

Crash through Workload Performance Boundaries with Azure Stack HCI

June 2021

H18771

White Paper

Abstract

This white paper shares results of performance testing on Dell EMC Integrated System for Microsoft Azure Stack HCI using validated AX-7525 nodes, in a system configured to meet the requirements of today's most demanding workloads.

Dell Technologies Solutions

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA June 2021, White Paper H18771.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

- Overview 4
- Sizing an Azure Stack HCI solution 6
- Lab performance testing results..... 9
- Maintaining peak performance on Azure Stack HCI 20
- Conclusion..... 22
- References 23

Overview

Business challenge

Resource-intensive workloads are the lifeblood of the digital economy. Traditional applications running in virtual machines (VMs) and modern cloud native applications running in containerized ecosystems drive underlying virtualization clusters to their limits. Many IT professionals support these business-critical services with antiquated systems that cannot deliver the end-user response times required to meet service-level agreements (SLAs). Even if the infrastructure is currently meeting expectations, inflexible configurations often prevent easily scaling-up or scaling-out to meet projected performance and capacity demands in the future.

Incorrect and suboptimal infrastructure designs and configurations also lead to performance degradation over the life of IT systems. Failure to test and validate components properly with the correct combination of BIOS, firmware, and operating system driver revisions can significantly impact IOPS, throughput, latency, CPU utilization, and more. IT administrators must also perform maintenance by routinely applying hardware and software updates to keep their systems secure and running at peak performance and availability. However, this task is often neglected due to risky and time-consuming manual update processes.

Once precisely tuned and optimized, highly performant systems need to be monitored carefully to ensure that they are meeting SLAs. Some IT operational tooling is complex and lacks the features and functionality required to address modernized cloud native applications, such as those running in Kubernetes clusters. These tools may not be updated frequently enough to keep up with the constant changes in enterprise IT environments.

Solution overview

Dell EMC Integrated System for Azure Stack HCI (Azure Stack HCI) is an ideal solution for organizations that are refreshing and modernizing their aging virtualization environments to support high-value, highly performant workloads. This all-in-one validated hyperconverged infrastructure (HCI) system includes full-stack life cycle management, native integration into Microsoft Azure, flexible consumption models, and solution-level enterprise support and services expertise.

Azure Stack HCI offers a broad portfolio of intelligently designed AX nodes that are workload-focused and validated with deliberately selected hardware components and BIOS, firmware, and driver revisions. High-performance architectures leverage technologies such as 25/100 GbE RDMA networking and all-NVMe storage configurations. Our engineering organization also validates scalable and switchless storage networking topologies using Dell EMC PowerSwitch network switches. We also provide prescriptive guidance on expanding clusters to accommodate future workload demand.

Dell EMC OpenManage Integration with Microsoft Windows Admin Center enables automatic cluster creation with the Dell EMC HCI Configuration Profile and one-click full-stack life cycle management using Cluster-Aware Updating. This ensures that Azure Stack HCI is deployed in a guided and consistent way and helps maintain the latest quality, security, and performance enhancements. The extensions and snap-ins provided

by Microsoft and Dell Technologies reduce risk, eliminate guesswork, and liberate IT administrators to focus on performance tuning and other high-value activities.

Azure Stack HCI's native integration with Microsoft Azure allows IT administrators to leverage feature-rich management and governance services such as Azure Monitor, Azure Security Center, and Azure Policy. These services are consistent with and complementary to Microsoft Windows Admin Center, which serves as the edge/local, always available management console. Microsoft continually adds new innovative functionality to Microsoft Azure services to ensure that even the most complex and demanding VM-based and containerized applications are properly monitored and performing at the levels required to meet SLAs.

In this white paper, we demonstrate how the architectural components of the Dell EMC Integrated System work together to deliver extraordinary value to high-performing workloads. We performed synthetic testing on an Azure Stack HCI cluster running four AX-7525 nodes, each populated with two AMD EPYC 7742 64-core processors, 24 NVMe drives (PCIe Gen 4), and 100 GbE RDMA networking. This configuration delivered the following results:

1. 5,727,985 IOPS at 1.3-millisecond latency using a 4 KB block size and 100% random-read I/O profile.
2. 700,256 IOPS at 9-millisecond latency using a 4 KB block size and 100% random-write I/O profile.
3. 105 GB/s and 8 GB/s throughput for 100% sequential-read and 100% sequential-write I/O profile, respectively, using 512k block sizes.
4. Even with two nodes failed, the cluster delivered 3,842,997 IOPS at under 1-millisecond latency using a 4 KB block size and 100% random-read I/O profile.

Audience

The audience for this white paper includes IT managers, system engineers, field consultants, and others with expertise in virtualization and cloud computing who support infrastructure running high-performing, mission-critical business applications and services.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#) or provide your comments by completing our [documentation survey](#).

Author: Michael Lamia

Contributors: Tom Smaistrila, Rizwan Ali


Note: For links to additional information about this solution, see the [Dell Technologies Info Hub for Dell EMC Integrated System for Azure Stack HCI](#).

Sizing an Azure Stack HCI solution

Before exploring the lab environment and performance testing results, we will first review the process of selecting and sizing an optimized cluster configuration. Dell EMC Integrated System for Microsoft Azure Stack HCI offers a broad range of validated AX node and network topologies to meet the requirements of nearly any use case and workload. Example scenarios include the following:

- For small Remote Branch Offices (ROBO), we may recommend a cluster with two AX nodes using lower core count processors, modest amounts of RAM, all-flash storage, and a [switchless storage network topology](#).
- For applications with significant storage capacity and expansion requirements, we may suggest a two-tier, hybrid storage configuration, and [scalable networking topology](#).
- For the industry’s most demanding workloads, we often select processors with the highest number of cores and fastest clock speeds, large memory footprints, NVMe drives, and 100 GbE RDMA networking.

Figure 1 illustrates our specification sheet at-a-glance. It depicts the component combinations included in our four AX node models.



	AX-640	AX-740xd	AX-6515	AX-7525
Processor	Intel Xeon Scalable Processors (Select Silver/Gold/Platinum options)		2 nd Gen AMD EPYC Processor	
Memory	96 GB to 1.5 TB	96 GB to 1.5 TB	64 GB to 1 TB	128 GB to 2 TB
Storage Configurations				
Min/Max Raw Storage	2.5 to 77 TB	8 to 192 TB	3 to 60 TB	6.4 to 184 TB
All Flash (All-NVMe)	✓	✓		✓
All Flash (SSD)	✓	✓	✓	
All Flash (NVMe + SSD)	✓	✓		✓
Hybrid (NVMe + HDD)	✓	✓		
Hybrid (SSD + HDD)	✓	✓		

Figure 1. Component options for AX node portfolio

Live Optics

Dell Technologies follows a consultative and systematic approach to helping IT professionals modernize their infrastructure. Before recommending any configurations like those mentioned above, we begin with a thorough analysis of currently running workloads. Dell EMC Live Optics is an industry-standard method of impartially documenting server and storage configuration and performance. By recording and exploring the current state of IT environments, we can help accelerate informed decisions and reduce the risk of under or overprovisioning the infrastructure.

An in-depth discussion of Live Optics is beyond the scope of this white paper. From a high-level, customers first engage with their Dell Technologies account team to have an account created on the [Live Optics portal](#). Then they download and run the Optical Prime data collector software to capture inventory and performance insights from current hosts

and VMs. Collectors can run from four hours to seven days, in offline or online mode, to create a local file of the results or stream the data securely to the associated online viewer profile. When data collection is complete, Live Optics visualizes the data using a variety of rich reports available in the portal, such as the examples depicted in Figures 2 and 3.

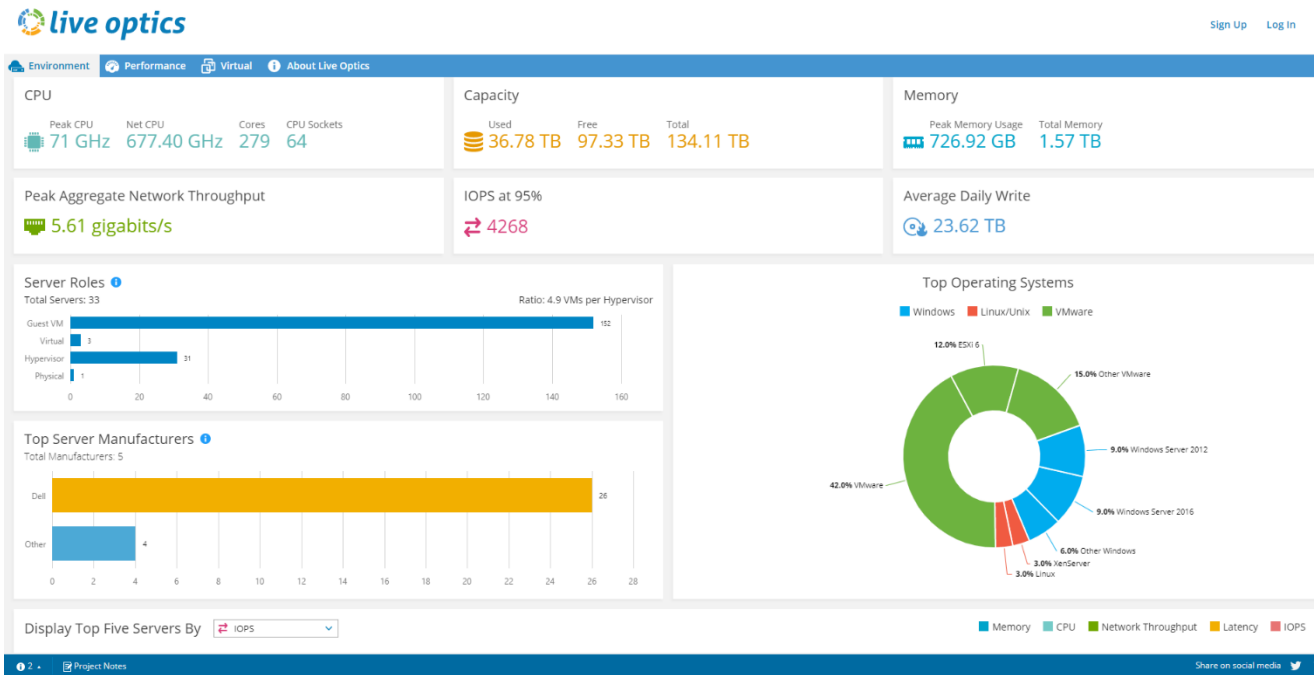


Figure 2. Demo environment report from Live Optics portal dashboard

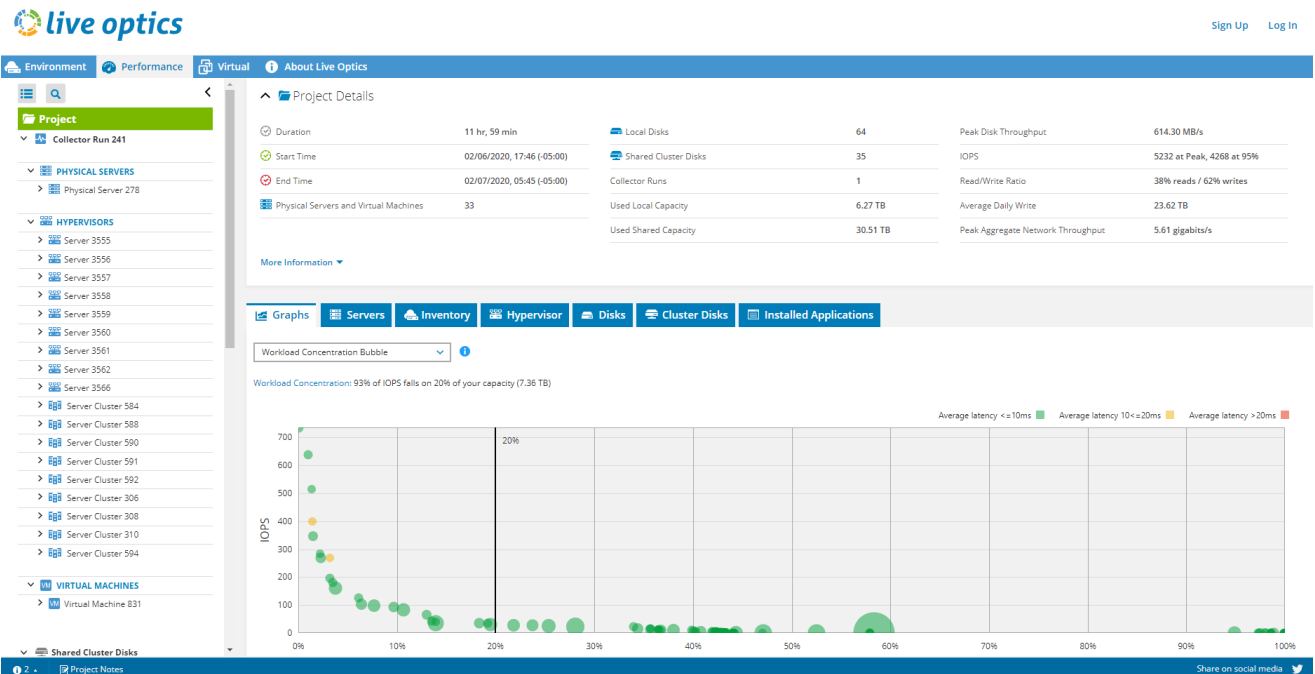


Figure 3. Demo performance report from Live Optics portal dashboard

Azure Stack HCI Sizer Tool

The Dell Technologies account team and customers use the data collected by Live Optics to influence selections in our exclusive Azure Stack HCI Sizer Tool. This tool integrates all the Azure Stack HCI design constraints and best practices from extensive engineering validation efforts. It provides peace-of-mind knowing that each potential configuration produced by the tool is an optimized and viable option. We work with our customers to generate multiple configurations for comparison. We then select the one that meets current demand and future growth projections within the allocated budget.

Note: The Azure Stack HCI Sizer Tool can only be accessed by Dell Technologies direct sales and its partners. Customers can contact their Dell Technologies or preferred channel partner to step through sample configuration options.

Lab performance testing results

In this section, we share the results of synthetic testing performed on Dell EMC Integrated System for Microsoft Azure Stack HCI running four AX-7525 nodes. We used [VMFleet](#), as many Microsoft customers and partners rely on this tool to help them stress-test their hyperconverged infrastructures. VMFleet consists of a set of PowerShell scripts that deploy VMs to a Hyper-V cluster and execute Microsoft's [DISKSPD](#) within those VMs to generate I/O. DISKSPD is very helpful for creating synthetic workloads to test an application's resource utilization before deployment.

Test lab setup

Figure 4 depicts the cluster's physical networking as built in the lab. We set up a scalable, fully converged network topology with management, VM, and RDMA storage traffic traversing two 100 GbE Mellanox CX-6 adapter ports. With this topology, Data Center Bridging (DCB) configuration was required on the S5232F-ON ToR switches. We also required Quality of Service (QoS) configuration in the host operating system. Figure 5 illustrates the virtual network configuration on the host operating system using Switch Embedded Teaming (SET).

Ancillary services required for cluster operations such as Active Directory, DNS, and a file server or cloud-based witness are outside the scope of this white paper. For information on all the steps required to deploy this cluster, refer to our [Deployment Guide](#) and [Network Integration and Host Network Configuration Options](#) guide.

Note: Out of band (OOB) connections from the iDRACs to a management switch are not included in the diagram.




Figure 4. Lab physical infrastructure

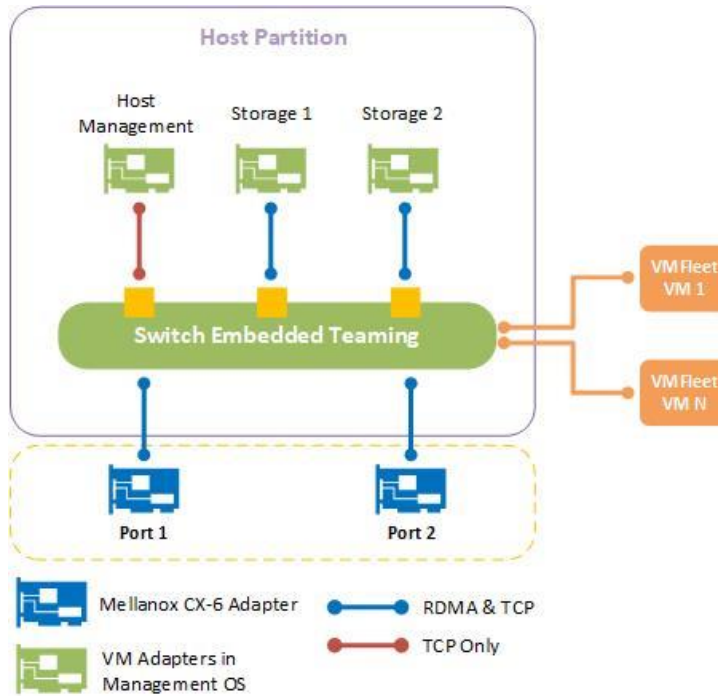


Figure 5. Virtual networking configuration

Tables 1 and 2 provide additional detail about the cluster configuration and AX-7525 specifications as built in the lab. We chose the three-way mirror resiliency type for the volumes we created with VMFleet because of its superior performance compared to erasure coding options in Storage Spaces Direct. The raw storage capacity of the cluster was a little over 150 TB. This resulted in just under 50 TB of usable storage capacity due to three-way mirroring's capacity efficiency of 33 percent. Microsoft and Dell Technologies also recommend leaving reserve capacity equivalent to the capacity of one drive per server.

Table 1. Cluster configuration

Cluster design elements	Description
Cluster node model	AX-7525
Number of cluster nodes	4
Network switch model	Dell EMC PowerSwitch S5232F-ON featuring 32 x 100 GbE QSFP28 ports
Number of ToR network switches	2
Network topology	Converged network topology
Volume resiliency	Three-way mirror
Usable storage capacity	Approximately 50 TB

Table 2. AX-7525 node specifications

Resources per cluster node	Description
CPU	Dual-socket AMD EPYC 7742 64-Core Processor
Memory	1 TB 3200 MHz DDR4 RAM
Storage controller for operating system	BOSS-S1 adapter card
Physical drives for operating system	2 x M.2 480 GB SATA drives configured as RAID 1
Physical drives for Storage Spaces Direct	24 NVMe drives (PCIe Gen 4)
Network adapter for management, VM, and storage traffic	1 x Mellanox ConnectX-6 DX Dual Port 40/100 GbE QSFP56 Adapter
Operating system	Microsoft Azure Stack HCI, version 20H2

VMFleet testing results

We ran the VMFleet tests under the following conditions:

- Healthy cluster running 64 VMs per node
- Healthy cluster running 32 VMs per node
- Degraded cluster with one node failed
- Degraded cluster with two nodes failed

Healthy cluster running 64 VMs per node

The following table presents the range of VMFleet and DISKSPD parameters used during our testing of the cluster in optimal health and running 64 VMFleet VMs per node.

Table 3. Test parameters for VMFleet and DISKSPD on a healthy cluster

VMFleet and DISKSPD parameters	Values
Number of VMs running per node	64
vCPUs per VM	2
Memory per VM	4 GB
VHDX size per VM	40 GB
VM operating system	Windows Server 2019
Data file size used in DISKSPD	10 GB
Cluster Shared Volume (CSV) in-memory read cache size	0
Block sizes	4 – 512 KB
Thread counts	1 – 2
Outstanding I/Os	2 – 32
Write percentages	0 – 100
I/O patterns	Random, Sequential

We first selected DISKSPD I/O profiles aimed at identifying the maximum IOPS and throughput thresholds of the cluster. By pushing the limits of the storage subsystem, we confirmed that the networking, compute, operating systems, and virtualization layer were configured correctly according to our [Deployment Guide](#) and [Network Integration and Host Network Configuration Options](#) guide. This also ensured that no misconfiguration occurred during initial deployment that could skew the results of I/O profiles that are more representative of real-world workloads. Our results are depicted in Table 4.

Table 4. VMFleet test results

I/O profile	Parameter values explained	Performance metrics
B4-T2-O32-W0-PR	Block size: 4k Thread count: 2 Outstanding I/O: 32 Write percentage: 0% I/O Pattern: 100% R andom read	IOPS: 5,727,985 Read latency: 1.3 milliseconds
B4-T2-O16-W100-PR	Block size: 4k Thread count: 2 Outstanding I/O: 16 Write percentage: 100% I/O Pattern: 100% R andom write	IOPS: 700,256 Write latency: 9 milliseconds
B512-T1-O8-W0-PSI	Block size: 512k Thread count: 1 Outstanding I/O: 8 Write percentage: 0% I/O Pattern: 100% S equential read	Throughput: 105 GB/s
B512-T1-O1-W100-PSI	Block size: 512k Thread count: 1 Outstanding I/O: 1 Write percentage: 100% I/O Pattern: 100% S equential write	Throughput: 8 GB/s

Healthy cluster running 32 VMs per node

To prepare for the other three test scenarios, we re-deployed VMFleet with 32 VMs running per node. We then stressed the storage subsystem using I/O profiles more reflective of the types of demanding workloads found in the modern enterprise. These applications had smaller block sizes, random I/O patterns, and a variety of read/write ratios.

The 8 KB block size and 30 percent random write I/O profile is typical of an Online Transactional Processing (OLTP) workload. We observed VMFleet driving over 1.6 million IOPS at slightly over 1-millisecond average latency with this I/O profile, indicating that the cluster has the potential to accelerate OLTP workloads and improve end-user response time to database applications. Even increasing the block size to 32 KB and increasing the write I/O to 50 percent produces over 400,000 IOPS at under 7-millisecond latency.

Figure 6 depicts the results of this workload testing.

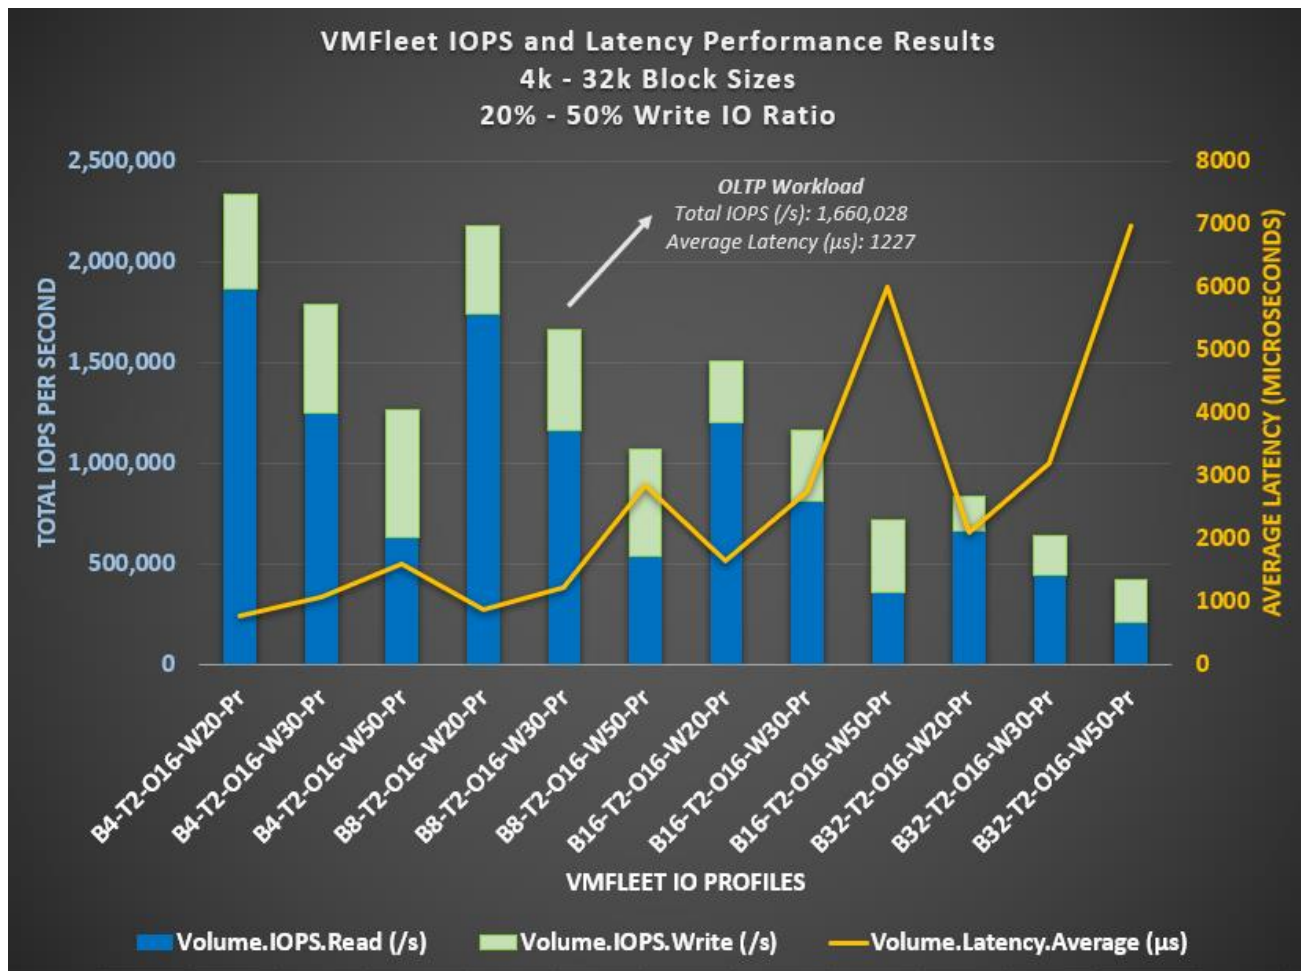


Figure 6. IOPS and latency results

Online Analytical Processing (OLAP) workloads focus on retrieval and analysis of large datasets. We selected larger block sizes and sequential I/O patterns to test workloads in this category. Throughput became the key performance indicator to analyze. The results shown in Figure 7 indicate an impressive sustained throughput that can greatly benefit this category of OLAP workloads. Other types of IT services that could benefit include file and application backups and streaming video.




Figure 7. Throughput results

Testing results on a degraded cluster

We simulated two different failure scenarios on the healthy cluster running 32 VMs per node. In the first scenario depicted in Figure 8, we shut down one node, which had the following impact:

- The cluster moved the CSV owned by Node 4 to Node 2. Node 2 became the new owner of that CSV.
- We live-migrated the failed VMs to Node 2 to ensure they were running on the same host that owned the CSV containing their virtual disks.




Figure 8. VMFleet testing after a single node failure

In the second scenario depicted in Figure 9, we shut down two nodes, which had the following impact:

- The cluster moved the CSVs owned by Nodes 3 and 4 to Nodes 1 and 2, respectively. Nodes 1 and 2 became the new owners of these volumes.
- We live-migrated the failed VMs to Nodes 1 and 2, ensuring that all VMs were running on the host that owned the CSV containing their virtual disks.

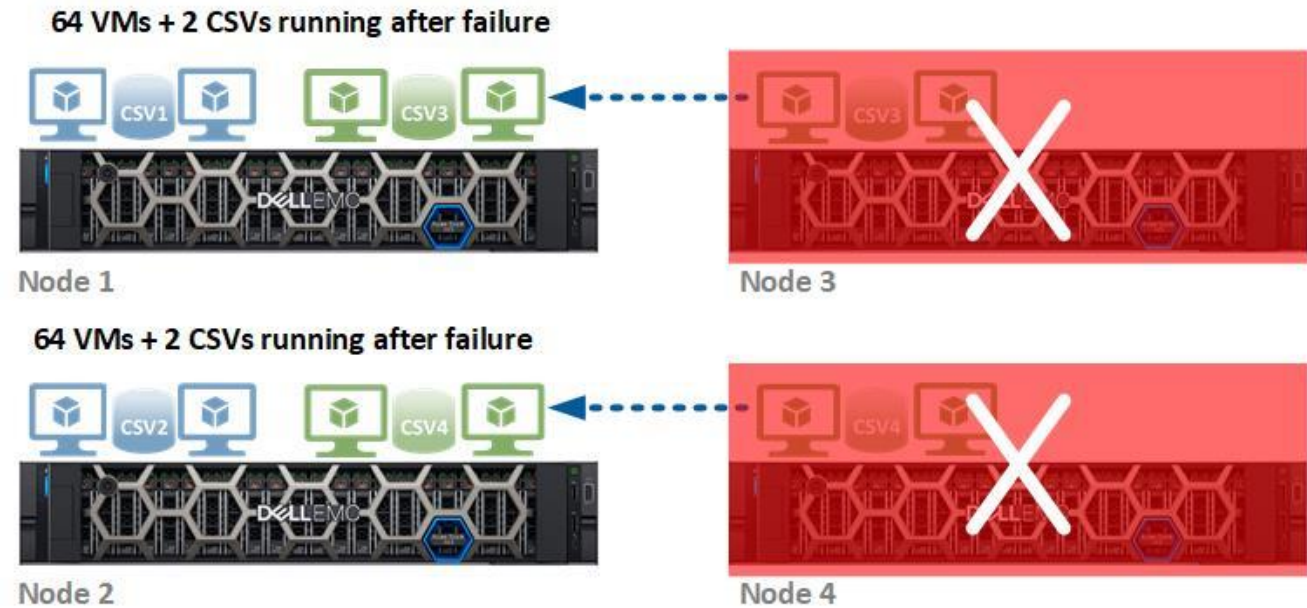


Figure 9. VMFleet testing after a two-node failure

According to the VMFleet test results, we continued to see outstanding performance even with the cluster in a degraded state. This allows system administrators extra time to recover from node failures without negatively impacting the end-user response times of their most critical applications. Table 5 compares the testing scenarios using the I/O profiles aimed at identifying the maximum thresholds.

Table 5. Performance comparisons

I/O profiles	Healthy cluster	One-node failure	Two-node failure
B4-T2-O32-W0-PR	IOPS: 4,856,796 Read latency: 378 microseconds	IOPS: 4,390,717 Read latency: 380 microseconds	IOPS: 3,842,997 Read latency: 262 microseconds
B4-T2-O16-W100-PR	IOPS: 753,886 Write latency: 3.2 milliseconds	IOPS: 482,715 Write latency: 5.7 milliseconds	IOPS: 330,176 Write latency: 11.4 milliseconds
B512-T1-O8-W0-PSI	Throughput: 91 GB/s	Throughput: 113 GB/s	Throughput: 77 GB/s
B512-T1-O1-W100-PSI	Throughput: 8 GB/s	Throughput: 6 GB/s	Throughput: 10 GB/s

Figure 10 shows the IOPS and latency results from the VMFleet tests. These tests used I/O profiles that are more representative of real-world applications with smaller block sizes, random I/O patterns, and a variety of read/write ratios. Notice that IOPS decrease and latency increases for the 4k and 8k block sizes, as one would expect. With the 32k block size, we observed the following behavior:

- Latency is less variable across the failure scenarios. We found an 8 percent decrease in latency after the single-node failure and a 14 percent decrease in latency after the two-node failure. This was primarily because write I/O did not need to be committed across as many nodes in the cluster. After the single-node failure, we observed that the network traffic between nodes decreased by more than 3 Gbps per node.
- IOPS **increased** by 20 to 30 percent after the two-node failure, depending on the percentage of writes in the I/O profile. There were two reasons for this behavior:
 - The 3-way mirrored volumes running on the healthy cluster became 2-way mirrored volumes running on the two remaining online nodes. With 33 percent fewer backend drive write IOPS, the overall drive write latency decreased, which drove higher read and write IOPS. However, this is only in cases when CPU was not the bottleneck.
 - Each of the two nodes remaining online were running double the VMs for a total of 64 VMs per node. The more VMs running on a virtualized host, the greater the potential for higher IOPS.

Note: After shutting down the nodes in the failure scenarios, we live-migrated the VMs to the same nodes that owned the CSVs hosting their virtual disks to avoid any I/O redirection, which would negatively affect performance.

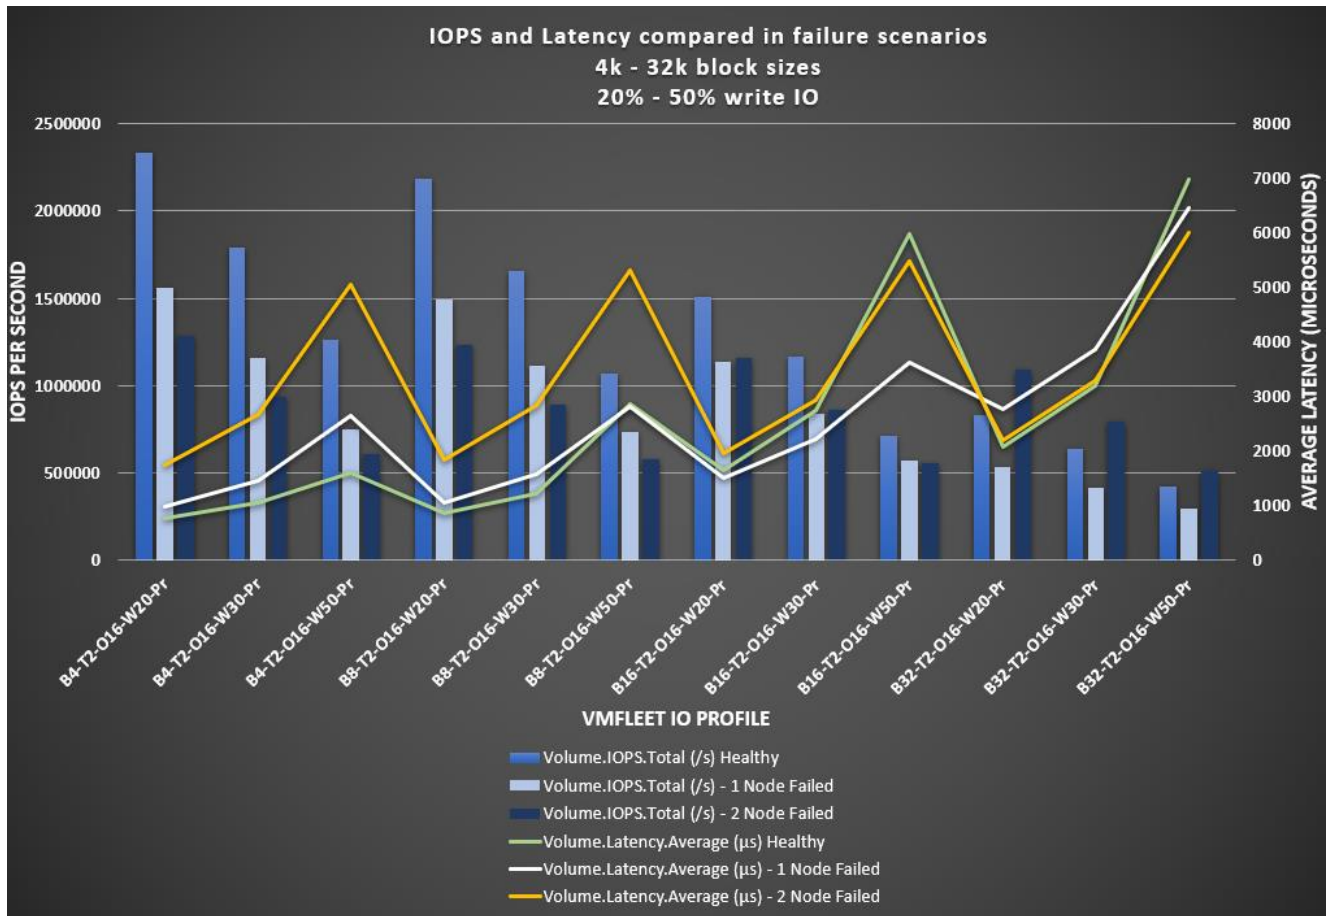


Figure 10. IOPS and latency results during failure scenarios

Figure 11 shows the throughput results from the VMFleet tests using I/O profiles with larger block sizes and sequential I/O patterns. Throughput decreased an average of 17 percent after the single-node failure. However, throughput **increased** by an average of 52 percent after the two-node failure for the same reasons as applied to Figure 10.

- The 3-way mirrored volumes running on the healthy cluster became 2-way mirrored volumes running on the two remaining online nodes. With 33 percent less backend drive write throughput, the overall drive write latency decreased, which drove higher read and write throughput. However, this is only in cases when CPU was not the bottleneck.
- Each of the two nodes remaining online were running double the VMs for a total of 64 VMs per node. The more VMs running on a virtualized host, the greater the potential for higher throughput.

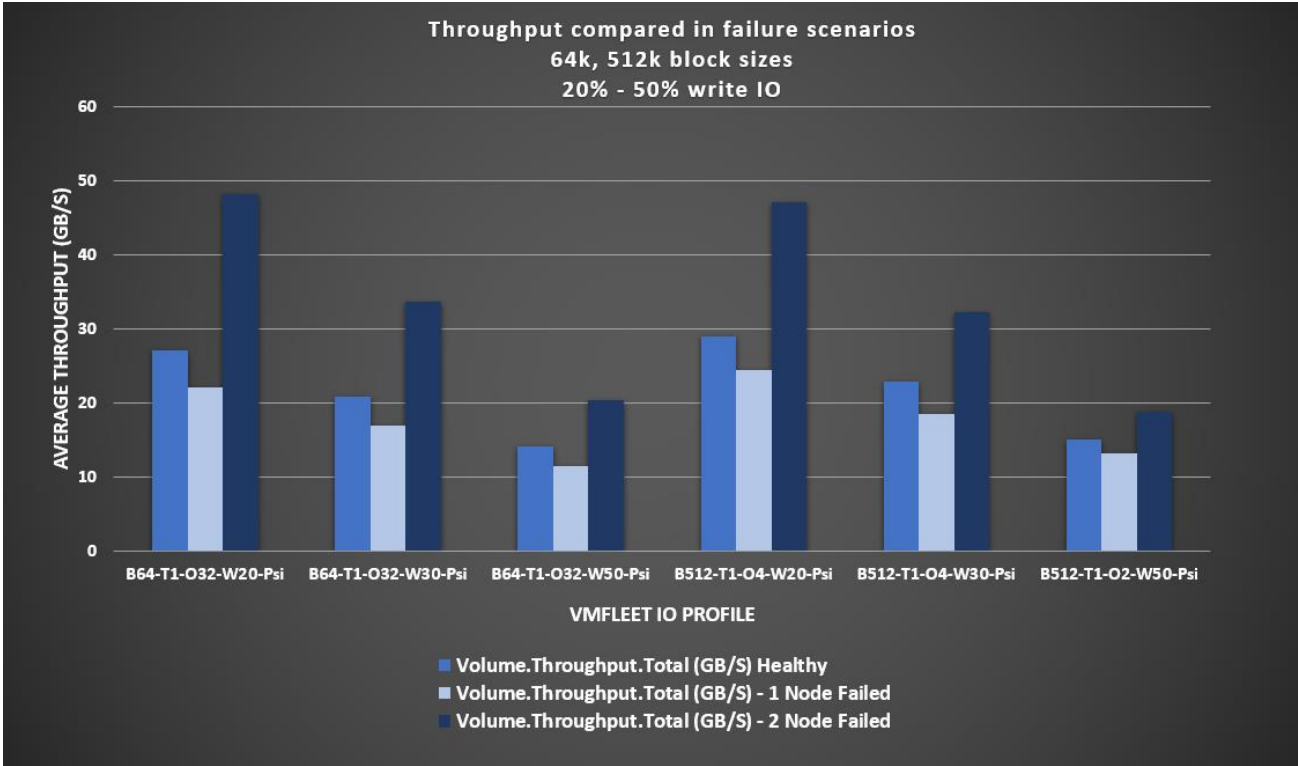


Figure 11. Throughput results during failure scenarios

Maintaining peak performance on Azure Stack HCI

To keep Dell EMC Integrated System for Microsoft Azure Stack HCI running at peak performance and reliability, IT administrators must take the appropriate steps to monitor and maintain it correctly. Dell Technologies provides the expertise to make Azure Stack HCI operations efficient and worry-free. Administrators will first want to familiarize themselves with the [Operations Guide – Managing and Monitoring the Solution Infrastructure Life Cycle](#). The following are some of the topics covered in this guide:

- Expanding clusters by scaling up and scaling out
- Recovering from AX node hardware and operating system failures
- Monitoring the hardware and software using Microsoft Windows Admin Center and Dell EMC OpenManage Integration for Microsoft Windows Admin Center
- Updating the operating system, BIOS, firmware, and drivers with a simple consolidated workflow in Windows Admin Center

Specialized expertise

For any Azure Stack HCI cluster expected to deliver high performance, administrators need to rely on specially trained experts who can help accelerate time to value and decrease mean time to resolve (MTTR) failures. Dell Technologies makes services simple, flexible, and worry free – from installation and configuration to comprehensive single-source support. Certified deployment engineers ensure accuracy and speed, reduce risk and downtime, and free IT staff to work on higher value priorities. One-stop cluster level support covers the hardware, the operating system, the hypervisor, and Storage Spaces Direct.

When it comes to expanding clusters, customers can choose to follow the prescriptive guidance in the Operations Guide or enlist the help of our support and services. In addition to improving the storage performance of the cluster, adding AX nodes increases storage capacity and provides more compute resources to the VMs. Administrators must verify hardware symmetry before attempting any expansion operations. This means that the processor model, HBA, and network adapter configuration in the new AX node must be identical to the existing nodes.

For the full list of requirements for scaling out the cluster to include additional nodes and scaling up existing nodes, see the Operations Guide.

Note: Operators should always consult the [Support Matrix](#) when investigating the addition of components to Azure Stack HCI clusters.

Life Cycle Management

Life Cycle Management (LCM) is a priority for every IT operations team. LCM is especially critical for Azure Stack HCI clusters running an organization's most resource-intensive, mission-critical workloads. This practice ensures that systems run at peak performance and availability, maintain a sound security posture, and benefit from newly released features and functionality.

Dell EMC OpenManage Integration with Microsoft Windows Admin Center makes the update process automated, repeatable, and predictable. The OpenManage Integration

extension includes a specialized snap-in to Microsoft's Cluster-Aware Updating extension. This innovative collaboration between Microsoft and Dell Technologies creates a single-click full stack update experience. A simple Windows Admin Center workflow orchestrates Azure Stack HCI operating system and AX node BIOS, firmware, and driver updates, requiring only a single reboot per node and resulting in no impact to running workloads in the VMs. Figure 12 depicts the initial steps in the updates workflow that check for AX node authenticity, correct licensing, and hardware symmetry.

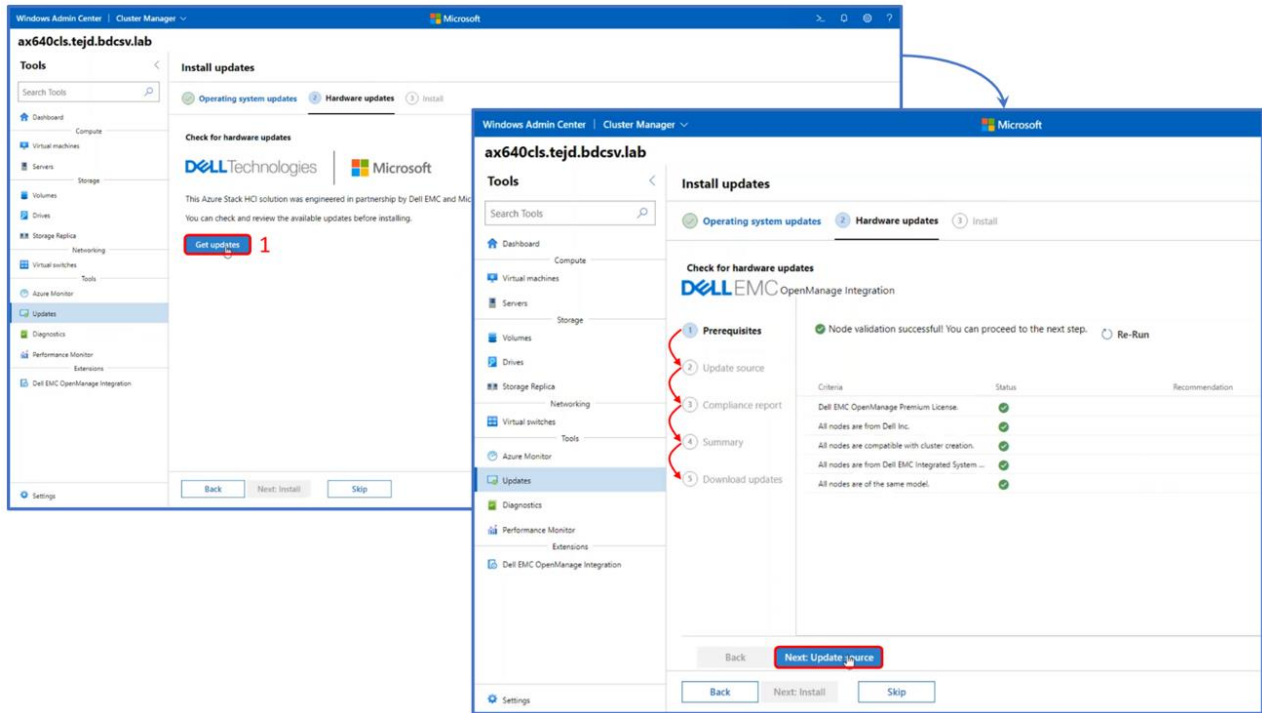


Figure 12. Dell EMC policies as part of updates workflow

For a detailed technical walkthrough of the full stack life cycle management with the Cluster-Aware Updating workflow in Windows Admin Center, please see the [OpenManage Integration Technical Walkthrough blog](#).

Monitoring with Microsoft Azure services

It is also critical to monitor a high-performing cluster to ensure it is meeting SLAs. Windows Admin Center with the Dell EMC OpenManage Integration extension provides local hardware and software monitoring in-band and completely agent-free. Since Azure Stack HCI is a service in Microsoft Azure, administrators can also take advantage of native integration with a host of Azure management and governance services such as Azure Monitor, Azure Security Center, and Azure Policy. This native integration is made possible in large part by Azure Arc.

Azure Arc enables administrators to simplify the management and governance of their on-premises workloads by projecting existing resources into Azure Resource Manager. Once onboarded to Azure Arc, VMs, Kubernetes clusters, and databases can be monitored and managed with the same familiar tools used for Azure resources. Tools include the Azure Portal and programmatic means such as PowerShell, Azure CLI, and the REST API.

Conclusion

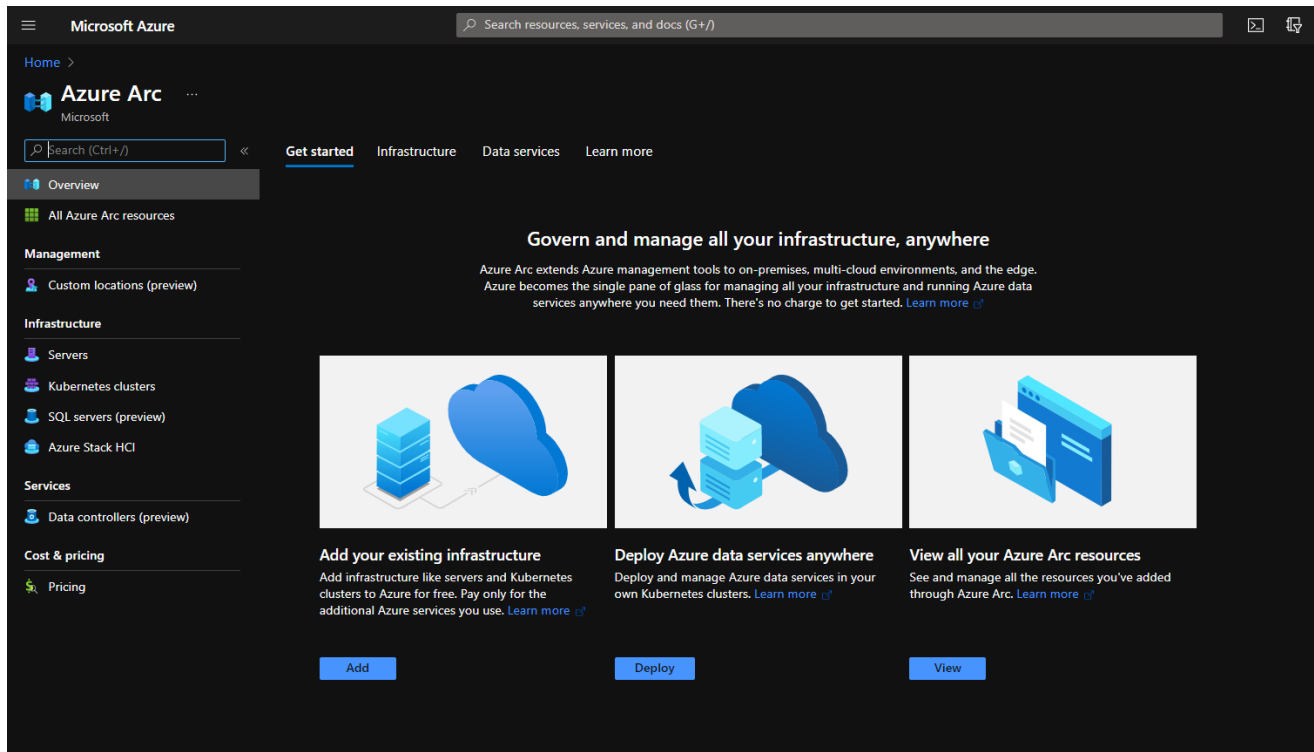


Figure 13. Azure Arc Overview blade in the Azure Portal

For more information about using Azure Arc to help manage and monitor our integrated system at scale, see the [Azure Arc documentation](#). See also the [Azure Monitor documentation](#) to learn more about monitoring high-performance Azure Stack HCI clusters.

Conclusion

Demanding applications require a modernized infrastructure that is intelligently designed and correctly sized to meet current performance SLAs and to scale easily to accommodate future growth.

In this white paper, we discussed the systematic and consultative approach Dell Technologies takes in the pre-sales process using Live Optics and our exclusive Azure Stack HCI Sizer Tool. We also shared synthetic testing results using VMFleet on a Dell EMC Integrated System for Microsoft Azure Stack HCI cluster running four AX-7525 nodes and two S5232F-ON switches. These results indicated that the model cluster we set up in the lab was capable of impressive performance – even running in a degraded state with two nodes failed. Finally, we highlighted best practices surrounding support, maintenance, life cycle management, and monitoring for resource-intensive infrastructures.

References

Dell Technologies documentation

The following Dell Technologies documentation provides additional information:

[Deployment Guide](#)

[Operations Guide](#)

[Network Integration Guide](#)

[Support Matrix for Microsoft HCI Solutions](#)

Microsoft documentation

The following Microsoft documentation provides additional and relevant information:

[Azure Stack HCI documentation](#)

[Azure Monitor documentation](#)

[Azure Arc documentation](#)