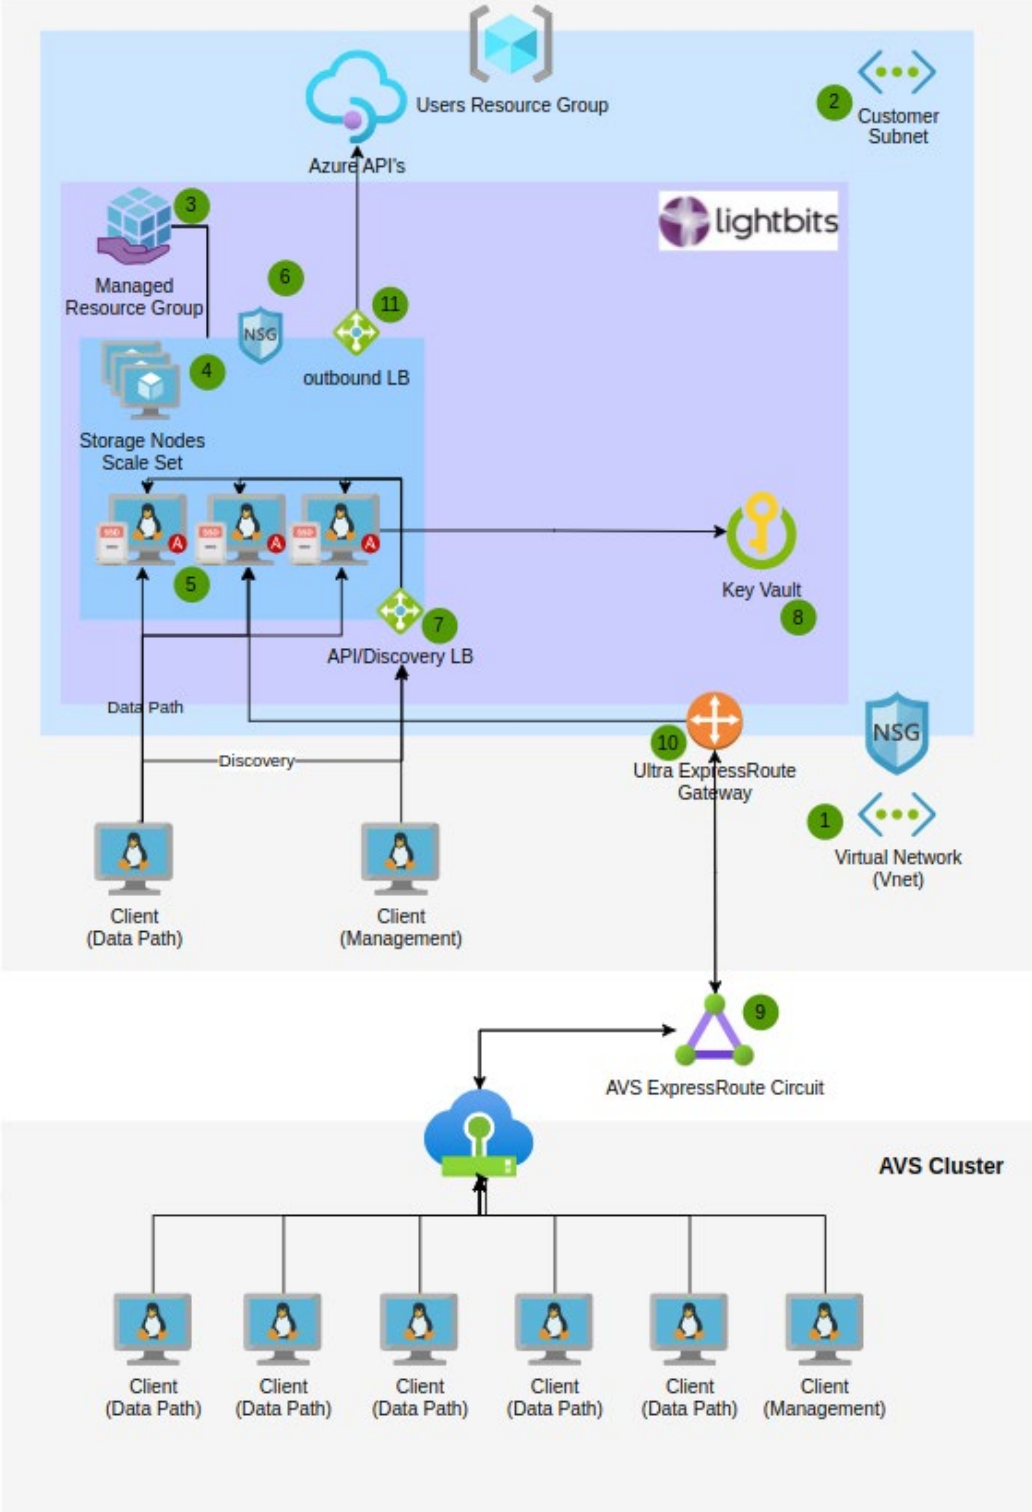

Figure 4 - Lightbits and AVS Architecture

© 2024 Lightbits Labs™

Gathering Information About the Deployment Environment

This section provides an example checklist to help you prepare for the installation.

Resource	Description
Subscription ID	The Azure subscription used by the AVS SDDC.
Resource Group	The Azure resource group used to house the AVS SDDC.
SDDC Name	The name of the AVS SDDC as shown in the Azure portal.
SDDC IP Address Range	The IP Address range used within your SDDC and the rest of your Azure environment.

Table 2 - Environment Information

Deploying Lightbits for AVS

NOTE: Users performing this action will require the permissions outlined in the Minimum Required Permissions article.

Navigate to the Marketplace Page

1. Access the Azure portal using the common link ms.portal.azure.com, and log in using your Azure credentials.
2. Once in the portal, use the search bar and search for Marketplace, and then click on the entry.

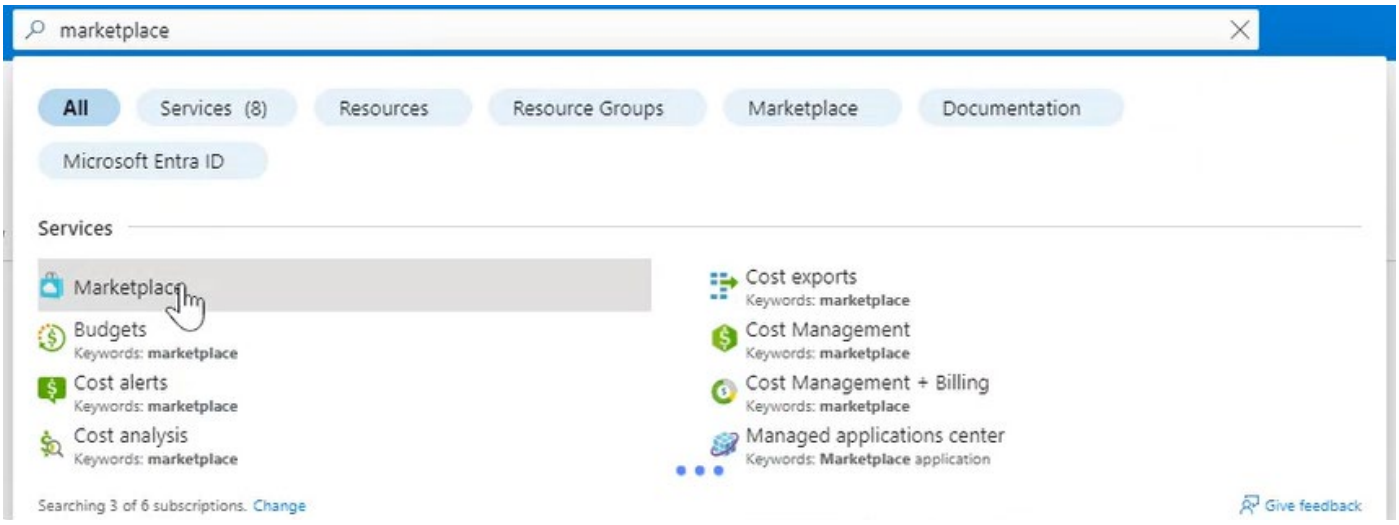

Figure 5 - Azure Marketplace Search

3. Inside the Marketplace blade, search for Lightbits and select a version higher than 3.6.1.

© 2024 Lightbits Labs™

Basics Tab

1. Inside the Basics tab, select:
   a. Subscription: This does not have to be the same subscription as your AVS SDDC.
   b. Resource Group: You can use a currently deployed resource group or create a new one.
   c. Region: Select the same region that the SDDC is deployed into for best performance.
   d. Multi-Az: Check this box to deploy the Lightbits cluster across all three availability zones in the region (where available).
   e. Availability Zone: If not deploying in Multi-Az, choose the zone that is the same as the AVS SDDC. This information can be found on the Private Cloud Overview page in the Number of hosts section.
   f. Application Name: Give the managed application a name.
   g. Managed Resource Group: Optionally give the managed application managed resource group a unique name to fit with the organizational naming conventions.
2. The output should look like the screenshot below.
3. Click Next.

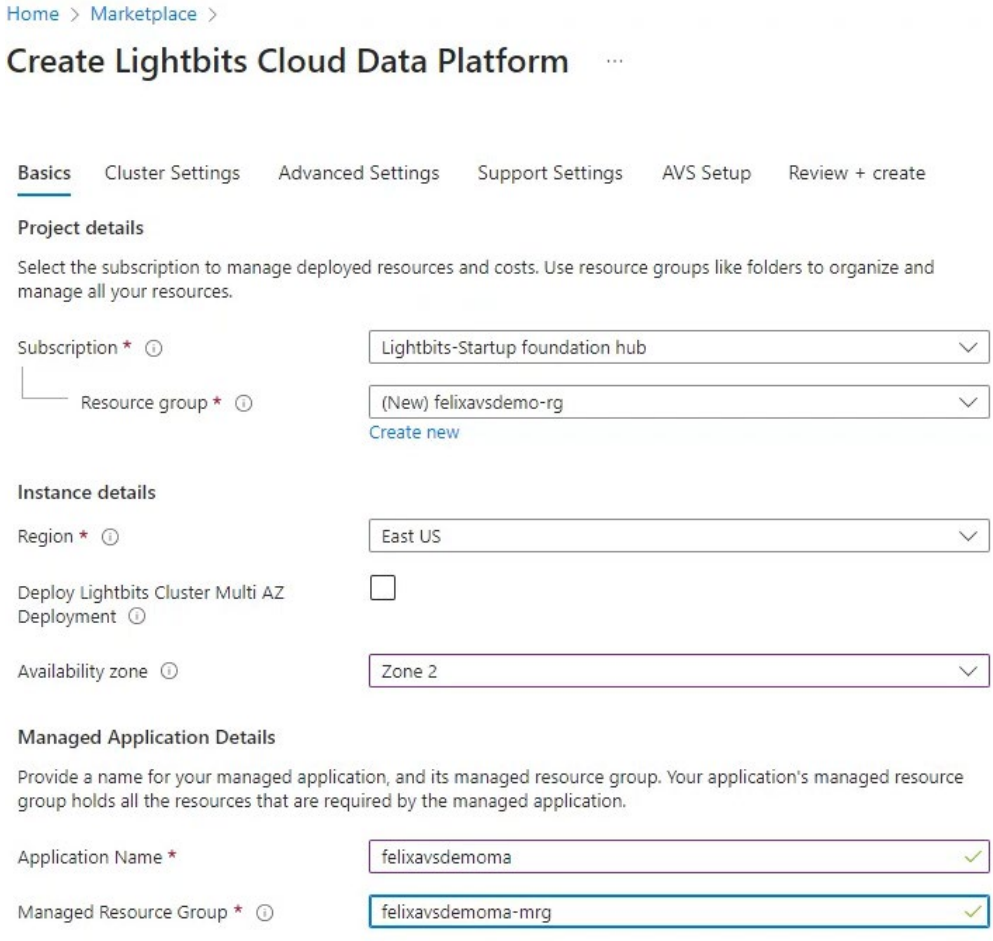

Figure 6 - Create Lightbits Basics

Cluster Settings Tab

1. Inside the Cluster Settingstab, select:
   a. Virtual Network ID: click Edit Network and change:
   b. Size: Choose the size of the Virtual Machines that will be a part of the Lightbits cluster. For help with sizing, please reach out to your Lightbits representative.
   c. Username: Choose a username that can be used to SSH into the Lightbits VMs.
   d. SSH Public Key Source: Generate a new key pair or select an existing key. Please note that this will need to be accessible for running the LBAVS commands.
   e. User Assigned Managed Identity: This can be left blank, since a new vNet is being created.
2. The output should look like the screenshot below.
3. Click Next.

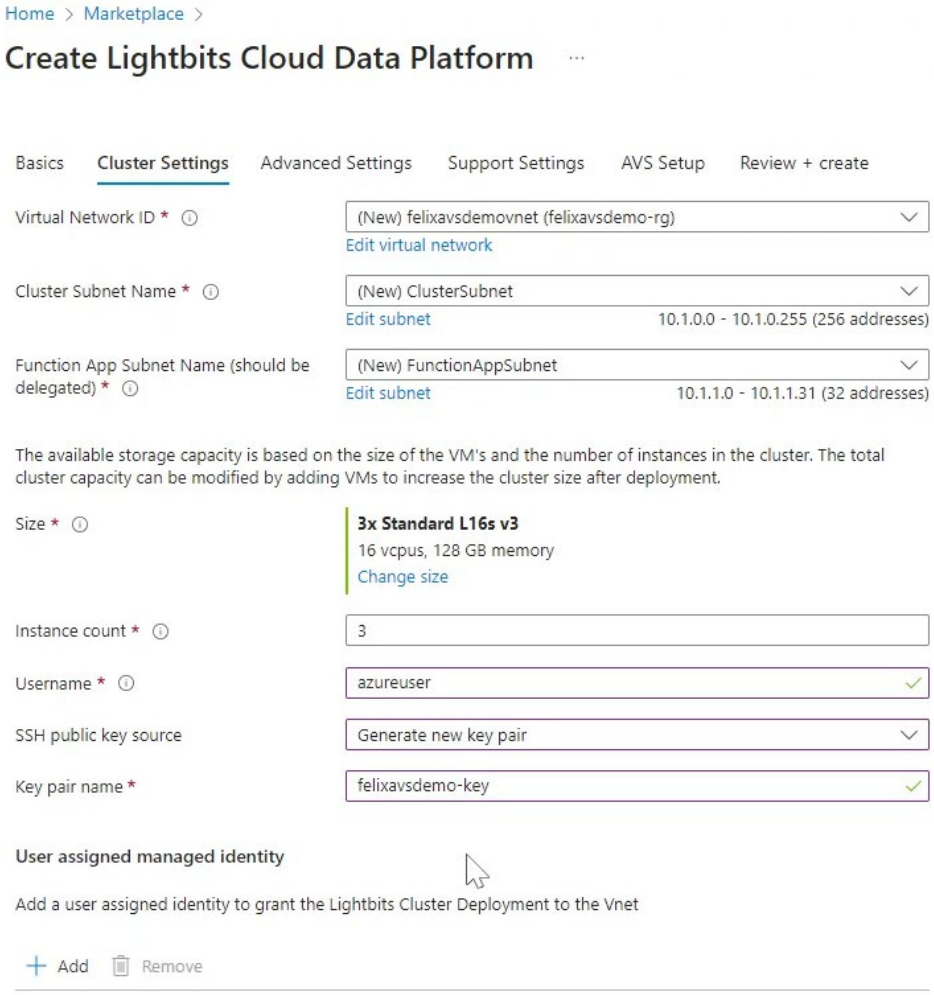

Figure 7 - Create Lightbits Cluster Settings

NOTE: For connecting directly to an AVS SDDC using the AVS Setup tab, a new vNet must be created. If you already have a vNet, then please follow the documentation about creating a cluster with an existing vNet, then manually connect the vNet to the SDDC.

Advanced Settings

1. Inside the Advanced Settings tab, select:
   a. Resource Prefix: Optionally change the prefix that will be applied to all resources deployed as part of the managed application, including the Virtual Machine Scale Set.
   b. Network Security Group Name: Optionally change the name of the generated Network Security Group that will be assigned to the managed application Virtual Machines.
2. The output should look like the screenshot below.
3. Click Next.

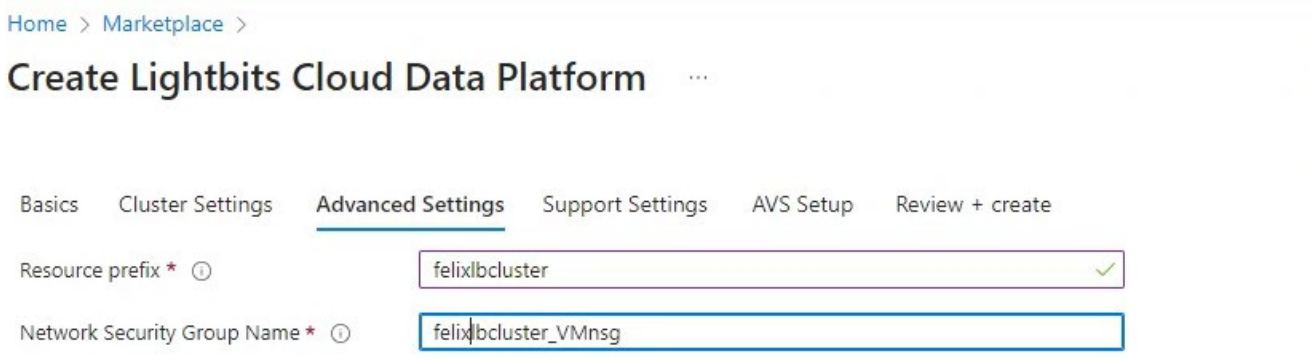

Figure 8 - Create Lightbits Advanced Settings

Support Settings

1. Inside the Support Settings tab:
   a. Check the box to allow the Lightbits Support team access to the cluster. This allows Lightbits to provide proactive support to the managed application. If this box is not checked, customers will have to provide Lightbits Support personnel with SSH access to the cluster, if required.
   b. Customer Name: Provide the name of a point of contact in the company in case Lightbits Support needs to contact you for troubleshooting or maintenance.
   c. Customer Email: Provide the email address of a point of contact in the company, in case Lightbits Support needs to contact you for troubleshooting or maintenance.
   d. Customer Phone Number: Provide the phone number of a point of contact in the company, in case Lightbits Support needs to contact you for troubleshooting or maintenance.

AVS Setup

1. Inside the AVS Setup tab:
   a. Check the box to create connectivity to the AVS SDDC during deployment. If this box is not checked, the express route, gateway, and connection must be made manually using the Azure documentation.

© 2024 Lightbits Labs™

b. Express Route Key: Generate a new ExpressRoute authorization key and place it into this section. Note that ExpressRoute authorization keys can only be used for a single connection at any time.

c. Express Route Address ID: Copy the ExpressRoute ID from the ExpressRoute section of the Connectivity blade of the AVS Private Cloud page. It will be in this format: /subscriptions/$SUBSCRIPTION/resourceGroups/$RESOURCE_GROUP/providers/Microsoft.Network/expressRouteCircuits/$CIRCUIT_NAME

d. Gateway Subnet: Provide a /27 address space that is within the vNet created in the Cluster Settings tab, but does not overlap with the other subnet ranges.

2. The output should look like the screenshot below.
3. Click Next.

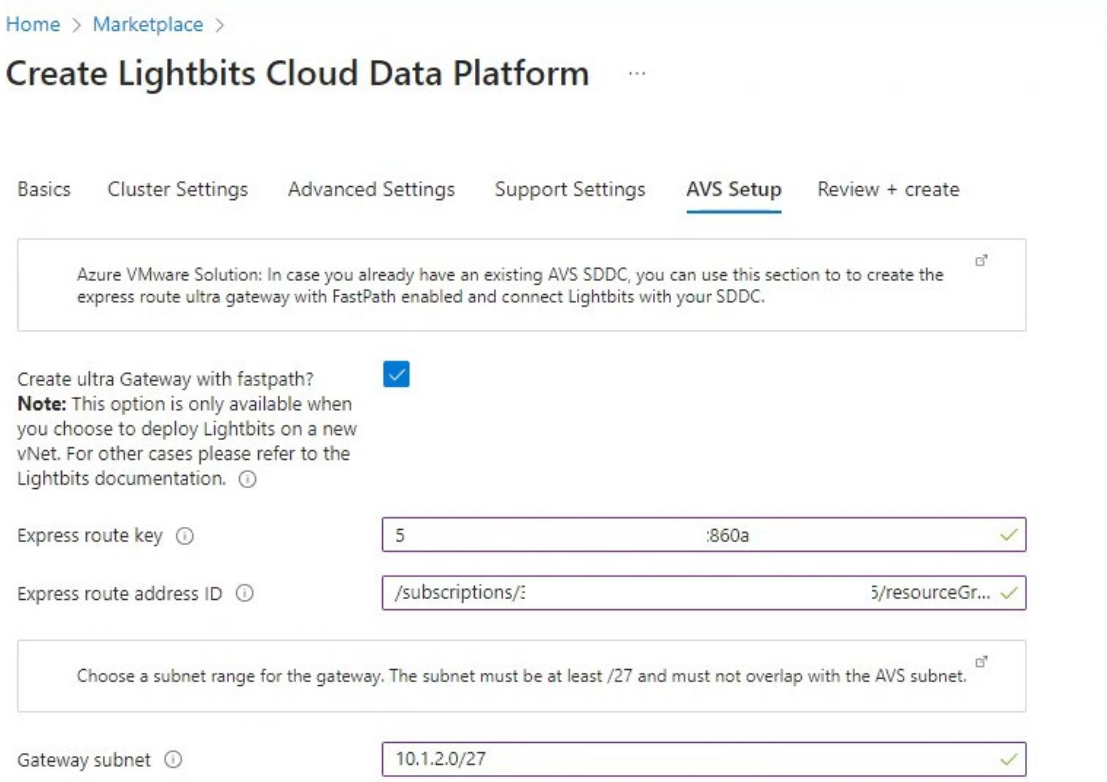

Figure 9 - Create Lightbits AVS Setup

Review

1. Check the Co-Admin Access Permission. This allows Lightbits access to the Managed Application Managed Resource Group when required.
2. Check that the information entered is correct, and click Create.
3. Wait for the deployment to complete. If the AVS Setup section is filled in, the deployment will take around 25 minutes to complete. If the AVS Setup section has not been filled in, the deployment will take around eight minutes to complete - but you will have to follow all of the steps outlined in the Azure documentation to connect the AVS SDDC to the Lightbits vNet.

© 2024 Lightbits Labs™

Configuring the LBAVS Managed Identity

Getting the Lightbits VMSS Managed Identity Information

1. Navigate to the Lightbits Managed Application Managed Resource Group by searching for the name given to the managed resource group during deployment in the Resource Groups blade.
2. Click on the Lightbits Virtual Machine Scale Set, and in the VMSS blade, click Identity.
3. Click on User assigned and the current user assigned managed identity will display. Note down the name of this managed identity. The name should look something like ${vmss_name}_ keyvaultidentity.
4. This will look something like the example from the screenshot below.

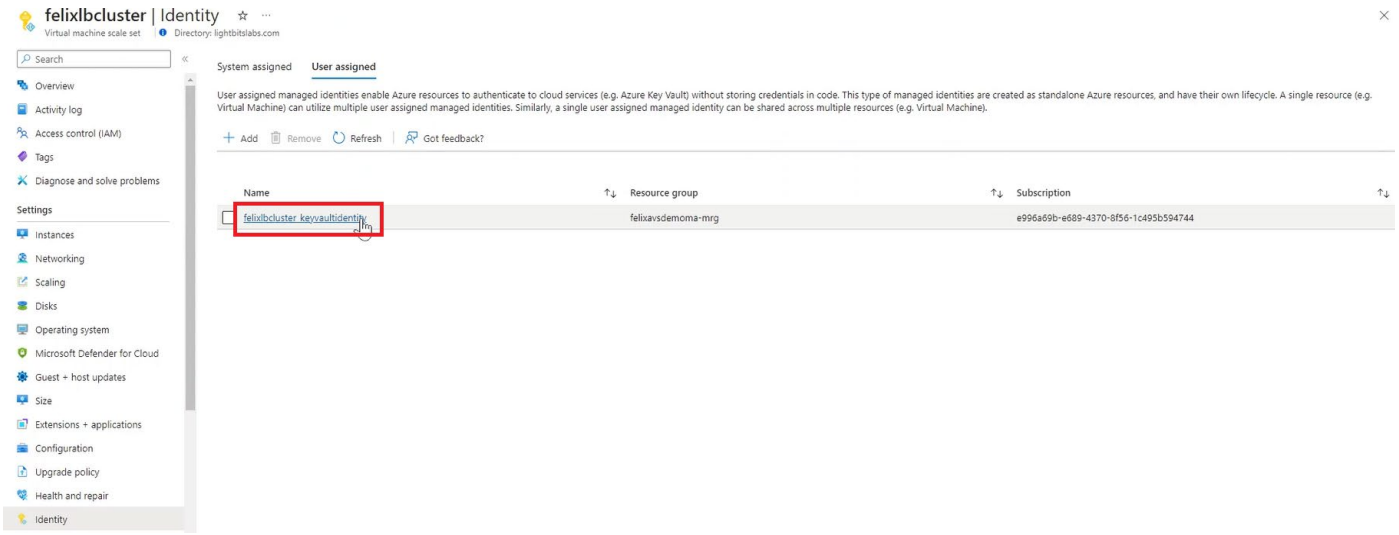

Figure 10 - Create Managed Identity

Assigning Permissions to the Managed Identity

1. Navigate to Resource groups and choose the resource group that contains the AVS SDDC.
2. Click Access control (IAM).

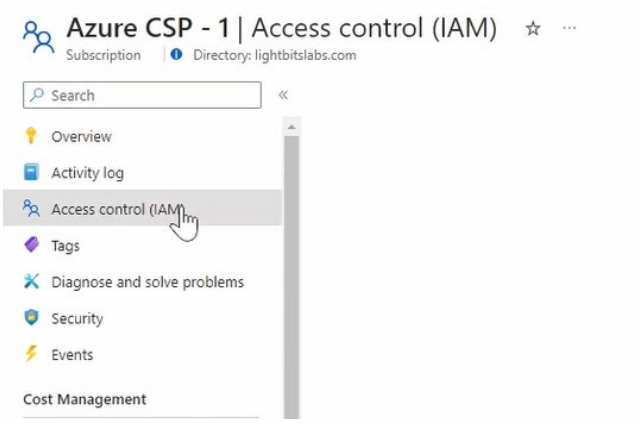

Figure 11 - Create Managed Identity - Permissions

3. Click +Add and Add role assignment.

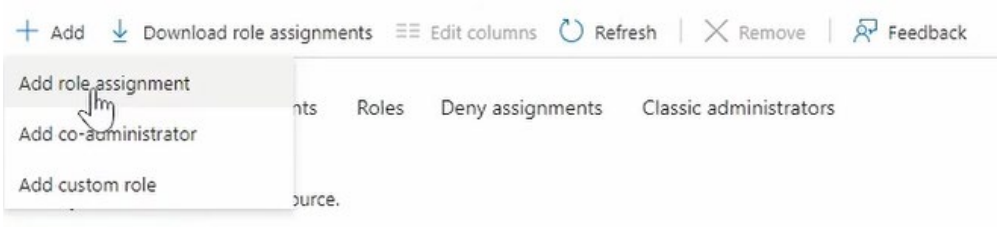

Figure 12 - Create Managed Identity - Add Role Assignment

4. Click Privileged administrator roles and choose Contributor.

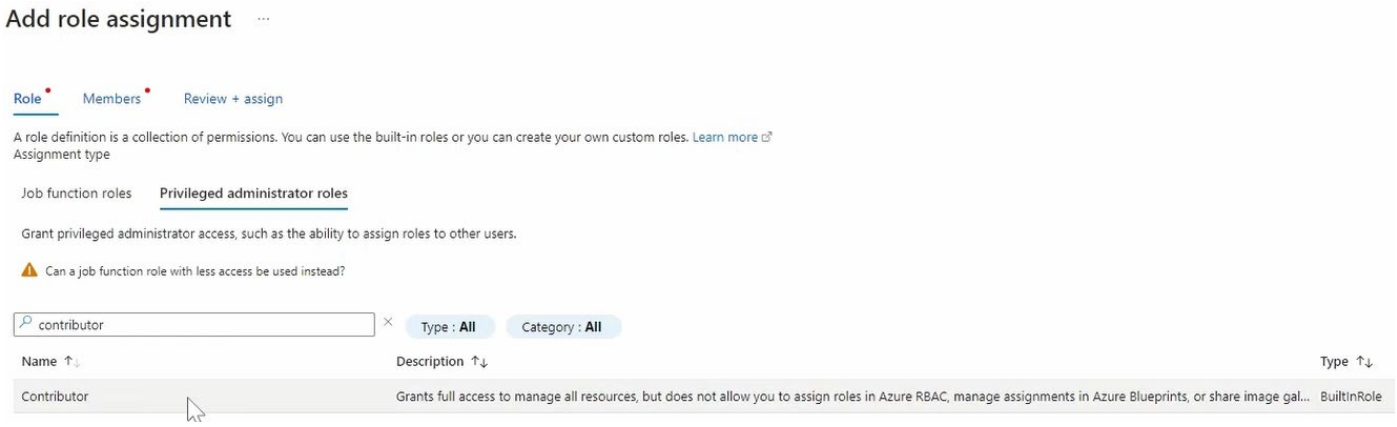

Figure 13 - Create Managed Identity - Select Contributor

5. Click Next. Choose Managed identity and +Select members, and then choose the managed identity that is already assigned to the Lightbits Managed Application VMSS. You may need to change the Subscription to the one that contains the Lightbits Managed Application.

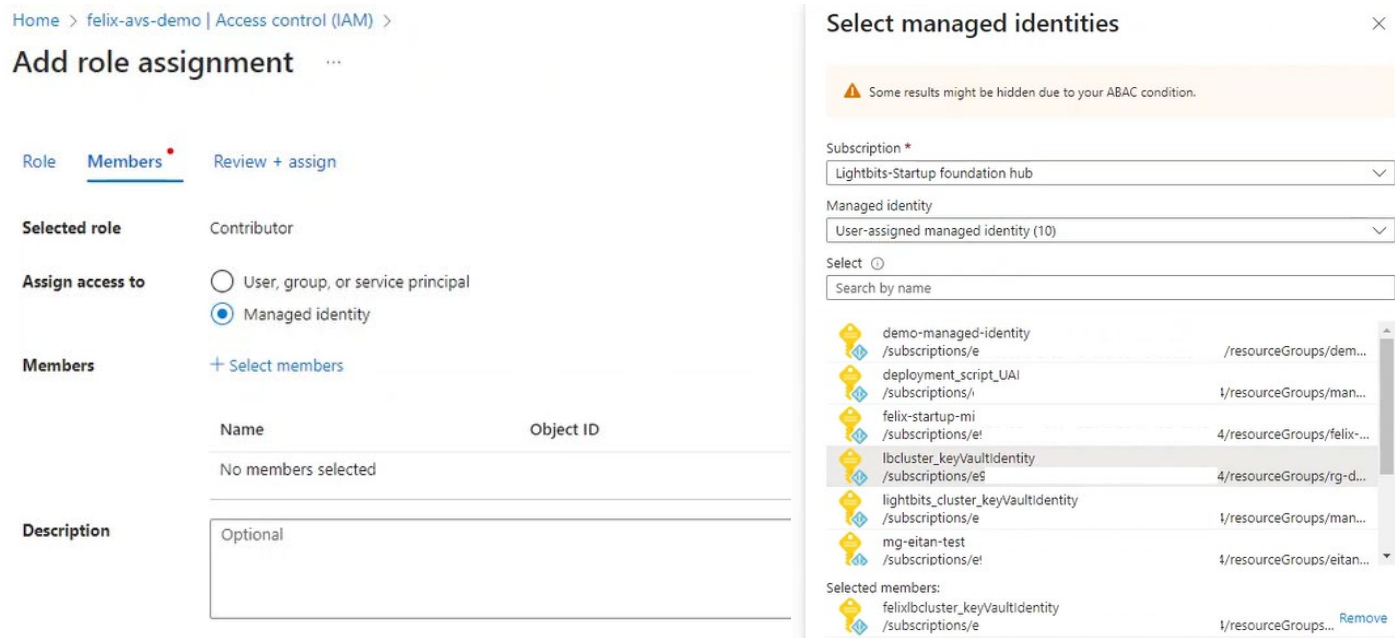

Figure 14 - Create Managed Identity - Select Managed Identity

Click Select and then Review + assign.

© 2024 Lightbits Labs™

Configuring the LBAVS Tool

Setting Up the LBAVS Configuration File

1. SSH into one of the Lightbits instances in the VMSS.
2. Navigate to /etc/lbavs/lbavs.yaml and open it with your favorite text editor such as vim or nano.
3. Copy the contents of the file below, replacing the following parameters:

   ${SUBSCRIPTION_ID}: The ID of the subscription where the AVS SDDC is installed.

   ${SDDC_NAME}: The name of the AVS private cloud that Lightbits will connect to.

   ${SDDC_RESOURCE_GROUP}: The name of the resource group that the AVS private cloud is deployed into.

   ${LIGHTBITS_MGMT_IP}: The IP address to access the Lightbits REST API. To get this IP address, see Find Private Management and Discovery Address in the Azure Managed Application Installation Guide.

   ${LIGHTBITS_JWT}: The JWT of the Lightbits cluster for authentication. To get the JWT, see Get the System JWT from Azure Portal.

```
Unset
azure-subscription-id: ${SUBSCRIPTION_ID}
avss:
  - name: azavs
    avs: ${SDDC_NAME}
    group: ${SDDC_RESOURCE_GROUP}
    useManagedIdentity: Yes
    managedIdentity:
      url: http://169.254.169.254/metadata/identity/oauth2/token
      apiVersion: “2018-02-01”
default-avs: azavs

lbcs:
  - name: lbc1
    endpoint: https://${LIGHTBITS_MGMT_IP}:443
    jwt: ${LIGHTBITS_JWT}
default-lbc: lbc1
```

NOTE: The appid and secret parameters can be left as the default and are only filled in when using a service principal to authenticate with Azure. user and pwd are only used for testing and should be left blank.

4. Save the file.

© 2024 Lightbits Labs™

Testing the Tool Setup

1. Create the log file to edit permissions: touch /var/log/lbavs.log
2. Change the permissions on that file to avoid having to run lbavs with sudo: chmod 755 /var/log/lbavs.log
3. When running the tool the first time, it is recommended to monitor the LBAVS logs. In another SSH session, run:

```
Unset
tail -f /var/log/lbavs.log
```

NOTE: The log file is created once the first LBAVS command has been run, so if this is the first time running the command, you may have to tail the file after running.

4. List datastores: the ${SDDC_CLUSTER_NAME} parameter is the name of the cluster in the AVS SDDC. By default, there is a cluster called Cluster-1:

```
Unset
lbavs datastore list -v ${SDDC_CLUSTER_NAME}
```

5. The response to list datastores should be No datastores

© 2024 Lightbits Labs™

Configuring AVS Networking

Creating NVMe/TCP Networks

Lightbits recommends two dedicated /25 virtual network address ranges outside of the address space of any network connected to the SDDC - for connectivity to the Lightbits cluster. During private preview, to create these networks, contact your Lightbits representative to work with Azure to enable this currently private feature.

Once the networks have been created, your SDDC will have two new VMKernel adapters per ESXi host, similar to the screenshot below.

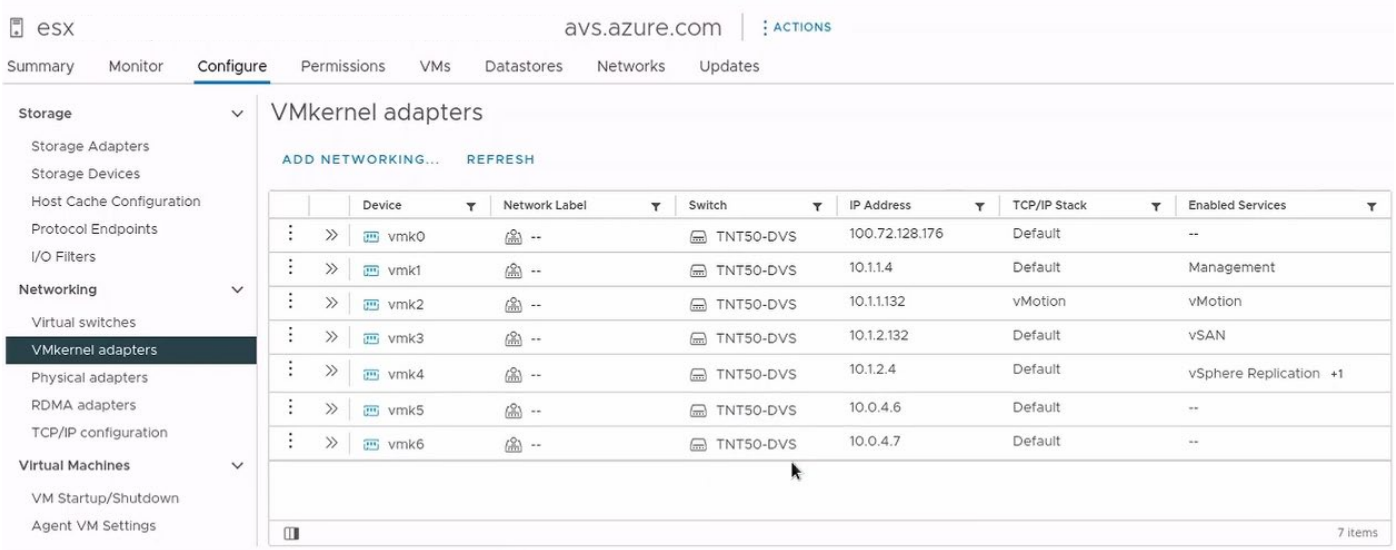

Figure 15 - AVS VMkernel Adaptors

Tagging VMKernel Adapters

1. Note down the VMKernel device names and ESXi hostnames for each host in the SDDC cluster.
2. For each VMKernel adapter and ESXi host, run the following command:

```
Unset
lbavs network tag -e ${ESXI_HOSTNAME} -n ${VMK_DEVICE_NAME}
```

3. Once complete, the VMKernel adapter should look like the below example. Note that NVMe over TCP is now an enabled service:

Figure 16 - AVS NVMe/TCP VMkernel Adaptor

© 2024 Lightbits Labs™

4. Repeat this process until each ESXi host has two VMKernel adapters tagged with NVMe over TCP.

Creating NVMe/TCP Storage Adapters

1. Note down the ESXi hostname and physical adapter names that use the DV switch. In this example, they are vmnic1 and vmnic2:

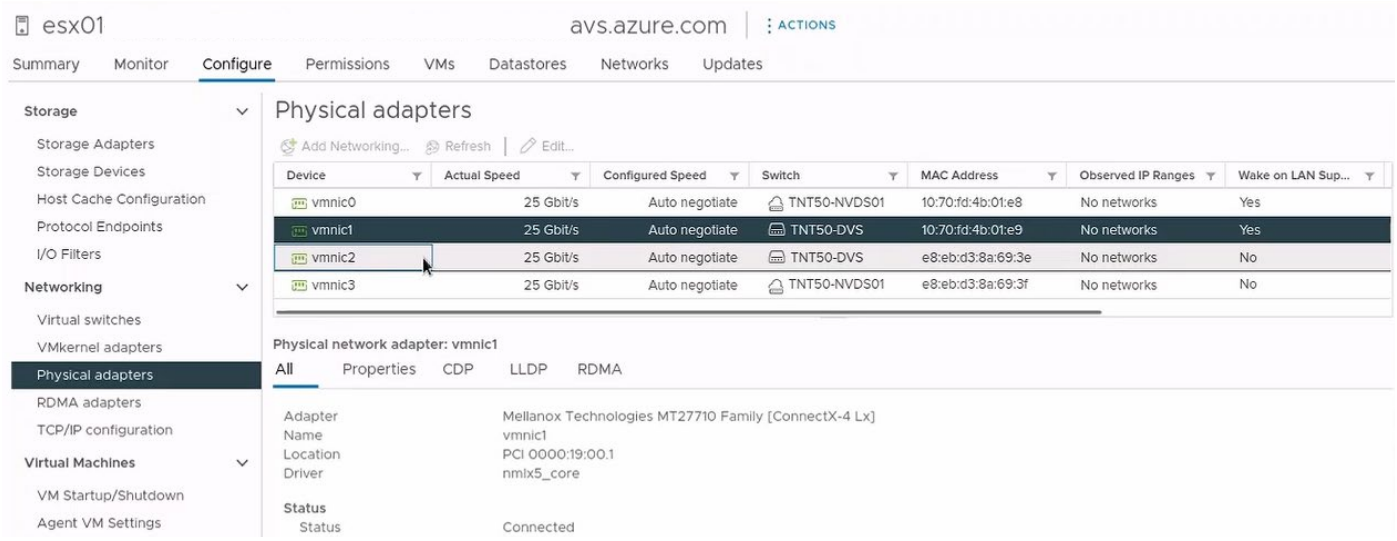

Figure 17 - AVS Physical Adaptors

2. For each physical adapter and ESXi host, run the following command:

```
Unset
lbavs network create -e ${ESXI_HOSTNAME} -n ${PHYSICAL_ADAPTER_NAME}
```

3. Repeat the process for each of the two physical adapters per ESXi host.
4. Once complete, there should be two new NVMe over TCP storage adapters per ESXi host:

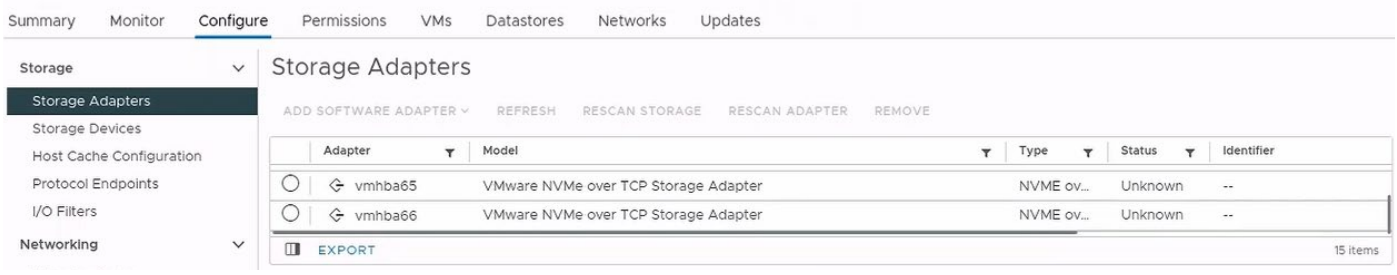

Figure 18 - AVS Storage Adaptors

© 2024 Lightbits Labs™

Connecting to the Lightbits Cluster

1. Once all VMKernel adapters and storage adapters have been created, you need to make a connection to the Lightbits storage cluster. To create this connection, run the following command:

```
Unset
lbavs network connect -v ${SDDC_CLUSTER_NAME}
```

2. This command needs to be run only once. LBAVS will initiate Run Commands to attach all the storage adapters to each Lightbits target node. This process will take some time to complete.
3. After the connect command has completed, you will see a number of controllers per storage adapter equivalent to the Lightbits cluster size. The example below shows a three-node Lightbits cluster connected to one of the storage adapters:

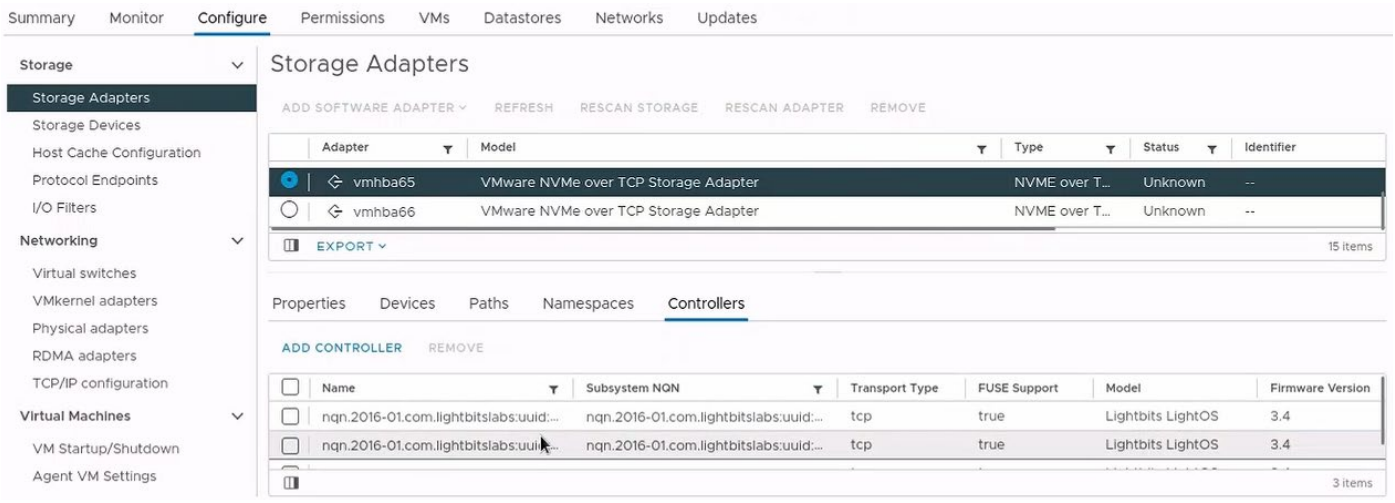

Figure 19 - AVS Lightbits Connections

© 2024 Lightbits Labs™

Create Datastores

A datastore in vSphere is presented by Lightbits as an NVMe/TCP block device before VMFS is overlaid onto the device. For optimal performance with Lightbits as well as following the best practices for running SQL Server on VMware, at least 1 datastore per Lightbits node should be created. This allows for SQL Server vmdks to be grouped and spread across multiple datastores.

To create a datastore, run the following command using the LBAVS client:

```
Unset

lbavs datastore create -n ${DATASTORE_NAME} -v ${SDDC_CLUSTER_NAME} -s ${SIZE} -r
${REPLICA_COUNT} -c ${COMPRESSION} -f ${FAILURE_DOMAIN}
```

Example command for creating a datastore:

```
Unset

lbavs datastore create -n lbds01 -v Cluster-1 -s 1TiB -r 3 -c true
```

A datastore will be created and LBAVS will output something similar to the below example:

Figure 20 - AVS Lightbits Datastore Created

Repeat this process for all Lightbits nodes in the cluster. For example, for a Lightbits cluster with 3 nodes, create 3 datastores. For a Lightbits cluster with 9 nodes, create 9 datastores.

If the Lightbits cluster is spread across multiple zones, ensure that the -f tag FAILURE_DOMAIN is set to the same zone as the AVS SDDC. If using multi-zone, only create as many volumes as there are nodes in the primary zone. The other zones will primarily be used for failures.

© 2024 Lightbits Labs™

Creating SQL Server VMs

The following describes (at a high-level) our Test Lab environment and the technical components, processes, configuration options used in the AVS infrastructure referenced in this document.

A single Azure VMware Solution SDDC, with three (3) ESXi Hosts running ESXi version 7, Update 3 (the current version available as at the time of our test and validation exercise).

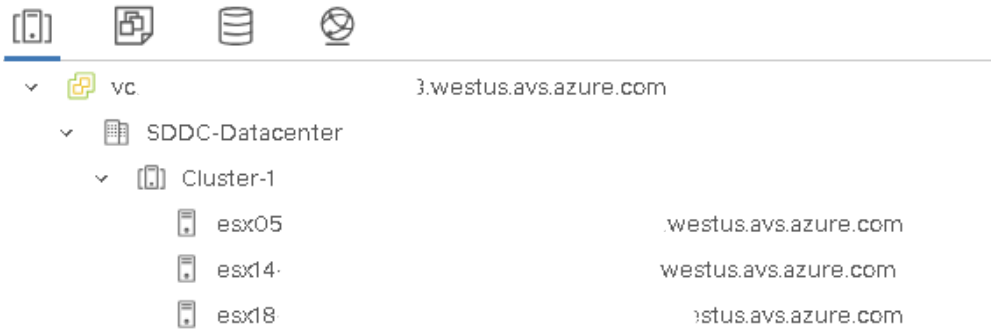

Figure 21 - Three ESXi Hosts in AVS Cluster

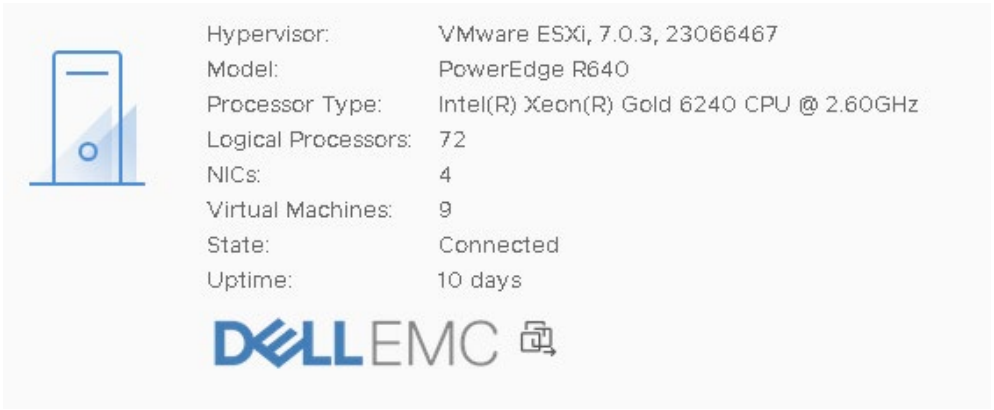

Figure 22 - ESXi Host Model and Hypervisor Version

© 2024 Lightbits Labs™

Each ESXi host's compute and storage capacity are as shown in the images below:

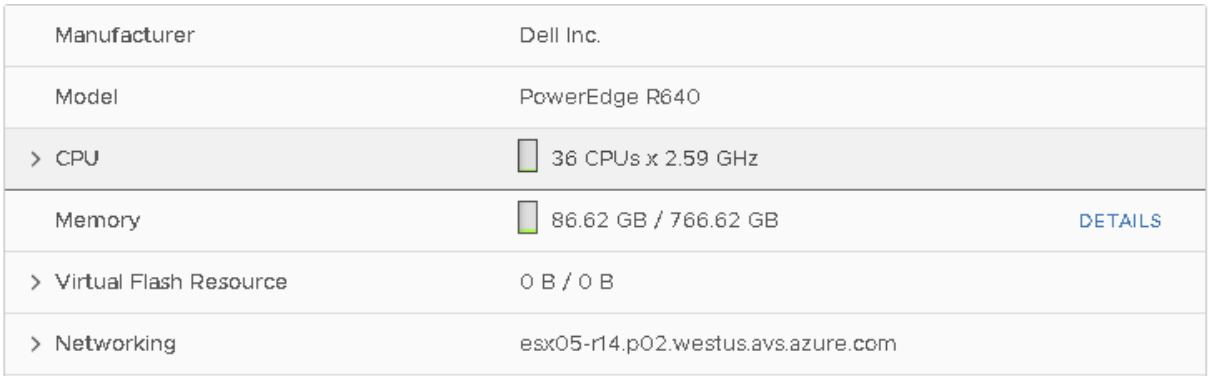

Figure 23 - ESXi CPU and Memory Capacity

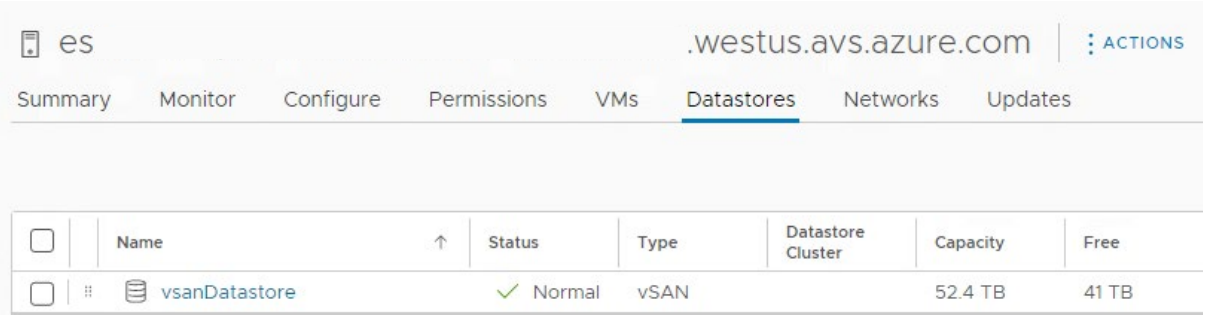

Figure 24 - ESXi Storage Capacity

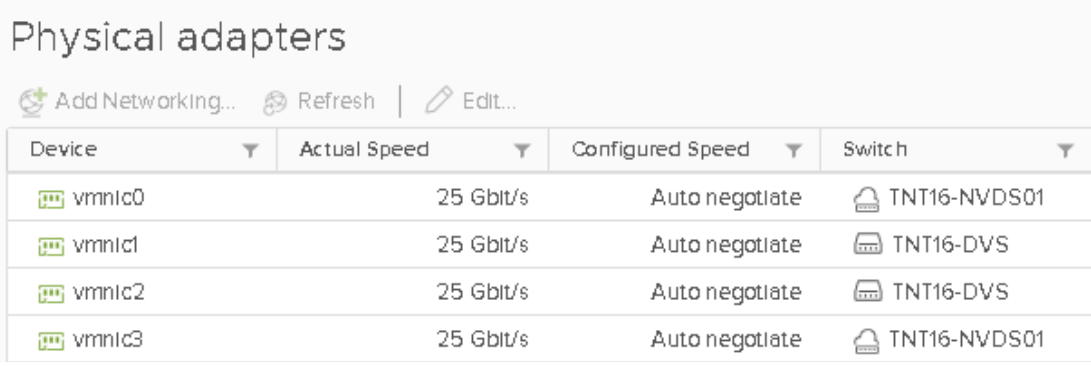

Figure 25 - ESXi Physical Network Capacity

© 2024 Lightbits Labs™

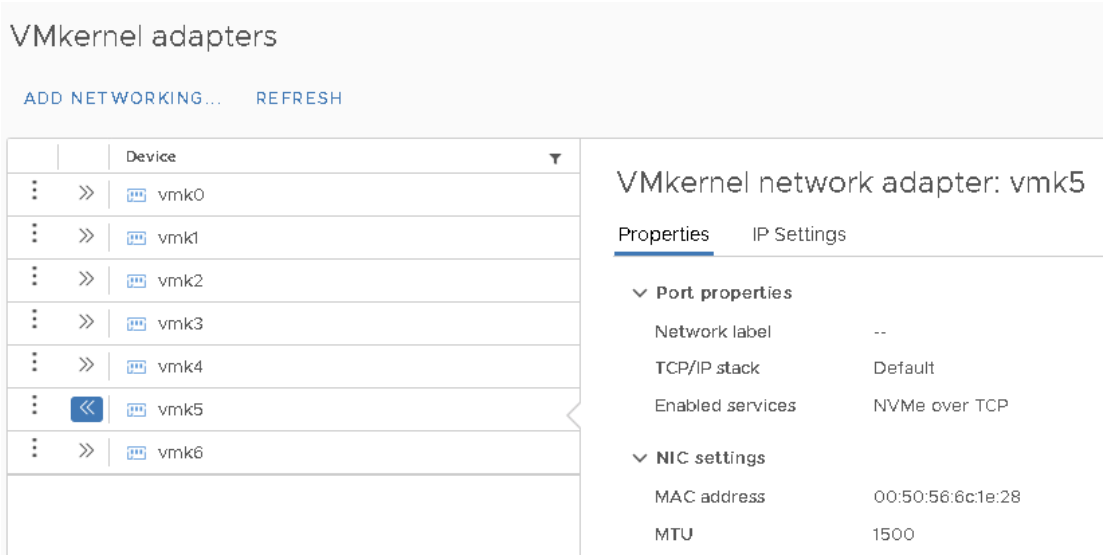

Figure 26 - ESXi NVMe VMKernel Adapters

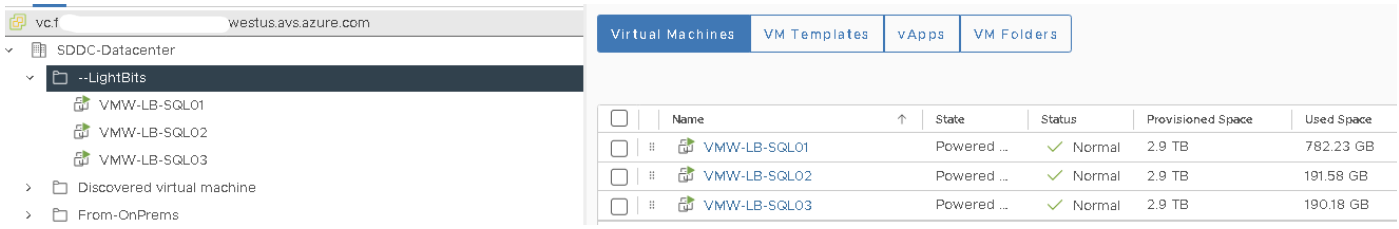

Figure 27 - Lightbits-Hosted Datastores

© 2024 Lightbits Labs™

SQL Server Client Configuration

There are three (3) VMs in our configuration, each running Windows Server 2022, as shown below:

Figure 28 - Windows Server VM

Each VM is allocated the following compute, virtual disks and SCSI controller resources:

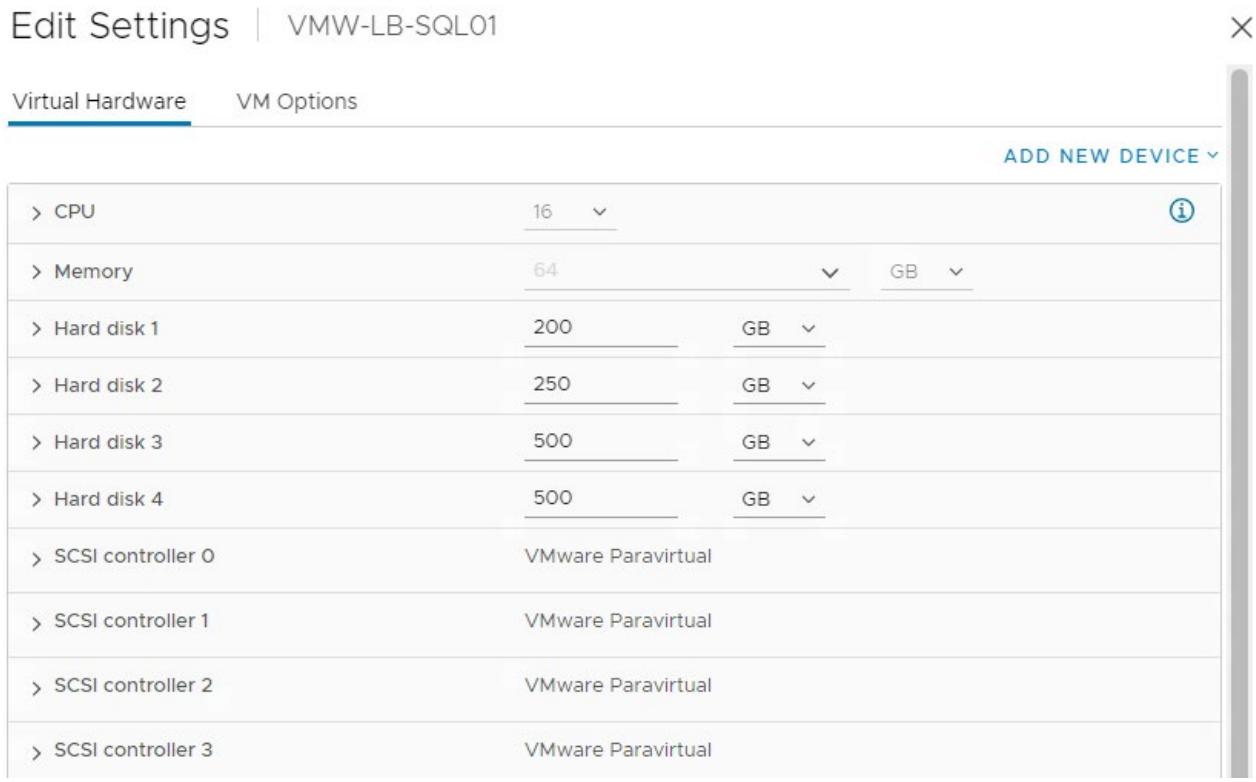

Figure 29 - Virtual Resource Allocation Per VM

© 2024 Lightbits Labs™

NOTE: To ensure parallelized and improved data IO throughput, each disk is connected to a separate virtual SCSI controller

The table below describes the size, sources and use of each allocated virtual disk.

Drive Letter	Capacity	Use	Source
C	200GB	OS Install	vSAN Datastore
D	250GB	SQL Data	Lightbits
L	500GB	SQL Logs	Lightbits
T	500GB	SQL TempDB and Backup	Lightbits

Table 3 - Disk Allocation and Use

The following is a view of the allocated disks within the Windows OS.

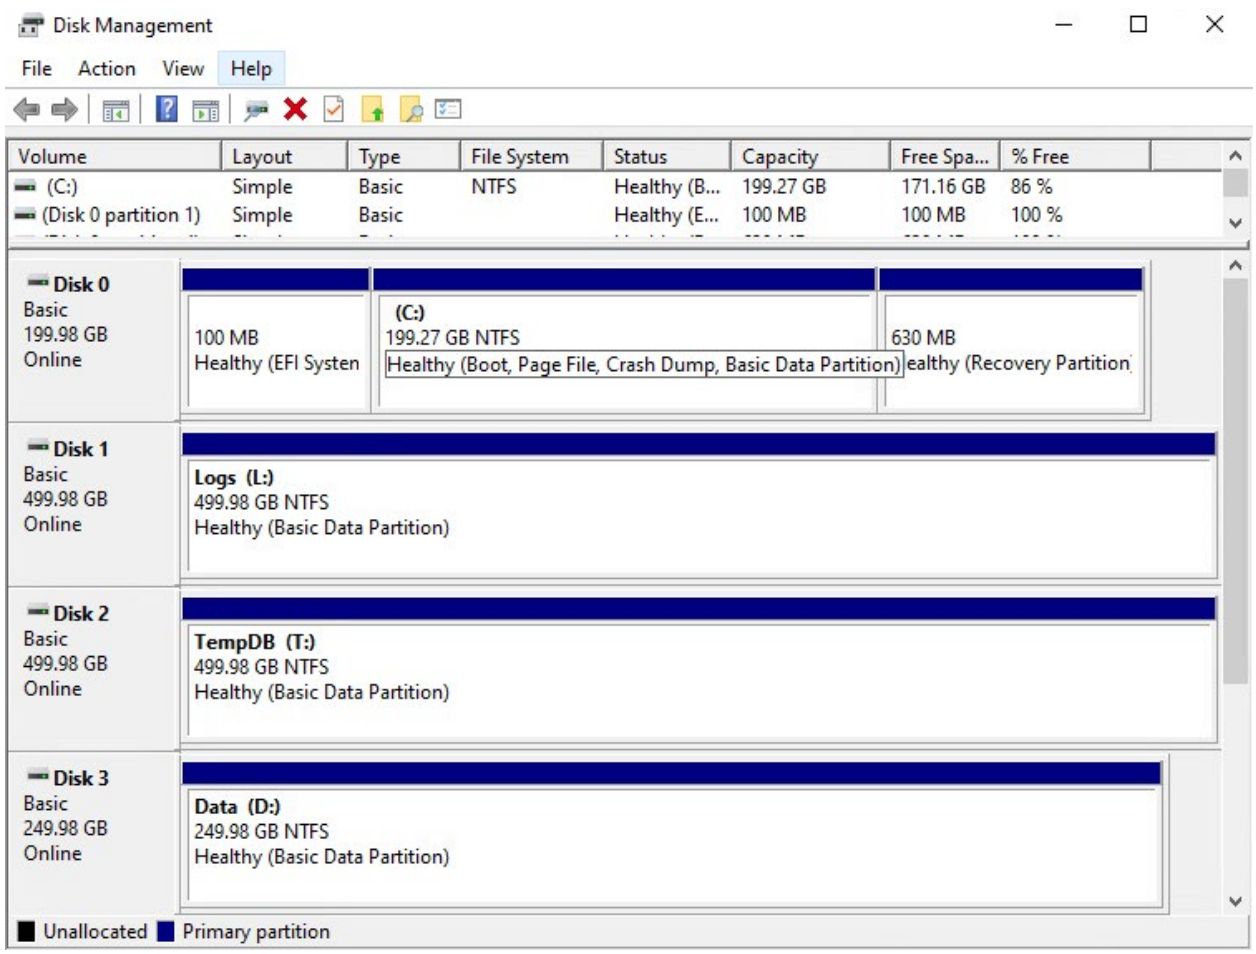

Figure 30 - In-Guest Disk Configuration

Database Setup

We installed SQL Server 2022 on each VM. Other than separating each SQL Server-related file type into their individual disk drives, no other optimization or performance-tuning exercise was performed on the installation for our initial test.

Figure 31 - Installed Microsoft SQL Server Version

© 2024 Lightbits Labs™

Performance Testing

Test Methodology

We imported an enlarged copy of the popular Adventure Works Sample database into each SQL Server instance on each VM. We then proceeded to perform a series of long-running insert operations into the sample database to observe the impact of these operations on the disk subsystem. We focused specifically on the following three metrics:

- Time taken to complete the insert operations
- Average disk IO throughput
- Disk response time (as observed in Windows)

Test Results

We have a general baseline and expectation with which to compare our observations, although a "performance bake-off" was not a focus of our exercise. Therefore, the following images are presented for academic purposes only and to enable you to use as input and guidance for your own internal testing as you evaluate the presented solution.

Records Insert Operation

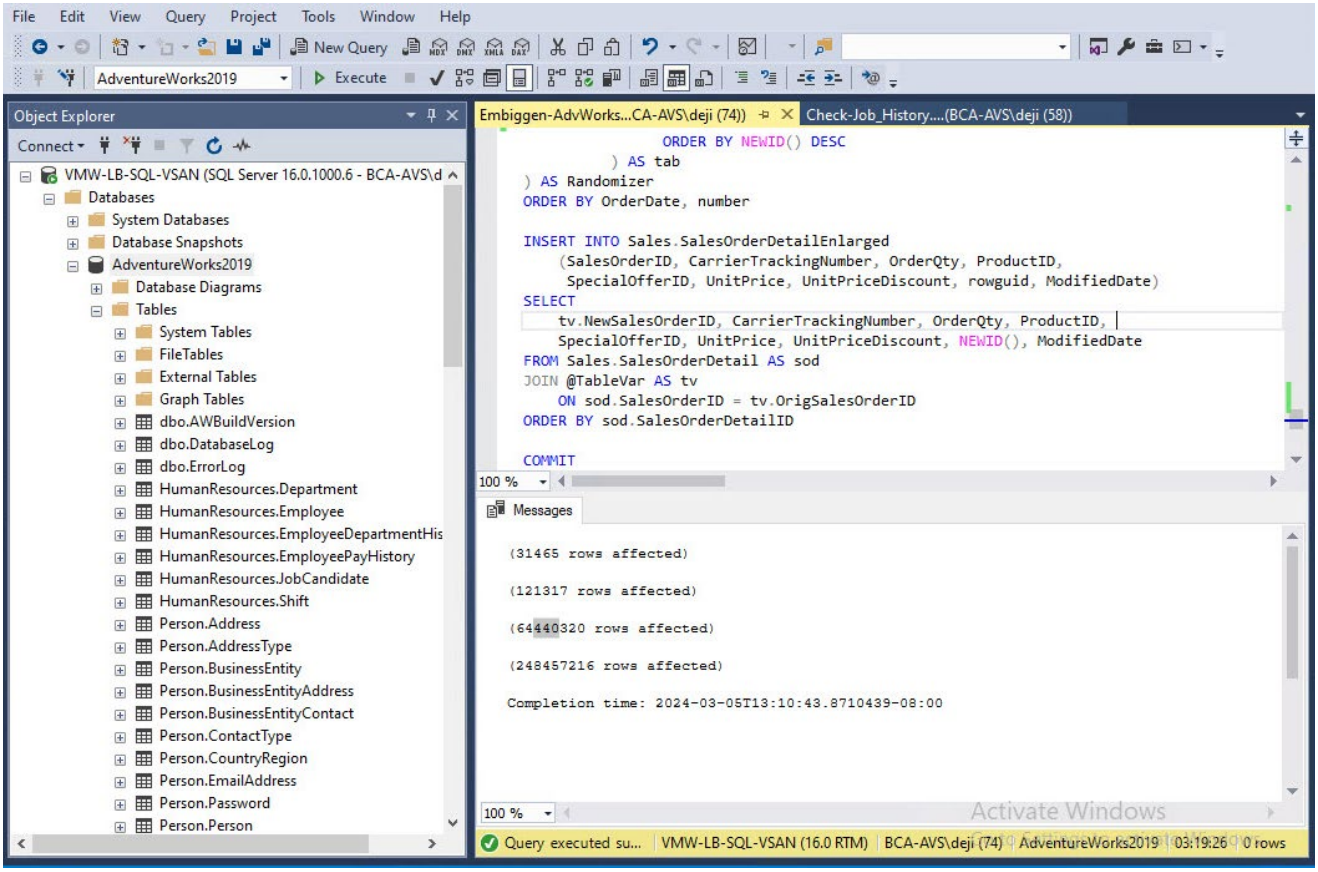

Figure 32 - Output From Records Insert Operation

Our database and logs restore operation returned the following results:

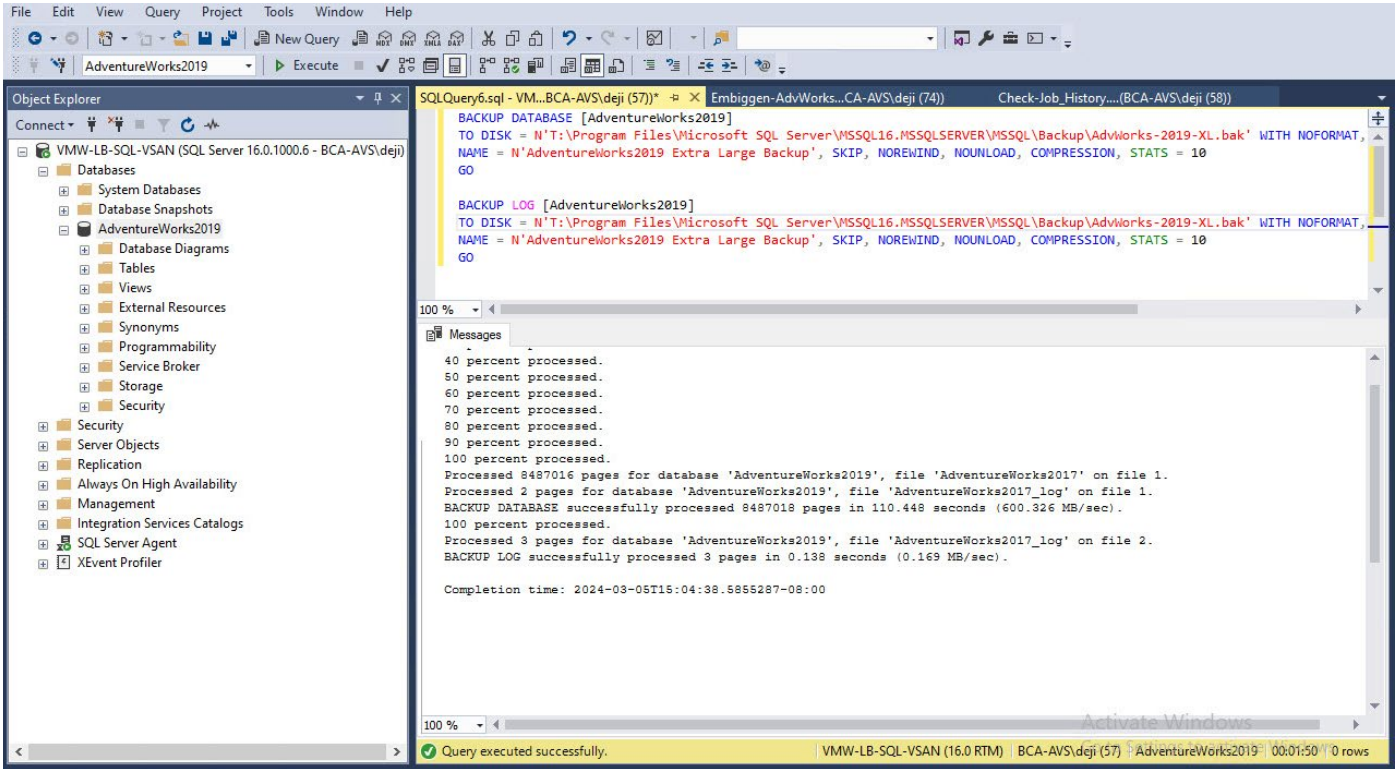

Figure 33 - Result of Database Backup and Restore Operation

Resource Monitor Metrics

For each of these test operations, we generally observed the following metrics in Windows Resource Monitor:

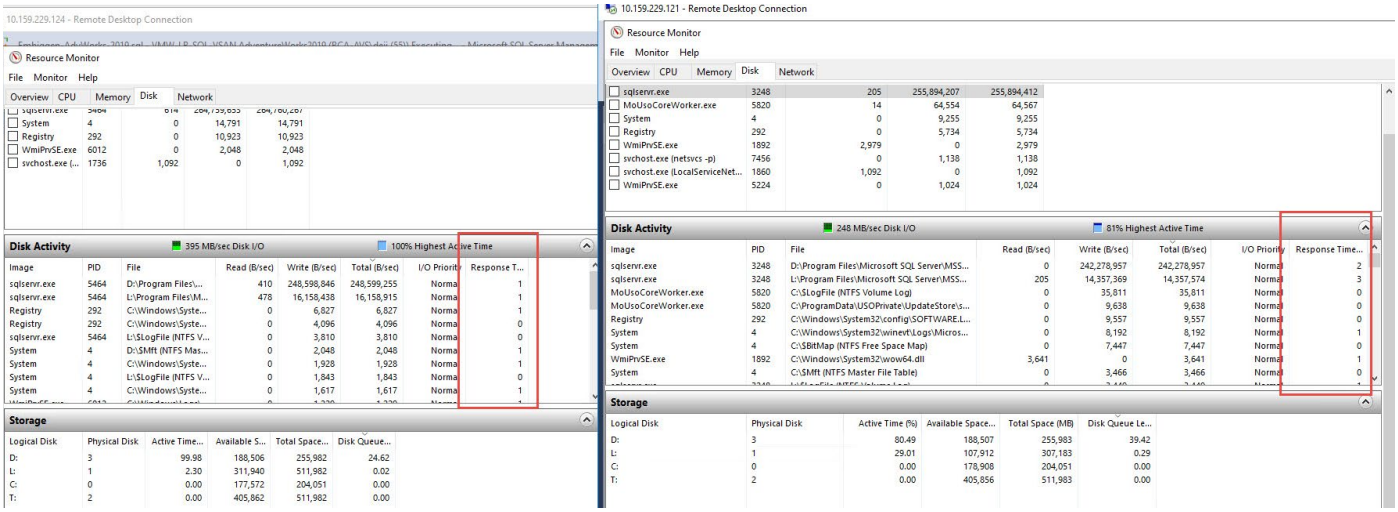

Figure 34 - Write Throughput and Latency During Data Inserts

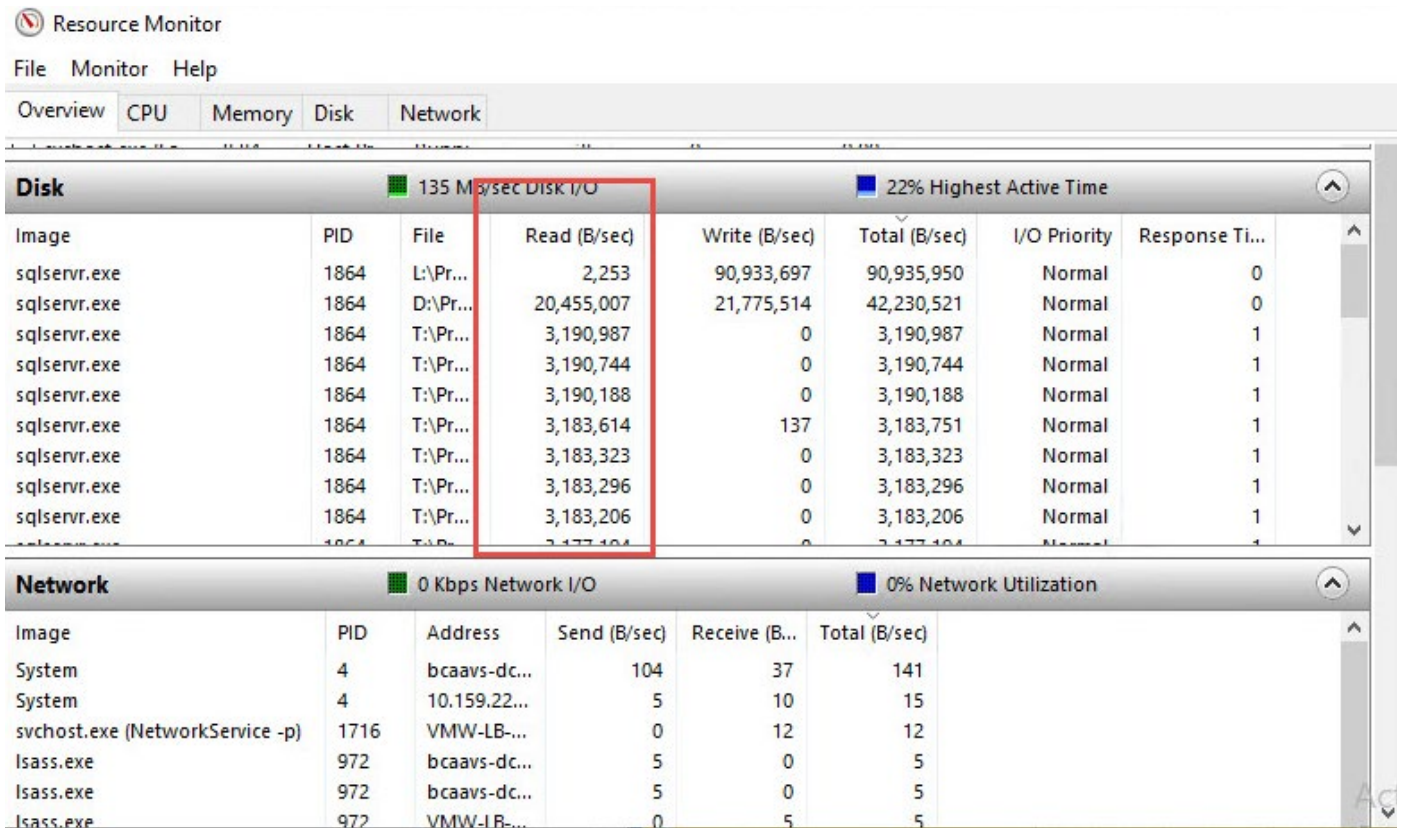

Figure 35 - Mixed Read/Write Throughput and Latency During Restore Operation

In addition to the above, we also conducted a series of synthetic storage IO performance tests using the popular Microsoft DiskSPD tool. As shown in these series of graph images, our tests simulated multiple usage scenarios with varying permutations of:

- File Size
- Read-Write Ratio
- Read-Write Type
- Block Size
- Outstanding IOs
- CPU Threads

We were focused on monitoring the following metrics:

- Latency
- CPU Loads
- IOPS
- Throughput

© 2024 Lightbits Labs™

Initial Test Results

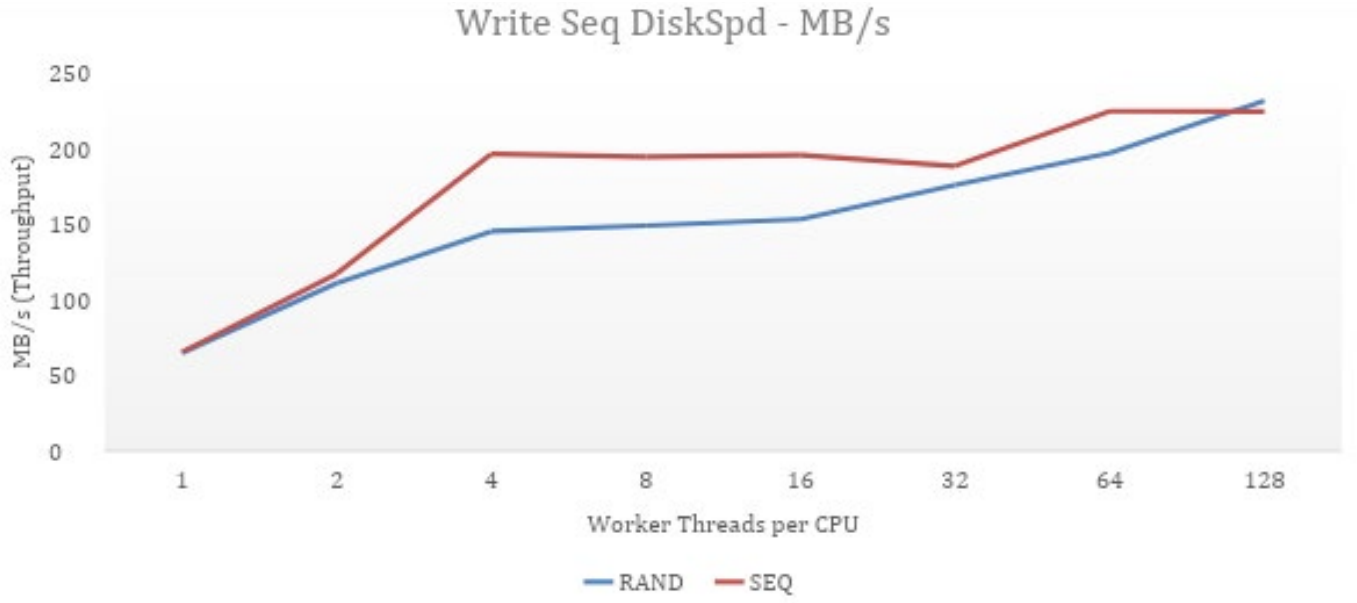

Figure 36 - Overall Write Throughput

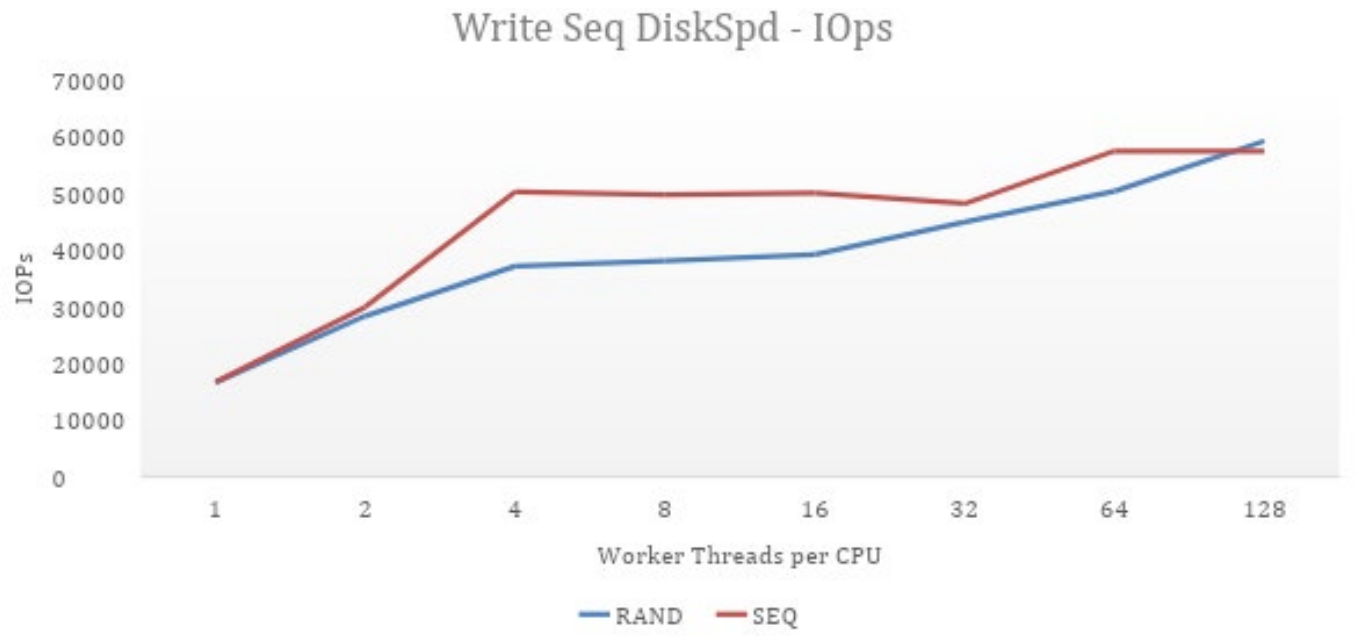

Figure 37 - Overall Write IOPs

© 2024 Lightbits Labs™

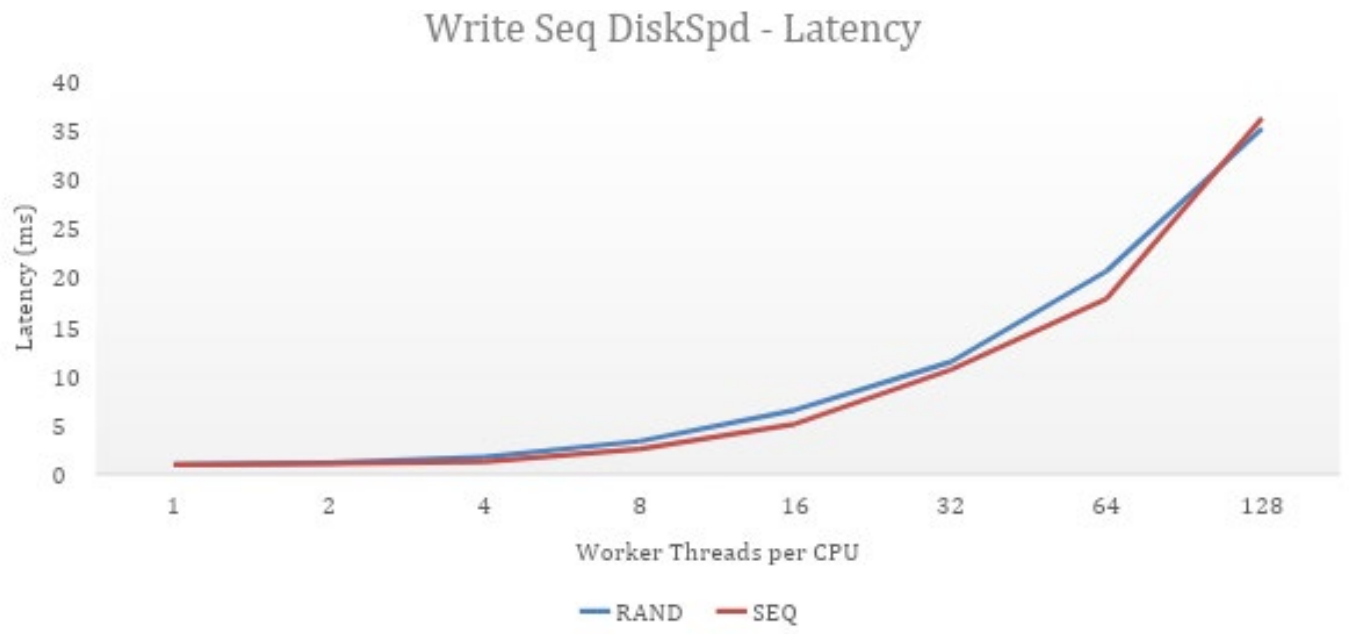

Figure 38 - Overall Write Latency

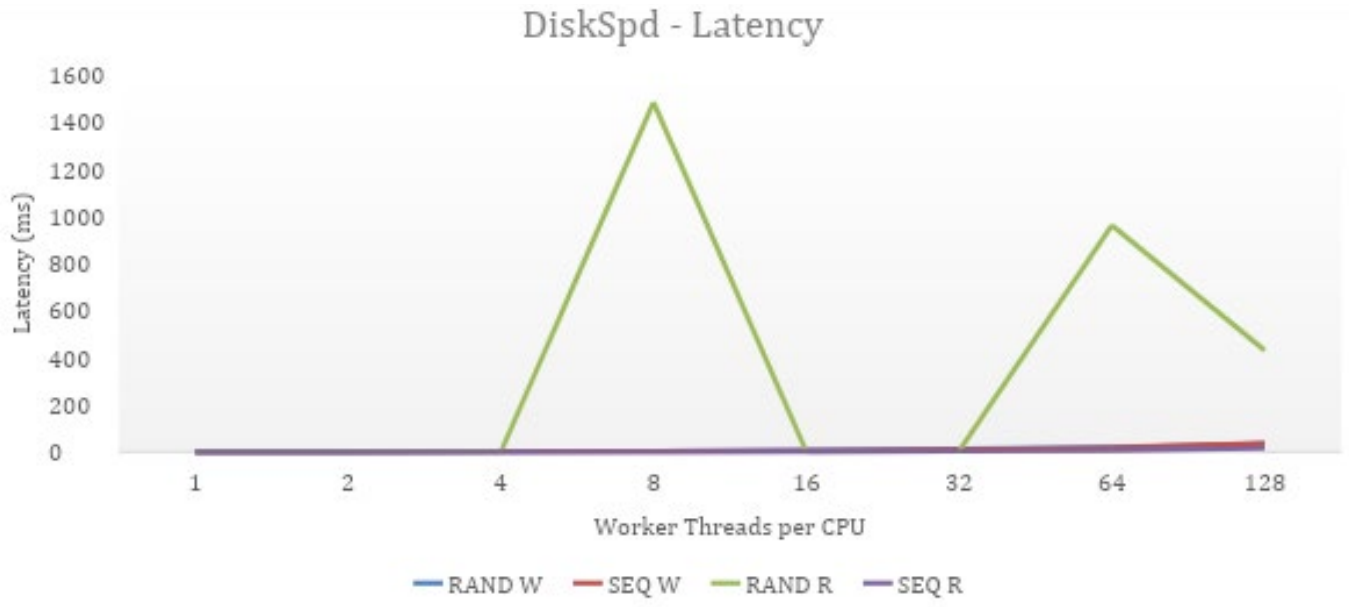

Figure 39 - Overall Disk Latency

© 2024 Lightbits Labs™

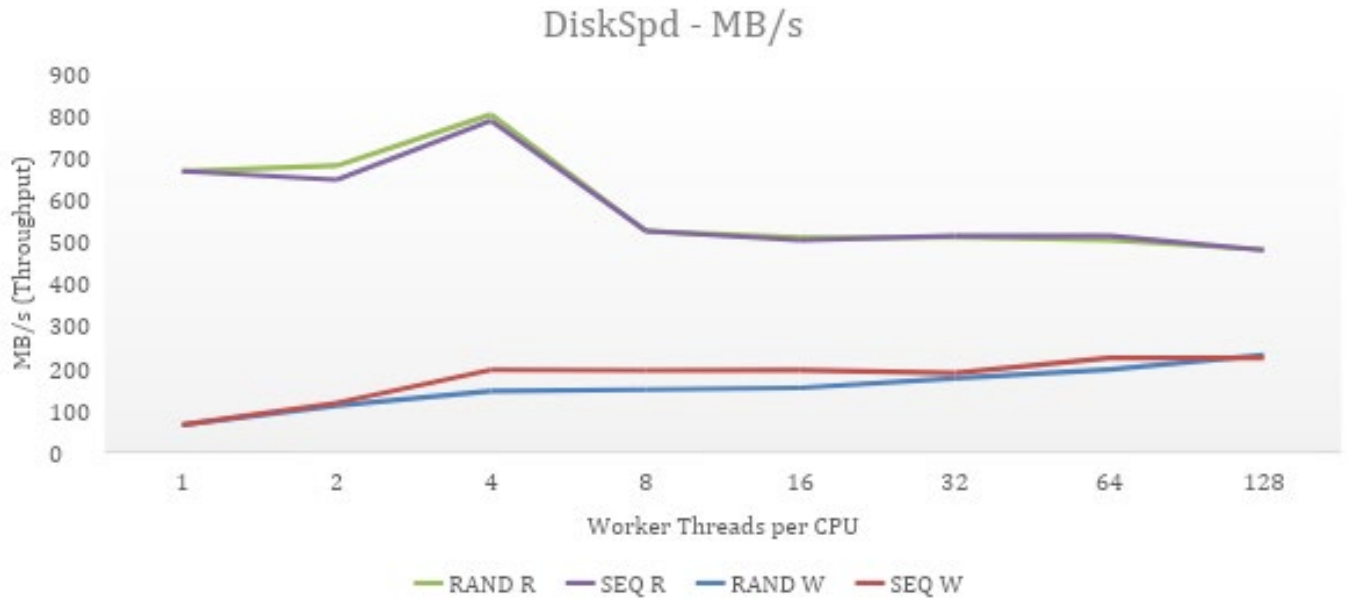

Figure 40 - Overall Disk Throughput

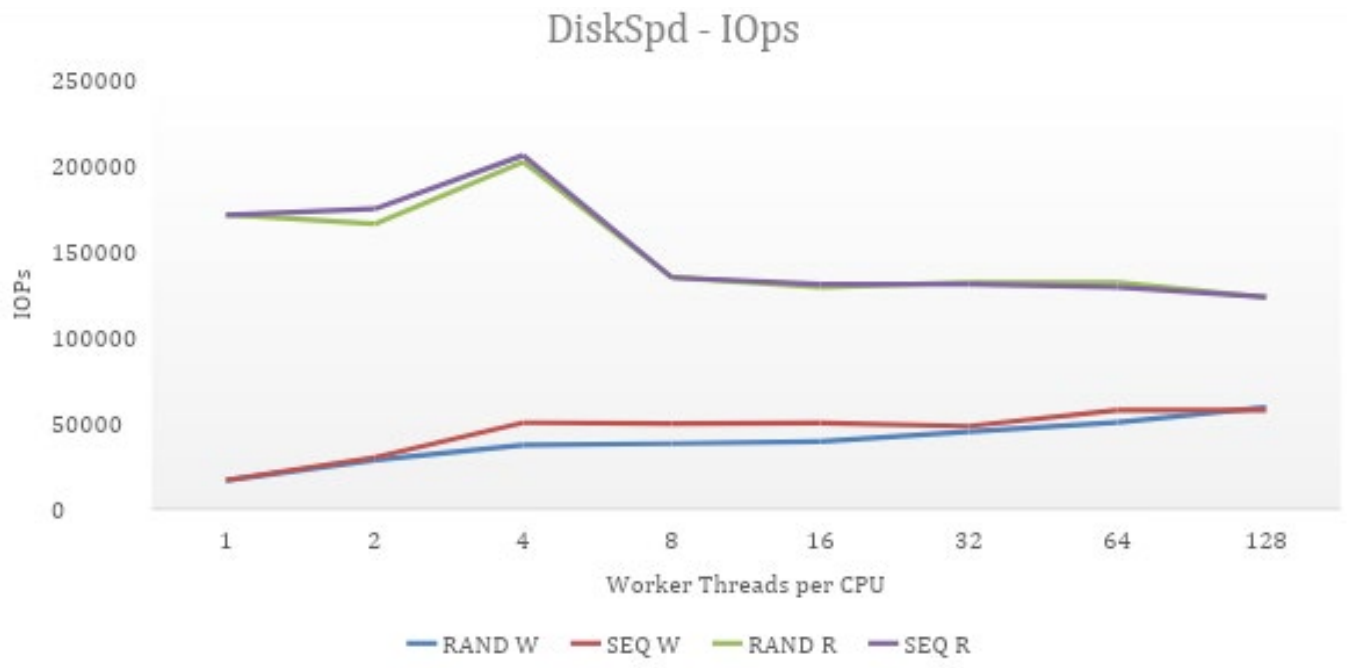

Figure 41 - Overall Disk IOPS

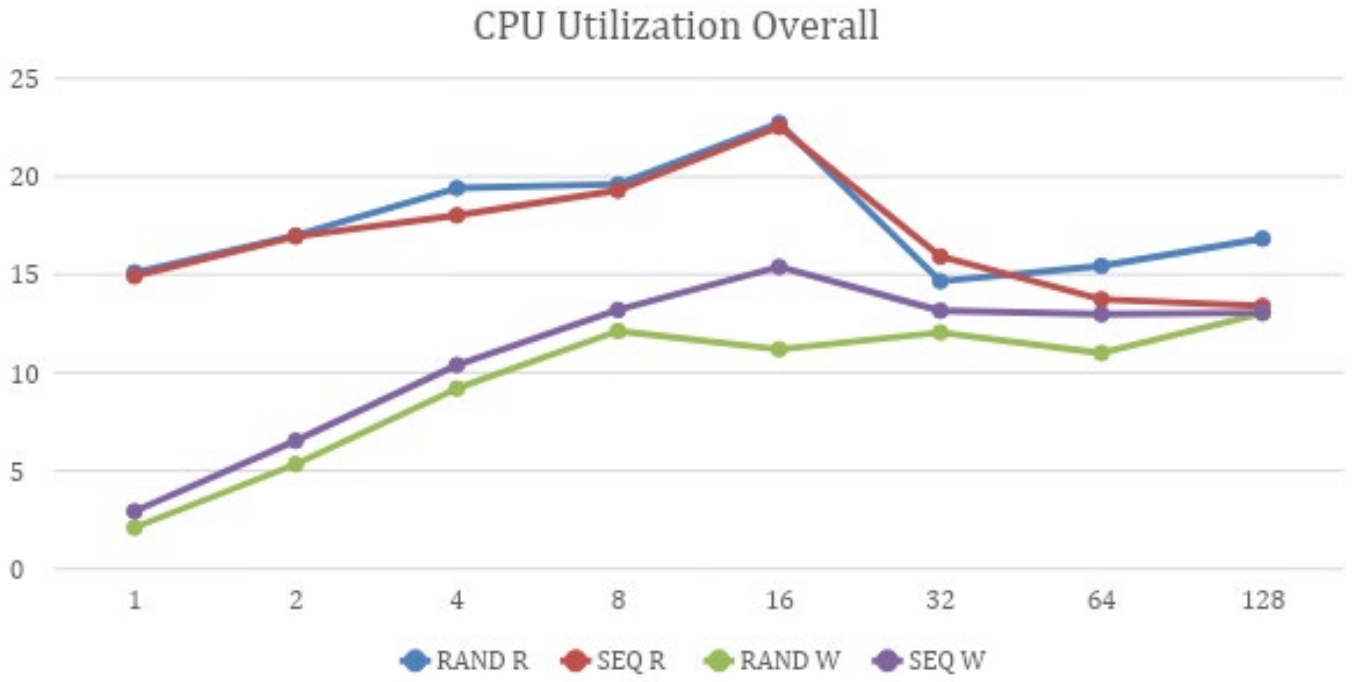

Figure 42 - Overall CPU Utilization

Optimized Test Results

Optimal storage IO performance is a critical metric for any enterprise-class, mission-critical SQL Server workload. This explains why almost every vSphere configuration best practices reference architecture/guide includes guidance on configurations that help improve IO throughputs and avoid bottleneck. Two such guides are: Architecting Microsoft SQL Server on VMware vSphere and Successfully Virtualizing Microsoft SQL Server for High Availability on Azure VMware Solutions.

As you have noticed, we created multiple VMDKs for our SQL Server VMs. We separated these VMDKs into multiple datastores and attached them to multiple vSCSI controllers. All of these choices are in line with VMware's standard (and indispensable) recommendations for optimizing storage performance for each VM by ensuring that IOs generated by the SQL Server application within an operating system have multiple parallel paths through which they traverse on their way to the storage subsystem.

Each device through which an IO travels has a finite capacity for how much IO it can service at any given point in time. This capacity is called the "queue depth" of that device. In absolute technical terms, a device's queue depth is the number of pending input/output (I/O) requests that a storage resource can handle at any one time. Anything beyond this number will be held in the device's queue until preceding IOs have been serviced. If IOs are held back for long as a result of competition for a device's limited capacity, SQL Server transactions experience degraded performance.

> "The depth of the queue of outstanding commands in the guest operating system SCSI driver can significantly impact disk performance. A queue depth that is too small, for example, limits the disk bandwidth that can be pushed through the virtual machine. See the driver-specific documentation for more information on how to adjust these settings"
>
> "In some cases large I/O requests issued by applications in a virtual machine can be split by the guest storage driver. Changing the guest operating system's registry settings to issue larger block size I/O requests can eliminate this splitting, thus enhancing performance. For additional information see VMware KB article 9645697"
>
> "If your storage subsystem uses 4KB native (4Kn) or 512B emulation (512e) drives, you can obtain the best storage performance if your workload issues mostly 4K-aligned I/Os. For more information on this subject, see VMware KB article 2091600 or the "Device Sector Formats" subsection of Viewing Storage Devices Available to an ESXi Host in the vSphere Storage Guide".
>
> — "Guest Operating System Storage Considerations" Performance Best Practices for VMware vSphere

© 2024 Lightbits Labs™

Based on the tuning recommendations above, we adjusted the VMware Paravirtual SCSI Adapter's (PVSCSI) queue depth in each of our VMs, as described in "Large-scale workloads with intensive I/O patterns might require queue depths significantly greater than Paravirtual SCSI default values (2053145)." This was done to obtain the performance metrics shown below:

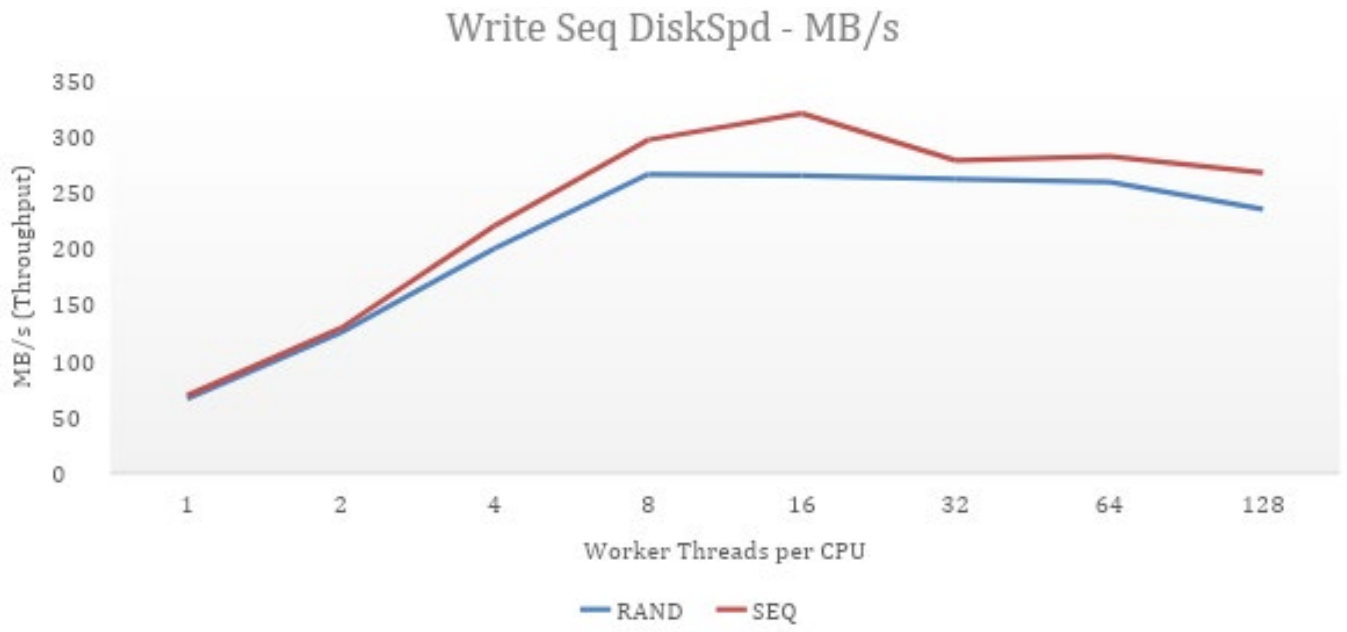

Figure 43 - Write Throughput - After Queue Depth Increase

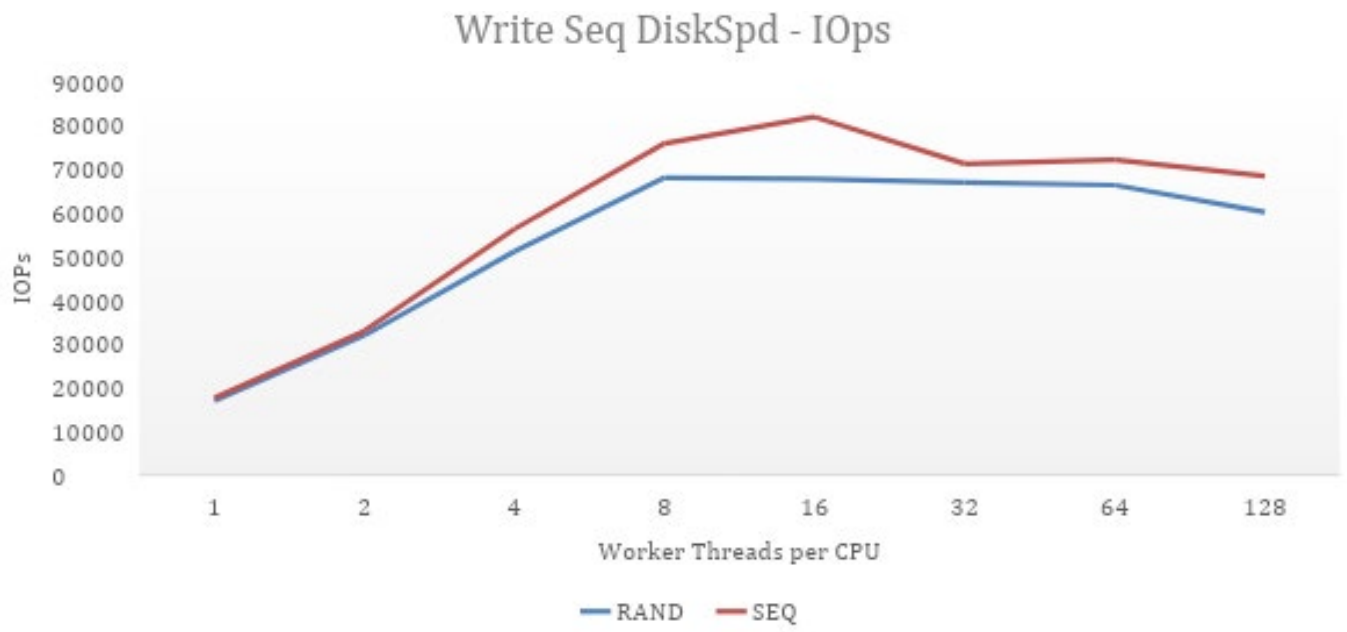

Figure 44 - Write IOPS - After Queue Depth Increase

© 2024 Lightbits Labs™

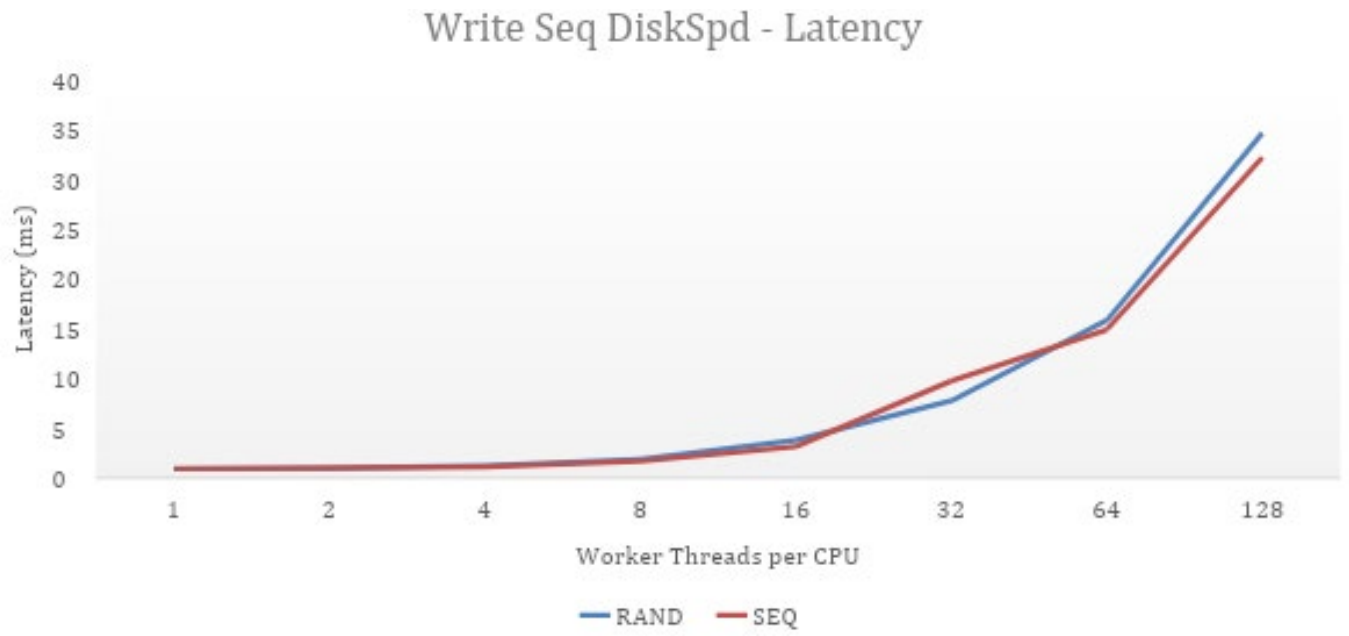

Figure 45 - Write Latency - After Queue Depth Increase

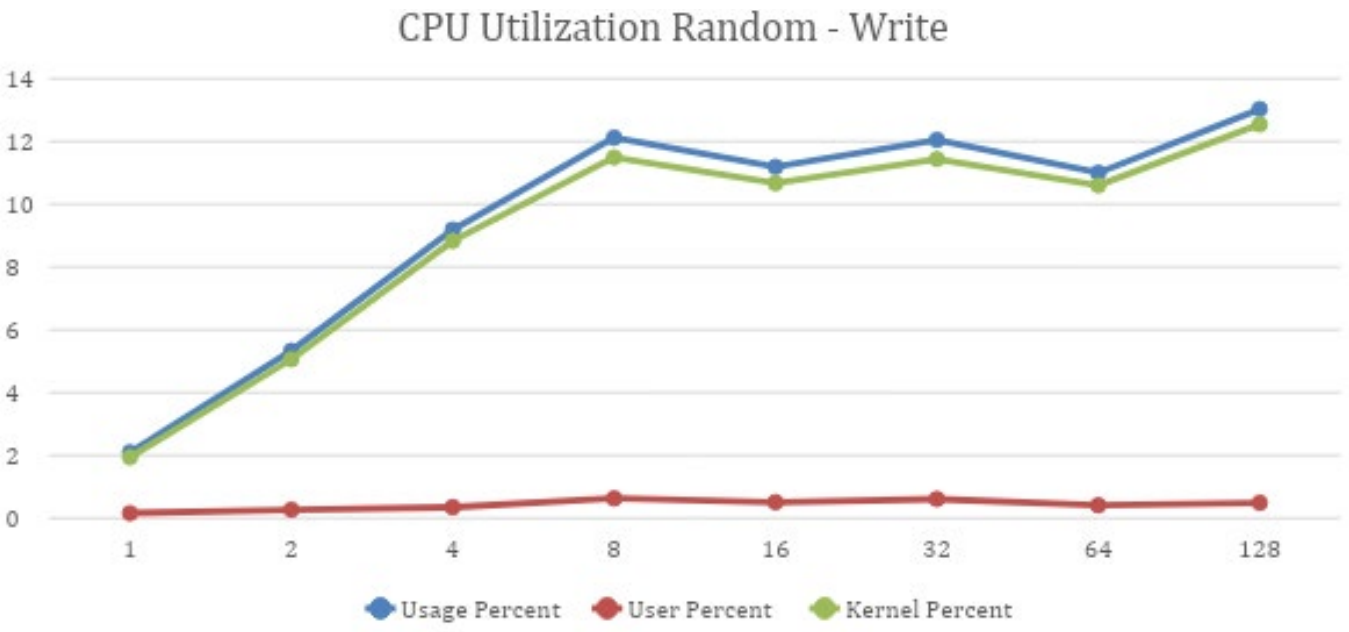

Figure 46 - Random Write CPU Utilization - After Queue Depth Increase

© 2024 Lightbits Labs™

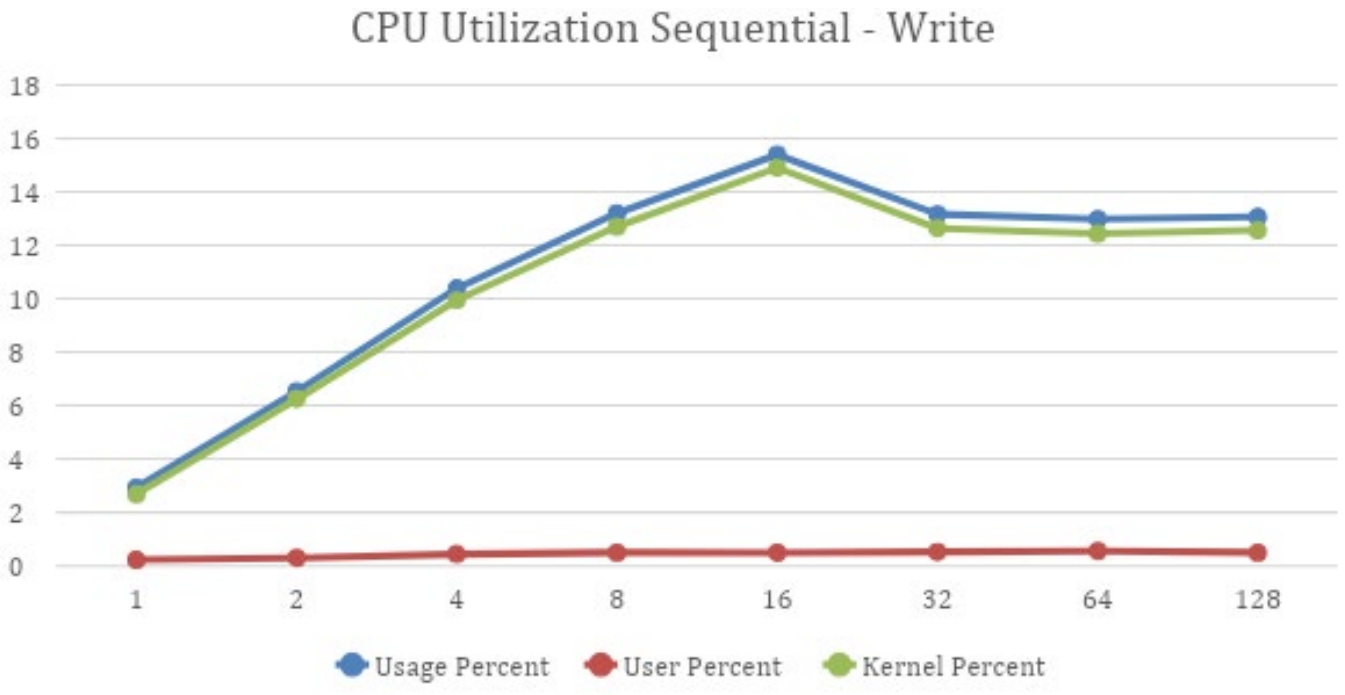

Figure 47 - Sequential Write CPU Utilization - After Queue Depth Increase

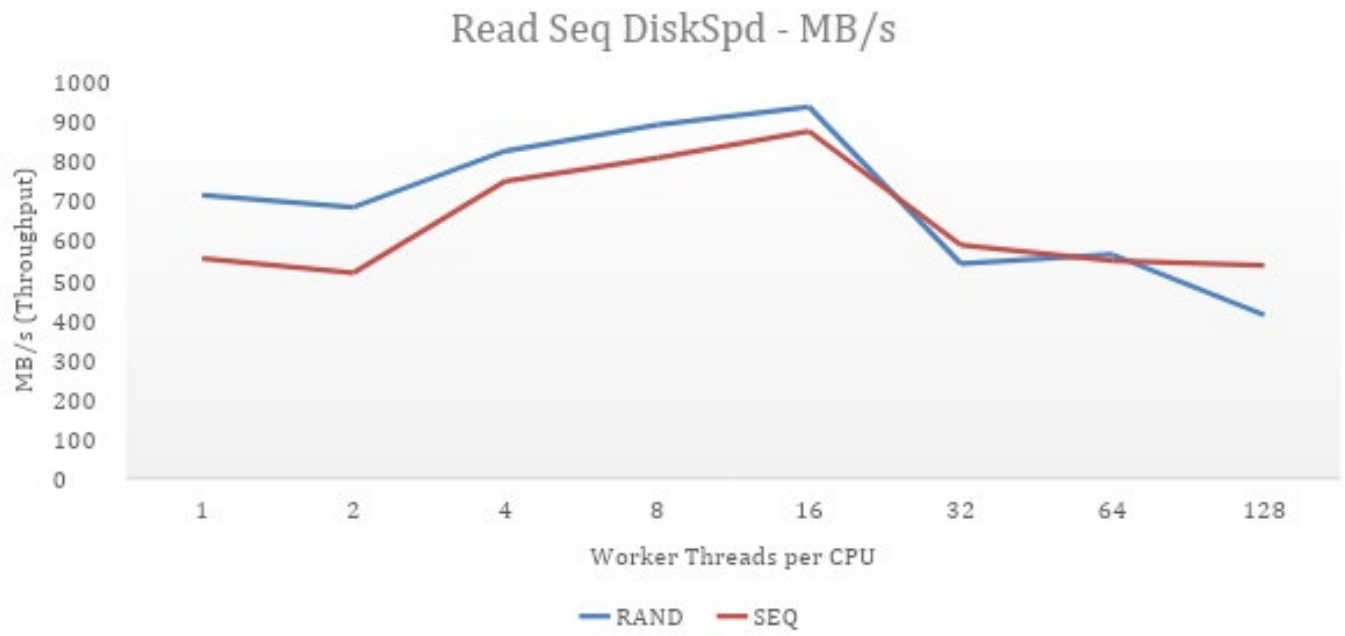

Figure 48 - Read Throughput - After Queue Depth Increase

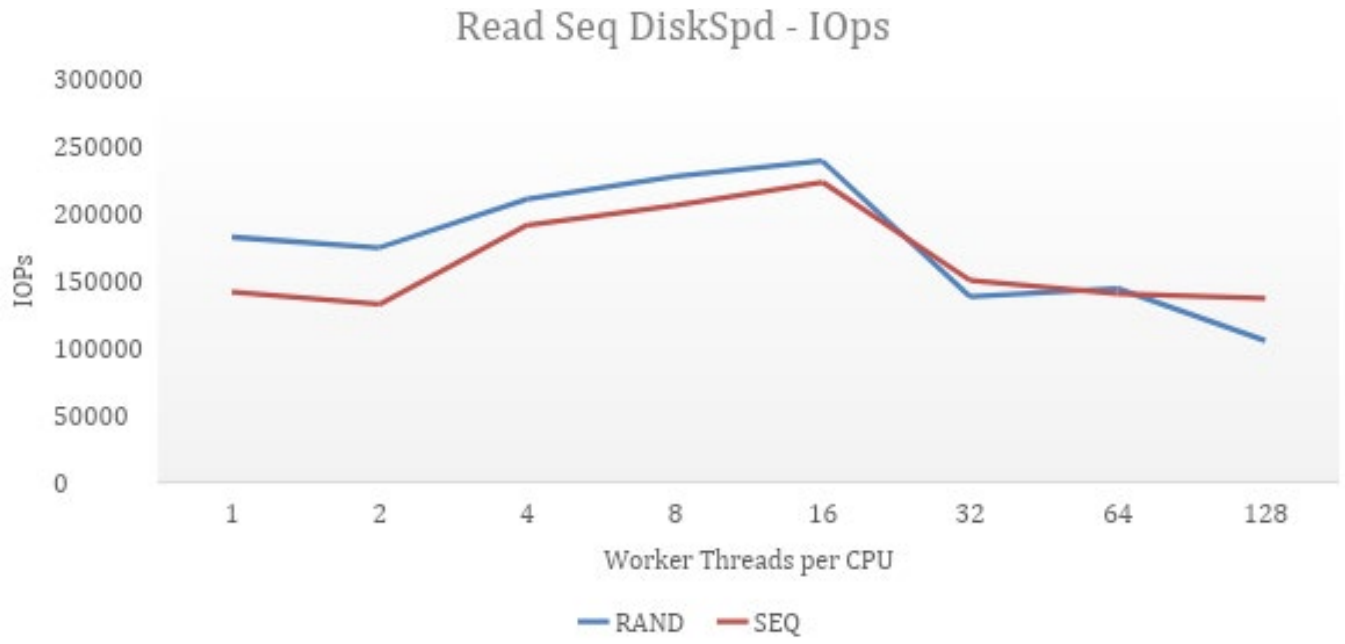

Figure 49 - Read IOPs - After Queue Depth Increase

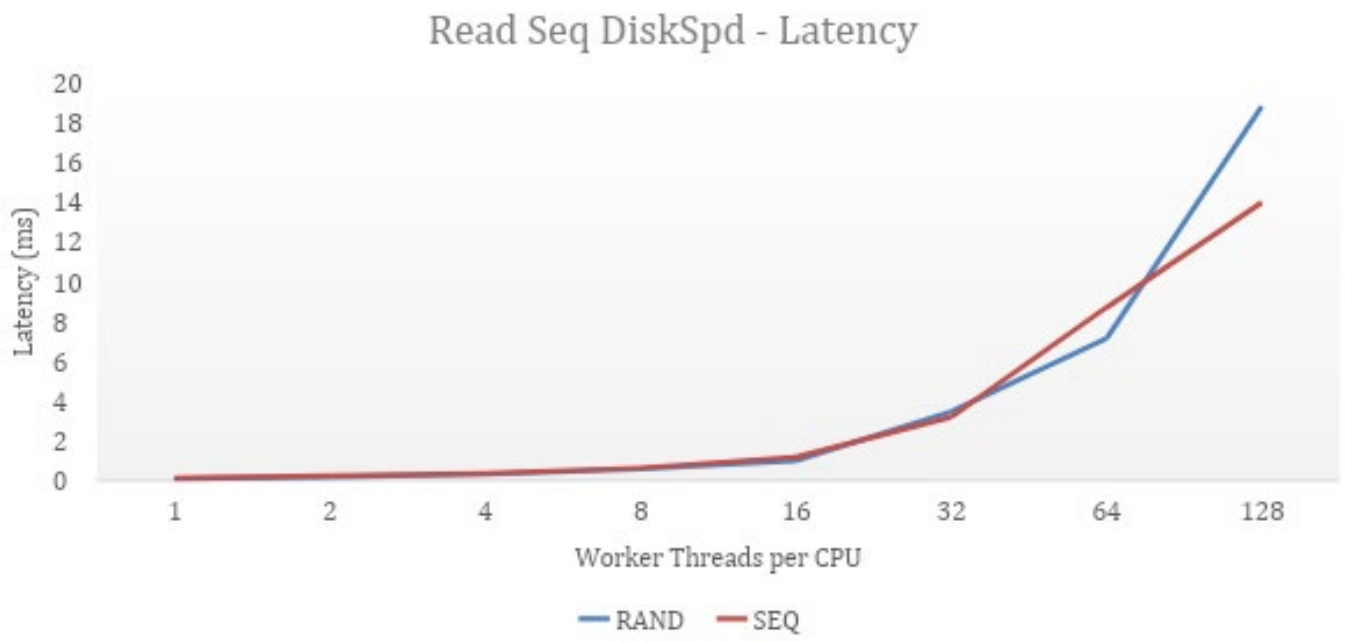

Figure 50 - Read Latency - After Queue Depth Increase

© 2024 Lightbits Labs™

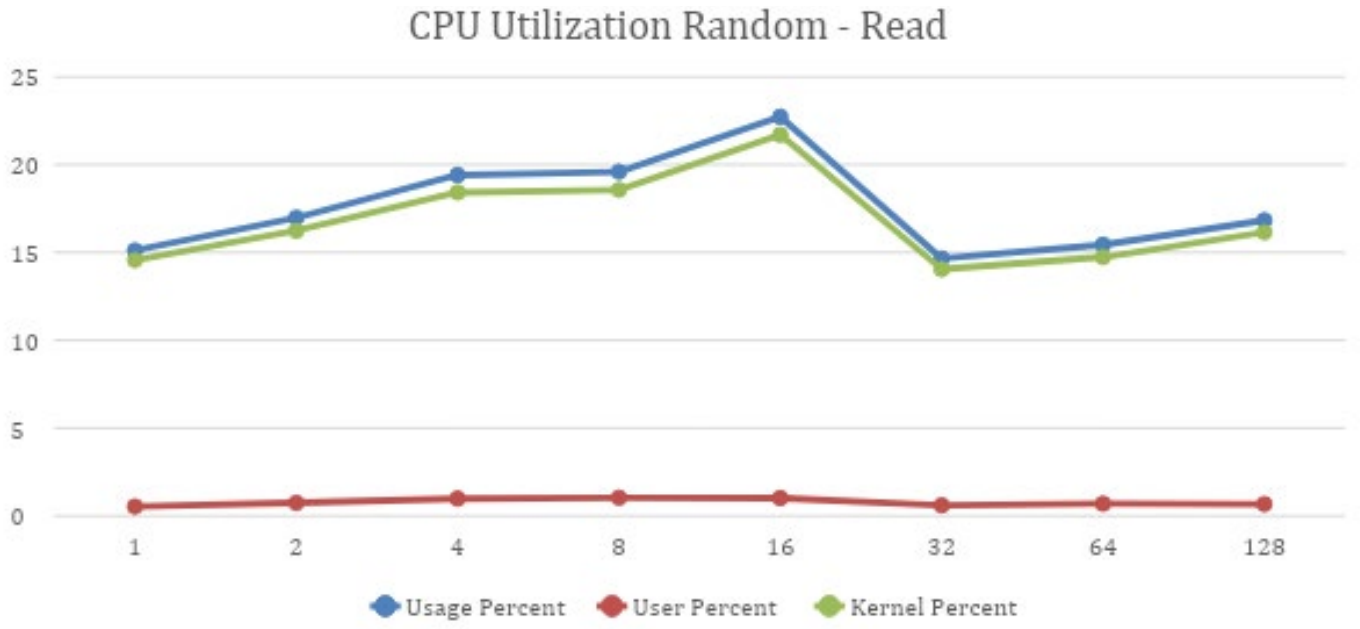

Figure 51 - Random Read CPU Utilization - After Queue Depth Increase

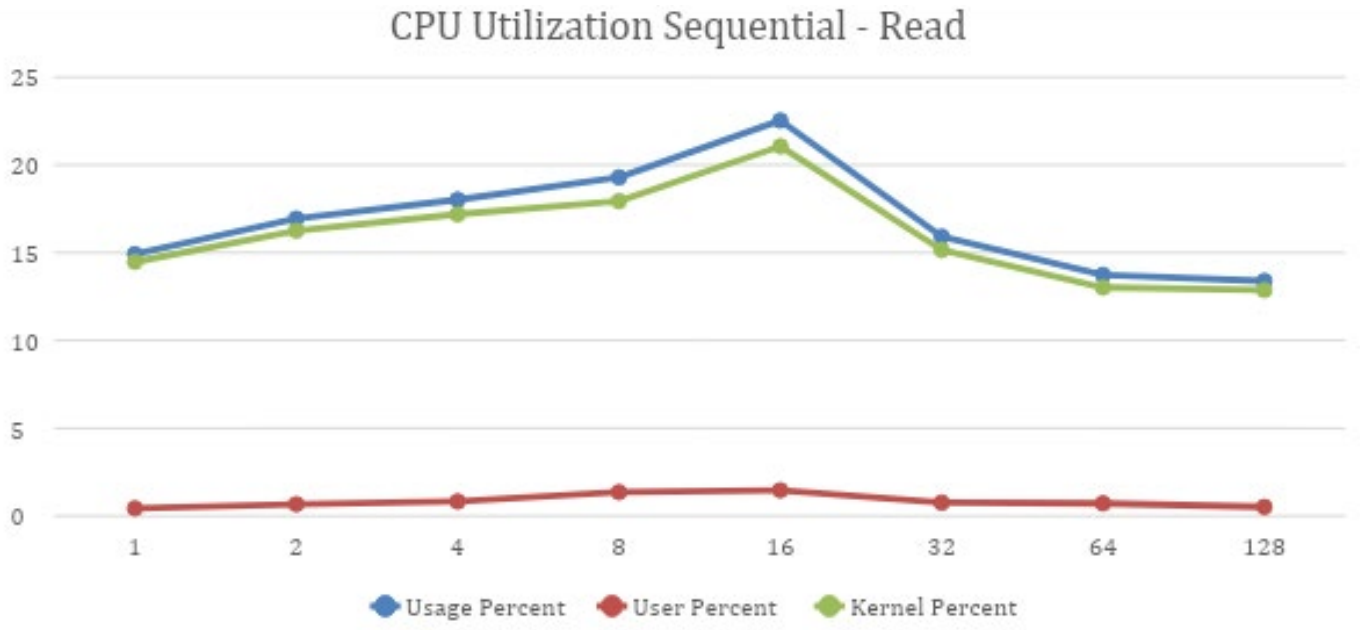

Figure 52 - Sequential Read CPU Utilization - After Queue Depth Increase

© 2024 Lightbits Labs™

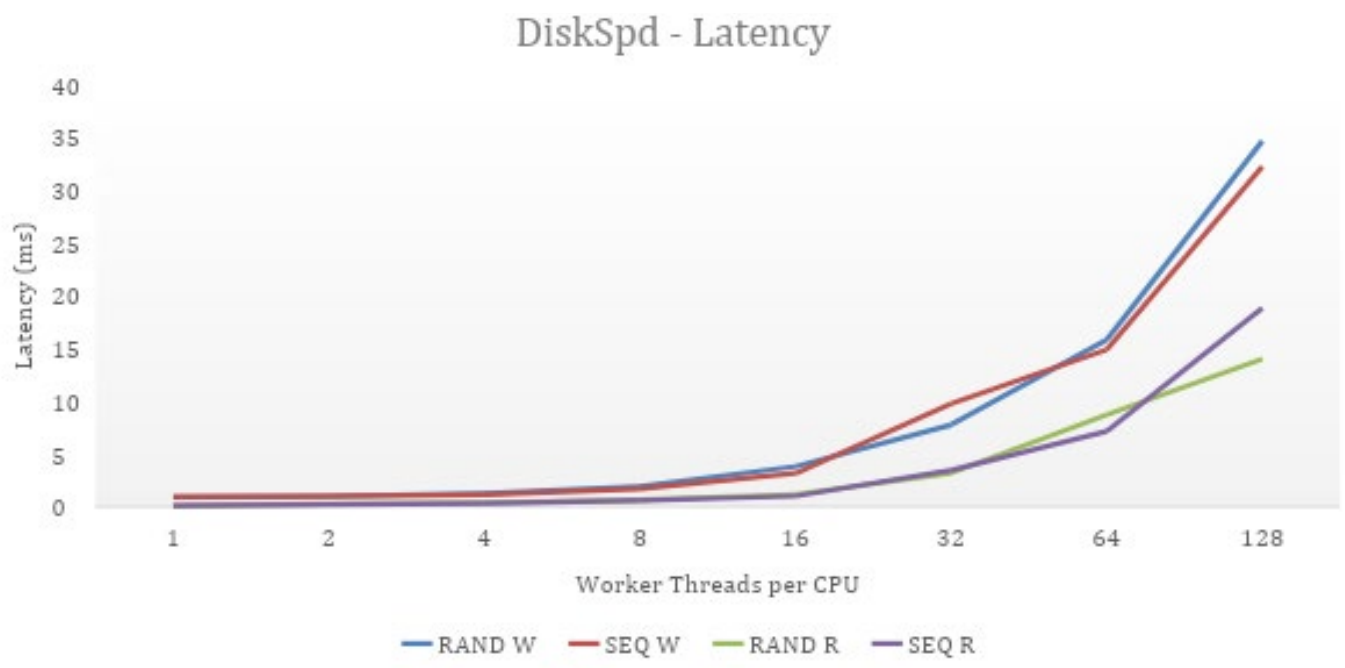

Figure 53 - Overall Latency - After Queue Depth Increase

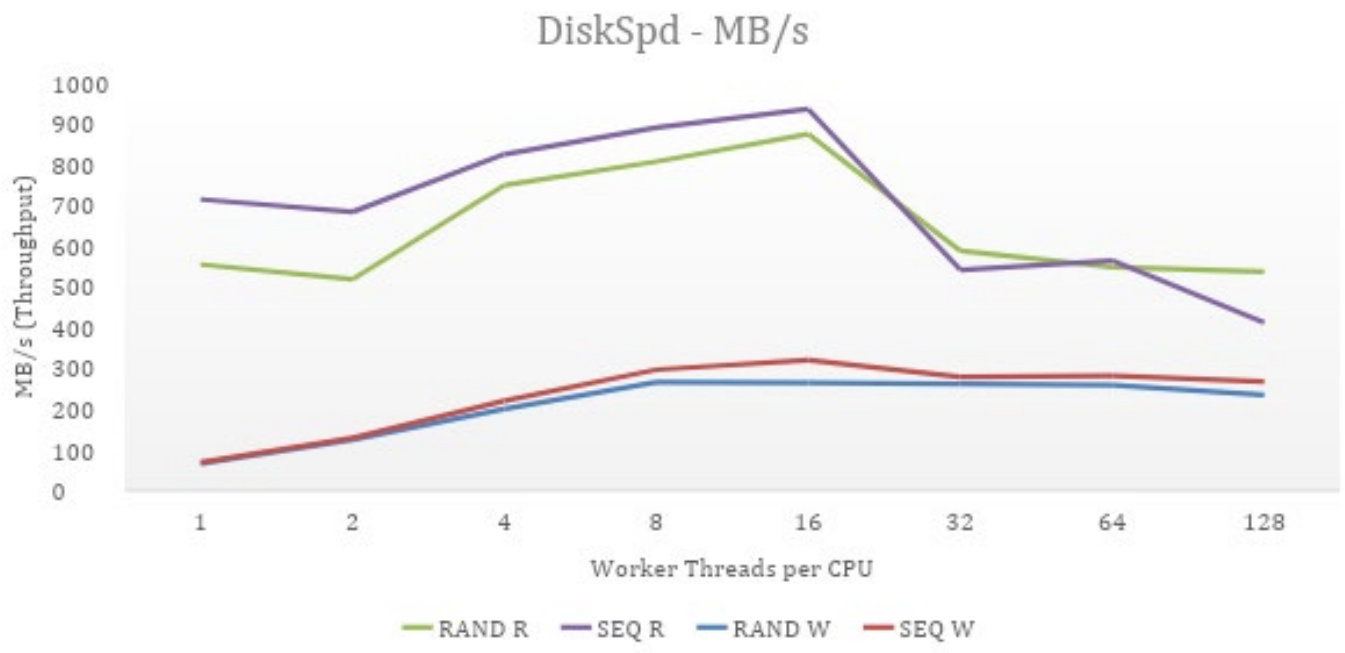

Figure 54 - Overall Throughput - After Queue Depth Increase

© 2024 Lightbits Labs™

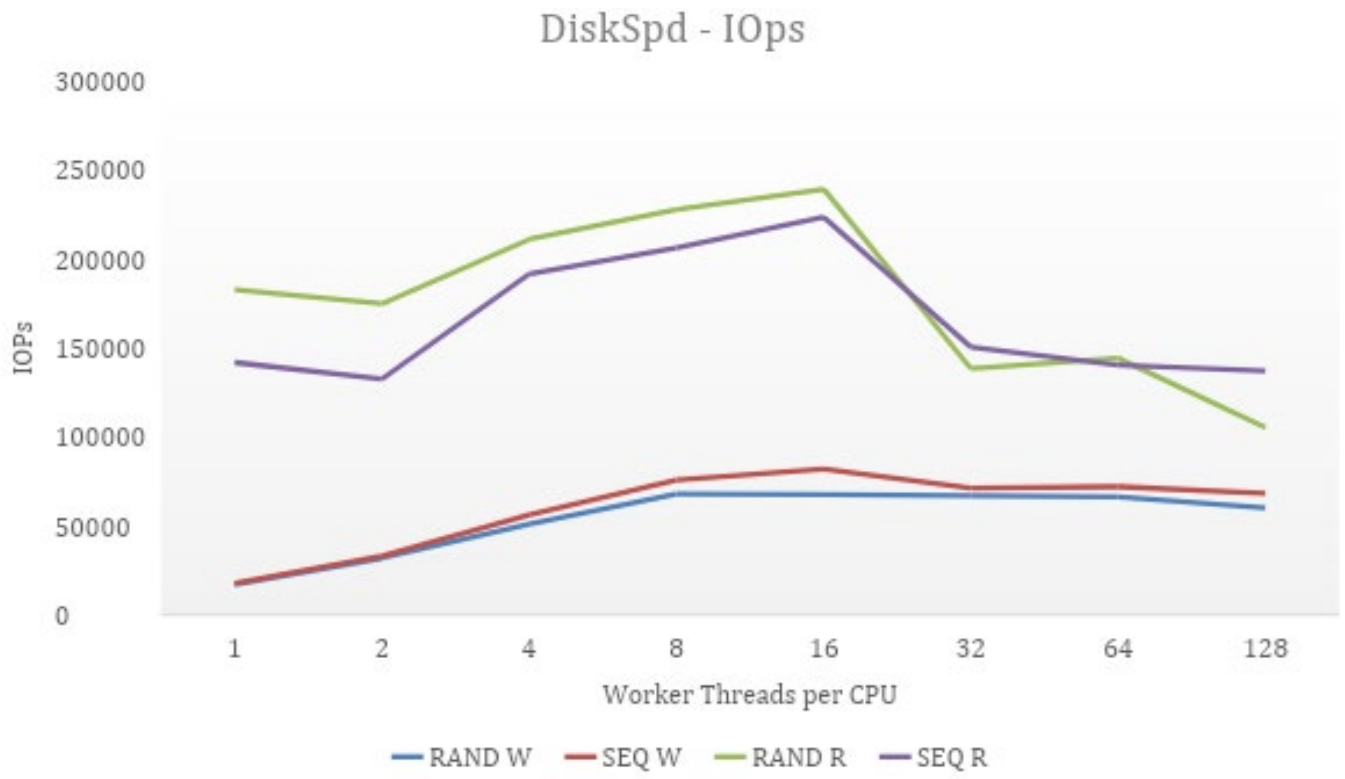

Figure 55 - Overall IOPs - After Queue Depth Increase

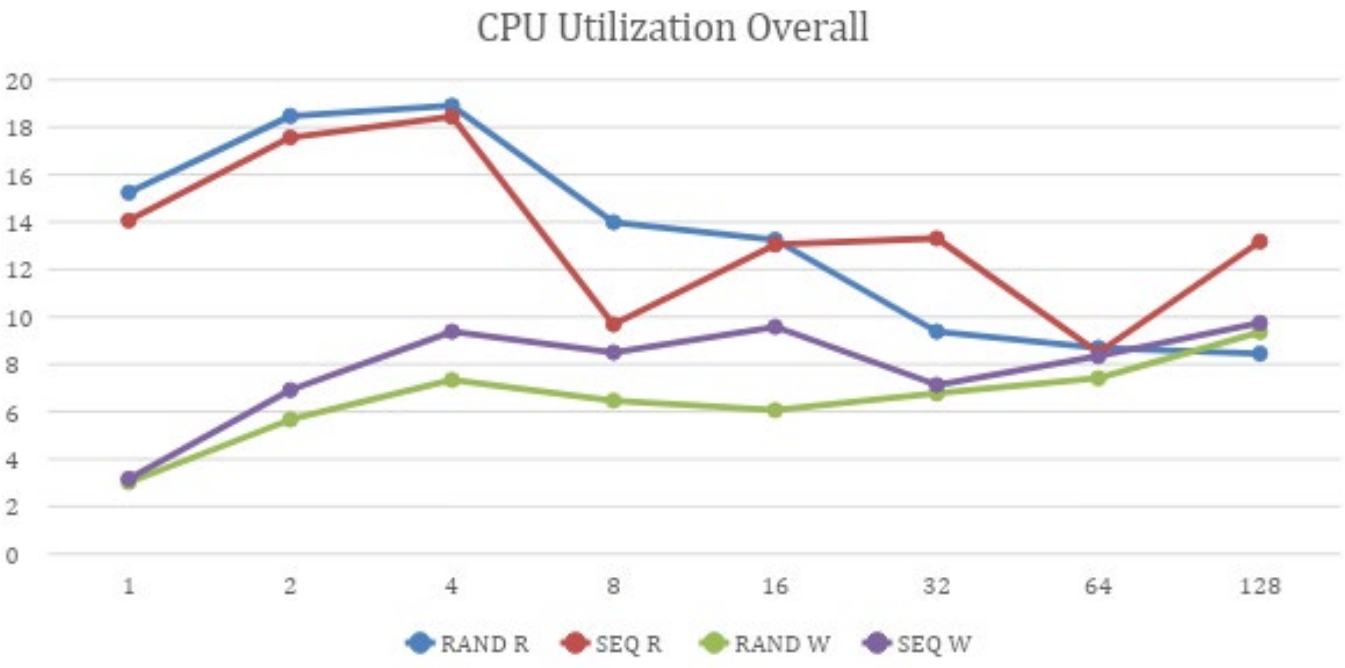

Figure 56 - Overall CPU Utilization - After Queue Depth Increase

© 2024 Lightbits Labs™

Other Benefits

As shown above in the performance testing and optimization graphs, Lightbits can provide the desired performance for SQL Server workloads on AVS, as well as other benefits above and beyond raw performance.

Cost

As mentioned in the introduction, one of the key motivations for using external disaggregated storage for AVS is to reduce the cost of deployments vs. adding ESXi hosts and causing underutilization of compute resources. Lightbits in Azure is deployed through the Marketplace and is charged relative to the VM size used for the cluster. This model enables you to start with a relatively small three-node cluster, and scale incrementally as your workloads do.

Due to being deployed as a set of VMs inside your Azure subscription, you can also take advantage of standard Azure discounts with reserved instances as well as any custom discounts applied to VMs from Azure.

As with vSAN, Lightbits can also compress data inline without sacrificing workload performance, which means that less data is stored on the underlying storage - a saving that is passed to the end user by requiring less storage instances. Due to Lightbits volumes (datastores) always being thin provisioned, administrators can over-provision storage assuming that data will be compressed. When the datastores fill up, Lightbits can automatically extend the cluster, adding more storage and balancing the datastores throughout the cluster to optimize capacity utilization.

This means that you only pay for the storage you need, when you need it - allowing for an optimized spend within Azure and significant cost savings.

Availability

Just as an AVS SDDC can span multiple Azure Availability Zones (AZs) in a region with a VMware stretched cluster, the Lightbits cluster can also be deployed across three zones for maximum availability. When this deployment is configured, each copy of the VMware datastore or Lightbits volume will be synchronously replicated across all three zones, ensuring that data is protected at all times. With this multi-zone configuration, only a minor write penalty is introduced based on the round-trip latency between the AZs. Read performance is unaffected if reading the volume from the local zone.

For customers who do not require zonal redundancy, Lightbits still replicates data synchronously between up to three nodes (VMs) in the cluster, enabling high availability of data within the cluster, even during a multiple VM failure event.

All upgrades to the cluster are non-disruptive and any node replacement due to VM failure is handled automatically by the Lightbits managed application, ensuring that data is always available for the applications.

If replicating data across AZs does not provide sufficient protection, Lightbits can utilize Azure Blob storage to back up data from the source Lightbits cluster and then restore it to either the same cluster or a different cluster. This allows for data protection at the regional level.

© 2024 Lightbits Labs™

Portability

The Lightbits software-defined storage solution in Azure is not only good for AVS workloads, but also for workloads running on Azure VMs or even Azure Kubernetes Service (AKS). With Lightbits, a single cluster can serve data to Azure IaaS, AKS, and AVS workloads simultaneously - enabling maximum utilization of a storage cluster and more workloads in Azure to take advantage of the resiliency, low cost, and high performance of Lightbits.

As well as being a supported storage solution for AVS, Lightbits is also recommended for Oracle workloads running on IaaS, even for the most performance-intensive databases.

The same Lightbits platform can be deployed across private and public clouds, allowing you to have the same experience and feature set no matter where your workloads reside. This flexibility is reflected in the Lightbits licensing, which can be portable across cloud environments. This ensures that the licenses are applied where they are needed, even if organizations are migrating data to and from the public cloud.

Development/Testing

In standard environments, development and test can be up to three times the size of production deployments. In order to correctly develop and integrate solutions, it is recommended to reflect the production architecture as closely as possible when building development/testing environments. Lightbits facilitates this process by allowing for thin snapshots and clones, allowing for production workloads to be snapshotted and used for development without taking up more space (only changed blocks) - while reflecting a real-world environment without risk to the source data.

For completely segregated environments, Lightbits on Azure allows administrators to take a snapshot from a source cluster, transfer it to Azure Blob for backup, and then restore the same dataset to a cluster in a different environment. This enables developers to have their own cluster to test with that is completely segregated from the production cluster, while still having production-like data to test with.

Lightbits Support 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. Your 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 perform their own monitoring, maintenance, and updates. The two models provide flexibility for you to choose the level of operational control, based on your IT preferences.

© 2024 Lightbits Labs™

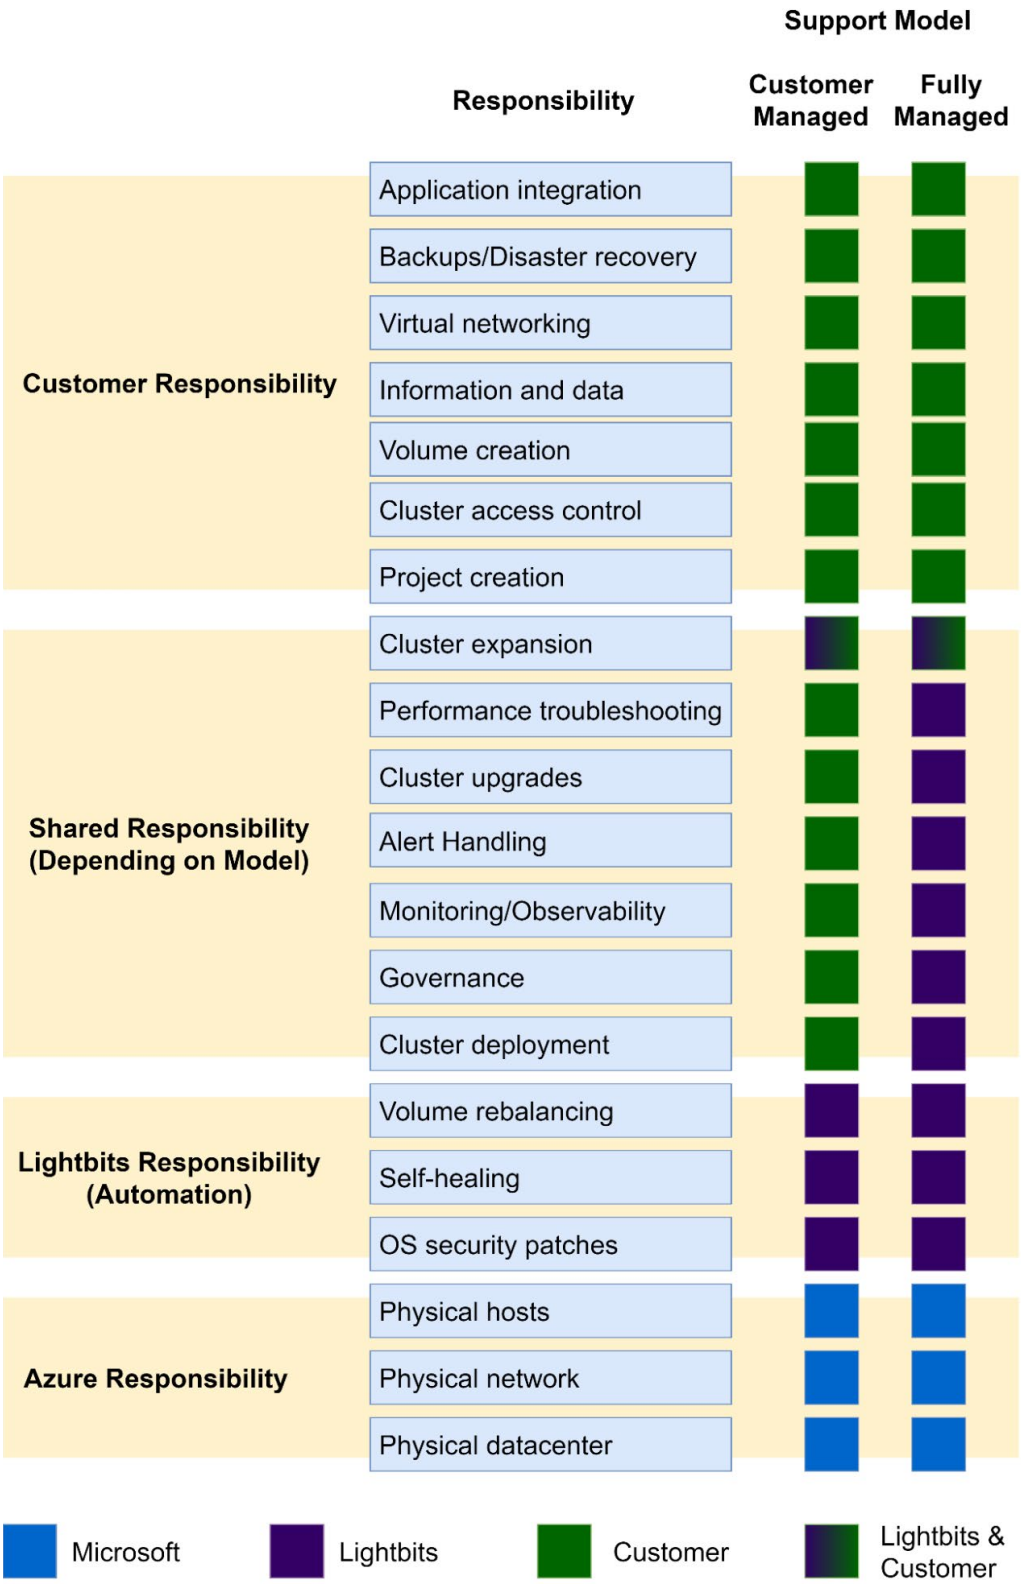

Figure 57 - Lightbits Managed Application Shared Responsibility Model

© 2024 Lightbits Labs™

Additional Support Information

To find out more about how our support models are implemented, Service Level Agreements (SLAs) and contacting Support, visit the Lightbits support documentation.

Request a Demo

To explore how Lightbits on Azure can set your cloud migrations up for success, request a demo today. We will review your specific cloud challenges, goals, and use cases, show you Lightbits in action, and discuss deployment options and pricing.

© 2024 Lightbits Labs™

About VMware by Broadcom

VMware is a Division of Broadcom Inc., a Delaware corporation headquartered in Palo Alto, California, is a global infrastructure technology leader built on more than 60 years of innovation, collaboration and engineering excellence.

VMware by Broadcom delivers software that unifies and streamlines hybrid cloud environments for the world's most complex organizations. By combining public-cloud scale and agility with private-cloud security and performance, we empower our customers to modernize, optimize and protect their apps and businesses everywhere.

Capable of deployment in the software-defined data center, across all clouds, in any app and out to the enterprise edge, VMware's unified software stack makes global enterprises more innovative, connected, resilient and secure.

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.

www.lightbitslabs.com

info@lightbitslabs.com

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

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

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.

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 non-infringement. 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. LBWP07/2024/05

COPYRIGHT© 2024 LIGHTBITS LABS LTD. ALL RIGHTS RESERVED

All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.