# Dell Integrated System for Microsoft Azure Stack HCI Tech Book

## A Hyperconverged Infrastructure System from Dell Technologies and Microsoft


**Abstract**

This document is a conceptual and architectural review of the Dell Integrated System for Microsoft Azure Stack HCI. The tech book first describes how hyperconverged infrastructure drives digital transformation and then focuses on the Azure Stack HCI system as a leading hyperconverged technology solution.



**Part Number: H18978**
**November 2023**


DELL Technologies

## Notes, cautions, and warnings

**NOTE:** A NOTE indicates important information that helps you make better use of your product.

**CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.**

**WARNING: A WARNING indicates a potential for property damage, personal injury, or death.**

# Contents

# 1

# Preface

**Topics:**



## Document description

This document is a conceptual and architectural review of the Dell Integrated System for Microsoft Azure Stack HCI. The tech book first describes how hyperconverged infrastructure drives digital transformation and then focuses on the Azure stack HCI system as a leading hyperconverged technology solution.

## Audience

This tech book is intended for Dell Technologies field personnel, partners, and customers involved in designing, acquiring, managing, or operating a Dell Integrated System for Microsoft Azure Stack HCI.

2

# Business Transformation and Data Center Modernization

**Topics:**



## Overview

Organizations of all sorts and sizes, across all vertical markets, and anywhere in the world are facing increased pressure to perform better. The introduction of new technologies results in consumers of goods and services demanding better products, and new ways of interacting with suppliers. This demand requires organizations to rethink their business models to revolutionize the customer experience. The ability to adapt and manage this business transformation and to achieve true agility is crucial for any organization to gain a competitive advantage, and to stay relevant. The requirement to focus on achieving the following outcomes has never been greater:

- Reduce costs and increase revenue
- Optimize go-to-market processes
- Improve customer experience and retention
- Gain industry share


**Figure 1. Business transformation**

A key enabler to gain customer mind share and achieve these outcomes is to adopt best in class technology solutions that enable organizations to develop modern applications, analyze buyer behavior and enhance customer interaction. The skill set required to excel as a Chief Information Officer (CIO) is no longer limited to technical knowledge but also includes deep business acumen to be a driving force on how technology can enable achieving the outcomes.

IT decision makers need to focus on digital transformation and the technology solutions they use must support the primary business goals of the organization. The IT environment must provide modern, secure, scalable, highly automated, and

continuously available services. These services must accommodate access anywhere at any time, all while technology budgets shift from the traditional IT organization to other business units. These challenges have driven cloud adoption and data center convergence that we have seen in recent years.

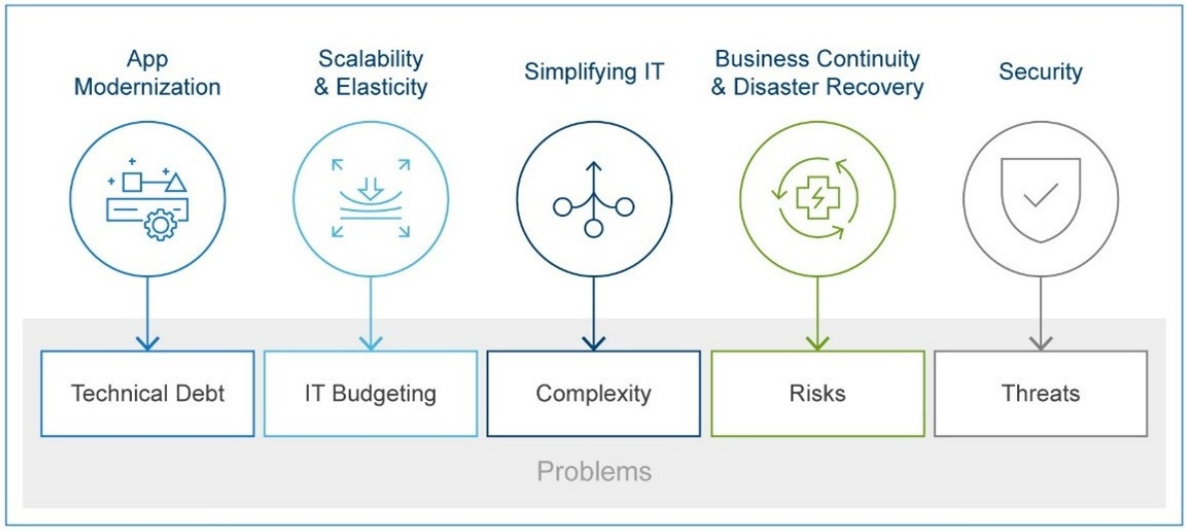


**Figure 2. IT transformation challenges**

# Technical debt

Thousands of applications, hundreds of servers, multiple cloud environments, several data centers and potentially numerous edge locations have introduced more complexity, risk, and cost. Add new application development methodologies, microservices, and containerization and it is easy to understand that traditional platforms and administration practices are not sufficient. This deficiency creates the need for data center modernization. Modern, secure, and scalable platforms with high degrees of automation for administration seamlessly integrate with a public cloud to create a true hybrid cloud experience. A hybrid cloud enables organizations to determine the location of their applications at any time and can make (or automate) that decision continuously. The decisions are based on:

- Economics – the cost of running applications in a specific location
- Physics – latency requirements for applications and proximity to data and users
- Land – regulations based on industry or government mandate for data or applications to be in a certain place

Further research shows that 72 percent of enterprises have a hybrid-first or private-first strategy [1] while 88 percent of cloud strategies include on-premises infrastructure [2].

Based on business requirements, IT professionals have the flexibility to choose the best location for their applications to run - either by running on-premises or moving workloads to the cloud. For the best of both worlds, applications can now be developed using a “build once, deploy anywhere" approach. As requirements change, workloads can be migrated to any hosting environment in an automated fashion without impacting business services.

# IT budgeting

The demand for IT resources like storage and compute is highly dynamic where needs increase or decrease depending on business demand. When “traditional” siloed infrastructures that lack this flexibility are used, there is enormous strain on the IT teams to keep up with business demand. The fear of not having enough resources to serve peak demand leads to overprovisioning and wasted resources that are unused. Elasticity and on-demand capabilities are the key to success to keep up with changing demand. This changing demand also requires different IT budgeting. Traditionally, IT budgets are set per year, purchasing and life cycle are determined for three to five years and are typically static. This does not align to the business demand for resources, nor does it align to the revenue streams of the business or cashflow of operations. For this reason many organizations have embraced the public cloud, believing that hyperscale computing like Amazon Web Services, Azure, and

1 Source: Everest Group Survey of 200 Enterprises. Used with permission.
2 Source: ESG Research Insights Paper, “Exploring Hybrid Cloud Adoption and the Complexity of Securing East-West Traffic” Commissioned by VMware, January 2020. Used with permission.

Google Cloud Platform are the only providers that can offer flexibility in scaling and payment. But there are many other ways to achieve this flexibility.

# Complexity

It has been common for traditional IT infrastructure purchases to be “best of breed” products for each technology. This choice leads to multivendor environments that are faced with interoperability issues, disparate management interfaces, and performance and support challenges. This complexity increases management overhead and time and inevitably leads to higher cost. Moving towards a limited number of vendors, or even a single vendor when applicable, that can provide a hybrid cloud operating model which provides a solution to these problems. Managing all resources from a unified, hyperconverged platform brings the following advantages:

- Reduced complexity with a complete validated platform
- Integrated life cycle management for storage, compute, and (potentially) networking
- High degree of automation to reduce human error and maximize application uptime
- Modern management tools that seamlessly integrate with the public cloud
- Management and maintenance capabilities in the public cloud to manage on-premises infrastructure at scale

# Risks

No company is exempt from the paralyzing consequences of catastrophes. A flexible and robust Business Continuity and Disaster Recovery (BCDR) plan that comprehensively considers people, processes, and technology systems is more essential than ever before. With the increased importance of IT resources in the business transformation, the dependency on the availability of these resources also increases. Business continuity and disaster recovery (DR) must be an integrated part of any IT infrastructure strategy. To address this challenge, IT organizations must adopt Hyperconverged Infrastructure (HCI) solutions that are integrated into a hybrid cloud operating model, which can be built from the ground up to simplify the implementation of high availability and DR for workloads running in virtual machines (VMs) and containers. This type of approach builds a strong base of resiliency in data centers and improves Recovery Point Objective (RPO) and shrinks Recovery Time Objective (RTO). Besides the simplicity of setup and maintenance, DR in a hybrid HCI infrastructure can provide automatic failover without the need for manual intervention. Using public cloud resources can be a cost-effective way to achieve DR, so it is imperative that HCI can integrate properly with the public cloud.

The best data protection strategies incorporate a combination of different technologies (deduplicated backup, archive, data replication, business continuity, and workload mobility) to deliver the right level of data protection for each business application. The following figure highlights that a reduced dataset holds the most valuable information and is protected with the most advanced technologies available. In contrast, more traditional protection tools can protect the rest of the data.

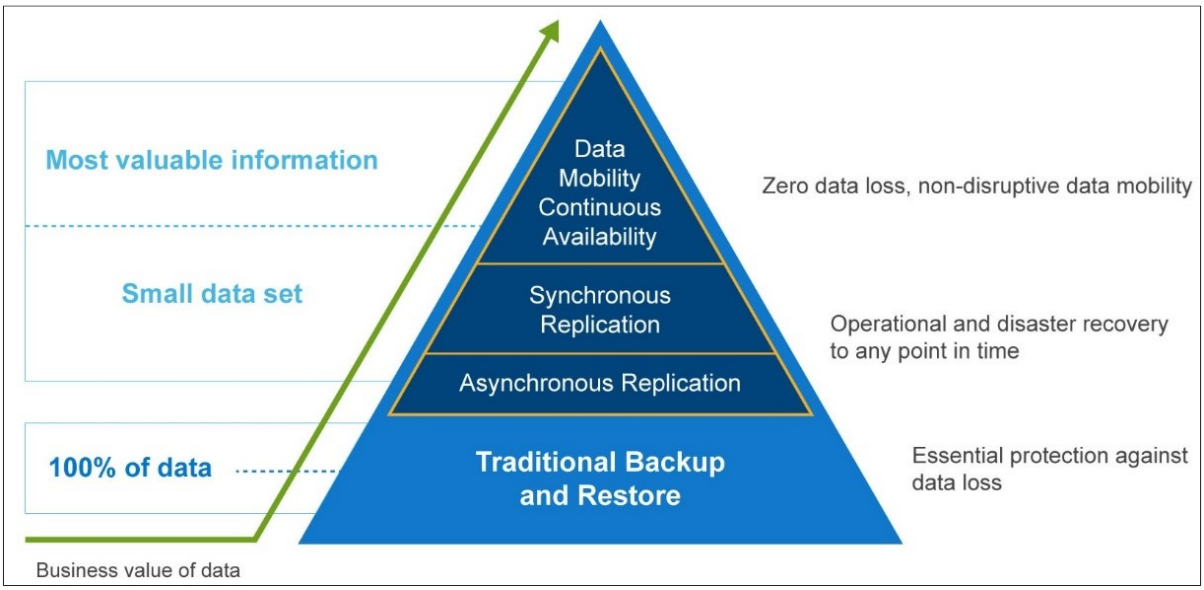


**Figure 3. Reduced datasets**

# Threats

Security has quickly risen to the top of CIO initiatives because of the rapid increase of threats that have surfaced in the past few years. Some examples include, unauthorized access into systems for financial gain and stealing information or damaging data, infrastructure and networks. These threats have made many people realize that security cannot be an afterthought but must be built from the ground up.

Some of the features that any hyperconverged infrastructure must include:

- Application of only properly signed firmware update packages, preventing vulnerabilities and data loss
- Use of secure management protocols such as RESTful APIs
- Lockdown mechanisms that prevent unauthorized changes to firmware and configuration
- Safe data erasure and encryption methodologies for data that resides on the systems

# Digital transformation and the rendezvous with the hybrid cloud

The investments that organizations make today will define the growth and performance over the next few years. Digital transformation happens both for consumers of goods and services as well as the suppliers of these goods and services. The four primary outcomes that companies are seeking from digital transformation are improved ways of customer engagement, empowering employees, optimizing operations and transformation of products.


**Figure 4. Four primary outcomes of digital transformation**

These digital transformation outcomes drive customers to hybrid cloud operating models and operations modernization through automation. To achieve these outcomes, Dell Technologies and Microsoft offer a unique Azure hybrid ecosystem that simplifies operations and enables innovation. The key questions that IT professionals are trying to answer are:

- How to simplify governance of resources and secure them seamlessly regardless of where they reside?
- How to bring cloud innovation to existing infrastructure?
- How to modernize on premises data centers to unlock the full potential of a hybrid cloud?

Dell Technologies and Microsoft designed a hybrid cloud solution to provide answers to these questions. Within the hybrid cloud strategy, important design considerations include:

- Some workloads cannot migrate to a public cloud due to data gravity (law of performance) and governance restrictions. The latter is common when dealing with personally identifiable information in industries like finance, health care and government
- Some workloads require unique low latency for faster AI capabilities (law of performance)
- Many companies want to leverage their current investments in infrastructure to have the best cost structure for their workloads (law of economics)

## According to analyst research

- 92 percent of organizations have both public and private cloud environments installed. (IDC's CloudPulse Q119, June 2019)

- 82 percent of organizations are already working with three or more cloud service providers outside their own organization. (ESG Research Insights Paper, “The Cloud Complexity Imperative,” February 2020)
- Data does not stay on one of these many clouds. Only 50 percent of all applications are expected to “stay in place” over the next year. (IDC’s CloudPulse Q119, June 2019)

# Accelerate your digital transformation journey

Dell Technologies can help organizations in their transformation journey by simplifying the deployment and management of their infrastructure with the introduction of HCI. A modern HCI platform is ideal for refreshing aging hardware, consolidation of virtual workloads and streamlining access to public cloud services. HCI is built on a software-centric design that tightly integrates compute, storage and virtualization resources in a single platform that is easy to manage and maintain. This infrastructure is enabled without sacrificing reliability and uptime while saving time and money spent on building and maintaining IT systems.

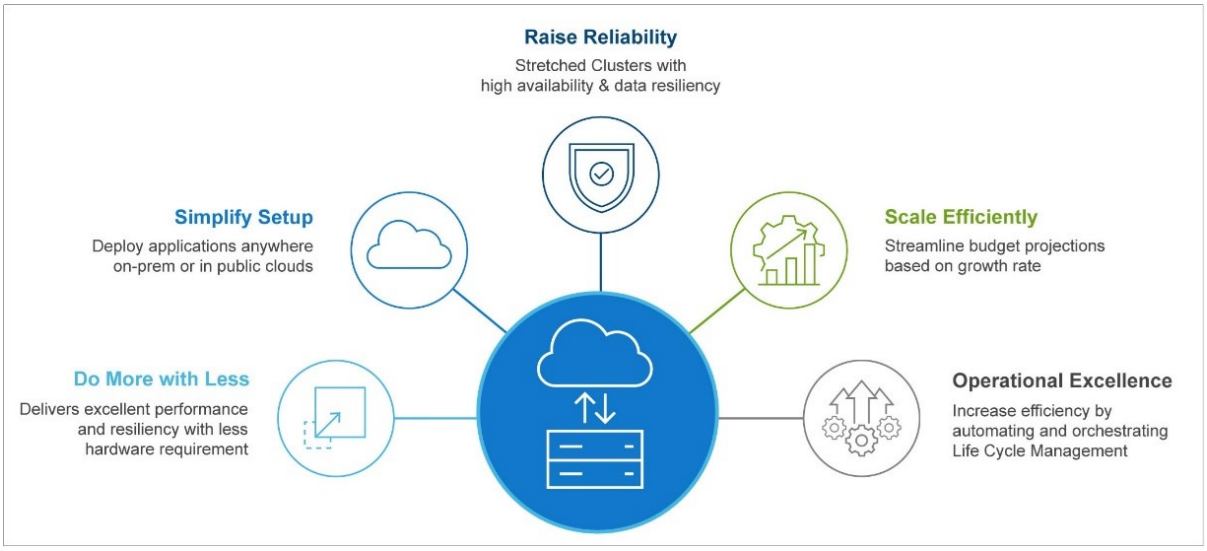


**Figure 5. Benefits of HCI**

Business goals and application requirements determine the performance that must be delivered from an HCI system. There are many possibilities when customers are equipped with an HCI appliance that can run their most critical workloads coupled with the ease of deploying applications anywhere on-premises or in a public cloud.

HCI solutions require less hardware to function when compared to three-tier architecture and provide the same great performance and resiliency as the traditional infrastructure. Organizations can realize a lower Total Cost of Ownership (TCO) with an immediate gain in productivity and reduced complexity.

HCI offers the ability to add both compute and storage capacity to the system by adding more cluster nodes. This ability is convenient for scaling resources as business needs grow. "Start with current needs and scale incrementally" is the approach used with HCI systems. Budget projections can be streamlined as organizations now know the exact number of nodes required to operate and can scale linearly and predictably depending on the business needs.

HCI provides a resilient foundation by featuring native DR for clusters combined with storage replication for volumes assuring business continuity in case if there is a catastrophic failure. Stretched clusters can be employed to spread workloads across two physical locations offering high availability and data resiliency. Automated failover recovers business services quickly, without the need of any manual intervention.

Dell Integrated System for Microsoft 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 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.

Dell Technologies is your single source for design, purchase, rack and stack, deployment, and comprehensive hardware and software support for Microsoft Azure Stack HCI solutions. Microsoft has consolidated all management tools into a single intuitive server management experience with Windows Admin Center (WAC) which modernizes and simplifies IT operations. Customers benefit from the Dell OpenManage Integration with Microsoft Windows Admin Center, which provides deep

hardware monitoring, inventory, and troubleshooting capabilities. It also has 1-Click Full Stack Lifecycle Management using Cluster-Aware Updating, automated cluster creation, dynamic CPU core management, and cluster expansion features.


**Figure 6. Dell leading the HCI market**

The Evolution of Datacenters to HCI and Hybrid Cloud chapter describes about how data centers have evolved over the years and how the hybrid cloud and HCI are becoming the new norm for modern infrastructures.

3

# Evolution of Data Centers to HCI and Hybrid Cloud

**Topics:**



## Overview

Never in the history of IT operations has the data center been more critical to the success of the global economy. Though increasing numbers of applications and data are moving to the edge, data centers remain the fulcrum of modern enterprises by hosting infrastructure at mass scale. These facilities can be corporate-owned and operated, secured in co-location centers, or fully managed by cloud service providers like Microsoft and Amazon. The rapid pace of technology innovation has required data center infrastructures to evolve to support the deluge of data, Internet of Things (IoT), and Artificial Intelligence (AI) and Machine Learning (ML). In this chapter, we discuss the journey from mainframes to modern hyperconverged systems running as part of a hybrid cloud computing operating model.

A data center is a physical facility in an organization that centralizes IT operations for hosting applications and critical data. The building blocks of a data center design include servers, storage, and networking components such as routers, switches, and firewalls that enable the delivery of shared applications and data. Facilities considerations must also be considered including the use of Uninterrupted Power Supplies (UPS) and Computer Room Air Conditioners (CRAC) to maintain environmental conditions vital for IT equipment operations.

Change is the new constant. The business climate, customer demographics, and target markets are changing radically, and it is evident that data centers are also evolving along with business needs. The following figure shows how data centers have evolved over the years:

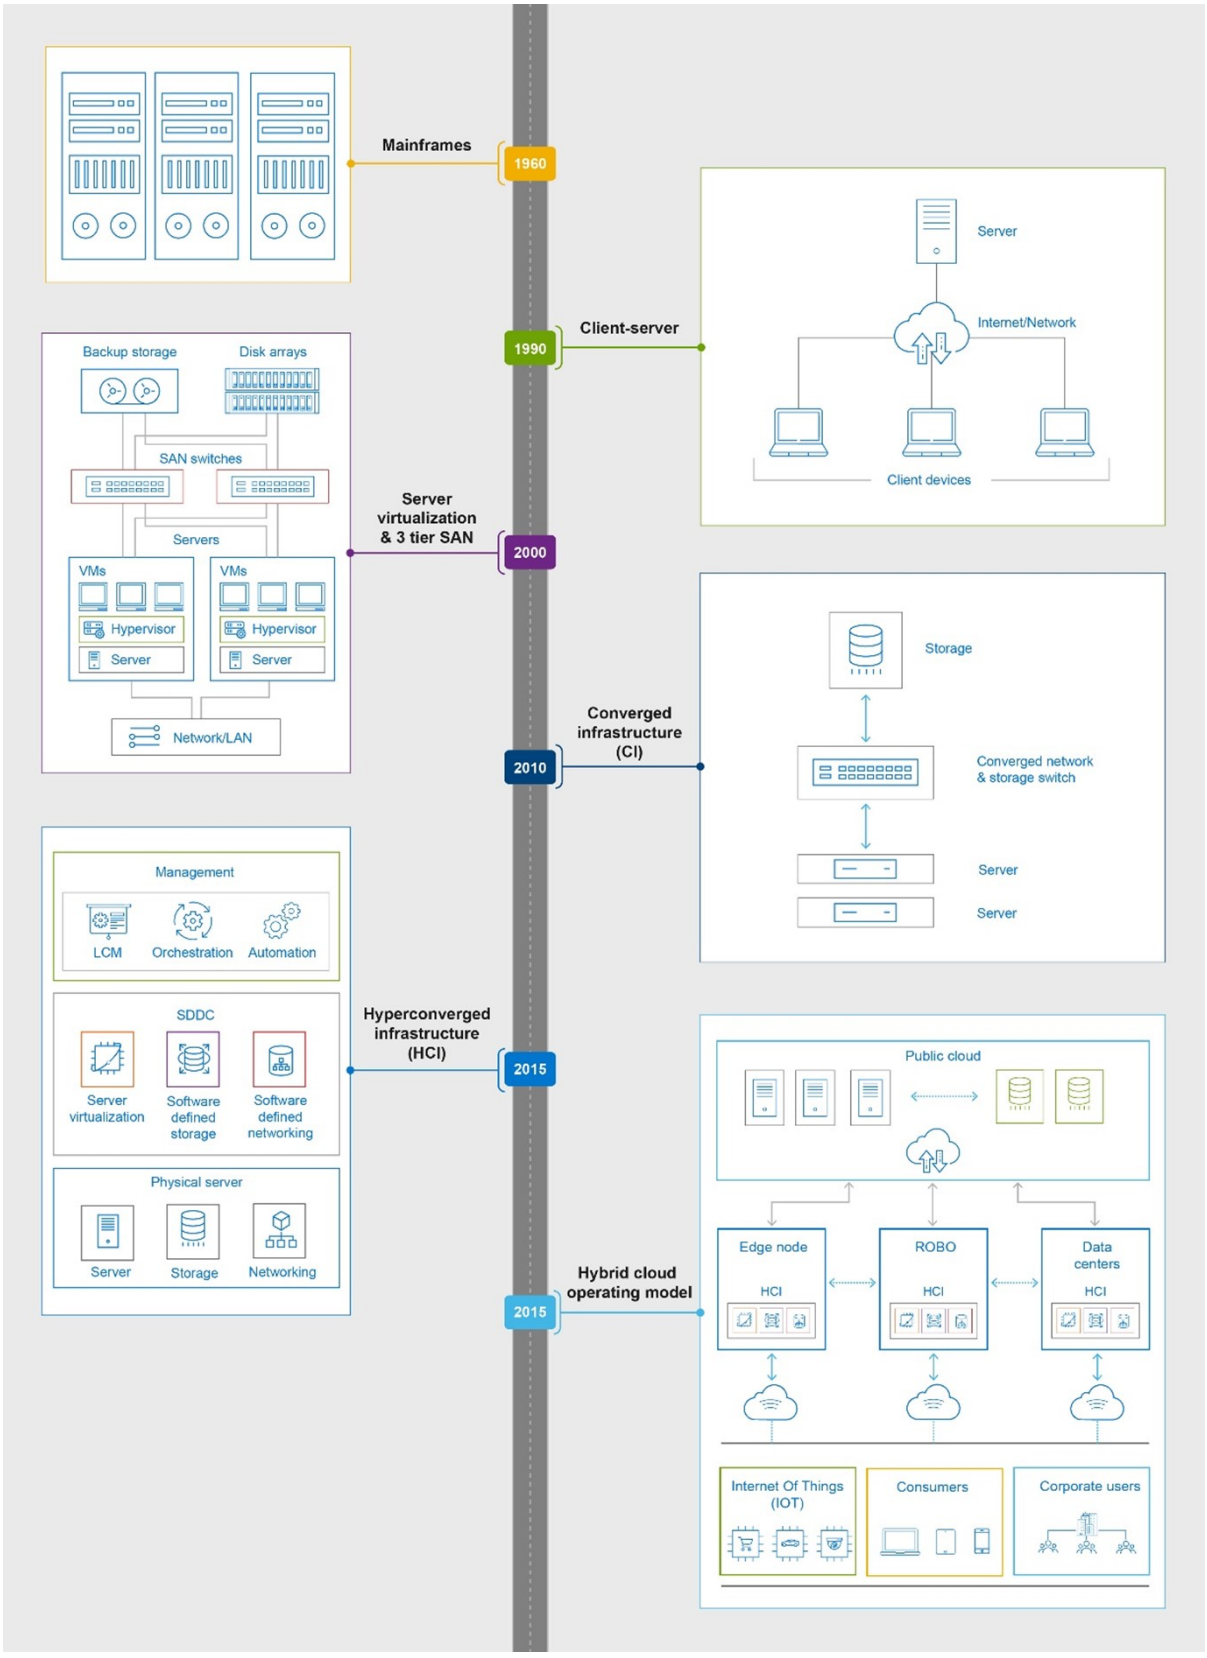


**Figure 7. Evolution of data centers**

# Mainframes – 1960s

Mainframe systems have been around for decades and remain operational today. When properly deployed, mainframes provide peak reliability, availability, serviceability, and security. They provide high throughput and process billions of transactions per day. Mainframes are still used today to run mission-critical applications in banking and financial institutions, airline, and retail industries.

Initial investment in mainframes was high because proprietary hardware and software made it expensive. IT operations required specially trained staff. Expertise to run present day mainframes is in increasingly short supply. Mainframes also consumed significant real estate and required a substantial investment in power and cooling to operate. These characteristics made mainframe technology only accessible to the largest organizations. To solve these challenges, the IT industry began moving to client/server architectures for hosting applications and data to lower costs and reduce complexity.

# Client-server architecture – 1990s

During the early rise of the Internet, client/server architectures started taking center stage. Servers intended for specific workloads were deployed. These servers were simpler to manage and more cost effective when compared to mainframes. Data centers became more prominent as companies started relying more on the large sets of data that led to growth of data center solution providers. The infrastructure decentralized the computing power and brought it closer to the developers and users. Applications became more UI based instead of command-line driven, which made it easier for business users to consume technology.

As this approach gained popularity and became ubiquitous, organizations used it to host a wide variety of back-end and front-end capabilities. Organizations raced to have a presence online with web applications such as Ecommerce engines to make it simpler for customers to obtain products and services. The more applications and data grew to support business functions, the more server, storage, and networking hardware was needed. Often, this equipment was designated for each department, different applications, and management services like backups, DR, and monitoring. The amount of power and cooling required increased substantially with the number of servers running in data centers. Organizations experienced situations in which a cluster of servers were not used to their full capacity. Low server consolidation led to expenses for unused space, power, and cooling. This scenario is known as “server sprawl” and the remedy was virtualization.

# Server virtualization and 3 tier SANs (iSCSI and FC) – 2000s

Virtualization was a breakthrough innovation in the world of data centers. Server virtualization introduced a software layer called the hypervisor that abstracted the CPU, storage, and memory of physical servers into pools of virtual resources. Virtual machines (VM) running on the hypervisor hosts ran their own operating systems and shared processor, memory, and storage. Applications could be hosted on multiple VMs residing on fewer physical servers -- enabling significant cost savings from consolidation and portability. Virtualization was appealing at a high level for consolidation, potential cost savings, simplified maintenance, and far more flexibility in backups and DR.

The classic three-tier architecture with Storage Area Networks (SANs) consisted of compute, storage, and networking that were disaggregated but connected. These architectures started gaining prominence due to their dedicated high-speed networks specifically purposed for data storage. SANs were built on Fiber Channel technology (FC) or on Internet Small Computer Systems Interface (iSCSI) with the latter being less expensive. SANs provided improved data security, highly responsive backups, increased scalability, and reliable disaster recovery.

Fairly complex to setup, configure and maintain, the SANs often required a substantial capital outlay and had their own share of problems, such as hardware compatibility, performance, and multivendor support issues. To address these challenges, Converged Infrastructures (CI) solutions were designed with highly integrated configurations.

# Converged infrastructure (CI) – 2010s

In-line with the SANs arrived the convergence of compute, storage, and networking on a pre-qualified turnkey appliance known as Converged Infrastructure (CI). It was known as, “data center in a box” and simplified all aspects of IT by seamlessly integrating all the system, network, storage, data protection and cloud management technologies into one engineered system. The deployment process was shortened with vendor validated solutions that reduced the time required to develop applications with trusted infrastructure and a single source of support.

Converged architecture contained a hardware-focused, building-block approach as the components were assembled on a siloed architecture sourced by multiple vendors. The configurations which were validated and delivered were rigid and had less room for hardware choices and incrementing resources later. Simplifying management of devices, automation, and Life Cycle Management (LCM) was quite a challenge. The progression towards improvising arose in the next generation of CIs known as HCIs.

# HCI – 2015s

The next advancement was towards virtualizing the three fundamental data center building blocks of server, storage, and networking to form Software Defined Data Centers (SDDCs). SDDCs resemble traditional data centers in terms of physical hardware but are characterized by a high degree of virtualization, abstraction, resource pooling, and automation. SDDCs eliminated the IT operational siloes and complexities by enabling a programmatic approach through virtualization.

HCI built on SDDC capabilities emerged as a foundation for running hybrid cloud IT operations. It uses software to assimilate core storage and compute functionality into a single highly virtualized solution. IT departments use HCI to lower total cost of ownership (TCO), increase performance and scalability of IT services, and improve overall productivity. Some HCI solutions can also seamlessly connect to cloud environments and leverage fully managed infrastructure and platform services. Designed with full-stack life cycle management, a fully productized HCI system can deliver efficient operations, flexible consumption models, and high-level enterprise capabilities.

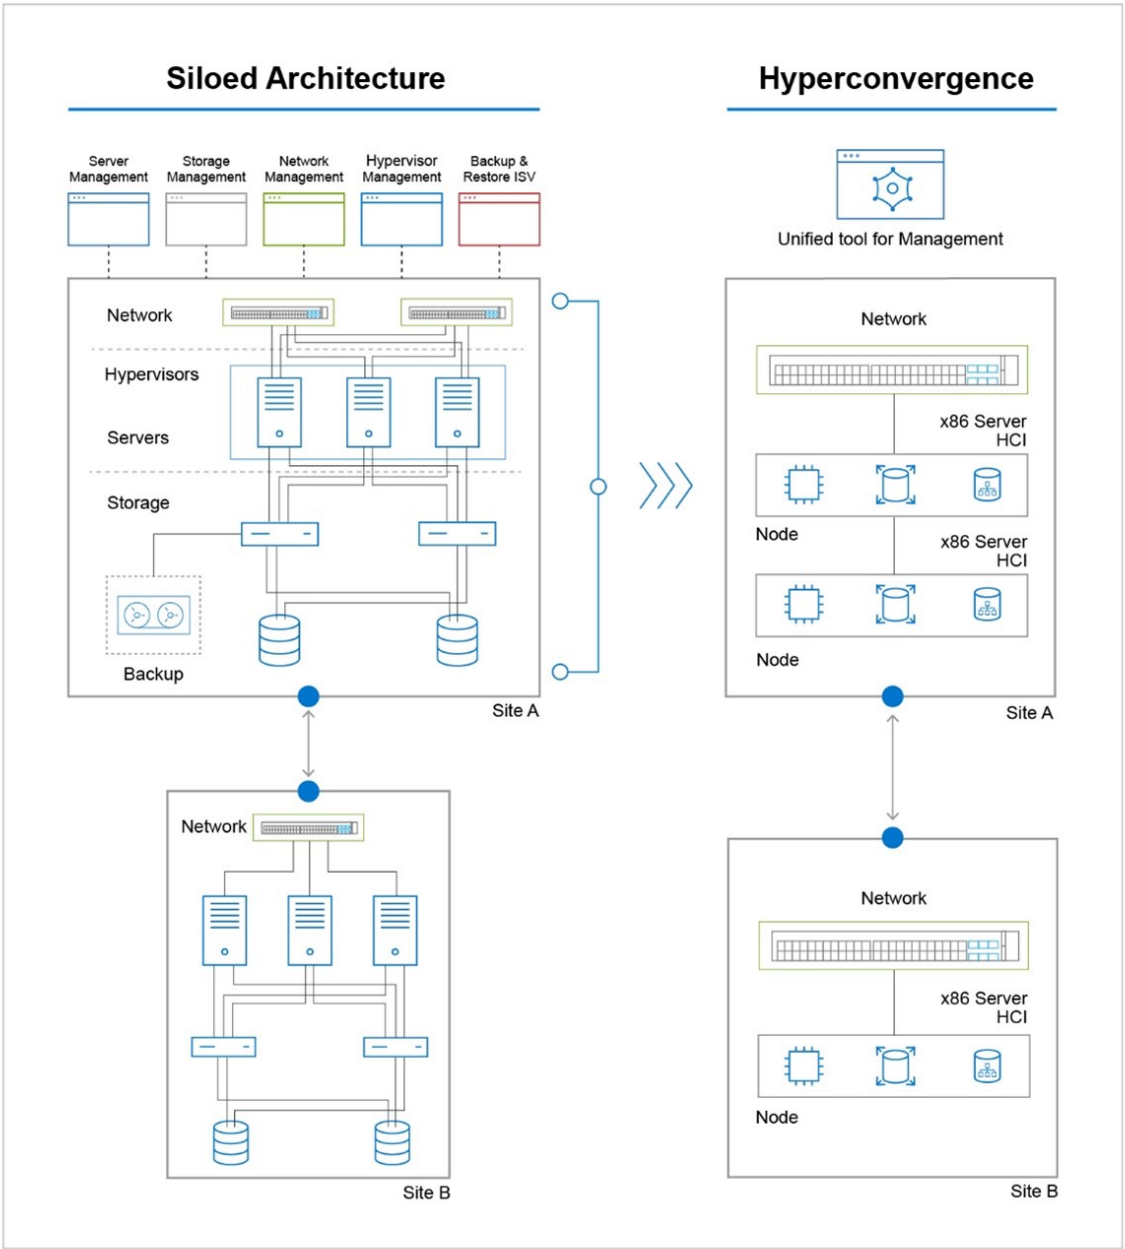


**Figure 8. Multi-tiered architecture collapsing to HCI nodes**

# Hybrid cloud operating model – 2015s

Hybrid cloud is a new model that enables end users to consume IT resources from any device, anywhere, and anytime. Hybrid cloud operating models allow companies to deploy their applications and host their data where it will most impact end users' response times and overall digital experience – from data center to edge to public cloud. Simplicity, flexibility, and on-demand consumption delivered by a public cloud provider, coupled with controlled, trusted, and secured on-premises platforms makes the hybrid cloud model the preferred strategy for most modern IT organizations.

The criticality of services deployed on infrastructure at the enterprise edge cannot be understated. The advent of Internet of Things (IoT) has ushered in the most transformational time in the human history. Consider self-driving cars that can learn and think, smart home security systems that can detect intrusion, or wearable health monitors that can predict medical emergencies. All of these are changing the way we live. Tremendous amounts of data are generated by IoT devices and fuel intelligent decision making from Artificial Intelligence and Machine Learning (AL/ML) systems. This deluge of data often requires computing power and storage capacity that resides near these devices to avoid latency in processing the information.

The A Closer Look at Microsoft Azure Stack HCI chapter provides a good insight on Microsoft Azure Stack HCI, a new HCI operating system that is delivered as an Azure service -- providing the latest security, performance, and feature updates with extended connection to Azure cloud services.

4

# A Closer Look at Microsoft Azure Stack HCI

**Topics:**



## Microsoft Azure Stack HCI – an overview

KEY TAKEAWAY

Introducing the Azure Stack HCI, the new addition to the Microsoft family. Delivered as an Azure service that integrates on-premise operations with the world wide capabilities of Azure, Azure Stack HCI allows you to discover and manage in-house IT infrastructure anywhere across the globe.
With Dell Integrated System for Microsoft Azure Stack HCI, you can modernize your infrastructure, consolidate virtualized workloads, and gain cloud efficiencies on-premises.

In December 2020, Microsoft announced the new Azure Stack HCI experience with features designed to transform traditional IT environments into highly agile hybrid cloud ecosystems. At its core is a purpose built hyperconverged infrastructure operating system delivered as an Azure service. Azure Stack HCI was deliberately created to deliver enterprise-class software defined capabilities to support any workload at the optimal price-performance ratio. IT organizations can elect to leverage familiar tooling, processes, and competencies for life cycle management. They can also natively integrate with Azure management, governance, data, and application services using Azure Arc.

The entire Azure Stack HCI architecture is inspired by Azure, using Hyper-V, Storage Spaces Direct, and Software Defined Networking technologies that have operated at massive scale in Azure for years. Traditional monolithic, cloud-native, microservice-based applications run securely, reliably, and at the desired performance levels within Windows and Linux-based virtual machines (VMs) and containers. A new feature exclusive to Azure Stack HCI is stretch clustering for disaster recovery. This stretch clustering solution provides automatic failover of VMs to restore production quickly and without the need for manual intervention.

Azure Stack HCI is hybrid by design. Each cluster is a unique resource manageable by Azure Resource Manager (ARM) using the Azure portal or through programmatic means such as Azure PowerShell and Azure CLI. Since it is a discrete Azure service, Azure Stack HCI always remains current with a rapid cadence of feature and security updates. It also uses an Azure subscription-based billing model at the host layer – charging a set fee per month for each core enabled on the cluster.

Microsoft Windows Admin Center is used as the always-on, edge/local management console. Complementary and consistent with Windows Admin Center is Azure Arc, which allows technologists to manage and operate their entire IT estate at scale from Azure. Azure Arc-enabled infrastructure technologies project non-Azure resources into Azure so they can be managed with services like Azure Monitor, Azure Security Center, Azure Backup, and Azure Site Recovery. With Azure Arc-enabled services

– including data and application services – software developers can build their applications in Azure and deploy them to Azure Kubernetes Service on Azure Stack HCI clusters located anywhere.

Two different deployment options are available for Azure Stack HCI. Customers can choose to obtain validated nodes based on standardized reference architectures or integrated systems. Integrated systems provide the fastest time to value, as they include the Azure Stack HCI operating system pre-installed in the factory with the hardware partners' firmware and drivers applied. At Dell Technologies, we take the integrated system approach and include capabilities that ensure success at each stage of the solution life cycle.

# Why choose Dell Technologies for Microsoft Azure Stack HCI

Dell Technologies has a distinct advantage over the competition because of our track record of delivering Microsoft solutions. With each evolutionary milestone in our journey for continuous innovation, we have learned and re-invested back into our Microsoft software-defined data center (SDDC) and HCI solutions. The following figure illustrates the SDDC and HCI milestones in the Dell Technologies partnership with Microsoft over the last several years.

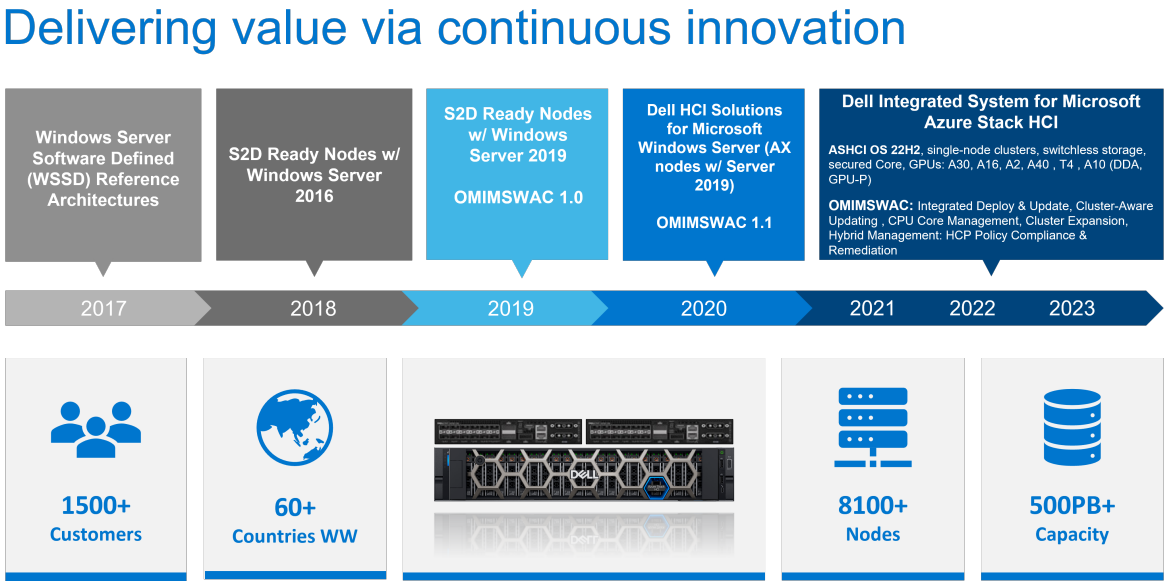


**Figure 9. Dell Technologies partnership with Microsoft**

We began this journey in 2017 when we certified Windows Server Software-Defined (WSSD) solutions on Dell PowerEdge 13th generation (13G) servers. The reference architectures explained how to implement Microsoft validated designs on PowerEdge servers, following engineering best practices for deployment and steady-state operations.

In 2018, we became the first solution provider to take a fully productized approach to running Microsoft software defined capabilities, including Hyper-V and Storage Spaces Direct (S2D). Our Dell S2D Ready Nodes based on Dell PowerEdge 14th generation (14G) servers were pre-configured with certified components and validated building blocks running Windows Server 2016. Dell engineering also validated network topologies using Dell PowerSwitch network switches. We vastly simplified ordering and reduced deployment risks while providing a streamlined and collaborative Customer Support experience.

The S2D Ready Nodes were refreshed in 2019 with the validation of Windows Server 2019 (End of Life June 30, 2023) on the supported platforms. We also introduced the 1.0 version of our Dell OpenManage Integration with Microsoft Windows Admin Center. The Dell OpenManage Integration extension added deep hardware monitoring, inventory, troubleshooting, and integrated iDRAC management for the S2D Ready Nodes.

More recently, our portfolio expanded in 2020 when we released the Dell HCI Solutions for Microsoft Server, based on Windows Server 2019 (End of Life June 30, 2023) running on our factory-ready, validated AX nodes from Dell Technologies. The AX nodes consisted of two Intel-based 14G models (AX-640 and AX-740xd, End of Life March 31, 2023) and a single AMD EPYC processor-based 15G model (AX-6515). The Dell OpenManage Integration with Windows Admin Center was updated to v1.1, which added life cycle management of BIOS, firmware, and drivers using Cluster-Aware Updating. Using this new feature, IT staff could apply hardware updates with confidence and save significant time, reducing manual steps by 82%[3] and reducing maintenance windows by 40%[4]. Clusters could be fully updated with only a single reboot per node in the cluster and no impact to the business services running in virtual machines.

After the release of the purpose-built Microsoft Azure Stack HCI operating system in December 2020, Dell Technologies released the Dell Integrated System for Microsoft Azure Stack HCI. This integrated system included AX nodes tested and validated to optimize the performance and resiliency of Azure Stack HCI clusters. By September 2021, the AX node portfolio consisted of the AX-650 and AX-750, based on Intel Xeon 3rd Gen scalable processors, and the AX-6515 and AX-7525, based on 2nd generation AMD EPYC processors.

In March 2022, Dell Technologies extended the support for all AX platforms to run Azure Stack HCI OS, version 21H2 and Windows Server Datacenter 2022.The Dell OpenManage Integration with Windows Admin Center was updated to v2.2.1 to support the new operating systems and added features such as intrinsic infrastructure security management with Microsoft Secured-Core sever and Dell Infrastructure lock. We also added support for Kernel Soft Reboot (KSR) for updates.

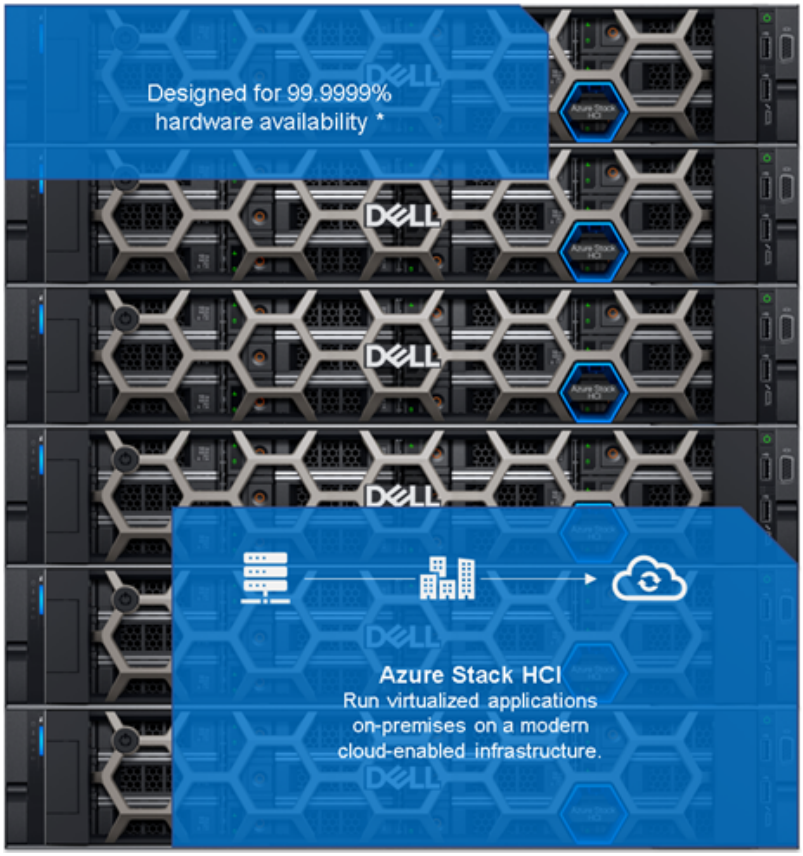


**Figure 10. Dell AX-740xd cluster**

Another version of the Dell OpenManage Integration for VMware vCenter, v2.3, was released in May 2022. This version added Dell HCI Configuration Profile Policies for Azure, which creates Azure Policy definitions for our hardware, operating system, and cluster-level recommended settings. Then, we provide compliance visibility and remediation for settings that deviate from our best practices.

The Dell OpenManage Integration extension also has numerous other features that continue to make IT administrators' lives easier and ultimately provide the business with improved response times and reliability for their mission critical applications. These features include automated cluster creation, one-click full stack life cycle management using Cluster Aware Updating, dynamic CPU core management, and cluster expansion preparation.

The second half of 2022 releases introduced support for NVIDIA A30 and A2 GPU adapters, single-node clusters and 20TB capacity 3.5" drives.

More recently, the 2023 quarterly releases added support for:

- Azure Stack HCI OS 22H2 (factory install and field support)
- New GPU options (NVIDIA A16, A40, and T4)

---

3 Internal Dell Technologies lab review of OpenManage integration with Windows Admin Center. Taken from "Dell Solutions for Microsoft Azure Stack HCI: Life Cycle Management Approach Comparison Benefits of Automating the Azure Stack HCI Hardware Update Process", December 2020

4 Internal Dell Technologies lab review of OpenManage integration with Windows Admin Center. Taken from "Dell Solutions for Microsoft Azure Stack HCI: Life Cycle Management Approach Comparison Benefits of Automating the Azure Stack HCI Hardware Update Process", December 2020

- New NIC options with Mellanox ConnectX-6 25 GbE (RoCE) and Intel E810 25/100 GbE (iWARP and RoCE) adapters
- Validation of Mellanox CX-6 25G OCP card for Factory Installation
- Validation of A10 GPU for After Point of Sale (APOS)
- Windows Server 2019 End of Life
- EOL for AX-640 and AX-740xd
- AX-6515 End of Life for CE countries 11/6/23

# What makes Dell Integrated System for Microsoft Azure Stack HCI different

The Dell Integrated System for Microsoft Azure Stack HCI goes far beyond a reference architecture or validated node by delivering efficient operations, flexible consumption models, and high-level enterprise expertise. Weaving simplicity into every aspect of the buyer's experience, we have made it increasingly easy for customers to order, deploy, maintain, and support their modernized, on-premises infrastructure. The following figure summarizes the unique and differentiating pillars of our integrated system.

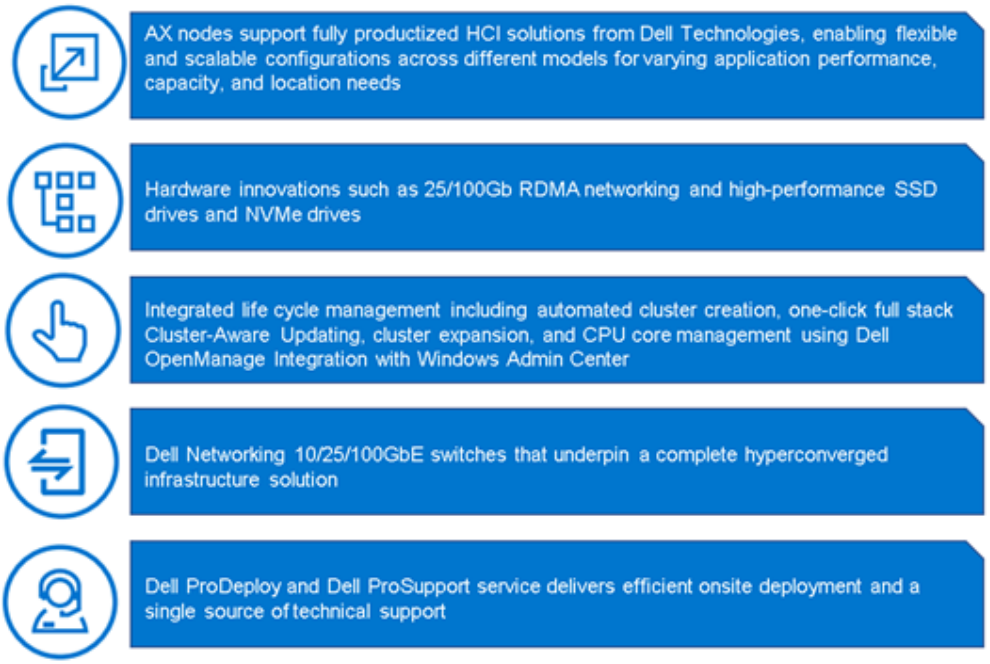


**Figure 11. Pillars of our integrated system**

Dell Integrated System for Microsoft Azure Stack HCI is an all-in-one fully productized HCI solution with flexible factory-ready AX nodes as the foundation. We have a broad portfolio of AX nodes that are intelligently and deliberately configured with a wide range of component options to meet the requirements of nearly any use case – from the smallest remote or branch office to the most demanding database workloads. The chassis, drive, processor, DIMM module, network adapter and their associated BIOS, firmware, and driver versions have been carefully selected and tested by the Dell Technologies engineering team to optimize the performance and resiliency of Azure Stack HCI. Our engineering has also validated networking topologies using Dell PowerSwitch network switches.

Azure Stack HCI is one of the best hyperconverged platforms in terms of performance. The combination of the intelligent design of the storage stack, combined with drives inside the nodes (closer to processor and memory vs external storage) and the use of RDMA protocols on networking lead to unparalleled disk IO and very low latency. In internal testing, Dell has observed over 1,000,000 IOPS per cluster node at sub-millisecond latencies. This is great for performance sensitive applications like SQL databases as well as other transactional systems. It also means that a desired performance level can be achieved with less expensive drives which saves money. Another added benefit is that disk performance is likely to no longer be a concern, which provides investment protection and added longevity of the system.

For workloads that demand peak performance, we have created higher-end configurations including 25/100GbE RDMA networking, Intel Optane persistent memory (Windows Server HCI only), and high-performance PCIe Gen 4 non-blocking NVMe drives. Dell PowerSwitch 10/25/100GbE network switches are fully qualified to work seamlessly within the Azure Stack HCI architecture. There are three deployment options for Dell Integrated System for Microsoft Azure Stack HCI:

- The Scalable model supports from 2 to 16 nodes in a cluster and uses top-of-rack switches for management and storage traffic networking.
- The Switchless storage usage model uses full mesh connections between the cluster nodes for storage traffic and supports from 2 to 4 nodes in a cluster.
- Single-node clusters are the simplest way to set up Azure Stack HCI. They are a great fit for remote, edge, or branch office projects because they are very sensitive to cost and space requirements.

**NOTE:** Dell Technologies recommend to setup data centers with 25 GbE network switches as the price per port is almost identical to 10 GbE network switches. This gives an advantage for IT organizations to achieve maximum performance with the right cost.

Arguably the most compelling aspect of our integrated system is our life cycle management capability.

- The Integrated Deploy and Update snap-in works with the Microsoft cluster creation extension to ensure a consistent, automated initial cluster creation experience on Day 1. Our snap-in guarantees adherence to our configuration best practices using Dell HCI Configuration Profile. The snap-in also updates BIOS, firmware, and drivers as part of the cluster creation workflow.
- When administrators are ready to apply updates, there is the one-click full stack life cycle management snap-in for the Microsoft Cluster-Aware Updating extension. This seamlessly orchestrates operating system, BIOS, firmware, and driver updates through a common Admin Center workflow.
- There is also dynamic CPU core management, which allows customers to right-size their cluster and maximize TCO by only enabling the cores needed to support running workloads.
- Finally, guided scale out cluster expansion helps customer eliminate guesswork and dramatically decrease time consumed preparing new nodes for cluster expansion.

Dell Technologies makes support 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 those higher value priorities. Our one-stop cluster level support covers the hardware, operating system, hypervisor, and Storage Spaces Direct software, whether you purchased your license from Dell Technologies or from Microsoft.

# Technical use cases and workloads

Azure Stack HCI is an ideal platform for a number of technical use cases and workloads. Dell Technologies takes a workload-focused approach when determining what configurations we choose to design and validate. Our supported configurations can be found in the Microsoft Azure Stack HCI Catalog. The following figure summarizes the technical use cases and workloads best suited for Azure Stack HCI.

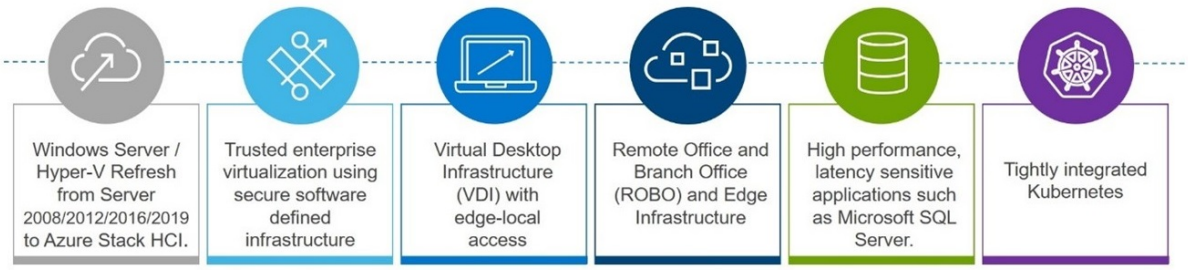


**Figure 12. Technical use cases**

## Windows server/ Hyper-V refresh

Many Microsoft-centric virtualization clusters running Windows Server and Hyper-V today are approaching End of Support (EOS) and End of Life (EOL). These clusters are running Windows Server 2012 or older on antiquated and inflexible infrastructure. Refreshing these environments is of paramount importance, and many customers would be best served adopting a modern HCI platform. They gain greater performance, reliability, and flexibility while also simplifying life cycle management and increasing operational efficiency.

## Trusted enterprise virtualization

With Azure Stack HCI, businesses can deploy trusted enterprise virtualization by providing a highly secure infrastructure for workloads through virtualization-based security (VBS). VBS relies on Hyper-V to implement a mechanism seen as Virtual Secure Mode (VSM) which creates a dedicated, isolated region of memory within its guest VMs. Through programming techniques, a designated security-sensitive operation can be carried out while blocking access to it from the host-OS, limiting potential vulnerability from kernel-based exploits.

## Virtual Desktop Infrastructure (VDI)

An Azure Stack HCI cluster is well suited for large-scale VDI deployments delivering user desktops with the help of virtual desktop brokers such Microsoft Remote Desktop Service (RDS) or Citrix Xen Desktop. Dell Technologies has also performed extensive testing with Azure Virtual Desktop, which will become generally available later in CY2022. Azure Stack HCI provides additional benefits by including centralized storage and enhanced security. It provides businesses a secure way to deliver client desktops on a wide range of devices without allowing users to store any data locally or upload any data from those local devices.

## Remote Office Branch Office (ROBO) and Edge Infrastructure

Azure Stack HCI meets the typical requirements for retail stores, branch offices, field sites, and other edge sites. With a low-level entry price point, customers can start with two-node clusters combined with a cloud witness or USB drive-based file share witness to set up their Remote Office Branch Office (ROBO) or edge locations. They can also avoid expensive network switch upgrades to support RDMA storage networking in ROBOs using a switchless storage networking topology. This topology relies on directly connecting Ethernet adapters that support RDMA between cluster nodes. From the management perspective, Azure Arc can be used to view and manage remote Azure Stack HCI deployments from Azure Resource Manager.

## Microsoft SQL Server

High performance, latency sensitive applications like Microsoft SQL Server greatly benefit from higher-end integrated system configurations. These configurations could consist of a single storage tier of all NVMe drives and 100GbE RDMA networking to maximize throughput and minimize latency. There is also a broad range of processor core counts in our portfolio for optimizing SQL Server licensing costs.

## Tightly Integrated Kubernetes

Azure Kubernetes Service on Azure Stack HCI (AKS-HCI) automates running containerized workloads at scale on-premises with fully integrated life cycle management and security. AKS-HCI is designed to be operated in a hybrid environment using code, consistent design patterns, and best practices from Azure Kubernetes Service (AKS). This means that software developers can start creating modern applications in Windows or Linux-based containers and run them on AKS or AKS-HCI without modifying code. With AKS-HCI in place, customers can use Azure Arc-enabled infrastructure and services to manage their entire IT estate from Azure and build cloud-native applications in Azure and deploy to AKS-HCI.

# Microsoft HCI Solutions from Dell Technologies - complete portfolio

Dell Technologies understands that there are certain situations where Dell Integrated System for Microsoft Azure Stack HCI may not meet an organization's requirements. That is why we will also continue to offer Dell HCI Solutions for Windows Server in our Microsoft HCI Solutions portfolio. The current integrated system comes pre-installed from the factory with Azure Stack HCI, version 22H2. The Windows Server HCI solution arrives with Windows Server datacenter 2022 OEM pre-installed with downgrade rights to Windows Server 2019 (EOL June 30, 2023) on AX-640 (EOL March 31, 2023), AX-740xd (EOL March 31, 2023), AX-6515, and AX-7525. After March 30, 2022 the new AX nodes will not be sold with Windows Server 2019 (EOL June 30, 2023) factory installed. These operating systems can also be manually installed. The following figure depicts the differences between the Windows Server operating system and Azure Stack HCI operating system.

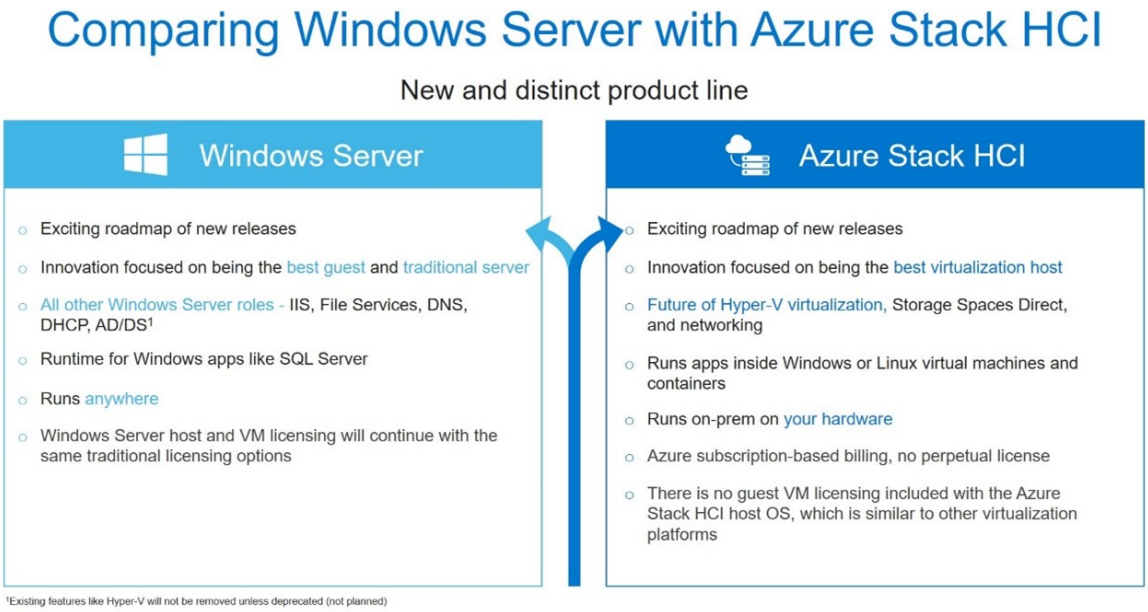


**Figure 13. Comparing Windows Server with Azure Stack HCI**

The Azure Stack HCI operating system is a new product line and is forked from the Windows Server operating system main development branch. All innovation and investment in the Azure Stack HCI operating system focus on supporting and further enhancing the software-defined capabilities of Hyper-V, Storage Spaces Direct, and Software Defined Networking. The Windows Server operating system remains a general-purpose operating system running the full complement of Windows Server roles and features such as IIS, DNS, DHCP, and Active Directory. The Windows Server operating system will also remain the runtime for the massive install base of Windows applications such as Microsoft SQL Server. Applications running on Azure Stack HCI will run inside of Windows or Linux-based VMs and containers.

Licensing is also different when comparing these operating systems. Windows Server host and VM licensing will continue to use Windows Server Data Center and Windows Server Standard licensing options. In contrast, Azure Stack HCI uses an Azure subscription-based billing model at the host layer. The cost at this layer is $10 per core per month. However, the Windows Server VMs running on the Azure Stack HCI platform still require Windows Server Data Center or Standard licensing.

In October 2022, during Ignite, Microsoft announced general availability for Azure Hybrid Benefit for Azure Stack HCI, which basically allows customers to bring their existing on-premises licenses to the cloud and significantly reduce the costs of running workloads in the cloud. Eligible customers with active Windows Server Software Assurance can now leverage Azure Hybrid Benefit to reduce costs on:

- Azure Stack HCI host fee: from $10/physical core/month to $0
- Windows Server unlimited virtualization on Azure Stack HCI: from $23/physical core/month to $0
- Windows Server unlimited virtualization on Azure Stack HCI: from $23/physical core/month to $0

This is a game changer for Azure Stack HCI, making it very cost effective for very high-level hybrid capabilities with Azure.

# Choosing the right Microsoft HCI solution

How does a business know when to choose Windows Server HCI compared with Azure Stack HCI? The following figure summarizes the high-level differences to consider when deciding which platform to choose to achieve specific business outcomes.

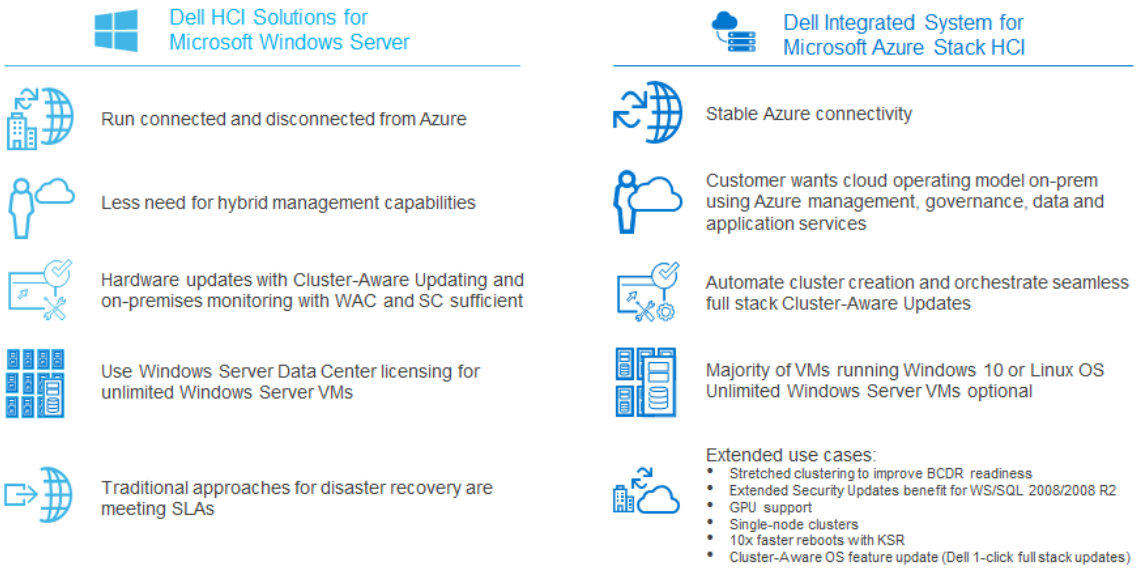


**Figure 14. Choosing the right Microsoft HCI solution**

The first question to address when considering one platform over another in the Microsoft HCI Solutions portfolio is: "Does the on-premises infrastructure have connectivity to Azure?". To remain fully functional, Azure Stack HCI requires connectivity to Azure and must communicate with the Azure subscription it is registered with once every 30 days. Windows Server HCI can run fully disconnected from Azure to protect highly secure and sensitive applications and data.

Management strategies will likely differ as well. Customers who are content with existing on-premises tools and processes may remain well served using Windows Server HCI. However, if customers prefer to move to a hybrid cloud operating model where they would like to use Azure for management and governance services for their non-Azure resources, Azure Stack HCI with Azure Arc is the ideal choice. Our Dell OpenManage Integration with Windows Admin Center also includes capabilities that are only supported for Azure Stack HCI, such as automated cluster creation and one-click full stack life cycle management using Cluster-Aware Updating.

As mentioned in the previous section, licensing approaches differ between these platforms. Some customers may prefer to continue using their Windows Server Data Center licensing on their Hyper-V hosts, which entitles them to run unlimited Windows Server VMs. Azure Stack HCI would be more appealing to organizations looking to transition to a more flexible cloud subscription-based approach. Though Windows Server VMs currently require Windows Server Standard or Data Center licensing in the Azure Stack HCI VMs, road maps indicate that future Windows Server editions will allow for subscription-based billing at the VM level.

Many organizations have legacy applications that can only run on older Operating Systems or Database Versions. This is why there is still a lot of Windows Server 2008, 2008 R2, 2012 and 2012 R2 in production. These versions have either reached or will reach the end of extended support in the near future. Microsoft offers Extended Security Updates for these operating systems and SQL server at no cost in Azure today, but that's not always a solution since a lot of these legacy applications simply won't run properly in Azure. Thankfully, customers can purchase Extended Security Updates for those OS and databases that need to remain on premises. Unfortunately, this is rather costly, which can become a burden on organizational budgets. If organizations choose to virtualize these older OS's (if not already) and decide to run these on Azure Stack HCI, these Extended Security Updates will be included at no additional cost. Instead of paying just for the security benefits, customers can now get a modern, more secure platform with all the benefits described in this document. This is a much better strategy and enables customers to really progress in modernizing their datacenters and adopting hybrid cloud while still supporting legacy Operating Systems and applications.

The Azure Stack HCI Architecture and Building Blocks chapter digs deeply into the building blocks of Azure Stack HCI, namely Hyper-V, SDS, and SDN, and its architecture in detail.

# 5

# Azure Stack HCI Architecture and Building Blocks

**Topics:**

# Overview

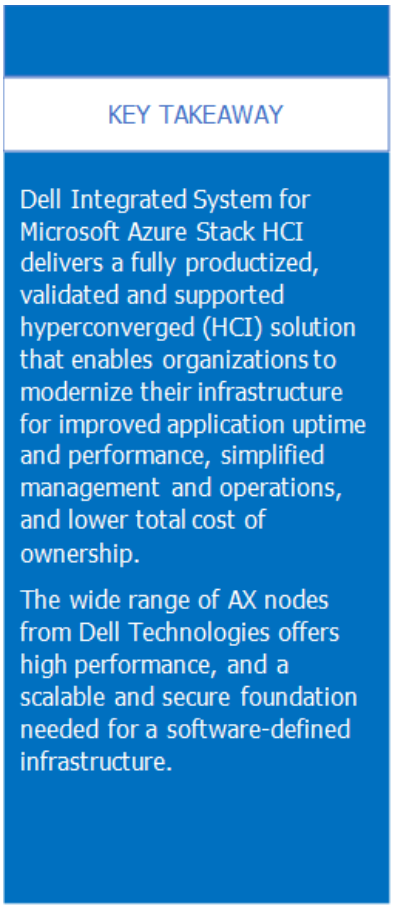


Azure Stack HCI delivers an effective blend of application performance and ease of management. This makes it a great solution for remote office/edge use cases, desktop virtualization such as Virtual Desktop Infrastructure (VDI), Microsoft SQL Server and other resource intensive workloads, and container orchestration platforms such as Kubernetes. It is structured to use Software Defined Data Center (SDDC) technology from Microsoft for the infrastructure elements, namely compute, storage, networking, and security. Effective for running existing apps with improved efficiency, it also lets you connect to Azure for hybrid cloud scenarios such as backup, cloud-based monitoring, and other Azure services.

Azure Stack HCI is factory installed on certified AX nodes from Dell Technologies, and supports familiar Windows Server Datacenter components and management tools such as:

- Hyper-V (SDC)
- Storage Spaces Direct (S2D)
- Software Defined Networking (SDN) (optional)
- Windows Admin Center
- Azure services
- PowerShell for scripting and automation

In this chapter, we will deep dive into the solution architecture of Azure Stack HCI with a walkthrough into its building blocks:

- Hardware ecosystem
- Elements of Software Defined Data Center (SDDC)
- Container platforms
- Hybrid Cloud scenarios
- Management tools

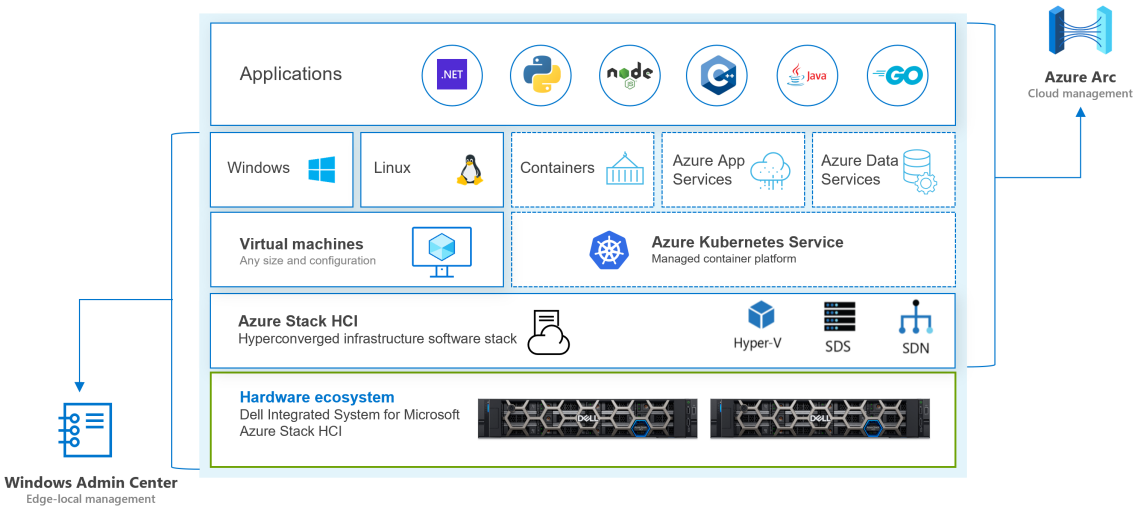


**Figure 15. Solution architecture of Azure Stack HCI**

# Hardware ecosystem

Central to Dell Integrated System for Microsoft Azure Stack HCI are the HCI solution-oriented AX nodes. These validated and certified AX nodes from Dell Technologies offer high performance, and the scalable and secure foundation needed for a software-defined infrastructure. Though they are built on the design of the popular Dell PowerEdge servers (Dell Technologies holds the No. 1 position in x86 servers from 2017 to 2019[5] AX nodes are deliberately configured and optimized by our engineering organization exclusively for Azure Stack HCI. AX nodes offer numerous configurations so that customers can choose the memory, storage, and compute capacities that align with their needs.

Dell Technologies has validated and certified an array of configurations across various node types in the solution catalog. The depth of offerings gives customers options and flexibility in their Azure Stack HCI deployments. When selecting nodes, one can choose either Intel Xeon Scalable processors or AMD EPYC processors. These can be configured with NVMe, SSD and HDD devices to create hybrid or all-flash nodes. The nodes can be networked by using the 2-node switchless clusters method or by using top-of-rack Dell PowerSwitch Data Center network switches for the storage network traffic. For detailed specifications for the distinct platforms the AX nodes offer, see Chapter 6: Dell Integrated System for Microsoft Azure Stack HCI Portfolio – AX Nodes.

Customers also get the benefit of simplified life cycle management (LCM) with the Dell OpenManage Integration with Microsoft Windows Admin Center (OMIMSWAC). This orchestrated LCM reduces the time required for tested and validated upgrades and allows IT teams to focus on higher value tasks.

# The hyperconverged infrastructure software stack

Hyperconverged infrastructure (HCI) unifies servers, storage, and networking into a distributed infrastructure platform, powered by intelligent software, to create the elements of a data center. It replaces the legacy infrastructure that consists of separate servers, storage networks, and storage arrays with an all-in-one integrated solution. HCI works on the principle of a software-defined data center (SDDC), in which infrastructure is virtualized and delivered “as-a-service” that includes:

- Software-defined compute – for virtualization and workload management
- Software-defined storage – for storage abstraction and data management
- Software-defined networking – for network virtualization and traffic management

Traditional enterprise applications can be supported in a more flexible and cost-effective manner because SDDC eliminates IT operational siloes and complexities by enabling a programmatic approach through virtualization.

5 PRESS RELEASE. "Dell Advances World's Top Selling Server Portfolio". Dell Technologies. https://corporate.delltechnologies.com/en-us/newsroom/announcements/2019/04/20190402-01.htm

# Hyper-V

Microsoft Hyper-V is the hypervisor powering Microsoft Azure at global scale. Its primary responsibility is to run virtual machines (VMs). VMs run in their own isolated and dedicated space, which means that more than one VM can run simultaneously on the same physical hardware. This is done by abstracting the physical hardware and dividing it into logical units, providing organizations the benefit of business agility and lower TCO (Total Cost of Ownership). The hypervisor uses a bare metal architecture and runs directly on the hardware platform. The guest operating systems run one level above the hypervisor and inside the virtual machine.

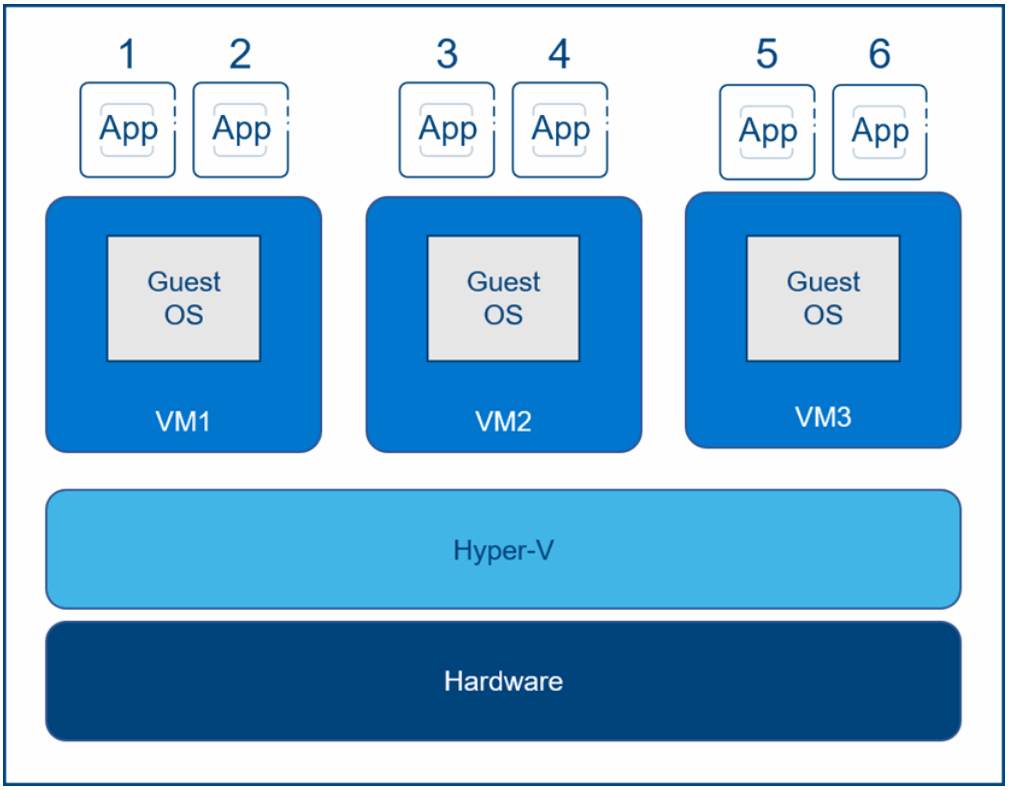


**Figure 16. Architecture diagram of Hyper-V**

Hyper-V is the foundation of the Microsoft Hybrid Cloud and is the first step in enabling the dynamic data center. Hyper-V is on the forefront in establishing and expanding a private cloud environment by providing a foundation for flexible and on-demand services such as creating and expanding VMs as utilization changes. By consolidating servers and workloads onto fewer and more powerful physical servers, Hyper-V helps to reduce the hardware footprint. Through integrated high availability, Hyper-V improves business continuity by minimizing the impact of downtimes.

Hyper-V offers many features each with its own set of benefits for Hyper-V administrators and users. Here are the top Hyper-V features and what they can do for an organization.

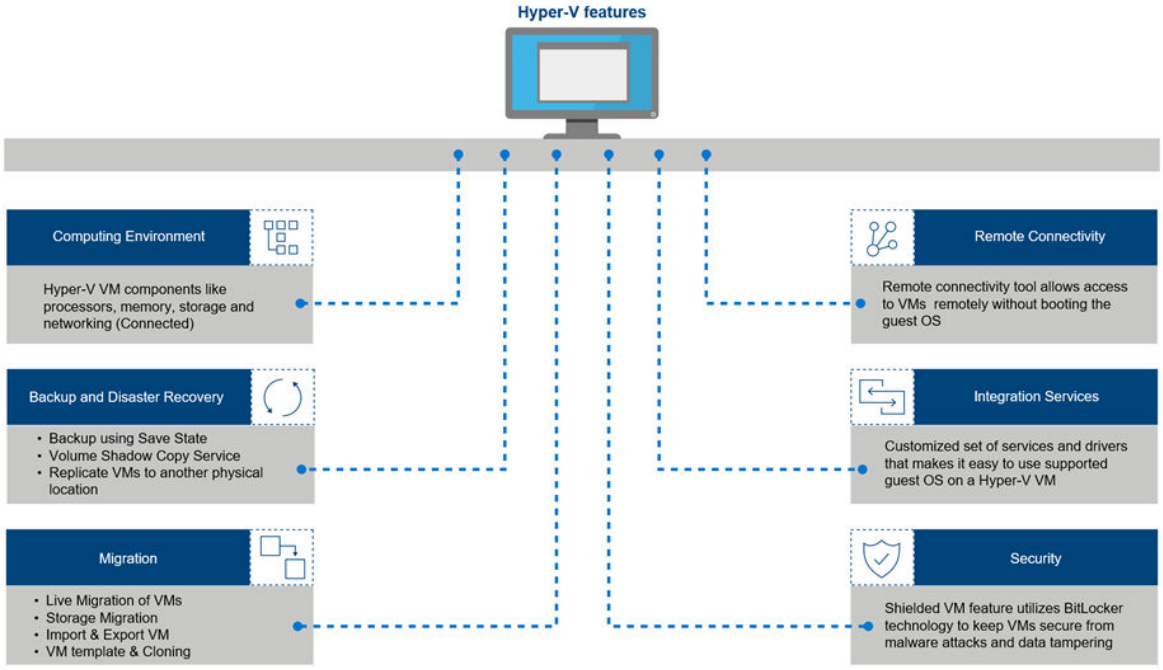


**Figure 17. Hyper-V features**

# Software-defined storage (SDS)

Software defined storage (SDS) is a significant building block of the Azure Stack HCI in which the software manages the storage hardware. SDS works on the concept of separating the management and presentation layer of the storage from the underlying physical hardware. It also eliminates the complex and error prone process of configuring Logical Unit Numbers (LUNs) over the Storage Area Networks (SANs) to deploy the virtualized workloads. Instead of using the proprietary and expensive technologies of SANs, customers can take advantage of this flexible and economical server-based solution by using local disks over remote file sharing protocols with high bandwidth and low latency networking.

SDS running inside the Azure Stack HCI enables customers to build their data centers with the following characteristics:

- Start as small as 2 servers with a switchless configuration. If scale up is in the plan, then start from 3-16 servers

**NOTE:** Scaling up from a 2-node cluster is not supported by Dell as it requires to reinstall the complete setup. For stretch clustering to be implemented, 4 or more nodes are required.

- Scales up to 384 drives and configures up to 3 PB of raw storage for each cluster
- Adds drives and memory to scale up
- Adds servers to scale out
- The storage pool automatically absorbs new drives making it hassle free for drive upgrade/replacement
- Offers better storage efficiency and performance at larger and distributed scale
- Simplifies procurement due to its converged design supports Ethernet, with no special hardware or cables required

**NOTE:** Dell Technologies recommends configuring the servers with all-flash single tier (SSD or NVMe) drives instead of hybrid drive configurations as it allows IT organizations to leverage benefits like reduced power consumption, physical rack space consumption, power, and cooling reduction. Flash drives have seen a rapid growth besides reduction in cost possess the ability to perform better and provide the same capacity as hybrid storage in a smaller footprint.

Storage Spaces Direct (S2D) evolved from Storage Spaces, first introduced with Windows Server 2012. While Storage Spaces was confined to pool resources from a single server and its direct attached storage (JBODs), iSCSI, and Fiber channel enclosures, the S2D was extended to pool storage across multiple servers. S2D is powered by features such as Failover Clustering, Cluster Shared Volume file system, Server Message block (SMB3), and Storage Spaces. The following figure shows the overall resource stack of S2D.

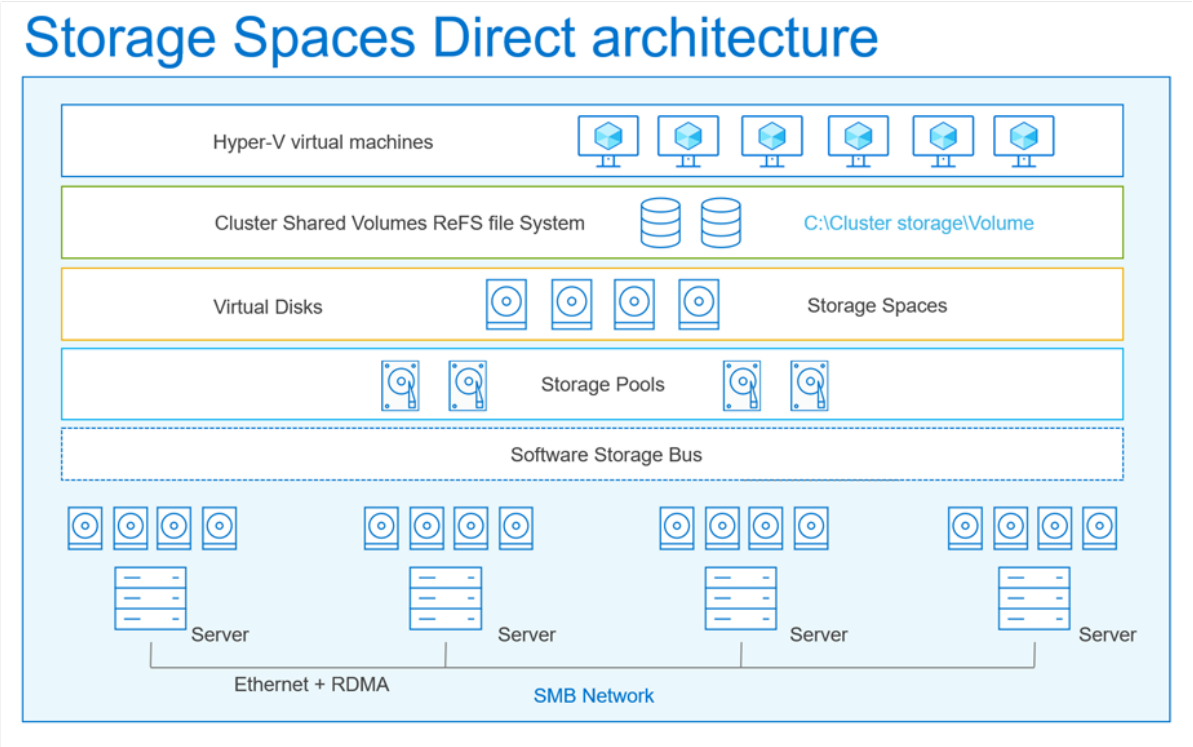


**Figure 18. Architecture of S2D**

**Storage pool** is a single logical disk entity that is formed by aggregating a pool of hard drives of any type and size. Storage spaces are virtual disks that are created from the available disk space provided by the storage pools. They are equivalent to LUNs in a SAN environment.

**Software Storage Bus** is an S2D component that establishes a software defined storage fabric that spans clusters. It enables all the servers to see each other's' local drives, providing a huge pool of storage space. It also enables a caching mechanism that provides server-side read/write caching that accelerates I/O and boosts throughput. The caching behavior is determined automatically by the type(s) of drives present in the server. The following table summarizes the different cache configurations and their behavior:

**Table 1. Caching configurations and their behavior**

| Configuration | Cache drives | Capacity drives | Cache behavior (default) |
|---|---|---|---|
| All NVMe | None (Optional: configure manually) | NVMe | Write-only (if configured) |
| All SSDs | None (Optional: configure manually) | SSD | Write-only (if configured) |
| NVMe + SSD | NVMe | SSD | Write-only |
| NVMe + HDD | NVMe | HDD | Read + Write |
| SSD + HDD | SSD | HDD | Read + Write |
| NVMe + SSD + HDD | NVMe | SSD + HDD | Read + Write for HDD, Write-only for SSD |

**Clusters** through Storage Spaces Direct use SMB3, including SMB Direct and SMB Multichannel, over Ethernet to communicate between servers.

- Server Message Block 3.x (SMB) is a network file sharing protocol that provides access to files over a traditional Ethernet network by means of the TCP/IP transport protocol.
- SMB Multichannel is part of the implementation of the SMB 3.x protocol, which significantly improves network performance and availability for devices by automatic configuration, increasing throughput, and fault tolerance.
- SMB Direct optimizes the use of remote direct memory access (RDMA) network adapters for SMB traffic, allowing them to function at full speed with low latency and low CPU utilization.

**Storage Spaces (Virtual Disks)** support the system by providing high resiliency and fault tolerance to disks using mirroring and parity. The virtual disks resiliency resembles the technology used in Redundant Array of Independent Disks (RAID).

**Cluster Shared Volumes** (CSV) is a clustered file system that enables multiple nodes of a failover cluster to concurrently read-from and write-to the same set of storage volumes. This is achieved by mapping CSV volumes to subdirectories within the C:\ClusterStorage\directory.Each cluster node then allows cluster nodes to access the same content through the same file system path. While each node can independently read from and write to individual files on a given volume, a single cluster node serves a special role of the CSV owner (that is, the coordinator) of that volume. A failover cluster automatically distributes CSV ownership between cluster nodes, however there is also a provision to assign a given volume to a specific owner.

**Resilient File System (ReFS)** is a filesystem purpose-built for virtualization that accelerates file operations such as creation, expansion, checkpoint merging, and integrated checksum to detect and correct bit errors. It also introduces tiers that rotate data between so-called "hot" and "cold" storage tiers in real-time based on usage.

# Planning volumes with SDS

The drives in a storage pool combine to form a volume to facilitate fault tolerance, scalability, and performance. Choosing a balance between all three depends on the type of workloads that run on those volumes. The following figure summarizes which workloads are a good fit for each resiliency-type, mapped with performance and efficiency.


| Resiliency | Fault tolerance | Storage efficiency | Minimum required fault domains | Workloads |
|---|---|---|---|---|
| Two-way mirror<br>Data is copied onto one disk | 1 | 50.0% | 2 | - Virtualized workloads<br>- Databases<br>- Other high-performance workloads |
| Three-way mirror<br>Data is copied on two disks. | 2 | 33.3% | 3 | - Virtualized workloads<br>- Databases<br>- Other high-performance workloads |
| Dual parity<br>Data is stored on two disks and keeps two bitwise parity symbols on other two disks<br>Fault-tolerance is same as Three-way mirror and storage efficiency gets better | 2 | 04 Servers : 50.0%<br>16 Servers : 80.0% | 4 | - Archival and backup<br>- Virtualized desktop Infrastructure |
| Mixed Accelerated Parity<br>A volume can be configured as part mirror & part parity<br>Writes land first in the mirrored portion and then gradually moved into the parity portion. | 2 | 04 Servers : 33.3%<br>16 Servers : 80.0% | 4 | - Archival and backup<br>- Virtualized desktop infrastructure |

**Figure 19. Different types of volumes offered by S2D**

Dell Technologies recommends using a 3-way mirror for volume considerations for most workloads that highly demand performant data transfers. The reason behind this is absolute greatest performance it offers amongst all other volume options. With a 33% usable capacity, the 3-way mirror offers a great resiliency which can withstand multiple failures at the same time besides providing access to data.

For more information about best practices designed for certain workloads, see Best practices and example configurations.

If there is a need for added resiliency to the volumes, they can be configured to use nested resiliency which is explained in the further section.

The following figure describes some fault tolerance scenarios and how Azure Stack HCI can help keep the organization running. In each case, a 4-node cluster keeps on running despite one or more drive/server failures.

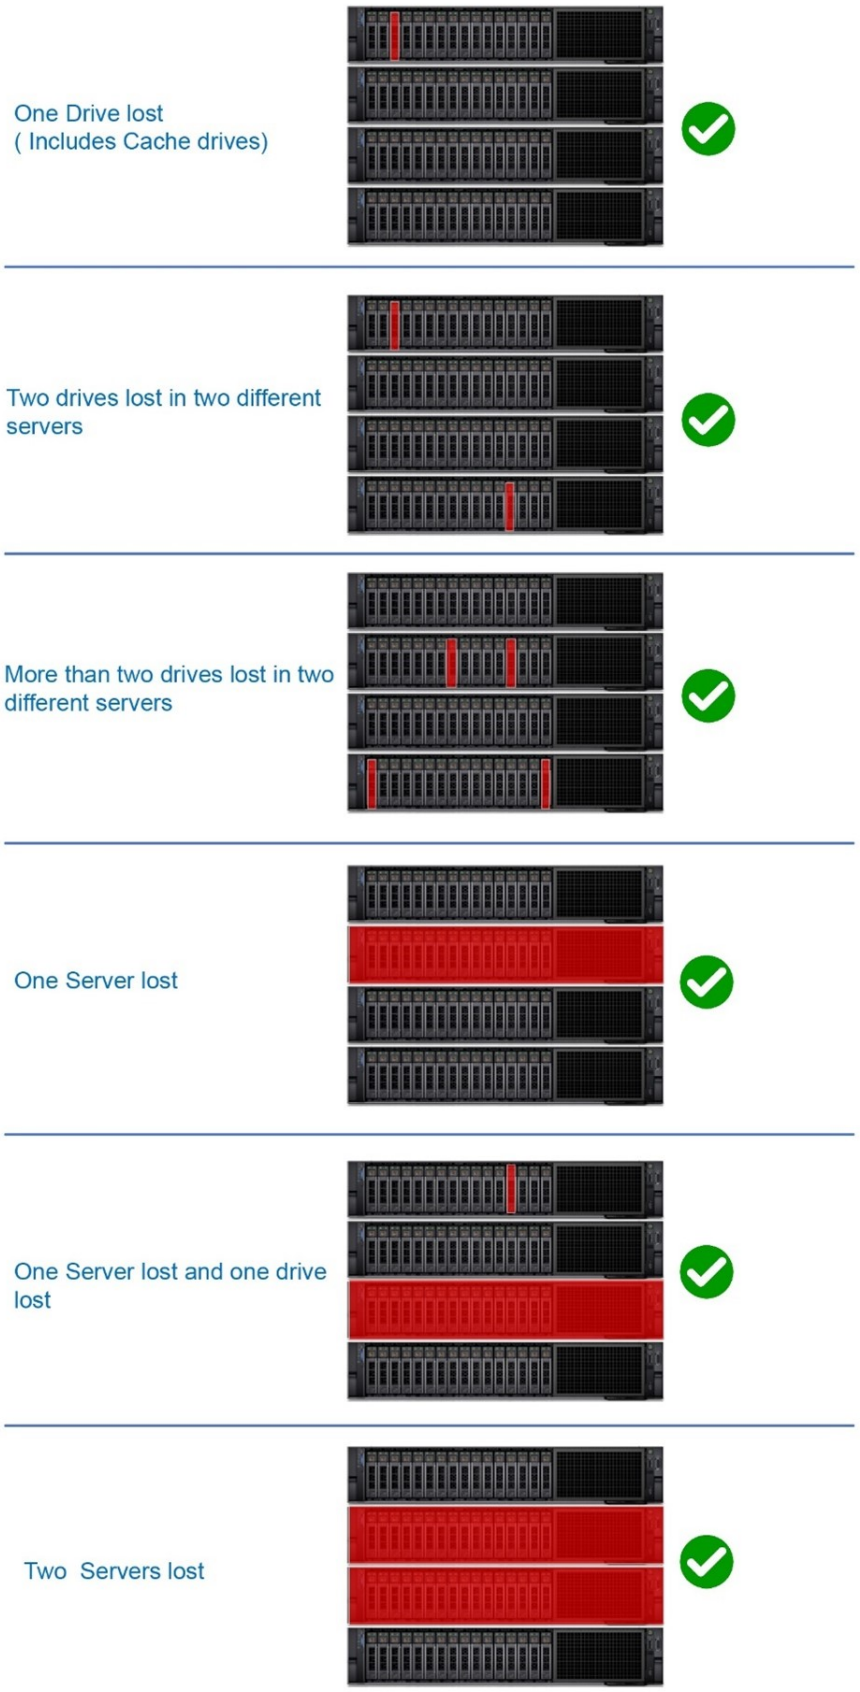


**Figure 20. Failure scenarios in which the SDS based HCI solution remains operational and online**

# Nested resiliency

SDS powered HCI brings in a feature for a two-node cluster called nested resiliency. This feature increases uptime for applications and virtual machines providing uninterrupted access to files and resources during multiple drive, volume, or server failures. Volumes created over nested resiliency can stay online even when multiple hardware failures occur at the same time, unlike with classic two-way mirroring. The only caveat is the lower capacity efficiency, meaning less usable storage space when compared to two-way mirroring.

The following two examples show that if two drives fail at the same time, or if a complete server goes down and a drive fails simultaneously, volumes running in a nested resiliency configuration stay online and accessible.

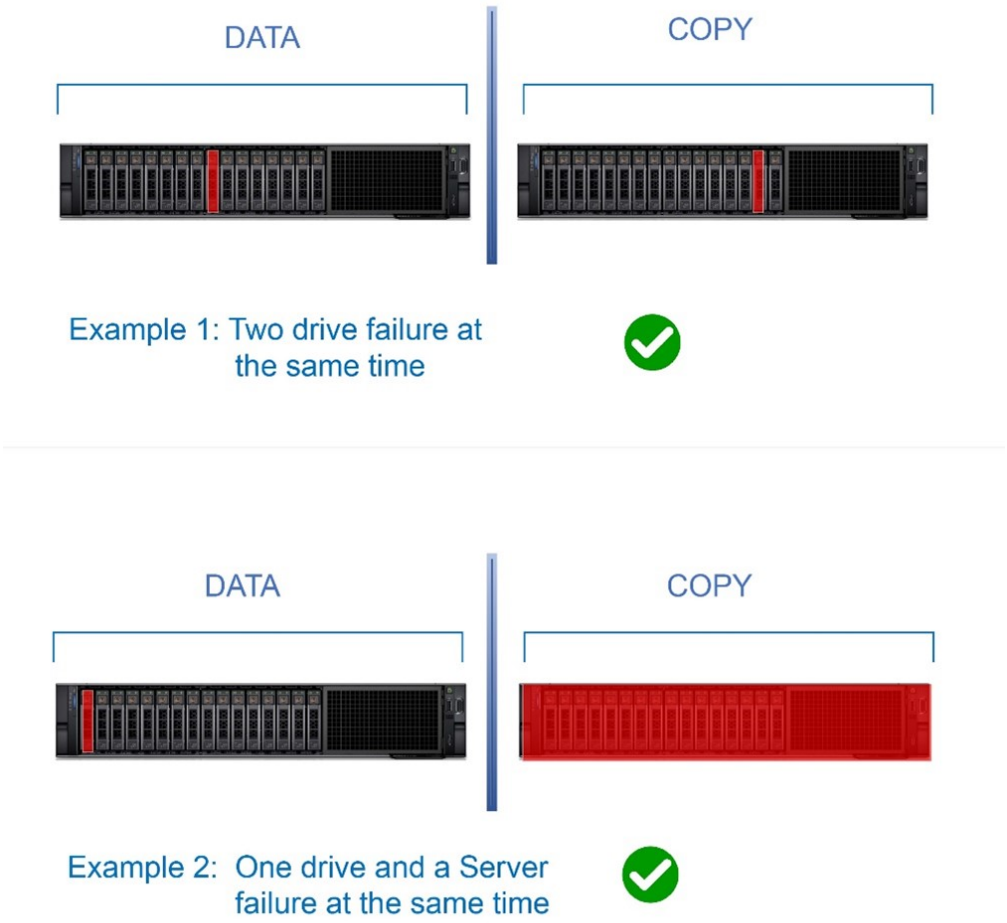


**Figure 21. Nested resiliency**

Two new resiliency options are included in Storage Spaces Direct, which uses software as a platform without the need for any hardware RAID.

**Nested two-way mirror**: This option works essentially as a four-way mirror with data getting written and copied, first within the server and data getting copied in between servers. This results in a volume that provides excellent performance with great resiliency.

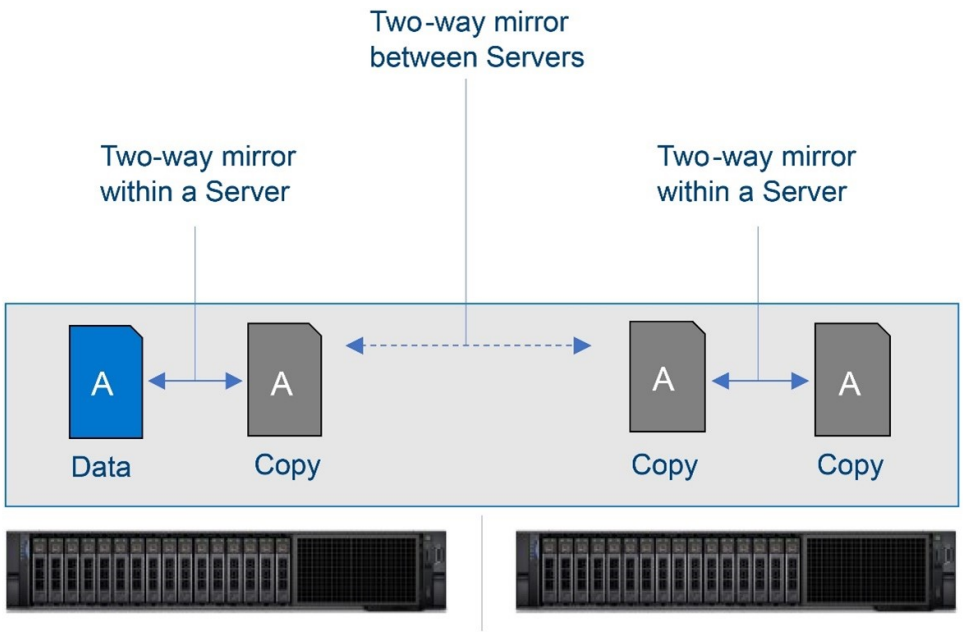


**Figure 22. Nested two-way mirror**

**Nested mirror-accelerated parity:** This option combines the nested two-way mirroring with nested parity. First, there is parity calculation performed by S2D, that is happening locally within each server and, at the next step the two-way mirroring is happening in between the servers. This can achieve higher capacity efficiency, depending upon the number of capacity drives in each server and the mix of mirror and parity specified per volume.

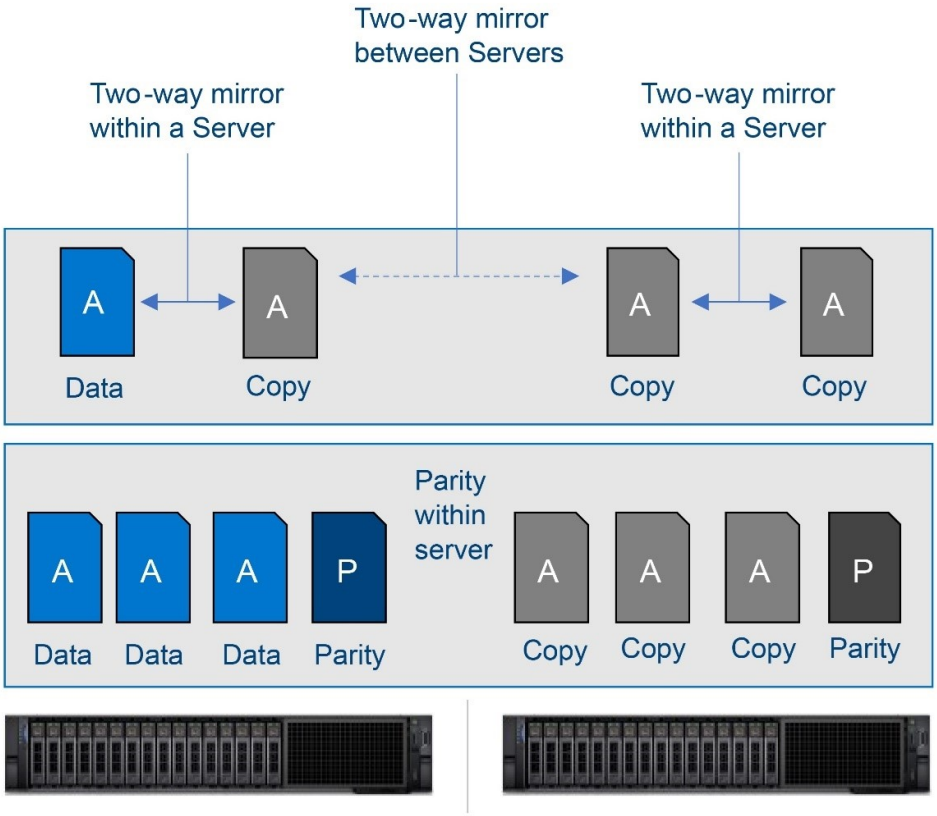


**Figure 23. Nested mirror-accelerated parity**

# Software defined networking (SDN)

Software defined networking (SDN) has made it possible for environments like Microsoft Azure, which operates in data centers across the globe, to perform tens of thousands of network changes every day. SDN provides a way to centrally configure and manage networks and network services, such as switching, routing, and load balancing in a data center. SDN can dynamically create, secure, and connect the network to meet the evolving needs of applications.

Virtual network elements such as Hyper-V Virtual Switch, Hyper-V Network Virtualization, Software Load Balancing, and RAS Gateway are integral elements of the SDN infrastructure. Existing SDN-compatible devices achieve deeper integration between workloads running in virtual networks and the physical network.

# Virtual switch

A virtual switch can connect virtual machines to networks that are external to a Hyper-V host, including an organization`s intranet or Internet. It is programmatically managed and can connect VMs to both virtual networks and physical networks. The switch is a software-based Ethernet network switch and is manageable in Hyper-V Manager and Windows Admin Center (WAC).

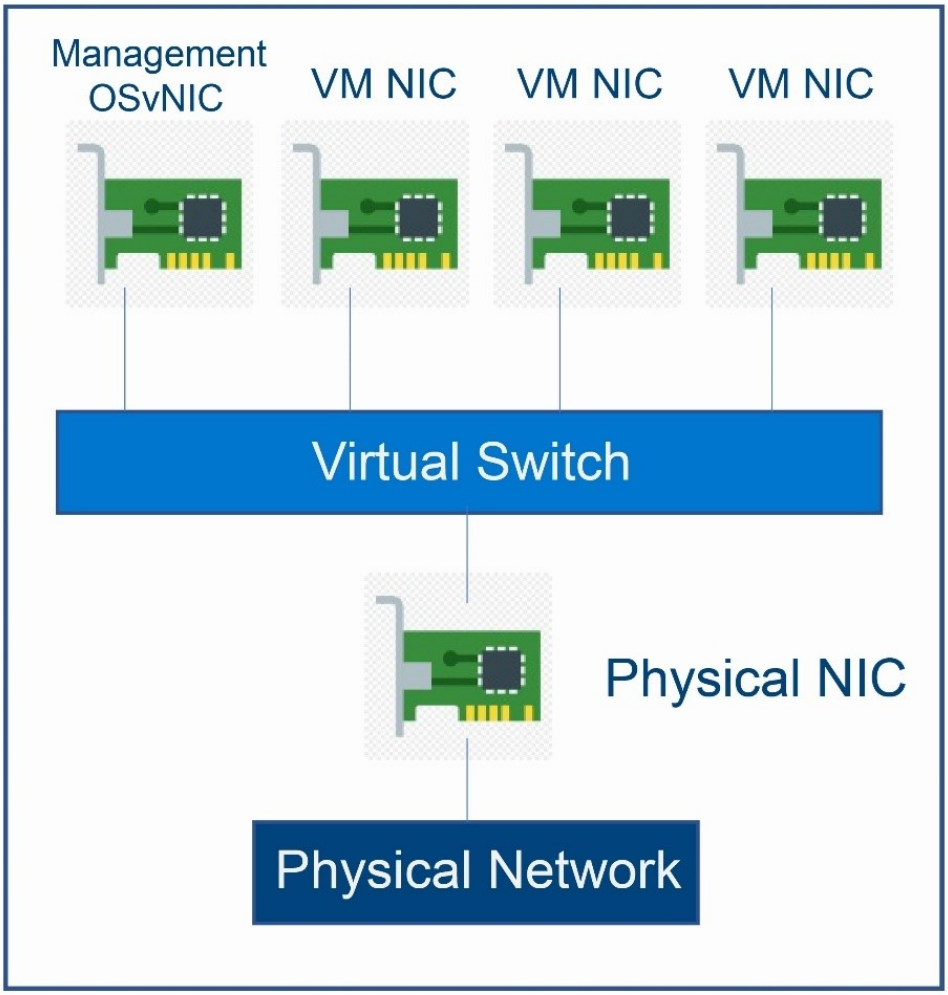


**Figure 24. Virtual networking**

Network virtualization works on the concept of shared network infrastructure, in which multiple virtual networks (potentially with overlapping IP addresses) run on the same physical network infrastructure. A Hyper-V virtual switch makes network management easy by allowing admins to configure and view VLANs, while also implementing security policies with Extended Port Control Lists.

# Network Controller

The Network Controller provides a centralized, programmable point of automation to manage, configure, monitor, and troubleshoot virtual network infrastructure in your data center. It is a highly scalable server role that uses Service Fabric to provide high availability. The Network Controller must be deployed on its own dedicated VMs.

Deploying Network Controller enables the following functionalities:

- Create and manage virtual networks and subnets. Connect virtual machines (VMs) to virtual subnets.
- Configure and manage micro-segmentation for VMs connected to virtual networks or traditional VLAN-based networks.
- Attach virtual appliances to your virtual networks.
- Configure Quality of Service (QoS) policies for VMs attached to virtual networks or traditional VLAN-based networks.

# Software Load Balancing

Software Load Balancer (SLB) can be used to evenly distribute customer network traffic among multiple VMs. It enables multiple servers to host the same workload, providing high availability and scalability. SLB uses Border Gateway Protocol to advertise virtual IP addresses to the physical network.

# Gateway

Gateways are used for routing network traffic between a virtual network and another network, either local or remote. Gateways can be used to:

- Create secure site-to-site IPsec connections between SDN virtual networks and external customer networks over the Internet.
- Create Generic Routing Encapsulation (GRE) connections between SDN virtual networks and external networks. The difference between site-to-site connections and GRE connections is that the latter is not an encrypted connection. For more information about GRE connectivity scenarios, see GRE Tunneling in Windows Server.
- Create Layer 3 connections between SDN virtual networks and external networks. In this case, the SDN gateway simply acts as a router between your virtual network and the external network.

Gateways use Border Gateway Protocol to advertise GRE endpoints and establish point-to-point connections. SDN deployment creates a default gateway pool that supports all connection types. Within this pool, you can specify how many gateways are reserved on standby in case an active gateway fails.

For more information, see Software Defined Networking (SDN) in Azure Stack HCI and Windows Server.

# Compute building blocks

## Virtual machine and containers

A Virtual Machine (VM) is essentially a computing environment created out of software that enables running programs just like a physical machine. It works on the concept of creating a "virtual" version of a computer, with dedicated amounts of CPU, memory, and storage that are borrowed from a physical server. The virtual machine is partitioned from the rest of the system, meaning that the software inside a VM cannot interfere with the host computer's primary operating system. The guest operating system runs the Virtual hardware and provides an isolated environment for running applications independently.

Containers in many ways represent a more granular, lower-overhead approach to virtualization than a virtual machine. Each container image is packaged up with application code, all dependencies, system tools, runtime, system libraries, and settings needed to run the containerized applications. They are plugged onto a container engine that sits atop the operating system infrastructure, to become containers at runtime. Containers are mainly used for cloud-native, distributed applications and to package legacy applications for increased portability and deployment simplicity.

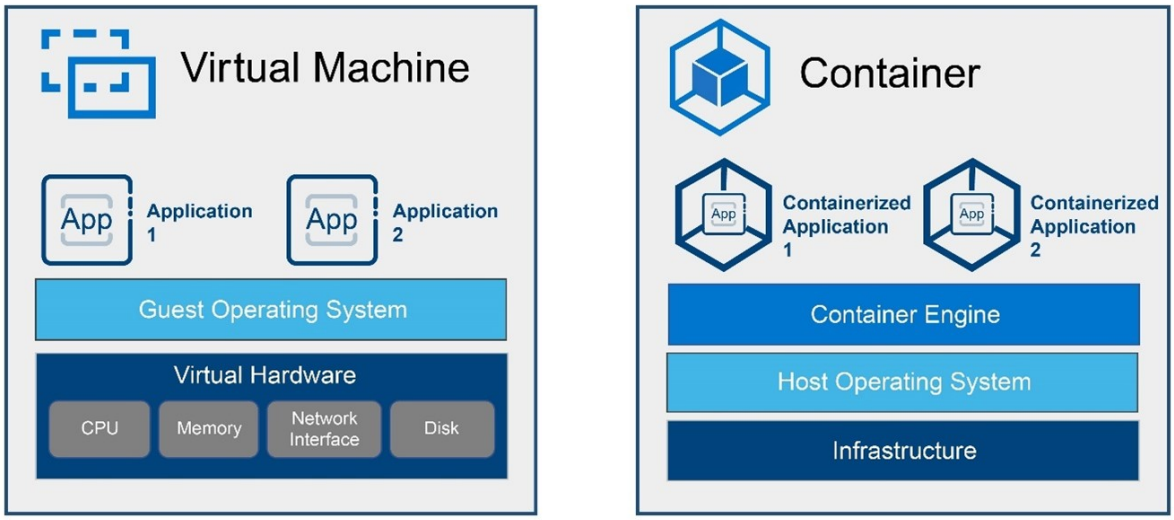


**Figure 25. Basic architecture of VMs and containers**

# Azure Kubernetes Service on Azure Stack HCI (AKS-HCI)

Azure Kubernetes Service on Azure Stack HCI (AKS-HCI) provides an effective platform consistency. Modern applications are increasingly getting built on the containerized approach, where microservices are packaged with their dependencies and configurations. Kubernetes, the core component of AKS-HCI, is open-source software for deploying and managing these containers at scale. As compute utilization increases, applications grow to span multiple containers that are deployed across multiple servers. Operating these applications at scale becomes more complex. To manage this complexity, Kubernetes provides an open-source API that determines how and where these containers will run.

Kubernetes [6] orchestrates with a cluster of VMs and schedules the containers to run on those VMs based on their available compute resources and the containers' resource requirements. Containers are then grouped into pods, the basic operational unit of Kubernetes. These pods scale based on the needs of the applications. Kubernetes also manages service delivery, load balancing, resource allocation, and scales based on utilization. It also keeps a check on the health of each individual resource and enables applications to self-heal automatically by restarting or replicating the containers.

Setting up and maintaining Kubernetes can be complex. AKS-HCI helps simplify setting up Kubernetes on-premises, making it faster to get started hosting Linux and Windows containers.

6 Kubernetes (pronounced "koo-ber-net-ees") is sometimes seen as "k8s" or "k-eights."

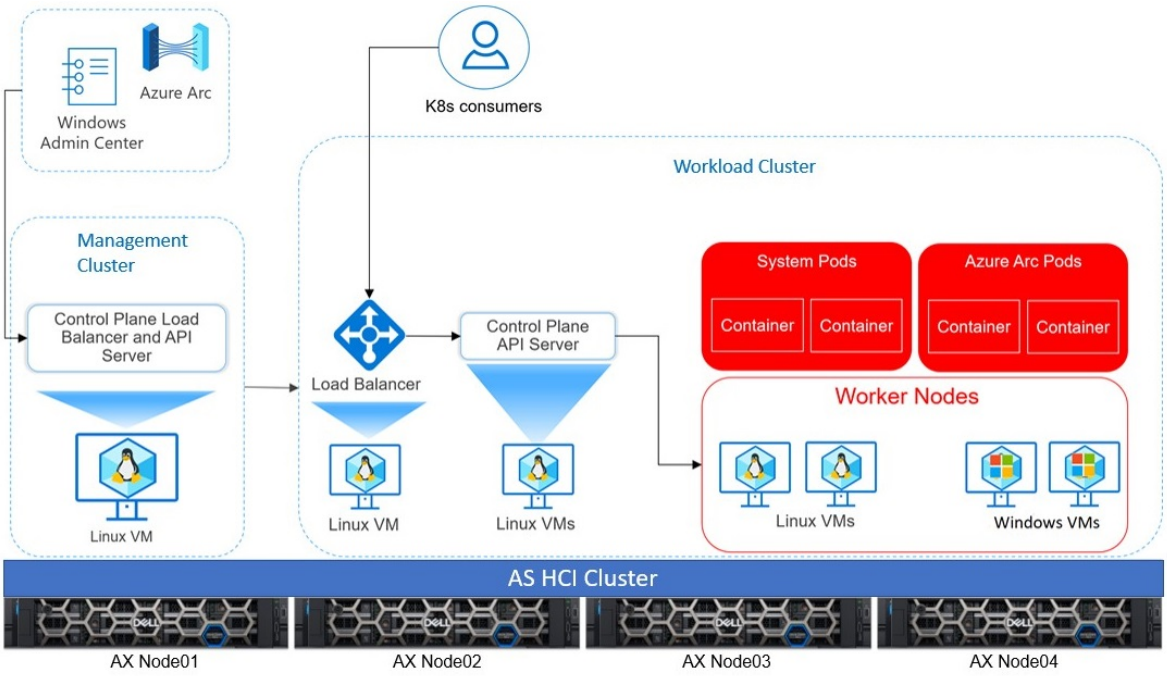


**Figure 26. AKS-HCI components**

**Windows Admin Center** and **PowerShell** are two options for managing the life cycle of Azure Kubernetes Service clusters on Azure Stack HCI. This figure shows the core infrastructure components of Kubernetes, which is divided into two main units based on their operating function. They are:

- **Management Cluster**: (also known as AKS host) provides the core orchestration mechanism and an interface for deploying and managing one or more clusters.
- **Workload Cluster**: comprises the target clusters where containerized applications are deployed.

**Management Cluster** is automatically created when the Azure Kubernetes Service cluster is created on Azure Stack HCI. It is mainly responsible for provisioning and managing workload clusters where workloads are designed to run.

Through the **API server**, the underlying Kubernetes APIs that are exposed provide the interaction with management tools, such as Windows Admin Center, PowerShell modules, or *kubectl*.

**kubectl** is the command line interface for Kubernetes clusters.

The **load balancer** is a single dedicated Linux VM with a load balancing rule for the API server of the management cluster.

The **Workload Cluster** is a highly available deployment of Kubernetes using Linux VMs, meant for running Kubernetes control plane components and running Linux and Windows Server-based containers. Multiple workload cluster(s) can be managed by one management cluster.

The **API server** allows interaction with the Kubernetes API and provides the interface for interaction of management tools, such as Windows Admin Center, PowerShell modules, or *kubectl*.

The **etcd** is a distributed key-value store that stores data required for life cycle management of the cluster. It stores the control plane state.

The **Load Balancer** is a Linux VM running HAProxy + KeepAlive to provide highly available load balanced services for the workload clusters deployed by the management cluster. In each workload cluster there is at least one load balancer VM defined with a load balancing rule.

A **Worker Node** is a VM that runs the Kubernetes node components and hosts the pods and services that make up the application workload.

A **pod** represents a single instance of the application running and typically has a 1:1 mapping associated with a container. In advanced scenarios, pods can contain multiple containers with shared resources.

A **deployment** defines the number of replicas (pods) to be created. The Kubernetes Scheduler ensures that if pods or nodes encounter problems, additional pods are scheduled to start on healthy nodes.

# Refactoring traditional applications to containers with AKS

The new default approach for running applications in the cloud-native environment is through containerization. Refactoring traditional existing applications to use a microservices architecture and running them within the Azure managed Kubernetes service (AKS) across containers, enables benefits such as lower cost, portability, efficiency, developer agility, scaling, and security. Transformation is simple because the application's source code requires no changes. The following figure shows the flow of containerizing applications to AKS cluster.

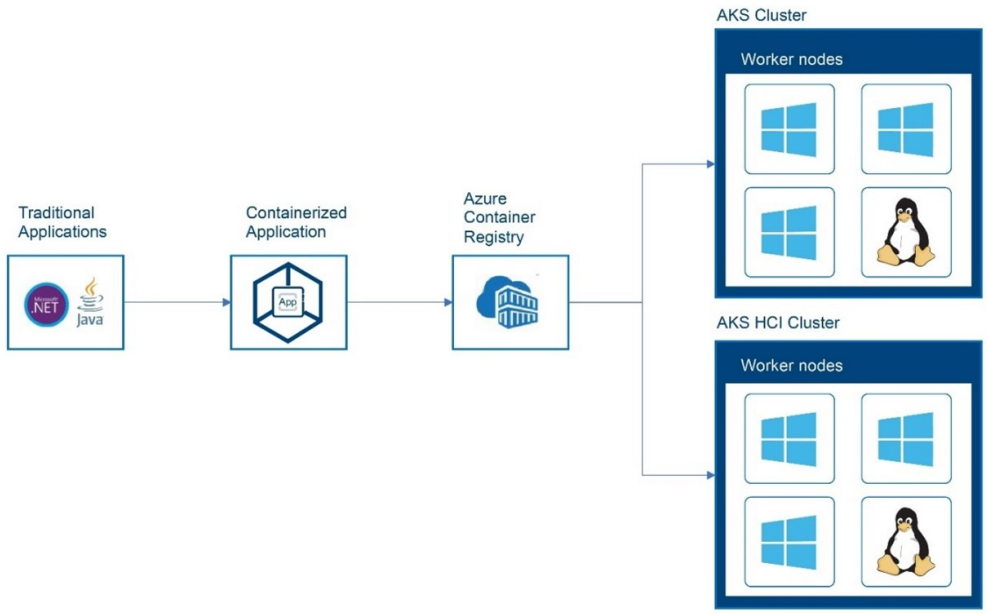


**Figure 27. Flow of application containerization to AKS cluster**

Steps to be performed for refactoring traditional applications to an AKS Cluster:

1. Prepare an application
   - Clone a sample application source from GitHub
   - Create a container image from the sample application source
   - Test the multi-container application in a local Docker environment
2. Push images to Azure Container Registry
   - Create an Azure Container Registry (ACR) instance
   - Tag a container image for ACR
   - Upload the image to ACR
   - View images in your registry
3. Deploy a Kubernetes Cluster
   - Deploy an AKS cluster on Azure Stack HCI
   - Install the Kubernetes CLI (kubectl)
   - Configure kubectl to connect to your workload cluster
4. Deploy an Application in Kubernetes
   - Update a Kubernetes manifest file
   - Run an application in Kubernetes
   - Test the application
5. Post-migration Activities
   - Scale an application
   - Update an application
   - Upgrade Kubernetes clusters

For details about moving applications to AKS-HCI, see Tutorial: Prepare an application for Azure Kubernetes Service on Azure Stack HCI.

# Managing the hybrid solution

Managing a hybrid solution is performed by using two complementary and consistent tools for Azure Stack HCI:

- **Windows Admin Center** as the locally deployed, always-available management console for managing servers, edge nodes, clusters, and HCI infrastructure.
- **Azure Arc** unlocks new hybrid scenarios by extending the Azure services and management to any infrastructure. The Azure control plane provides a platform to manage the multi-cloud and on-prem IT resources by projecting non-Azure resources in Azure. This means that all existing and new IT resources across the entire IT estate can be managed consistently and at-scale, wherever they reside, from Azure.

# Windows Admin Center

Dell Technologies is constantly improving the Dell OpenManage Integration with Microsoft Windows Admin Center (OMIMSWAC). This integrated tool is a lightweight, browser-based application that provides troubleshooting and managing of individual windows servers, clusters, and hyper-converged infrastructure. It enables streamlined life cycle management at the server level for:

- PowerEdge servers running Microsoft Windows Server
- Clusters based on AX nodes as part of the Microsoft Azure Stack HCI and Microsoft Windows Server HCI
- Storage Spaces Direct Ready Nodes as part of the Dell HCI Solutions for Microsoft Window Server and Hyper-V
- Failover clusters based on PowerEdge servers running on-premises and in Azure hybrid environments

This feature rich solution simplifies an IT administrators' tasks, such as:

- Creating clusters using automated cluster creation
- Performing a one-click full stack Lifecycle management (LCM)
- Cluster expansion
- Dynamic CPU core management
- Infrastructure compliance visibility and remediation with Dell HCI Configuration Profile Policies for Azure
- Intrinsic infrastructure security management with Microsoft Secured-core Server and Dell Infrastructure Lock
- Viewing the inventory of PowerEdge servers and Dell Solutions for Microsoft Azure Stack HCI Ready Node clusters
- Gaining a unified view of health, hardware, and firmware inventory details about device components
- Viewing iDRAC information about the PowerEdge server, clusters, and HCI, for out-of-band management, and to launch the iDRAC console from within Windows Admin Center
- Troubleshooting issues with a dashboard view of system health and status
- Accessing support for Dell online catalogs and Dell Repository Manager
- Cluster-Aware Updating (CAU) to update BIOS, drivers, firmware, and systems management applications with zero workload impact for Azure Stack HCI clusters and other cluster types built with PowerEdge servers
- Microsoft HCI Solutions from Dell Technologies certified badging ensures that you are using validated AX node configurations
- System Update and CAU using Dell online and offline catalogs supporting connected, offline, and edge scenarios

# Azure Arc

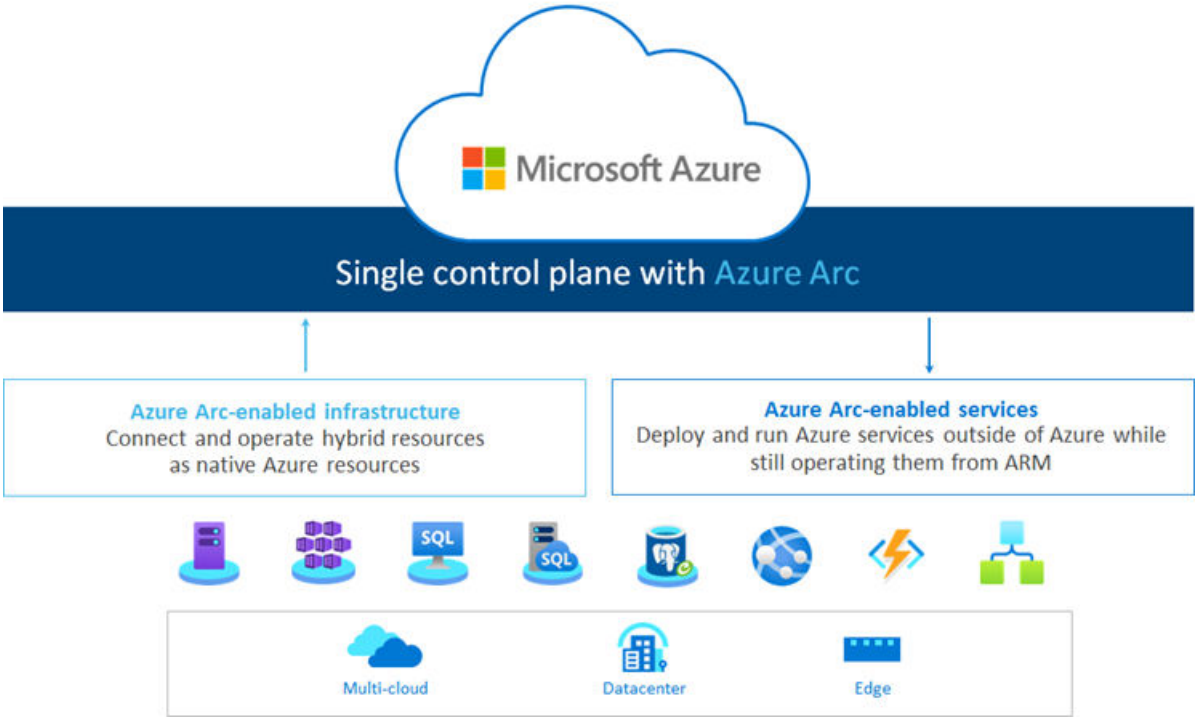


**Figure 28. Azure Arc**

A lot of organizations have resources spread across multiple datacenters and other public clouds -- all managed by different management tool sets. Azure Arc bridges these distributed environments across on-premises, edge, and multi-cloud into a unified central management plane in Azure. The resources need not be migrated or moved to a common directory service to be managed, they simply stay where they are. This means that users can now manage and operate all their existing and new IT resources consistently and at-scale, wherever they reside, from Azure. After they are under Azure Arc management, resources such as Windows and Linux servers, SQL server instances, and Kubernetes clusters can be organized, governed, and secured centrally at scale. This also includes running Azure services anywhere on the infrastructure to benefit from the cloud.

To unpack this a bit more, **Azure Arc-enabled infrastructure** enables you to connect your resources, which live outside of Azure today, and operate them as if they were native Azure resources, using the same management tools and services that Azure provides.

With **Azure Arc-enabled services**, you have the flexibility to deploy fully managed Azure services anywhere – on-premises or in other public clouds -- so you can take advantage of cloud benefits everywhere, such as scalability, fast deployment, and always up-to-date cloud innovation. Importantly, you can initiate and manage these deployments right from the Azure Portal. Arc enabled infrastructure makes it possible to manage the multi-cloud data center and edge devices under a single plane of glass.

## Azure app services and data services

**Azure App Service** is an HTTP-based service for hosting web applications, REST APIs, and mobile back ends. Customers can develop applications using their favorite language, such as .NET, .NET Core, Java, Ruby, Node.js, PHP, or Python. Applications run and scale with ease on both Windows and Linux-based environments.

App Service adds the power and features of Microsoft Azure to the application, such as security, load balancing, autoscaling, and automated management. It also adds DevOps capabilities such as continuous deployment (from Azure DevOps, GitHub, Docker Hub, and other sources), package management, staging environments, custom domain, and TLS/SSL certificates. With App Service, customers pay only for the Azure compute resources used.

Using **Azure Arc-enabled data services,** IT administrators or DBAs first deploy a data controller to a workload cluster on AKS-HCI. To take advantage of the functionality of Azure Arc-enabled SQL Managed Instance and Azure Arc-enabled PostgreSQL Hyperscale on-premises, the data controller orchestrates and automates the provisioning, scaling, updating, monitoring, backup, high availability, and disaster recovery of the database instances.

This service provides the elasticity, instant provisioning, and always current feature set IT administrators, DBAs, and software developers require to solve their database management challenges. Seasoned DBAs can use their existing knowledge and toolsets to manage and maintain the instances and databases.

The Dell Integrated System for Microsoft Azure Stack HCI Portfolio – AX Nodes chapter provides some technical insight into Dell AX node offerings as part of flexible and scalable infrastructure configurations.

6

# Dell Integrated System for Microsoft Azure Stack HCI Portfolio – AX Nodes

**Topics:**

- A modern platform for Microsoft environments
- Powering the AX nodes, the Intel, and AMD processor
- Dell integrated systems with security built-in

## A modern platform for Microsoft environments

KEY TAKEAWAY

Dell Integrated System for Microsoft Azure Stack HCI delivers a fully productized, validated and supported hyperconverged (HCI) solution that enables organizations to modernize their infrastructure for improved application uptime and performance, simplified management and operations, and lower total cost of ownership.

The wide range of AX nodes from Dell Technologies offers high performance, and a scalable and secure foundation needed for a software-defined infrastructure.

AX nodes from Dell Technologies are intelligently designed based on the industry-leading Dell PowerEdge server architecture and are engineering validated and certified to run Azure Stack HCI. Dell Integrated System for Microsoft Azure Stack HCI offers flexible and scalable infrastructure configurations across different models for varying application performance, capacity, and application proximity requirements. In the design phase, the Dell team deliberately selects each configuration's components. Thereafter, only tested updates to the BIOS, firmware, and driver revisions are applied to the system, for performance and resiliency.

The modular architecture enables a predictable pay-as-you-grow approach which aligns to changing business needs and user demands. The flexible configuration options help customers choose performance, capacity and graphics acceleration required for their business environments. Cluster sizes can range from as few as two nodes and grow to a maximum of 16 nodes. To expand beyond 16 nodes, IT staff can leverage cluster sets and benefit from Azure-like Fault Domains and Availability Sets.

The form factors of the AX nodes range from one to two rack units (RU) in width and varying depth. There are different node configurations allowing for single socket or dual socket multi-core processors from Intel or AMD. Hardware innovations such as 100GbE RDMA networking and high-performance PCIe Gen 4 NVMe drives also are available. Storage configurations can be all-flash with SSD or NVMe drives or hybrid with a mix of SSD and NVMe or flash and spinning drives (HDDs).

The productized AX nodes come with Azure Stack HCI OS pre-installed at the factory. Dell OpenManage Integration with Microsoft Windows Admin Center offers full support for all AX nodes and provides 1-click full stack lifecycle management (LCM) with cluster-aware updating, automated cluster creation, cluster expansion, and dynamic CPU core management.

Through the Microsoft Cloud Solutions Provider (CSP) program from Dell Technologies, we can be your single-source of procurement and support for the entire infrastructure stack, including Azure Stack HCI software subscription, reducing time-to-value, and simplifying on-going support activities.

Organizations can benefit largely from the following with AX nodes:

## Scalable and flexible configurations options

Dell offers a broad range of configuration options for the hardware where business require volumes at enterprise scale or at small scale catering to SMBs. The datacenter can be setup starting with just two nodes, and then to scale out easily as business needs grow we would three nodes and more. AX nodes are also capable of scaling up, so you can right-size the cluster for current demand and add more performance and capacity when needed. However, Dell Technologies' guidance is that all nodes in an Azure Stack HCI cluster maintain homogeneity - i.e. identically configured with all the same components and BIOS, firmware, and driver updates. All configurations designed, validated, and assured to deliver balanced performance and capacity for every possible technical use case and workloads.

## Experience operational efficiency

Dell OpenManage Integration with Microsoft Windows Admin Center delivers full-stack lifecycle management enabled by Cluster-Aware Updating, easy cluster expansions and CPU core management. Updates occur with only a single reboot per node in the cluster, and there is no interruption to workloads running in the VMs. Automated cluster creation helps accelerate the process and reduce the risk of human error in deployment. The built-in disaster recovery with stretch clustering provides automatic failover to restore production quickly, and without the need for manual intervention, ensuring the system stays up and running. Unified view of system health, hardware monitoring, and firmware inventory information of the device components makes it easy for troubleshooting issues.

## Get world class integrated services and support

This intelligently integrated foundation from Dell Technologies delivers simple but proven enterprise-class deployment, maintenance, and support. Services have been designed to be simple, flexible and worry free—from installation and configuration to single-source support. Certified deployment engineers ensure accuracy and speed, while reducing risk and downtime. One-stop cluster-level support covers the hardware, the operating system, hypervisor, and Storage Spaces Direct software. Dell support-offerings helps customers get ahead of problems before they happen utilizing the proactive, predictive, and automated Secure Connect Gateway for faster resolution and reporting.


**Figure 29. Dell Integrated System for Microsoft Azure Stack HCI- AX nodes**

Dell Technologies offers multiple configuration options of hardware that have been validated and guaranteed to deliver the optimal balance of performance and capacity to address a broad set of Azure Stack HCI use cases and workloads. Azure Stack HCI is now available on 15G configurations with the latest generation of both Intel and AMD CPUs, delivering breakthrough performance and density for use cases that include Edge, ROBO and Core datacenter locations. The table below describes in more detail six different models of AX nodes.

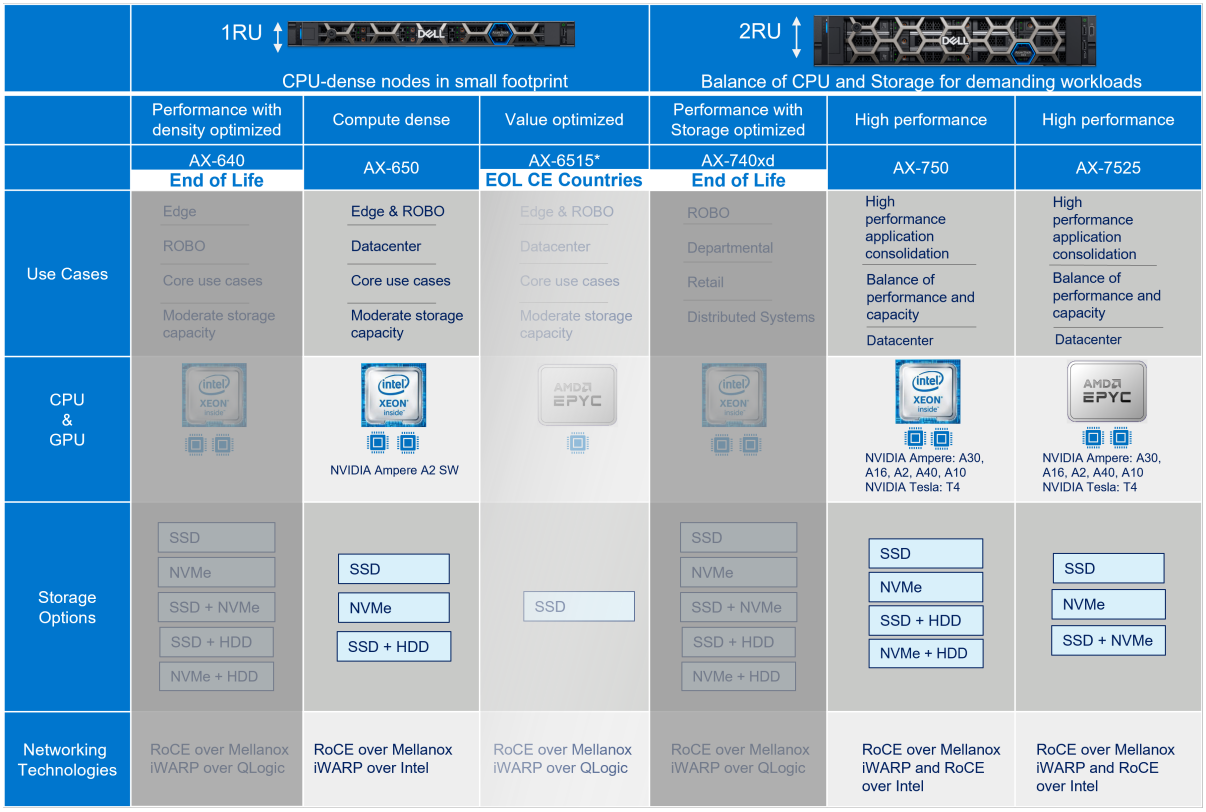


**Figure 30. Six models of AX nodes**

The AX-640 (EOL March 31, 2023) and AX-740xd (EOL March 31, 2023) are built on the 14th generation (14G) of Dell server platform powered by 2nd generation Intel Xeon scalable processors. The 15th generation(15G) of Dell Servers consist of the AMD EPYC processor-based AX-6515 and AX-7525 followed by the 3rd Generation of Intel Xeon scalable processor-based AX-650 and AX-750.

Each AX node includes one Boot Optimized Storage Solution (BOSS) card: either BOSS or BOSS-S2. This enables a RAID solution for booting the HCI OS on the AX nodes.

Stretch clustering is validated for all-flash configurations on all models and is applicable only for servers running Azure Stack HCI.

Dell Technologies recommends configuring the servers with all-flash single tier (SSD or NVMe) drives instead of hybrid drive configurations as it allows IT organizations to leverage benefits like reduced power consumption, physical rack space consumption, power and cooling reduction. Flash drives have seen a rapid growth besides reduction in cost possess the ability to perform better and provide the same capacity as hybrid storage in a smaller footprint.

**NOTE:** This is the list of affected territories by the AX-6515 End Of Life for CE countries: Aland Island (Finland), Acores or Azores (Portugal), Albania, Andorra, Austria, Belgium, Bosnia, Bulgaria, Canarias or Canary Islands (Spain), Corse Island (France), Croatia, Cyprus, Czech Republic, Denmark, Estonia, Elba (Italy), Faroe Islands, Finland, France, French Guiana or Guyana (France), Germany, Gibraltar, Greece, Guadeloupe (France), Guernsey, Herzegovina, Holy See (Vatican City), Hungary, Iceland, Ireland, Isle of Man, Italy, Jan Mayen, Jersey, Latvia, Liechtenstein, Lithuania, Luxembourg, Maderia (Portugal), Malta, Martinique, Mayotte, Moldova, Monaco, Montenegro, Netherlands, Norway, Poland, Portugal, Reunion (France), Romania, Saint Martin (France), San Marino, Sardinia (Italy), Serbia, Sicily (Italy), Slovakia, Slovenia, Spain, Svalbard, Sweden, Switzerland, Turkey.

The following figure shows the different supported configurations.

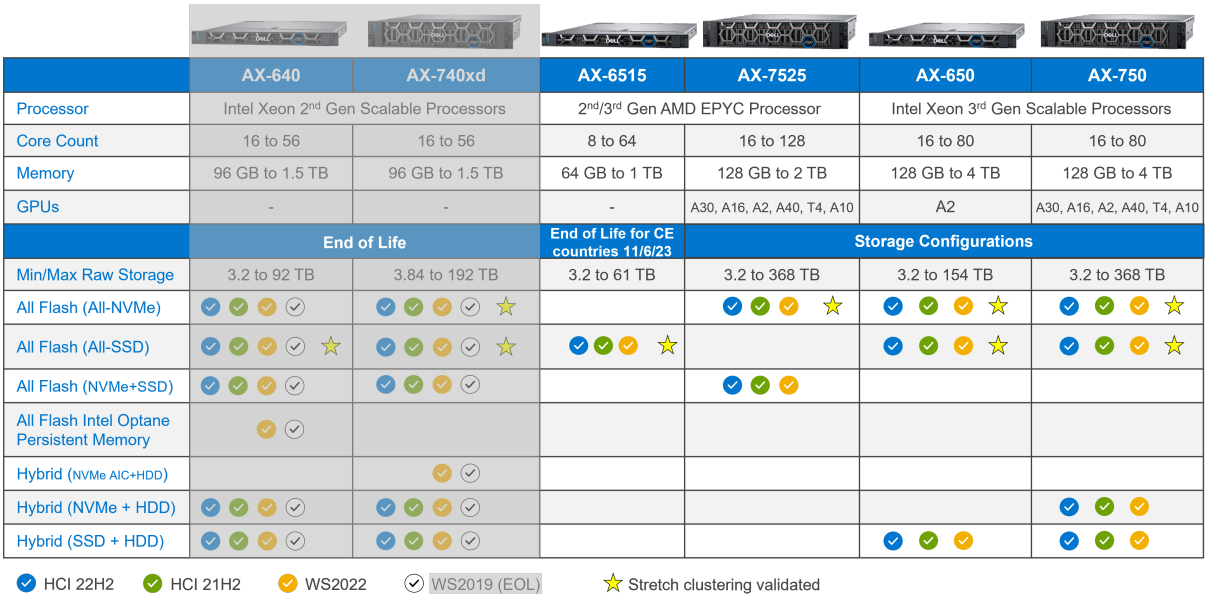


**Figure 31. Intelligently designed portfolio**

For detailed specifications on the distinct platforms, see the Dell Integrated System for Microsoft Azure Stack HCI Specification sheet.

# Powering the AX nodes, the Intel, and AMD processor

## Intel® Xeon® scalable processors:

Intel® Xeon® Scalable platforms are powerful infrastructure that represents an evolutionary leap forward in agility and scalability. Disruptive by design, it sets a new benchmark in platform convergence and capabilities across compute, storage, memory, network, and security. An innovative approach to platform design in Intel® Xeon® Scalable processors unlocks the power of scalable performance for today's data centers and communications networks—from the smallest workloads to the most mission-critical applications.

Intel innovation is driving the modernization and hybrid cloud transformation of the traditional enterprise data center. Migrating to the newest generation of high-performing and energy-efficient Intel-based hardware tunes a data center for highly optimized performance across a broad set of enterprise workloads while lowering costs and improving resource utilization. Over time, evolving to a software-defined infrastructure (SDI) across all the critical domains of the data center (compute/storage/network) will deliver critical automation, orchestration, and telemetry capabilities to help businesses unlock the full capabilities of multi-cloud computing. With modern, industry-standard Intel® servers and technologies that run on software-defined infrastructure, customers can seamlessly manage an environment that supports development and delivery of cloud-native applications and mission-critical workloads on secure private clouds, while also integrating with public clouds, many of which already run-on Intel® architecture.

## AMD EPYC™ processors:

AMD EPYC™ are a new breed of server processors which have set a higher standard for data centers. AMD has brought many features first to the market with the processor generation:

- First with 7 nm technology enabling higher transistor density and energy efficiency
- First with PCIe 4.0 delivering 128 lanes to double the I/O performance over PCIe 3.0
- First with 64 cores (128 threads) in a single socket


AMD EPYC™ has been engineered for data centers that rely on CPU performance. From oil and gas exploration, to in-memory databases, to big data analytics to production rendering to standard data center applications, highly parallel workloads have more cores to work with. Traditional CPUs typically must scale up to a 2-socket server to overcome an imbalance of resources. With AMD EPYC™, 1-socket servers satisfy many workload needs, helping increase density and reduce capital, power, and cooling expenses.

AX nodes are optimized for the 2nd Gen/3rd Gen AMD EYPC processors taking advantage of the additional cores, faster and additional memory channels, and of PCIe 4.0 for faster I/O.

# Dell integrated systems with security built-in

As servers become more critical in a software-defined data center architecture, server security becomes the foundation of overall enterprise security. Servers must emphasize security at both the hardware and firmware level by leveraging an immutable Root-of-Trust that can be used to verify subsequent operations within the server. Dell has established a chain of trust that extends throughout the server life cycle, from deployment through maintenance to decommissioning. Below is the list of security features integrated to the AX nodes.

## Cyber resilient architecture

Dell 14th and 15th generation AX node servers feature an enhanced Cyber Resilient Architecture that provides a hardened server design to Protect, Detect, and Recover from cyberattacks. Some of the key aspects of this architecture are:


**Figure 32. Cyber resilient architecture**

# Protect

Dell servers uses an immutable, silicon-based Root-of-Trust to cryptographically attest to the integrity of BIOS, iDRAC, and other critical firmware. This Root-of-trust is based on onetime programmable, read-only public keys that provide protection against malware tampering. In contrast to Security Laggards, Dell works with extensively vetted silicon chip manufacturers to customize the chip and build in this root of trust technology.

Secure Boot checks the cryptographic signatures of UEFI drivers and other code loaded prior to the operating system running.

Dell servers use digital signatures on firmware updates to assure that only authentic firmware is running on the server platform. iDRAC will scan firmware updates and compare their signatures to what is expected using the silicon-based Root-of-Trust. Any firmware package that fails validation is aborted and an error message is logged into the Lifecycle Log (LCL) to alert IT administrators.

Dynamic System Lockdown which can be enabled without a server reboot by an IT administrator prevents users with lesser privileges from making changes to the server. By enabling lockdown mode, users can prevent configuration drift in their data centers when using Dell tools and agents, and protect against malicious attacks against embedded firmware when using Dell Update Packages.

TPM can also be used to enable the BitLocker™ hard drive encryption feature to address threats of data theft or exposure from lost, stolen, or inappropriately decommissioned systems.

Enterprise Key Management delivers a central key management solution to manage data-at-rest across the organization.

Security-Enhanced Linux operating system (SELinux) operates at the core kernel level on the iDRAC and does not need any input or configuration from users. SELinux logs security messages when an attack is detected. These log messages indicate when and how an attacker tried to break into the system.

Physical I/O ports such as USB inputs can be dynamically disabled using iDRAC. This permits the disablement of these ports for production use but also temporarily grants access for crash cart debugging without rebooting the server.

Shielded VMs are part of the core hypervisor and are protected against inspection, theft, and tampering from malware running on a Hyper-V host as well as the fabric admins administering it.

# Detect

Lifecycle log is a collection of events that occur in a server over a period. Lifecycle log provides a description of events with timestamps, severity, user ID or source, recommended actions, and other technical information that could come handy for tracking or alerting purposes.

iDRAC provides the capability to configure different event alerts and actions to be performed when a particular Lifecycle Logs event occurs.

Dell servers are provided with hardware intrusion detection and logging feature, with detection working even when no AC power is available. Sensors on the chassis detect when anyone opens or tampers with the chassis, even during transit. Servers that have been opened while in transit generate an entry in the iDRAC Lifecycle log after power is supplied.

# Recover

Dell servers include two types of recovery:

- BIOS recovery
- Rapid operating system (OS)

These features enable rapid recovery from corrupted BIOS or operating system images. In both cases, a special storage area is hidden from run-time software (BIOS, operating system, device firmware, so on). These storage areas contain pristine images that can be used as alternatives to the compromised primary software.

It is recommended to keep firmware updated to ensure servers have the latest features and security updates. However, there may need to rollback an update or install an earlier version if any issues are encountered after an update. Firmware Rollback to the previous version, is also verified against its signature.

At the end of a system's life cycle, it either must be retired or repurposed. The goal of System Erase is to erase sensitive data and settings from the server storage devices and server non-volatile stores such as caches and logs so that no confidential information unintentionally leaks. It is a utility in Lifecycle Controller that is designed to erase logs, configuration data, storage data, cache, and any embedded apps.

# Secured-core server

## Windows hardware lab kit

At Dell , hardware devices and drivers are tested as part of the Windows Hardware Compatibility Program using the Microsoft test framework known Windows Hardware Lab Kit or Windows HLK .This is done to ensure that the system which is getting developed is certified as compatible with Windows Server operating systems starting from Windows Server 2016.

## UEFI secure boot

Industry-standard UEFI (Unified Extensible Firmware Interface) Secure Boot checks the cryptographic signatures of UEFI drivers and other code loaded prior to the operating system running ensuring only authorized firmware and operating system bootloaders are initialized during the boot process.

## Trusted Platform Module (TPM)

TPM can be used to perform public key cryptographic functions, compute hash functions, generate, manage, securely store keys, and do attestation. Attestation and remote attestation solutions can use the TPM technology to take measurements at boot time of a server's hardware, hypervisor, BIOS, and operating system, and compare them in a cryptographically secure manner against base measurements stored in the TPM. If they are not identical, the server identity may have been compromised and system administrators can disable and disconnect the server either locally or remotely.

## VBS and HVCI

Virtualization-based security (VBS) and Windows Hypervisor Code Integrity (HVCI) service create a secure, hardware-isolated environment that effectively isolates memory and critical components to prevent attacks and unauthorized access to critical parts of the operating system.

## Dynamic Root of Trust for Measurement (DRTM)

DRTM is a technology which lets the server boot initially into untrusted code, but shortly after that launches the system into a trusted state by taking control of all CPUs and forcing them down a well-known and measured code path. This has the benefit of allowing untrusted early UEFI code to boot the hypervisor, but then being able to securely transition into a trusted

and measured state. The AX nodes based on intel processor comes with the Intel® Trusted Execution Technology (Intel® TXT) whereas the AMD based platforms are with the SKINIT (Secure Init and Jump with Attestation) instruction.

## DMA boot protection

Through the DMA Protection feature (also known as Direct Memory Access Protection), the operating system and the system firmware are protected against malicious and unintended Direct Memory Access (DMA) attacks for all DMA-capable devices (including M.2 PCIe slots) during the boot process and operating system runtime.

## Secure default password for iDRAC

The AX nodes are shipped with a unique, factory-generated iDRAC password to provide additional security. They are made available on the pull-out Service Tag on the front of the chassis, adjacent to the server asset label. Users who choose to use this default option must note this password and may use it to log in to iDRAC for the first time, rather than using a universal default password. For security purposes, Dell Technologies strongly recommends changing the default password.

The Networking Topologies for Azure Stack HCI Solutions chapter encompasses on various configurations of AX nodes to form the primary compute cluster that is deployed as HCI.

7

# Networking Topologies for Azure Stack HCI Solutions

**Topics:**

# Overview

**KEY TAKEAWAY**

Networking the Dell Integrated Systems for Microsoft Azure Stack HCI is the crucial component of the solution without which the Storage Spaces Direct and Hyper-V comes to a standstill. In this chapter, we simplify the network itself and look at the deployment options available for connecting the HCI nodes.

Microsoft HCI Solutions from Dell Technologies includes various configurations of AX nodes that power the primary compute cluster that is deployed as HCI. HCI uses a flexible solution architecture rather than a fixed component design.

Hyperconverged Infrastructure demands enterprise-class network architectures that are resilient and highly performant. The networking fabric must provide a better and more efficient way of transferring data across compute nodes in the data center. This is to cater to the demands for more bandwidth, more processing power, faster access to critical applications, and lossless data transfers.

Designing the lanes for Azure Stack HCI network traffic is an inherent requirement for the complete hybrid solution. Having the right network design to interconnect the HCI cluster nodes and provide access to workloads is imperative to maximize the benefits of the complete solution.

The network traffic can be categorized as follows during the design considerations:

- Compute traffic: Traffic to and from virtual machines
- Storage traffic: Traffic for Storage Spaces Direct (S2D) using Server Message Block (SMB)
- Management traffic: Traffic meant for cluster management, such as Active Directory, Remote desktop, Windows Admin Center and PowerShell

Dell Integrated System for Microsoft Azure Stack HCI includes a broad array of validated AX node configurations to meet the demands of any workload. Each component in these nodes is deliberately chosen and installed with the optimal firmware and driver revisions from the curated Dell HCI solution catalog. This is especially important when selecting the ideal network adapter that will address the I/O requirements of an application. The AX nodes are designed to work with Mellanox and QLogic network adapters. Our integrated system goes a step further by including Dell PowerSwitch network switches that have been validated to optimize the Azure Stack HCI experience. Azure Stack HCI can be deployed using one of two network topologies: Scalable or Switchless Storage. The following figure shows the network integration options available based on the Dell Integrated System for Microsoft Azure Stack HCI.

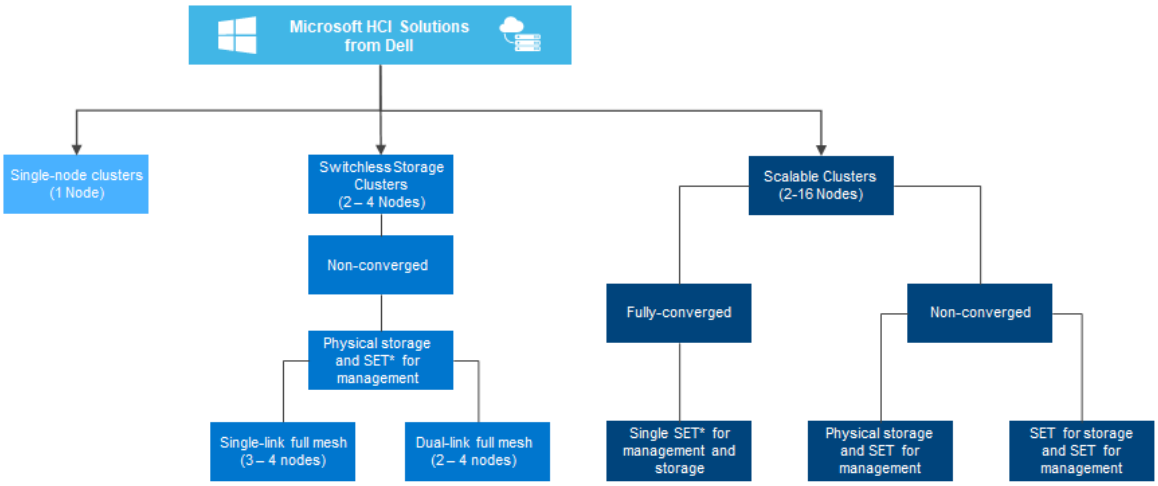


**Figure 33. Network integration options available on Dell Integrated System for Microsoft Azure Stack HCI**

# Scalable

This option supports 2 to 16 AX nodes in a cluster and uses top-of-rack (ToR) switches for management and storage traffic networking. These architectures support both fully converged and non-converged network topologies. Host management and storage networking using each approach can be implemented in different ways. The scalable model supports:

- Fully converged single Switch Embedded Teaming (SET) for storage and management
- Non-converged physical and SET for storage and management, respectively
- Non-converged all-SET for storage and management

**NOTE:** Dell Technologies recommends to setup data centers with 25 GbE network switches as the price per port is almost identical to 10 GbE network switches. This gives an advantage for IT organizations to achieve maximum performance with the right cost.

# Switchless storage

This option supports 2 to 4 AX nodes and can be deployed without consuming high-speed ToR switch ports for storage traffic, thereby reducing the initial investment. This model uses full mesh connections between the cluster nodes for RDMA storage traffic. The switchless storage networking model supports only a non-converged architecture implemented as a single-link or dual-link full mesh topology. These topologies are described in detail in the following sections.

# Switch Embedded Teaming (SET)

SET is a NIC teaming solution that you can use in environments that include Hyper-V and the Software Defined Networking (SDN) stack in Windows Server operating system and Azure Stack HCI operating system. It enables grouping several physical network adapters (from one to eight) into one or more virtual network adapters in a Hyper-V environment. These virtual network adapters provide fast performance and fault tolerance in the event of a network adapter failure. The main advantage of SET compared to NIC teaming, either of which can be enabled in the operating system, is the RDMA convergence and the RDMA virtualization for host. In this way, storage traffic with SMB Direct can be converged with other traffic. SET reduces the need for more physical NICs by doing the same job with less administrative overhead and network complexity.

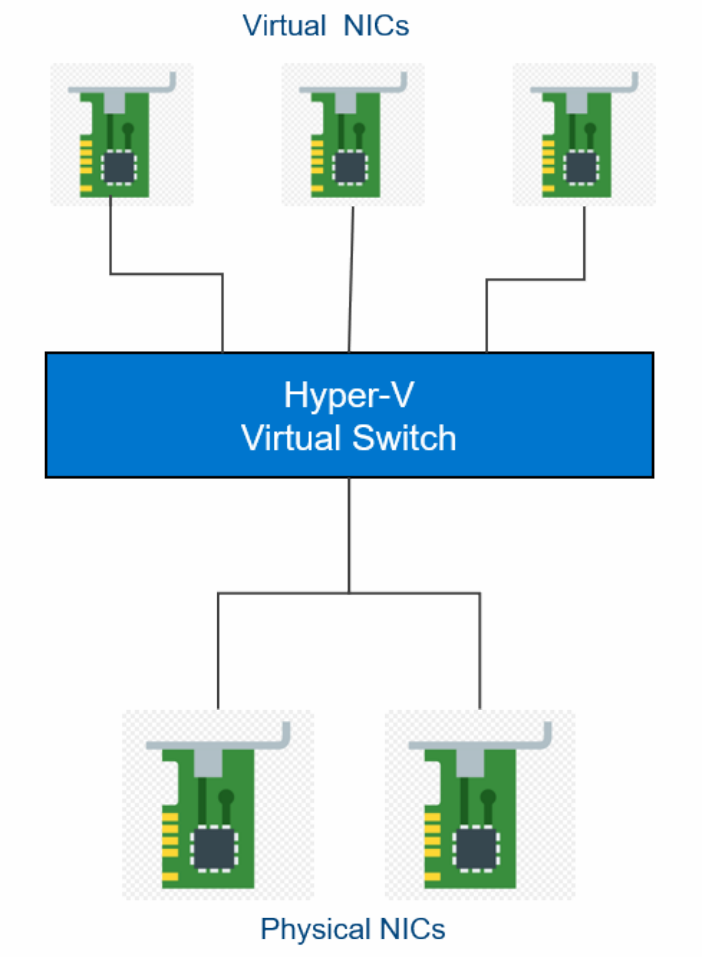


**Figure 34. Switch embedded teaming**

# Remote Direct Memory Access

Remote Direct Memory Access (RDMA) is a technology that allows servers on a network to exchange data in main memory, bypassing the operating system and the CPU resources. It enables direct transfer of data in and out of a server by implementing a transport protocol on the NICs. The technology uses zero-copy networking to read data directly from the main memory of one system and write that data directly to the main memory of another system. This improves performance and throughput by freeing up server resources. It has proven useful in applications that require fast and massively parallel high-performance computing (HPC) clusters and data center networks.

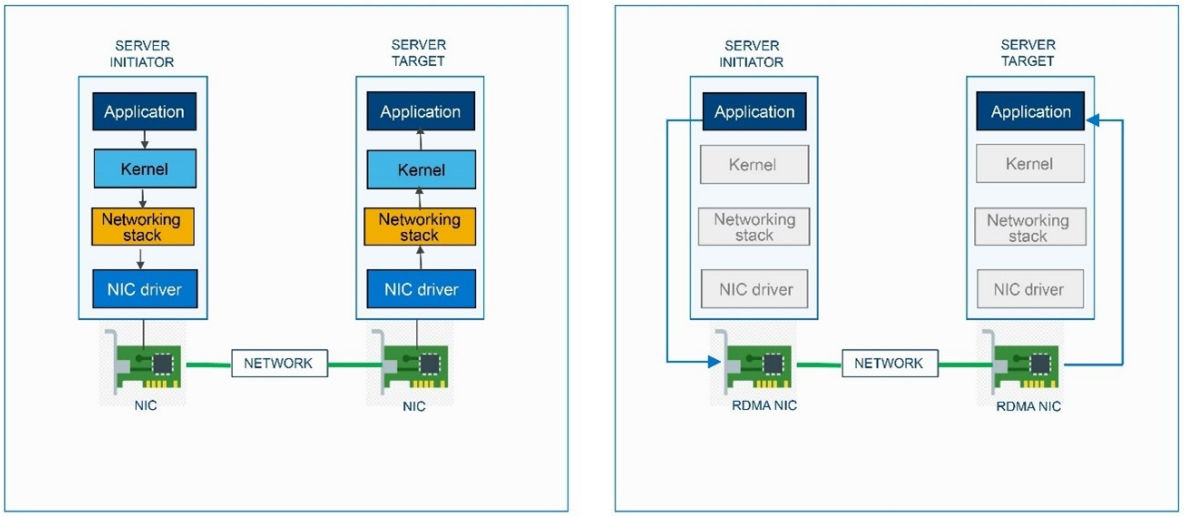


**Figure 35. Standard network connection vs Remote Direct Memory Access (RDMA) network connection**

# Scalable architecture for Dell Solutions for Azure Stack HCI

The Dell Integrated System for Microsoft Azure Stack HCI scalable architectures support from 2 to 16 nodes in a cluster. The scalable architectures support fully converged and non-converged network topologies. TheDell Integrated System for Microsoft Azure Stack HCI consists of two top-of-rack network switches that provide high availability by connecting the nodes to two separate network switches and implementing Switch Embedded Teaming (SET). SET can be configured either as a single interface or configured separately for storage and management traffic along with physical network adapters.

When using Remote Direct Memory Access over Converged Ethernet (RoCE), configure data center bridging (DCB) in both fully converged and non-converged topologies. DCB provides enhancements to the Ethernet protocol, which improves the functionality of data center networks. To take advantage of the Mellanox RoCE network adapters, Priority Flow Control (PFC) and Enhanced Transmission Selection (ETS) are required. PFC and ETS must be configured on all nodes and all network switches interconnecting the nodes. The scalable network topology allows for future cluster expansions as business needs grow.

# Fully converged network configuration

In the fully converged network topology, all storage ports from the server are connected to the same network fabric. Within the host operating system, the NIC ports are used for both storage and management/VM traffic.

The following network diagrams depict the integration of the fully converged architecture.

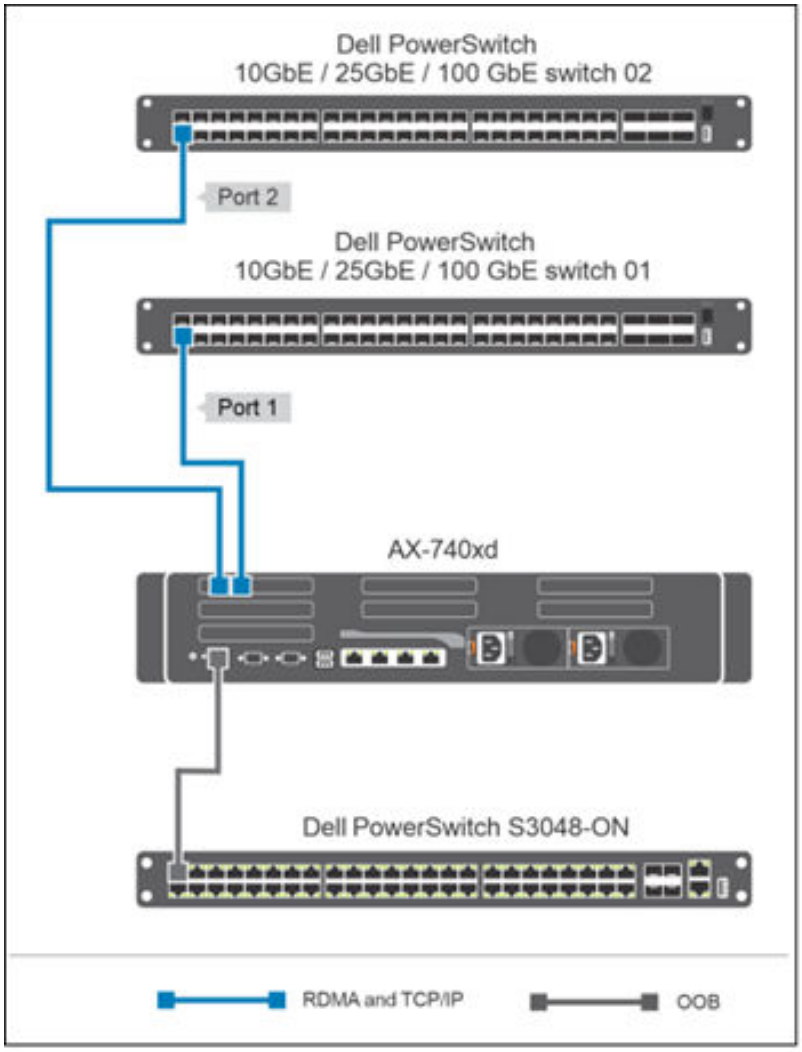


**Figure 36. Fully converged network topology with two NIC ports**

# Non-converged network configuration

In non-converged network configurations, the storage traffic is separated from the management/VM traffic using dedicated storage network adapters. The storage traffic can either be on the physical ports or implemented as virtual adapters in the host operating system that is connected to a switch embedded team (SET). The following figure shows this network integration option.

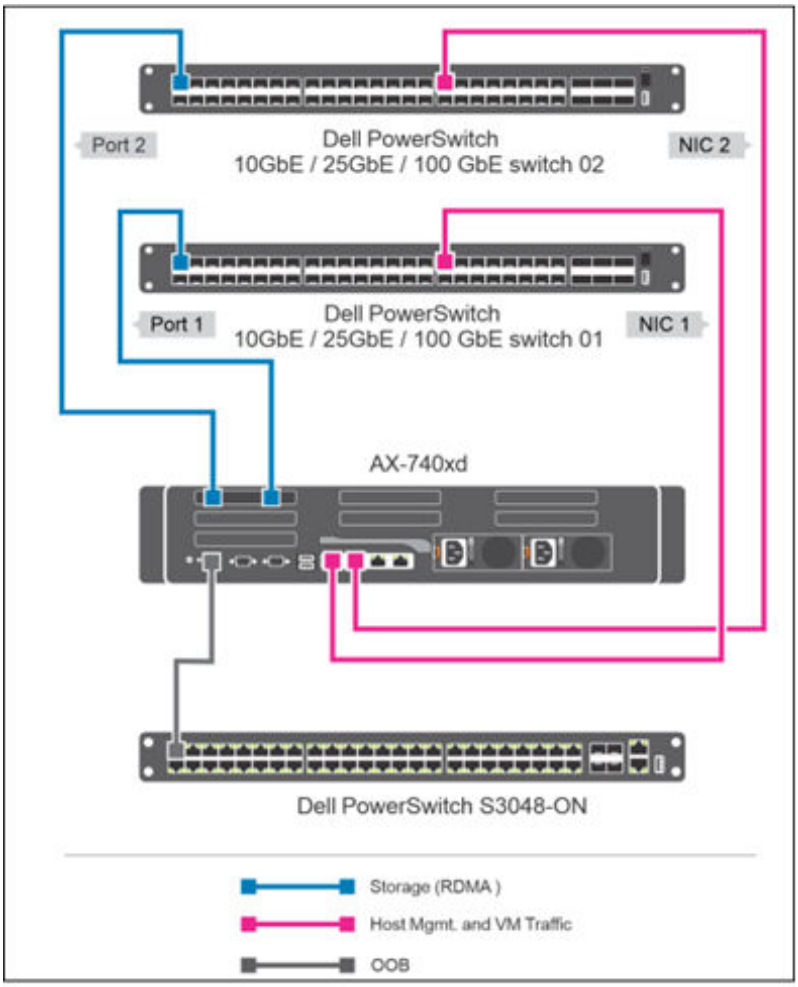


**Figure 37. Network integration in a non-converged configuration with two NIC ports**

# Switchless storage networking

Dell Integrated System for Microsoft Azure Stack HCI requires a reliable, high bandwidth, low latency network connection between each clustered node for storage network traffic. Dell Technologies recommends switchless storage to avoid the extra expense of high-speed network switches in locations such as ROBOs (Remote Office Branch Office).

Take cluster size into careful consideration before deploying a full mesh switchless Storage Spaces Direct cluster. Although it is possible, expanding a switchless cluster by adding a node may require installing additional network adapter cards. Creating the additional networks within Windows Server 2019 (EOL June 30, 2023) while the cluster remains running is not advised and the PowerShell scripts in the guide are not designed for expansion purposes. Dell Technologies has not validated expanding switchless clusters. To expand a switchless cluster, Dell Technologies recommends deploying the cluster from scratch, to properly follow cabling procedures and to prevent any operating system misconfiguration.

Microsoft announced support for 2-node switchless clusters for Windows Server Azure Stack HCI and Failover Clustering. Most recently, Microsoft has extended this support to three and four node clusters, using a full mesh storage network topology. A full mesh interconnect requires direct network cable connections between every node of the cluster. These direct connections can be configured either using single or dual-link full mesh topologies. For redundancy and performance purposes, dual-link direct network connections between every node of the cluster are recommended.

A full mesh switchless Storage Spaces Direct cluster provides these advantages:

- Data center environments that have an existing 1 GbE network switch infrastructure reduce costs that are associated with upgrading to 10 GbE (or faster) switches, because 10 GbE is a storage network minimum requirement of Storage Spaces Direct.
- Data center environments that have an existing 10 GbE network switch infrastructure would gain increased network throughput for storage traffic, because there are 25 GbE network adapter connections between each clustered node.

  **NOTE:** Management, VM, and any other external network traffic will still require connectivity to a network switch.

# Single-link full mesh topology

The following figure illustrates an example of a 3-node cluster that is configured with a single-link full mesh storage network interconnect. External network traffic uses the LOM/rNDC/OCP integrated network adapter, while storage network traffic uses a 10/25 GbE PCIe adapter per server node. The management network offers redundancy in the event of a storage network connectivity (cable or port) failure.

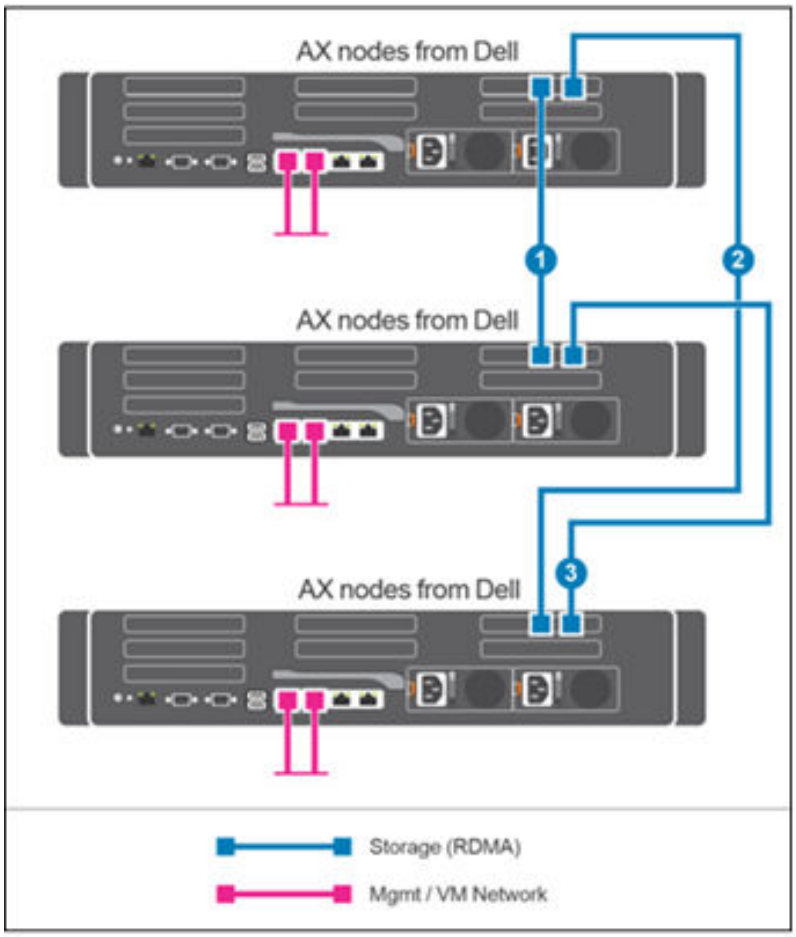


**Figure 38. 3-node single-link full mesh**

**NOTE:** Single-link full mesh topologies rely on the management network for storage and live migration redundancy purposes. The management network is also shared with VM traffic. During failure scenarios, and high throughput network load, contention between these three traffic types can degrade the cluster. For this reason, Dell Technologies recommends a dual-link full mesh network interconnect for high-performance storage configurations and workloads.

# Dual-link full mesh topology

The following figure indicates the two storage networks required to create a dual-link full mesh interconnect between two clustered nodes (also known as back-to-back). The blue dots indicate the recommended cabling order, beginning with the PCIe adapter in the lowest slot #, Port 1. After cabling storage, connect all nodes, LOM/rNDC/OCP, Ports 1 and 2 to a management/VM network.

# 2-node back-to-back connected architecture

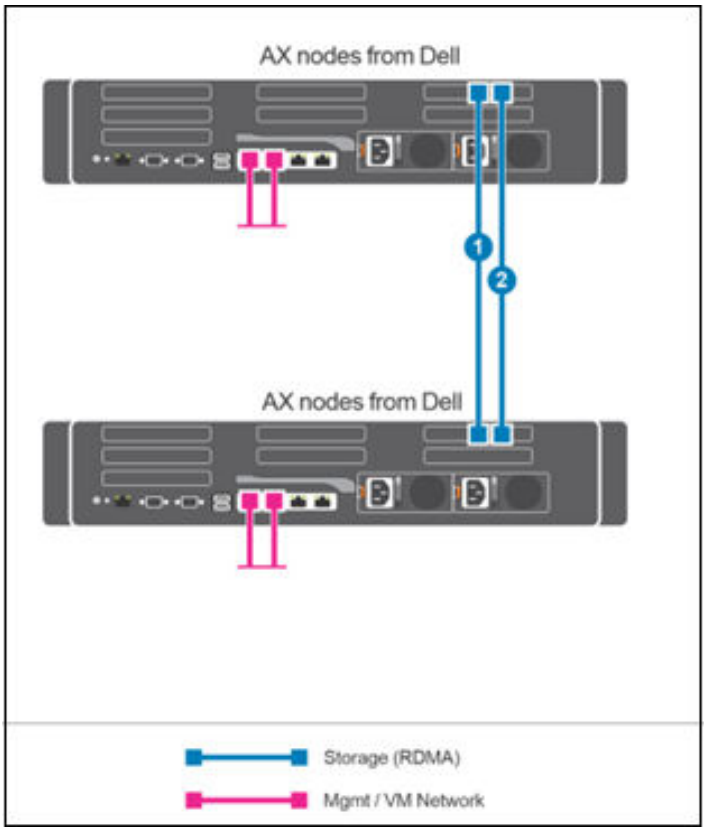


**Figure 39. 2-node dual-link full mesh cabling**

# 3-node single-link full mesh switchless interconnect

The following figure illustrates an example of a 3-node cluster that is configured with a dual-link full mesh storage network interconnect. External network traffic uses the LOM/rNDC/OCP integrated network adapter, while storage network traffic uses two 10/25 GbE PCIe adapters per server node. In this case, each node has redundant storage network connectivity to every other node in the cluster.

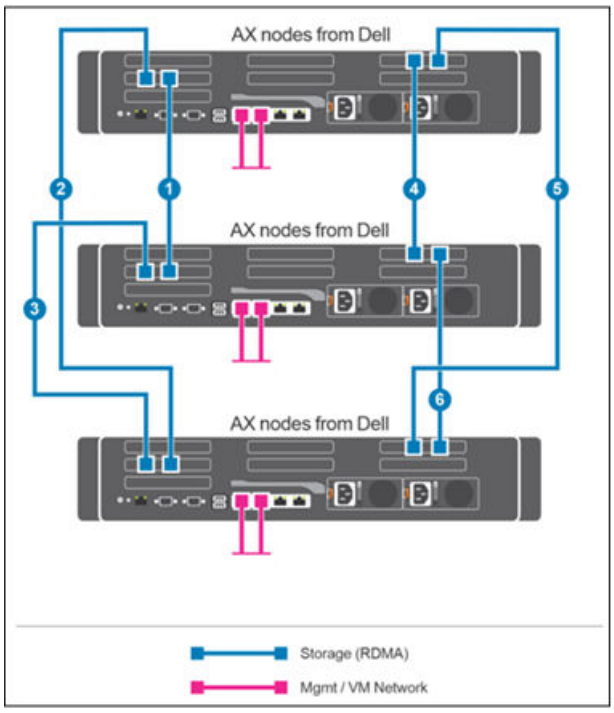


**Figure 40. 3-node dual-link full mesh topology**

For more information about network configuration options, see Dell Solutions for Microsoft Azure Stack HCI - Network Integration and Host Network Configuration Options.

# Single-node clusters

Single-node clusters are a unique opportunity for remote, edge or branch projects, very sensitive to costs and may tolerate the resiliency of a single server. Single-node clusters support the same workloads as multi-node Azure Stack HCI clusters and behave very similarly.

PowerShell has to be utilized for initial deployment and initially there is limited Windows Admin Center support. Follow these considerations and requirements when using Azure Stack HCI on a single server.

# PowerSwitch model options for scalable infrastructure

Dell PowerSwitch network switches are designed for flexibility and high performance for today's demanding modern workloads and applications. They are tested and validated to work across all possible configurations of the Dell Integrated System for Microsoft Azure Stack HCI. However, while choosing the right switch for the data center, we must consider the number of ports required taking into account any future cluster scale up operations. These switches are fully supported as part of the Windows Server HCI and Azure Stack HCI architecture.

# iWARP vs RoCE for RDMA

Traditional TCP/IP networks offer adequate bandwidth for many workloads, although it comes with some limitations, including CPU overhead and higher than wanted round-trip latency. The RDMA protocol reduces CPU overhead and latency by skipping through layers of the stack, resulting in rapid data transfers.

Two network implementations of RDMA, the Internet Wide Area RDMA protocol (iWARP) and RDMA over Converged Ethernet (RoCE), are embedded in the network adapter hardware, offloading nearly all the work from the CPU to offer faster networking:

- iWARP implements RDMA over IP networks using TCP, making it ideal for organizations that want to use RDMA over their existing IP network infrastructure without any specialized hardware. iWARP requires no additional configuration at the ToR switches for its implementation.
- RoCE v2 uses UDP and requires QoS (Quality of Service) to ensure packet delivery. It relies on an Ethernet network configured to use Layer 2 Priority Flow Control (PFC) or Layer 3 DSCP PFC to minimize congestive packet loss.

| Topic | iWARP (Internet Wide Area RDMA Protocol) | RoCE (RDMA over Convergent Ethernet) |
|---|---|---|
| Ease of deployment | No special network configuration (simple to install/configure) | Requires network configuration (more complex to install/configure) |
| QoS | Not needed | PAUSE frames, or DCB |
| Routable | Yes | Yes with RoCE V2 |
| Latency | Comparable to RoCE except for HPC applications | Best |
| Throughput | Similar to RoCE | Best |
| Vendor support | QLogic, Intel | Mellanox and Intel |
| Deployment size recommendation | Data center or cluster ideal for remote communication (storage replica) | May be more appropriate for clusters |
| Interoperable between vendors | Not supported | Yes - RoCE from two different vendors should not reside in the same server |

See Reference Guide: Switch Configurations – RoCE for details about the switch configurations for all RoCE-based deployments for Microsoft Azure Stack HCI Solutions from Dell Technologies.

See Reference Guide: Switch Configurations - iWARP for details about the switch configurations for all iWARP-based deployments for Microsoft Azure Stack HCI Solutions from Dell Technologies.

# Dell Fabric Design Center

Dell Fabric Design Center (FDC) is a cloud-based user interface that allows the user to design solution-specific data center fabrics and integrated turnkey solutions. It generates logical and physical network views, Bills of Materials, and Cabling diagrams from the details of clusters that are intended for the business. This can be used in conjunction with Live Optics and the Microsoft HCI Solutions from Dell Technologies sizer tool in the early design phases.

**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.

The Stretch Clustering with Azure Stack HCI chapter signifies how stretched clusters enable an organization to create an active workload balanced data center. Stretched clusters deliver workload mobility, reduce downtime, and provide disaster avoidance and load balancing capabilities.

8

# Stretch Clustering with Azure Stack HCI

**Topics:**



## Azure Stack HCI stretch clustering

KEY TAKEAWAY

Stretched clusters, a native feature of Azure Stack HCI, provides disaster recovery and business continuity to clusters in a data center without the need for any third party software. It allows spanning Azure Stack HCI systems between multiple sites using the Storage replica feature and automatic VM failover.
This characteristic of built-in resilience of Azure Stack HCI, combined with the ability to connect to Azure services, helps to achieve a hybrid hyper-converged system making it a complete, stable, and reliable solution to host business critical workloads.

As organizations modernize their infrastructure and application architectures, they rely more heavily on the digital assets that support their business operations. Because these workloads and data form the lifeblood of businesses, a flexible and robust DR solution is crucial.

Resiliency has become the watchword for organizations facing an array of threats, in the form of natural disasters, major IT failures, or cyberattacks. A Business Continuity and Disaster Recovery (BCDR) plan ensures an organization`s ability to remain operational after an adverse event has occurred. However, implementing a DR solution can be costly, complex, and time consuming to design and maintain. It also taxes already-strained IT staff and budgets. In addition, aging infrastructure often requires manual intervention during a DR scenario, which can be error-prone and can lead to unacceptable recovery time objectives (RTOs) and recovery point objectives (RPOs).

Dell Integrated System for Microsoft Azure Stack HCI addresses these challenges by providing a native DR solution through automated failovers to restore production quickly without manual intervention.

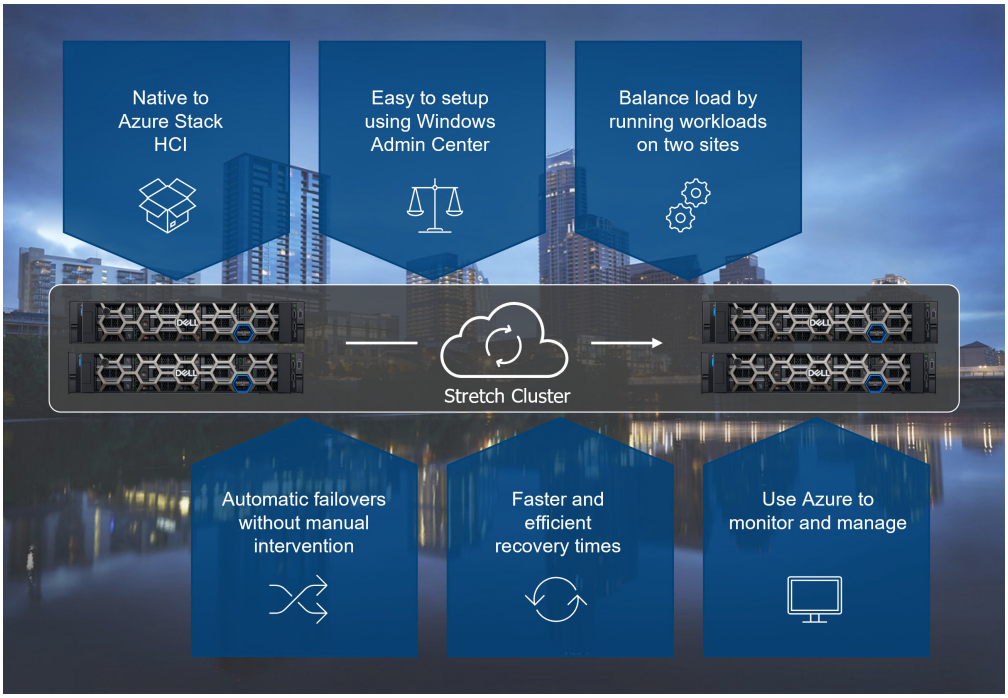


**Figure 41. Stretch cluster in Azure Stack HCI**

Grouping two or more servers in a single failover domain forms a cluster. By contrast, a stretched cluster is formed by extending the connectivity of servers to two different failover domains at multiple sites. This enables system administrators to increase resiliency for their IT infrastructure.

The following figure shows a stretch cluster that is implemented with servers in New York City and London. Both sites are connected to each other over the network. If one site fails, all the workloads automatically fail over to the secondary site, ensuring business continuity. This failover occurs without the need for an IT administrator to take action.

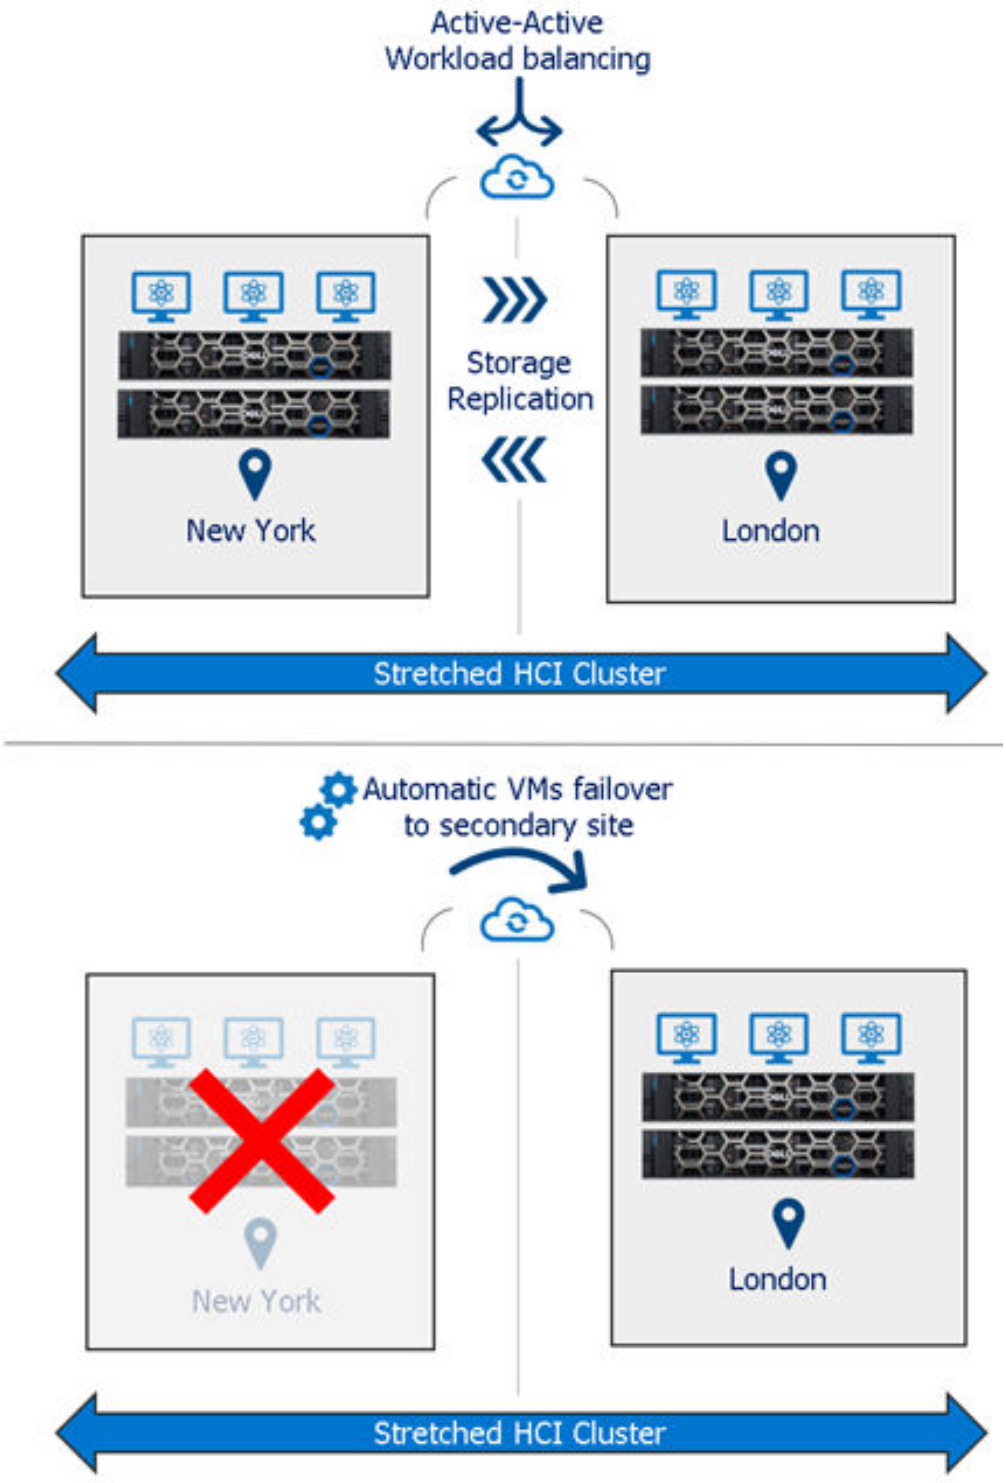


**Figure 42. Business continuity with Azure Stack HCI clusters**

Azure Stack HCI features integrate stretch clustering capabilities that use storage replica for volume replication. Stretch clustering enables organizations to split a single HCI cluster across two locations that might be in different rooms or buildings, cities, or even different geographical regions for DR.

# Comparison of traditional and stretched clusters

Stretch clustering with Azure Stack HCI offers great benefits for certain workloads. The following table provides a brief summary for how an organization can benefit from stretch clusters when compared to traditional clusters.

In a traditional cluster, we have the option to create a 2 or 3-way mirror confined to a single site. By contrast, in a stretched cluster each volume on the primary site will require a replica volume on the secondary site and log volumes on both sites. This has profound implications on storage efficiency and performance as we end up with 4x to 6x data resiliency for the data across the two sites.

**Table 2. Traditional cluster compared to stretched cluster**

| Traditional | Stretched Cluster |
|---|---|
| Single Site | Two sites |
| Site failure will cause complete outage | VMs restart on secondary site in case of site outage - Automated failover |
| Single Storage Pool | Two Storage Pools |
| 2x to 3x data resiliency | 4x to 6x data resiliency |
| No performance impact | Performance impact can occur due to inter-site throughput and latency<br>Additional writes to volumes and their respective log volumes on secondary site can cause overhead |
| Complete data loss in case of a calamity | Zero data loss for synchronous replication<br>Based on RPO for asynchronous |

# Business continuity planning

A thorough application rationalization and dependency-mapping effort must always precede conversations about recovery of IT infrastructure. Technologists and application owners must perform a business impact analysis to determine each application's criticality to business operations and to prioritize application recovery if a major service disruption occurs. As part of this analysis, business owners of the services must define their RTO RPO.

- **Recovery Point Objective** (RPO) is the amount of data loss that an application can sustain after a major service outage before causing an unwanted impact on the business. RPO is also measured in time because it is the point in time from which the data must be recovered. RPO focuses on the data or transactions of an IT system.
- **Recovery Time Objective** (RTO) is the duration of time in which the business process must be restored after a disaster occurs. The organization is only able to tolerate the degradation of a service or process for this period before the degradation causes serious consequences.

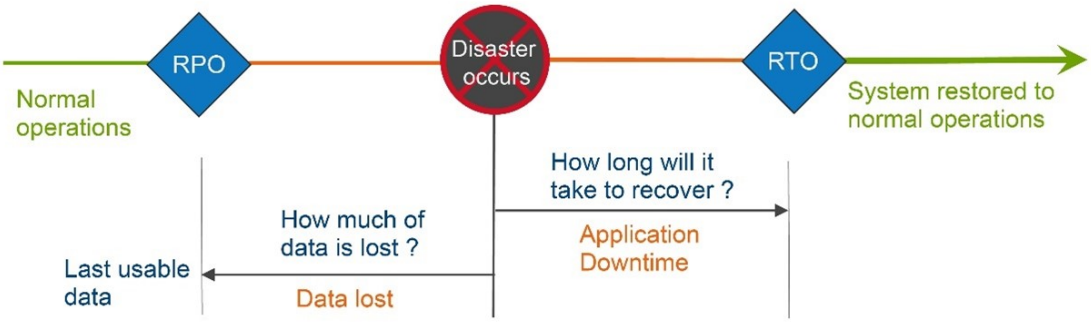


**Figure 43. RPO and RTO**

# Stretch clustering topology

Dell Technologies has validated the integrated system for Azure Stack HCI Stretch Cluster by creating a detailed and prescriptive reference architecture. The Reference Architecture—Dell Integrated System for Microsoft Azure Stack HCI:

Stretched Cluster Deployment provides an overview of the Microsoft Azure Stack HCI operating system and guidance about how to deploy stretched clusters in your environment. A robust set of configurations and different models enables you to customize your infrastructure for application performance, capacity, or deployment location requirements.

For a better understanding of stretch clustering, the following figure shows an example of an eight-node stretched cluster spanned across two different sites, Site A and Site B. Each site has its storage pool of disks, on top of which the volumes are created. Volumes are stretched across the two sites to provide better resiliency.

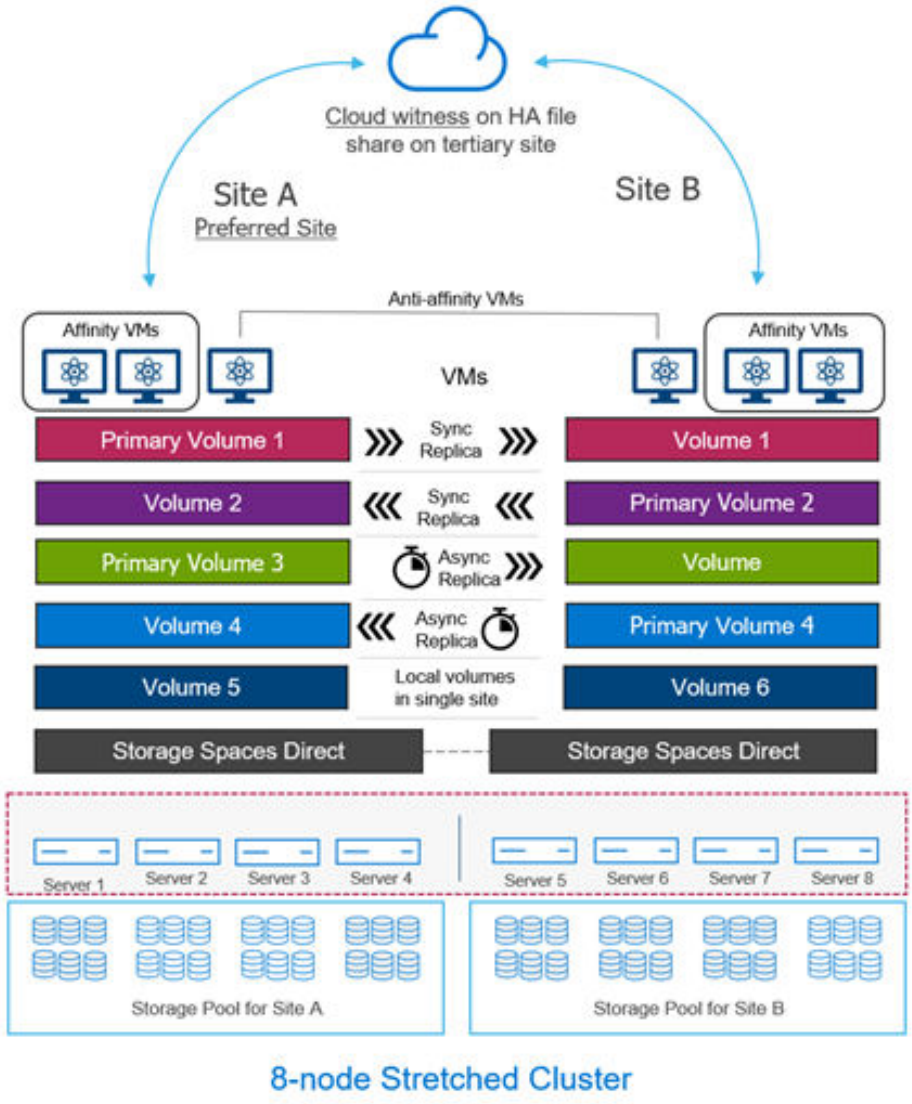


**Figure 44. Stretch clustering topology**

# Servers

Each site must have the same number of servers with the same hardware configuration. These servers are the nodes that are a part of the stretch cluster. To implement a robust configuration, you can choose to have a minimum of four nodes (two nodes per site) and a maximum of 16 nodes (eight nodes per site) across both the sites.

# Storage pools

S2D provides the software-defined storage layer for Azure Stack HCI. A stretched cluster environment has two storage pools, one per site. The Azure Stack HCI operating system can stretch volumes across sites making it appear as though there is only

one volume. The primary volume is accessible from the nodes at that site. The secondary volume at the other site is meant for standby and is brought online when the primary volume goes offline.

## Volumes

The servers residing in different sites replicate volumes either synchronously or asynchronously. With the synchronous approach, writes to persistent storage are replicated to both sites before they are acknowledged. With the asynchronous approach, writes are acknowledged when they are persisted to one site, and then replicated to the other site moments later, based on the chosen RPO. Local volumes that operate in the boundaries of a single site can be created.

## Setup configuration

A stretched cluster can be set up as either active/active or active/passive. In an active/active setup, both sites actively run the VMs or applications, therefore the replication occurs bidirectionally. In an active/passive setup, the passive site is always dormant unless there is a failure or planned downtime, waiting for an automatic failover from the active site.

## Preferred site configuration

In a stretched cluster, designate one site as the preferred site to define a location on which all resources will run. The other site becomes the secondary site. The system uses the preferred site if there is a loss of network connection between the two active sites, the preferred site is the one that remains operational. These designations also ensure that VMs and volumes come up on the preferred site after a cold start or after network connectivity issues are resolved.

## Affinity and anti-affinity VMs

Azure Stack HCI enables you to define the affinity and anti-affinity controls for VMs in a cluster. An affinity rule ensures that two VMs reside in the same site to ensure locality. For example, the web tier of an application and its database are in the same zone. An anti-affinity rule ensures that two VMs stay on separate sites, for example when hosting two domain controllers on separate sites.

## Cloud witness

To orchestrate a quorum in a stretched cluster, a tie-breaker vote, often called a witness, is required. A dynamic quorum ensures that preferred sites survive if network connectivity fails. Using a dynamic quorum, weighting is decreased from the passive (replicated) site first to ensure that the preferred site survives if all other things are equal. In addition, server nodes are pruned from the passive site first during regrouping after events such as an asymmetric network connectivity failure. During a quorum split across two sites, if the cluster witness cannot be contacted, the preferred site is automatically elected. The server nodes in the passive site then drop out of cluster membership. This scenario allows the cluster to survive a simultaneous 50 percent loss of votes. The preferred site can also be configured at the cluster role or group level. In this case, a different preferred site can be configured for each VM group. This enables a site to be active and preferred for specific VMs.

# Azure Stack HCI stretch clustering scenarios and concepts

In general, stretch clustering on Azure Stack HCI is an ideal DR solution for the following scenarios:

- Introducing automatic failover with orchestration for recovery of a web-based application's front-end server tier after a disaster renders one hosting location unavailable.
- Distributing primary and secondary instances of infrastructure core services, such as Microsoft Active Directory, across two physical locations.
- Hosting applications with lower write I/O performance characteristics.
- Running file-system-based services and other business services that can tolerate being hosted on crash-consistent volumes. For database workloads such as Microsoft SQL Server, which often cannot sustain the loss of even a single transaction, using application-layer recoverability solutions such as SQL Always-On might be more appropriate.
- Stretch clusters are set up either as active/passive or active/active.

# Active/passive setup

In an active/passive setup, business consumers actively connect to a preferred site for their application and workload resources. The infrastructure at the secondary site essentially remains idle until a failover occurs. When a failover occurs, active workloads are restarted at the secondary site, and business consumers connect to their applications running on that secondary site.

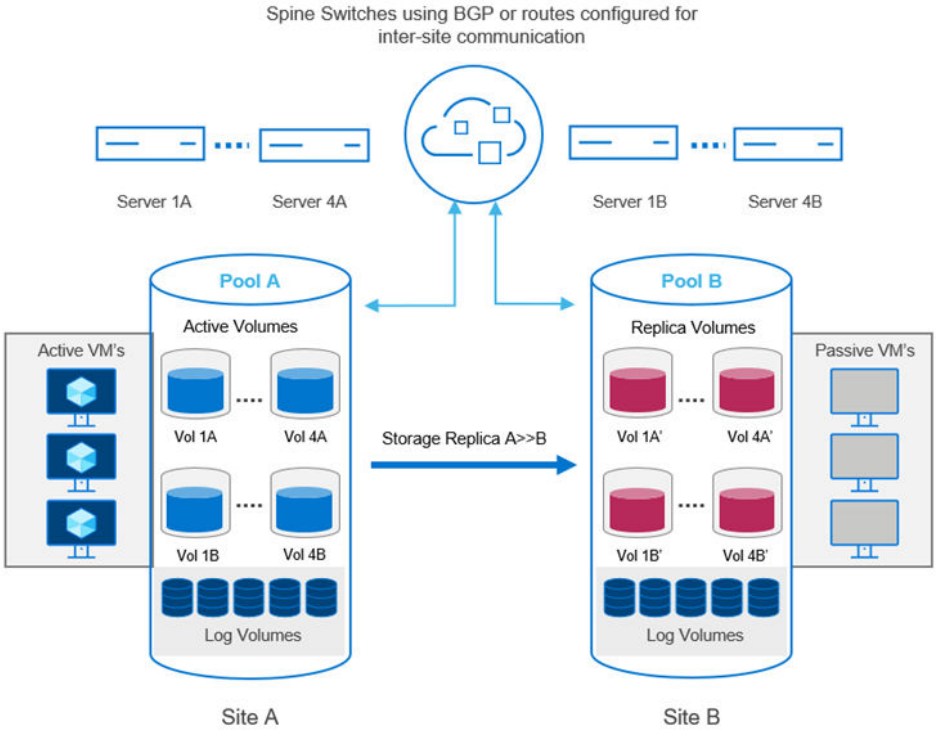


**Figure 45. Active-passive stretch clustering architecture**

# Active/Active

In an active/active setup, replication occurs bidirectionally from either site. Business consumers connect to active applications in both sites at any given time. This setup tends to be a more efficient use of an organization's investment in infrastructure because resources in both sites are being used.

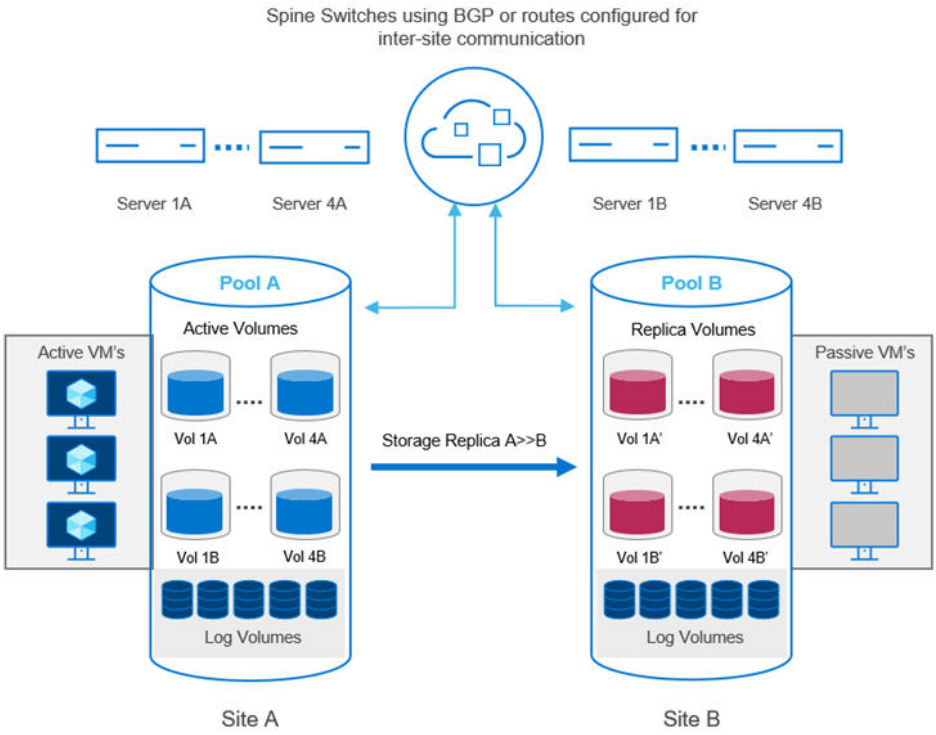


**Figure 46. Active/active stretch clustering architecture**

The Storage Replica service supports synchronous and asynchronous replication:

## Synchronous replication

Mirrors data across low-latency networks with crash consistent volumes to help ensure that no data loss occurs at the file system level during a failover. This replication technique enables workloads to come back online automatically. With synchronous replication, data blocks are written to log files on both sites before being committed.

## Asynchronous replication

Mirrors data across greater distances over networks with higher latencies. This strategy does not guarantee that both sites have identical copies of the data at the time of failure. This replication technique does not enable workloads to be brought back online automatically after a failover. With asynchronous replication, the remote node accepts blocks of replicated data and lazily acknowledges back to the source copy.

# Stretch clustering best practices with Azure Stack HCI

Stretch clustering included with Azure Stack HCI, must be configured on validated hardware found in the Azure Stack HCI Catalog. The engineering team at Dell Technologies has deliberately designed and validated the Dell Integrated System for Microsoft Azure Stack HCI to optimize the HCI experience. The Dell Integrated System for Microsoft Azure Stack HCI: Stretched Cluster Deployment Reference Architecture Guide is based on detailed results from extensive testing of stretch clustering in the labs.

Critical design considerations and best practices in this reference architecture include:

- The cluster size must be a minimum of four nodes and a maximum of 16 nodes. Both sites must be running the same number of AX nodes in the cluster, and these nodes must have identical hardware configurations.

- Consider using three-way mirror volumes for best performance. Capacity requirements must be carefully weighed considering the raw storage used and the storage efficiency of mirror volumes in a stretch clustering configuration. A three-way mirror results in 25 to 33 percent storage efficiency.
- Latency between sites is a significant consideration in a stretch clustering environment. The networks connecting the two sites must have enough bandwidth to accommodate the data rate of change and to contain write I/O. Although we make no specific recommendations on distance between sites, note the following guidelines:
  - For both synchronous and asynchronous replication, aim for less than 200-millisecond roundtrip latency between the AX nodes and the witness. If the witness is a highly available file share, it must be created at a tertiary site and not at either site hosting the stretch cluster nodes. The witness could also be an Azure cloud witness.
  - For synchronous replication, aim for an average roundtrip latency of 5 milliseconds or less between the AX nodes in Site 1 and the AX nodes in Site 2.
  - For asynchronous replication, which has no specific latency requirements, the most important consideration is correct RPO configuration.
- I/O-intensive workloads are not good candidates for stretch clustering environments. Both synchronous and asynchronous replication involve high I/O to the underlying hard drives during write operations, which can result in a significant performance impact.
- Storage Replica is not a backup and restore solution. Because it is a general purpose, storage-agnostic replication engine, use Storage Replica along with backup and recovery software capabilities. Back up only the active volumes, not the replica, or log volumes.
- Configure preferred sites in a stretched cluster to define a location to run all resources. Such configuration ensures that VMs and volumes become available on the preferred site after a cold start or after network connectivity issues.5B

# Validated network topology

## Basic configuration

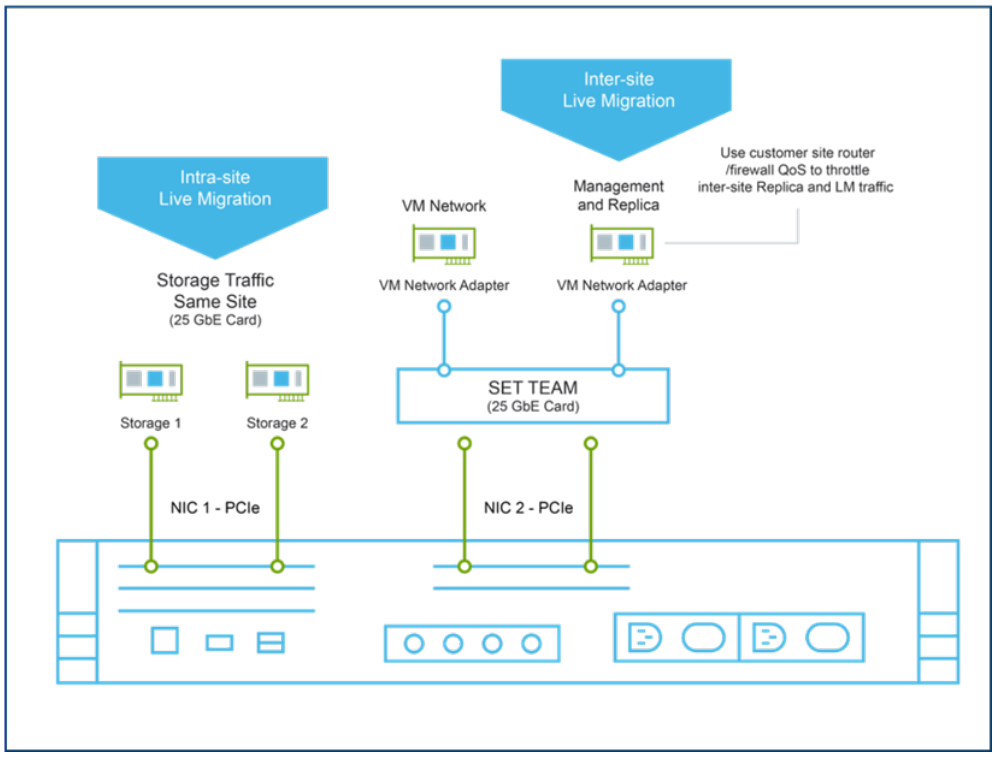


**Figure 47. Network topology for a stretched cluster (basic)**

This section describes the host network configuration and network cards that are required to configure a basic stretched cluster. The purpose of this topology is to keep the host and inter-site configuration simple with little or no change to a standard stand-alone cluster networking architecture.

This topology uses two 25 GbE NICs for each host on both sites. One NIC is dedicated to intra-site storage traffic, similar to a stand-alone Storage Spaces Direct environment. The second NIC is used for management, compute, and Storage Replica traffic. To ensure management traffic is not bottlenecked due to high traffic on the Replica network, the customer network team should throttle traffic between the two sites using QoS rules. It is recommended that the network bandwidth is throttled to 50 percent of the capacity of the total number of network cards supporting the management NIC team.

## High throughput configuration

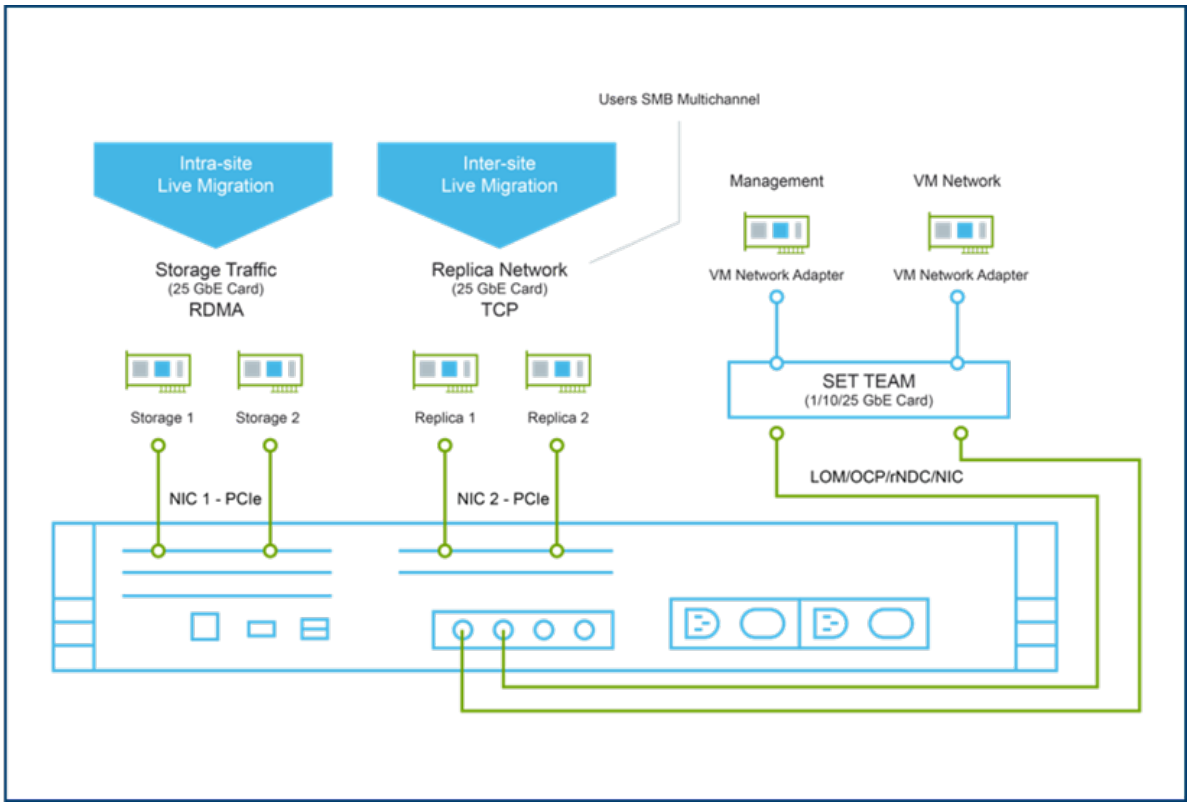


**Figure 48. Network topology for a stretched cluster (advanced)**

This topology uses two 25 GbE NICs for each host to configure a high throughput stretched cluster and one rNDC for management traffic. One NIC is dedicated for intra-site RDMA traffic, similar to a stand-alone Storage Spaces Direct environment. The second NIC is used for replica traffic. SMB Multichannel is used to distribute traffic evenly across both replica adapters to increase network performance and availability. SMB Multichannel enables the use of multiple network connections simultaneously and facilitates the aggregation of network bandwidth and network fault tolerance when multiple paths are available. For more information, see Manage SMB Multichannel.

The Set-SRNetworkConstraint cmdlet is used to ensure that replica traffic flows only through the dedicated interfaces and not through the management interface. Run this cmdlet once for each volume.

## Server and site affinity rules for VMs

PowerShell commands can be used to create affinity and anti-affinity rules for VMs in a cluster. An affinity rule is one that establishes a relationship between two or more resource groups or roles, such as VMs, to keep them together in an Azure Stack HCI cluster. An anti-affinity rule does the opposite, keeping specified resource groups apart from each other.

You can use storage affinity rules to keep a VM and its associated Virtual Hard Disk v2 (VHDX) on a Cluster Shared Volume (CSV) on the same cluster node. This ensures CSV redirection does not occur and keeps application performance at optimal levels. For more information, see Storage affinity rules.

For several proof-of-concept test scenarios, see the white paper Adding Flexibility to DR Plans with Stretch Clustering for Azure Stack HCI.

The Hybrid Cloud Management chapter covers how we can manage and monitor the hybrid cloud infrastructure from a single pane of glass.

# 9

# Managing the Hybrid Cloud

**Topics:**

# Managing the hybrid cloud

KEY TAKEAWAY

Windows Admin Center (WAC) is the future of remote server management, designed to modernize and simplify the IT administrative experience. Consolidating different management tools into a single intuitive remote server management experience, the Windows Admin Center integrated with Dell OpenManage enables IT professionals to operate and manage their AX nodes for their Azure Stack HCI solutions.

Azure Arc extends Azure services and management to any infrastructure. With Arc, IT administrators can now consistently build, deploy, operate, and manage all your workloads running on-premises traditional nodes, physical or virtual environments, and cloud-native and edge applications, all using a consistent control plane.

Modernizing data centers with the Hybrid cloud framework provides the ability to improve scalability and performance, lower total cost of ownership (TCO), and make it easier to manage and secure the IT infrastructure. Azure Stack HCI optimizes on-premises hardware by leveraging virtualization and integrating with key Azure hybrid services. As shown in the following diagram, Azure Stack HCI brings a cloud-centric operational model by connecting to Azure hybrid services through Azure Arc.

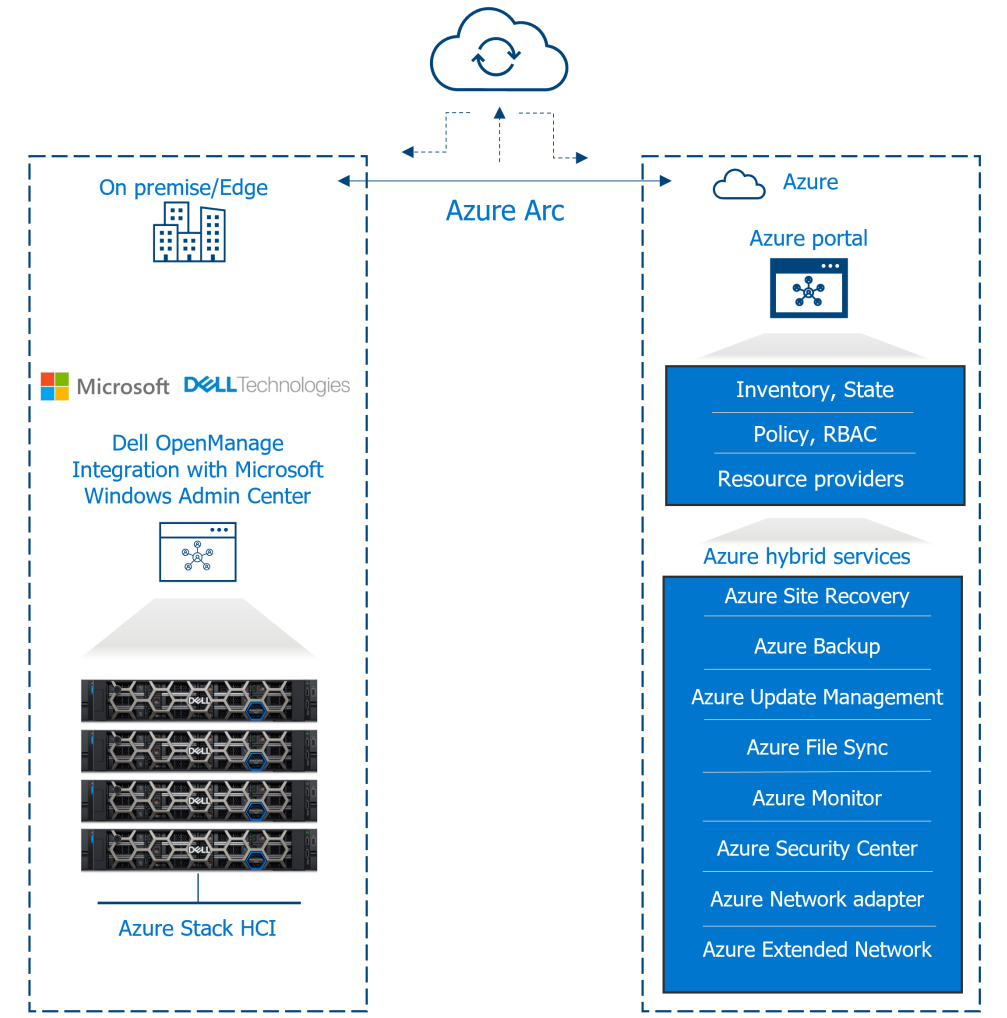


**Figure 49. Simple setup of Azure Stack HCI connecting to Azure hybrid services through Azure Arc**

Windows Admin Center (WAC) is a simplified management tool for managing Windows Servers, clusters and edge/HCI nodes residing on-premises. It brings in a lot of value being a remote management console that can have an upside in a modern IT environment. By connecting to your Windows server from anywhere, WAC can be installed on a Windows 10 or 11 workstation, a single-server node, or on a virtual machine. It contains core tools for troubleshooting, configuration, management, and maintenance for Windows Server, Windows Client, Software-Defined Storage (SDS), Software-Defined Network (SDN), Microsoft Hyper-V Server, and more. WAC also enables customers to take advantage of cloud services by allowing them to connect their HCI/Windows Server nodes to Azure Hybrid services. For a complete list of all Azure hybrid services integration with Windows Admin Center, see the official documentation from Microsoft: Connecting Windows Server to Azure hybrid services.

WAC encompasses a lot of common scenarios, compared to traditional Microsoft Management Console (MMC) in-box tools. It is presently a complementary tool to RSAT (Remote Server Administration Tools) and might become the exclusive administrative tool. Roles such as Active Directory, DHCP, DNS, have been added to WAC. For a detailed look at what tools are included with Windows Admin Center, see the Microsoft documentation Manage Servers with Windows Admin Center.

As more HCI nodes are deployed, it becomes increasingly difficult to manage the whole IT estate, in addition to the existing traditional infrastructure and multi-cloud resources. Also, edge scenarios are poised to create even thicker server sprawls. Azure Arc extends Azure`s common control plane to manage servers and Kubernetes clusters across on-premises, multi-cloud, and edge. Azure Arc empowers customers to organize, govern, secure, and manage resources by delivering a consistent multi-cloud and on-premises management platform, all from a single pane of glass.

OpenManage Integrations for Microsoft System Center support native management of PowerEdge servers and Dell Solutions for Microsoft Azure Stack HCI. Console-based offerings are made available for integration with Microsoft System Center Operations Manager, System Center Configuration Manager and System Center Virtual Machine Manager.

# Dell OpenManage Integration with Microsoft Windows Admin Center (OMIMSWAC)

Microsoft Windows Admin Center, first released in April 2018, is a locally deployed, light weight, browser-based management tool that lets IT administrators manage Windows servers with or without Azure connectivity or dependency. It gives complete control over all aspects of server infrastructure and is useful in managing servers in private networks and edge locations that are not connected to the Internet.

WAC was built with extensibility in mind from the beginning, with the ability for Microsoft and third-party developers to build tools and solutions to simplify and automate IT administrative tasks. Microsoft offers an SDK that enables developers to build their own tools for Windows Admin Center. Microsoft is constantly improving Windows Admin Center by adding new features and updates regularly. WAC sends out an alert whenever there is an update available.

Dell Technologies has made WAC even better, in conjunction with OpenManage Integration with Microsoft Windows Admin Center. The integration with WAC adds another level of management, visibility, and control to the resources in the data center. The extension is designed to communicate with the cluster nodes in-band and completely agent-free using the OS-to-iDRAC passthrough and Redfish technology. Its first release in 2019 featured the hardware and firmware inventory, real-time health monitoring, iDRAC integrated management, and troubleshooting tools. Since then, based on customer feedback, Dell Technologies has improved the feature set by adding capabilities such as Full Stack Cluster-Aware Updating, CPU core management, Cluster node expansion, and much more.

The OMIMSWAC solution brings the following benefits to organizations:

- Streamlines operational efficiency and offers great flexibility by directly monitoring and managing servers and cluster environments.
- Enables organizations to make informed and rapid decisions impacting the availability of virtualized, software defined, and cloud infrastructure.
- Saves IT attended time by bringing automation to Cluster Creation workflow which uses a consistent, repeatable, and guided process. Dell HCI Configuration Profile (HCP) embedded as part of cluster creation workflow ensures an error free, hardware compatibility and symmetry checks for nodes in the cluster.
- Simplifies Life Cycle Management operations with 1-Click Full Stack Cluster-Aware Updating which updates the operating system, firmware, drivers, and BIOS of the AX nodes in a single workflow with zero impact to workloads running on clusters.
- Eliminates time consuming guess work during a cluster expansion process by allowing to select only compatible AX nodes with HCI configuration profile checks.
- Reduces Azure subscription cost by allowing IT administrators to size and modify the CPU cores based on workload performance in a cluster aware fashion.
- Proactively secures HCI infrastructure from malicious threats and inadvertent changes using Secure-Core and Infrastructure Lockdown features.
- Prevents configuration drift and provides infrastructure compliance visibility and remediation with Dell HCI Configuration Profile Policies for Azure.

For the latest and the greatest set of features, see Dell OpenManage Integration with Microsoft Windows Admin Center key features.

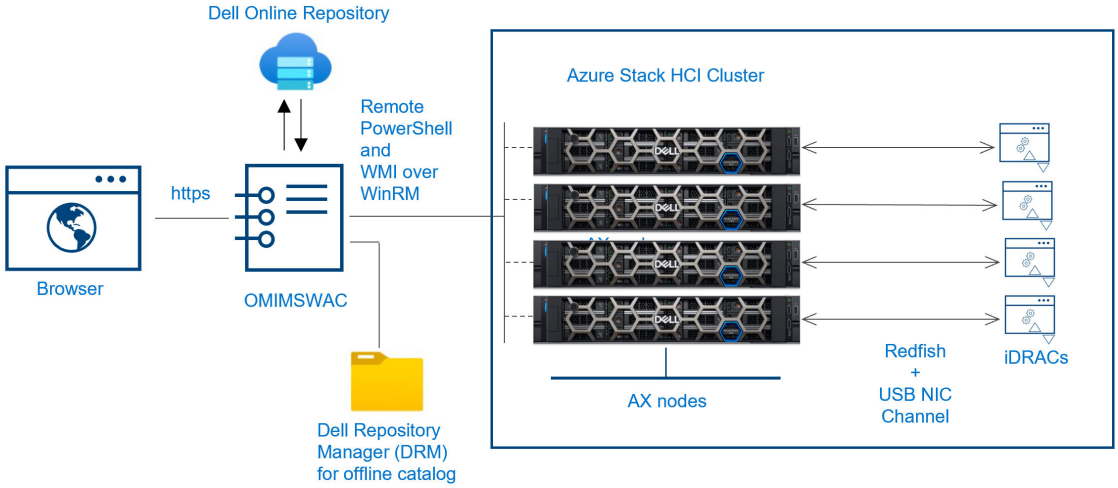


**Figure 50. Block diagram of OMIMSWAC**

Integrated Dell Remote Access Controller (iDRAC) is designed for secure local and remote server management and to help IT administrators deploy, update, and monitor Dell PowerEdge servers and AX nodes. The iDRACs are embedded onto each server and communicate regularly with the server hardware.

OMIMSWAC is installed as an extension for Windows Admin Center on a supported system. It can then be accessed from anywhere over a browser with an IP address or a URL. OMIMSWAC communicates to individual servers through remote PowerShell and WMI (Windows Management Instrumentation) over WinRM (Windows Remote Management) which in turn connects to iDRACs over the in-band USB NIC channel to collect hardware health and inventory. For BIOS, firmware, and driver updates, OMIMSWAC fetches the update packages and utilities from an online repository and runs them on the individual nodes. Customers also have an option to update the system when not connected to the Internet, through the offline catalog provided by Dell Repository Manager (DRM). Redfish is the DMTF's industry standard protocol which provides a RESTful interface for server management which is both human readable and machine capable. It leverages common Internet and web services standards to expose information directly to the modern tool chain.

# Important features of OMIMSWAC

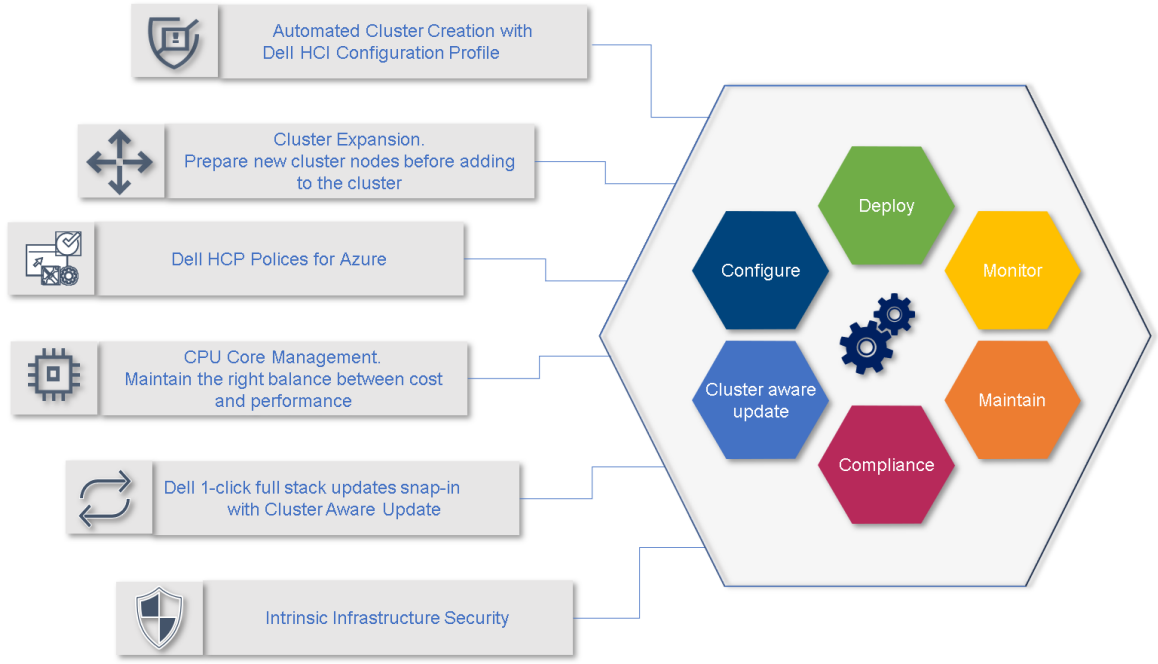


**Figure 51. Features of OMIMSWAC**

# Automated cluster creation with Dell HCI configuration profile

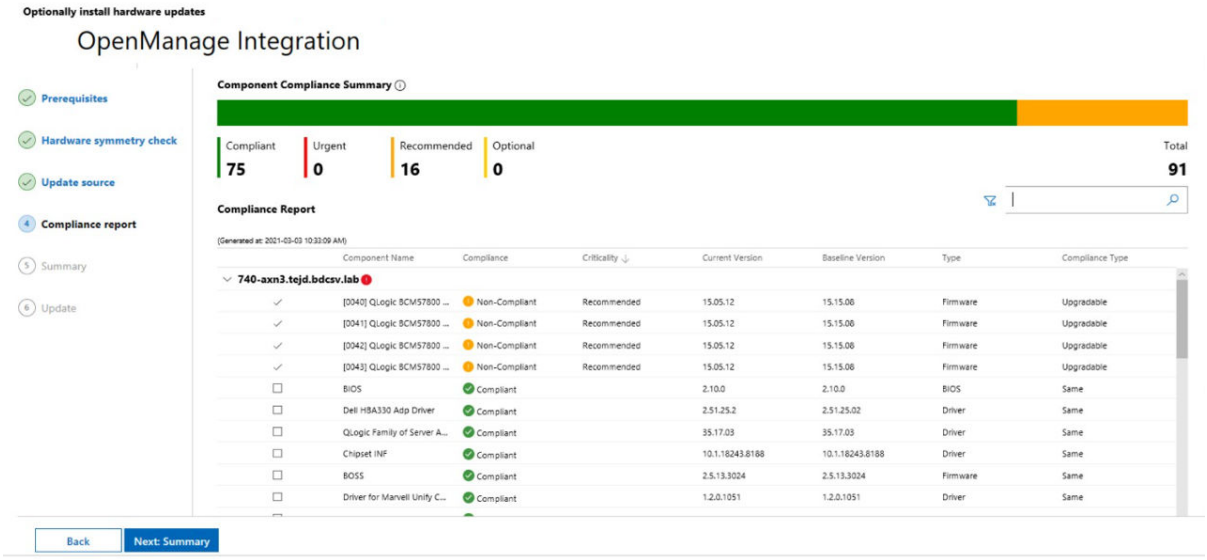


**Figure 52. Automated cluster creation workflow**

The Microsoft Azure Stack HCI operating system comes pre-installed with the AX nodes from the Dell factory. After the initial rack and stack activities, a simple deployment workflow in OMIMSWAC will automatically create the cluster. The Auto Cluster Creation feature uses OpenManage Integration snap-in that enables update of AX nodes while creating the cluster. This helps in hardware configurations compliance of selected nodes with Dell recommended hardware configurations. For optimal cluster performance, the OpenManage Integration performs a solution symmetry check to ensure the homogeneity of all cluster nodes before the deployment is initiated. This integration can also apply hardware solution updates, such as for BIOS, firmware, and drivers, at the same time as operating system updates, as part of cluster creation -- all with just a single reboot. To summarize the benefits of automated cluster creation, here are a few points:

- Save time and reduce risk associated with manual deployment efforts.
- Optimal cluster configuration using a consistent, repeatable, and guided process.
- Hardware compatibility and symmetry checks ensure that hardware configurations are validated against Dell HCI Configuration Profile.
- Perform hardware updates for connected or edge/disconnected nodes.
- Compliance report of operating system, BIOS, firmware, and drivers.
- Remediation before cluster build.

# Cluster expansion

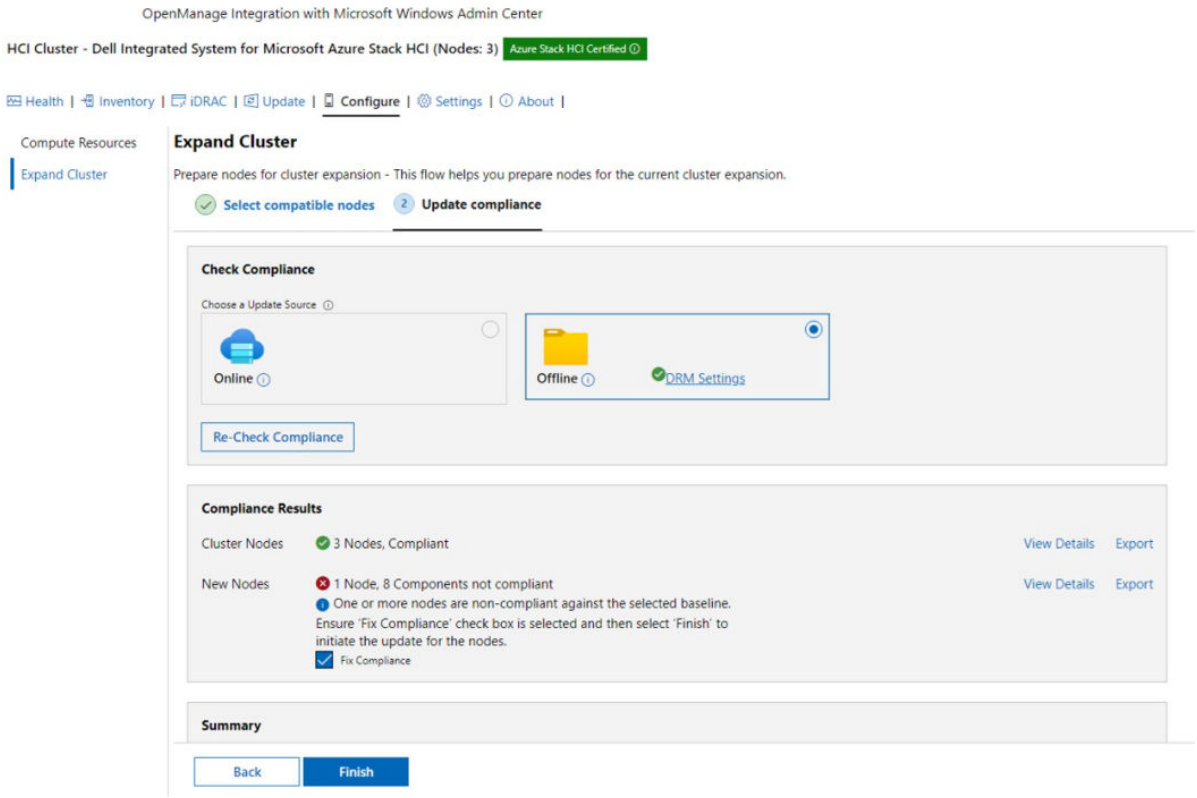


**Figure 53. Cluster expansion workflow**

Cluster Expansion is another significant value add which helps prepare the cluster nodes before adding them to the HCI cluster. OMIMSWAC readies the cluster nodes before adding them, but expanding a cluster requires some steps recommended by Microsoft. For details about creating and expanding clusters, see HCI Operations Guide—Managing and Monitoring the Solution Infrastructure Life Cycle. The new node to be added is checked for high-level compatibility, licensing, and HCI Configuration Profile to ensure it can be added successfully and consistently. Following these checks, a compliance check will be performed to ensure the new nodes' BIOS, firmware, and driver versions are consistent with those running on the cluster. If in case the new node is non-compliant the workflow will guide through an update process where the node is updated followed by a reboot.

# CPU core management

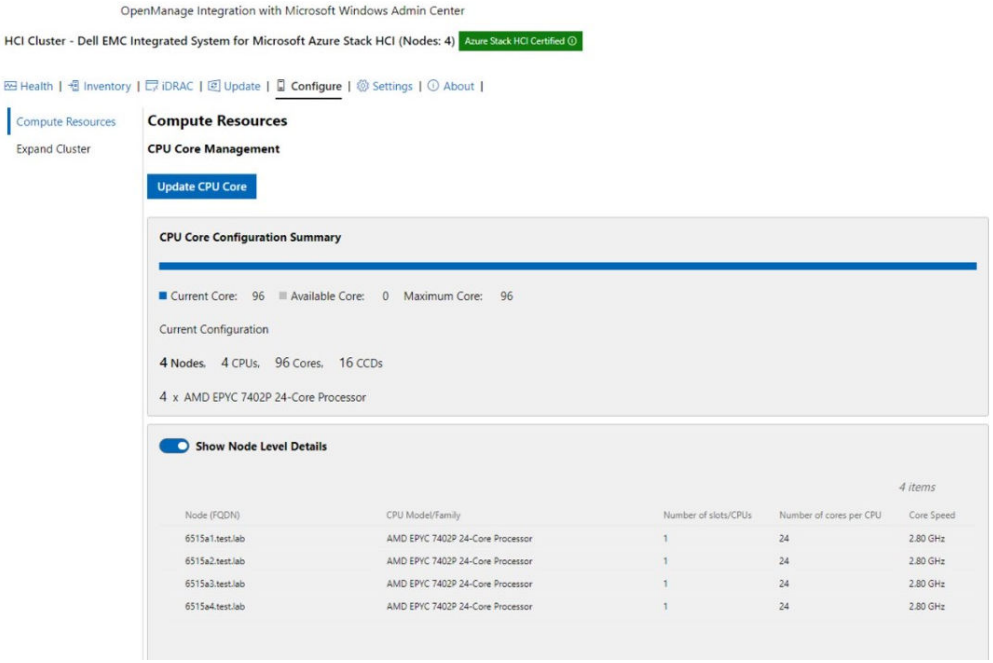


**Figure 54. CPU core management workflow**

CPU core management feature right sizes the cluster for workloads which creates the best Return on Investment (ROI). The Azure Stack HCI operating system is billed based on subscription which is on a per core basis for on-premises servers. This can become expensive because many servers are equipped with high core count processors. OMIMSWAC can help lowering this expense and optimizing the Total Cost of Ownership (TCO) by modifying the CPU core count dynamically. The tool provides a summary of current CPU core configuration along with available cores for expansion or reduction. Users can adjust and simultaneously apply the CPU core count BIOS settings without leaving the OpenManage Integration extension with a reboot-now or on a next-reboot option. Because this happens through Cluster-Aware Updating, the mandatory reboot of servers will not interrupt the running workloads.

# 1-Click full stack LCM using cluster-aware updating

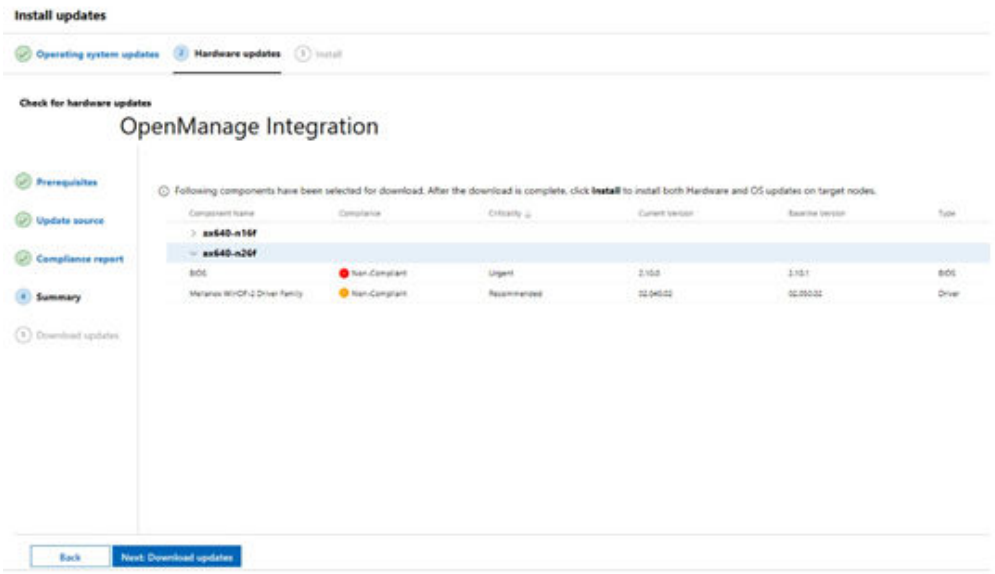


**Figure 55. 1-Click full stack LCM workflow**

As with any production system, it is critical to apply operating system and hardware updates routinely to protect businesses from security vulnerabilities and to take advantage of innovative new features and performance and reliability enhancements. Microsoft releases quality and security updates for the Azure Stack HCI operating system, and Dell releases updates for BIOS, firmware, and drivers regularly. To keep the system up to date, Dell OpenManage Extension installs a snap-in to Microsoft's Cluster-Aware Updating extension. This produces a simple consolidated workflow that orchestrates the updates of the OS and the hardware. All these updates only require a single reboot per node in the cluster, with no impact to running workloads. Below is a summary of benefits and reduced risks:

- Single workflow orchestrates operating system, BIOS, firmware, and driver updates.
- One reboot per cluster node at a time with no interruption to workloads.
- Node validation checks ensure cluster update readiness.
- Dell Repository Manager created solution catalog for edge/disconnected scenarios.
- Fully automated experience with the online solution catalog.
- Compliance report and remediation.

## Infrastructure lock for AX nodes

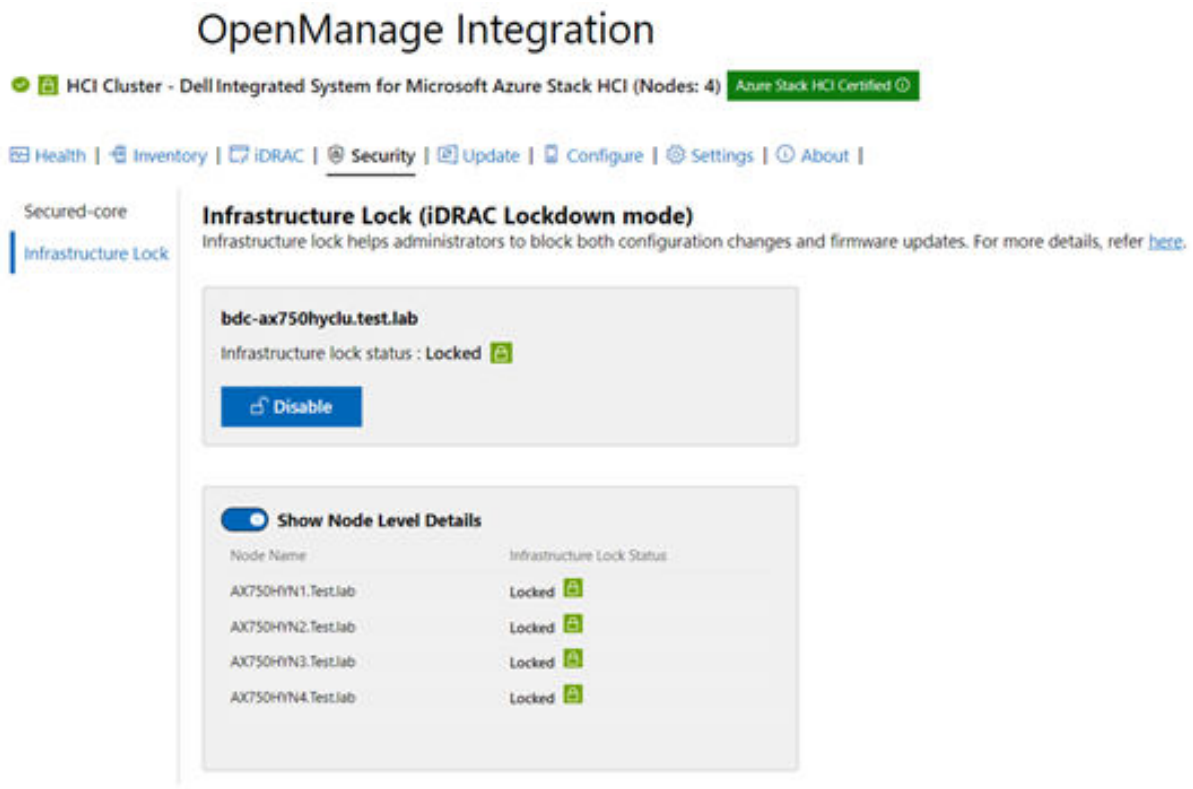


**Figure 56. Infrastructure lock workflow**

The AX nodes are protected from unintended changes that occur in a data center for various reasons by Infrastructure lock built into the tool. With this lock turned ON, the feature prevents the infrastructure from any configuration changes or updates to be done to the BIOS, iDRAC, firmware, and drivers. The tool facilitates these setting changes across all the nodes in the cluster. It also allows administrators to disable the infrastructure lock for HCI and failover clusters for the intended configuration changes, and update the cluster, and enable it back automatically after the operation is finished.

## Secured core

This feature allows you to enable Secured-Core server for the infrastructure (failover clusters) running on Azure Stack HCI and Windows Server 2022.

- In OpenManage Integration extension, this feature enables advanced hardware protection of the Secured-Core server consisting of hardware root-of-trust and firmware protection at the BIOS level in a cluster-aware manner.
- In Windows Admin Center Security extension, this feature enables dependent operating system security features in failover clusters.

- OpenManage Integration leverages the Windows Admin Center security extension tool to display the secured core status of clusters and cluster nodes at the operating system level.
- With the Secured-Core server feature in OpenManage Integration extension, when combined with the operating system capabilities enabled by Microsoft WAC security tool extension, your infrastructure is fully protected against current and future threats.

**NOTE:** Dell Technologies and Microsoft recommend enabling Secured-Core server for Azure Stack HCI and Windows Server HCI 2022 that includes both the Dell Infrastructure and Microsoft Operating System security features to protect the infrastructure from external threats. The Secured feature is supported on 15th generation of AX node servers with AMD Milan or Intel Ice Lake processors running Azure Stack HCI version 21H2 or Windows Server 2022.

## Cluster monitoring and management

The Dell OpenManage Integration extension's home page is carefully designed to give an overall health status of the cluster in a single view. Its monitoring, inventory, and troubleshooting capabilities include:

- Real-time health status
- Overall cluster status
- Status by node
- Component level health
- iDRAC integrated management
- Hardware and firmware inventory and compliance
- Support for traditional or new dark theme

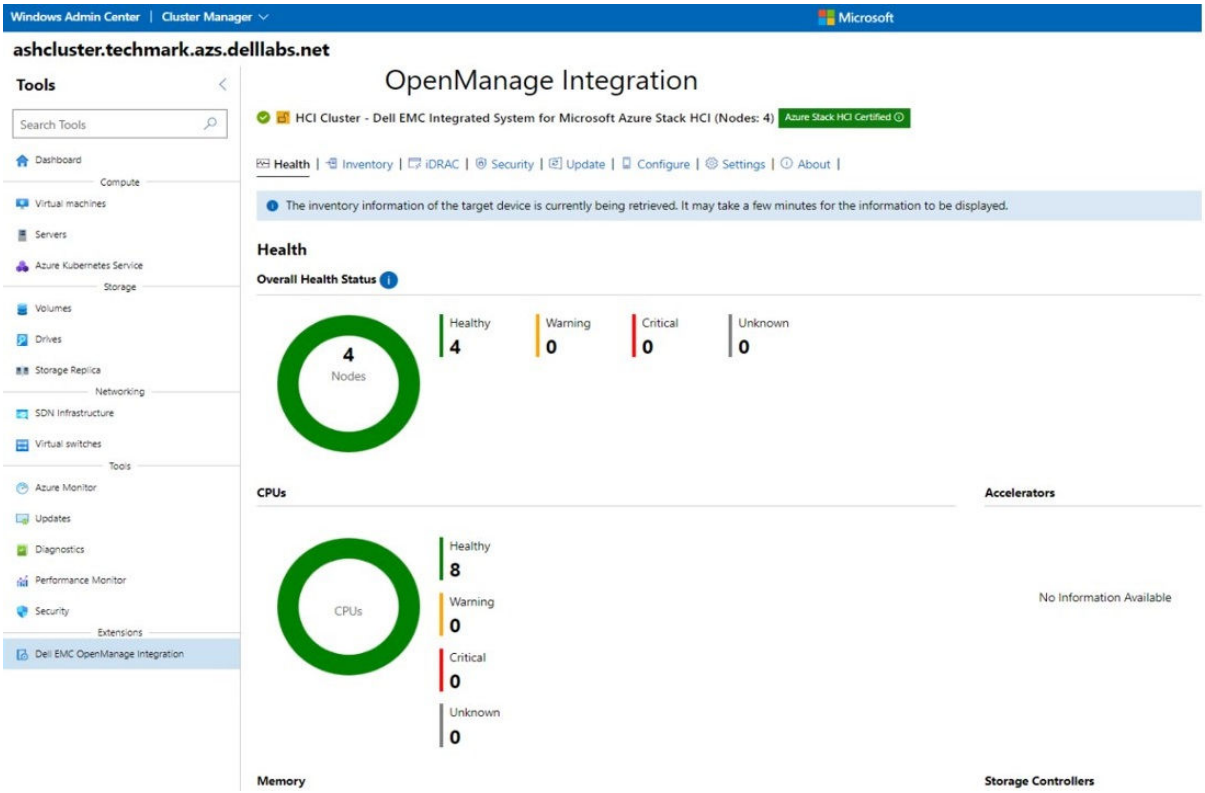


**Figure 57. Dell OpenManage integration with Windows Admin Center – overall health status page**

# Kernel Soft Reboot

Microsoft introduced Kernel Soft Reboot (KSR) with Azure Stack HCI, version 21H2 (or later) which is a premium feature available on all Azure Stack HCI integrated systems that allows faster server reboots by bypassing the full reboot BIOS/ Firmware initialization. Currently, KSR only applies for updates that do not require a cold firmware/BIOS initialization. KSR is disabled by default, and it is important to note that KSR needs to be disabled for our WAC extensions to work.

# Dell HCI Configuration Profile (HCP) policies for Azure

Dell HCI Configuration Profile (HCP) is the specification from Dell which captures the best practices and recommended configuration for Azure Stack HCI and Windows-based HCI solutions from Dell to get better resiliency and performance with the HCI solutions.

This feature functions at the cluster level and is supported for clusters running with Azure Stack HCI OS 21H2 (or later) and pre enabled Windows Server 2022.

IT admins can manage Integrated System environments through two different approaches:

- At-scale through the Azure portal leveraging the Azure Arc portfolio of technologies.
- Locally on-premises using Windows Admin Center.

By using a single Dell HCP policy definition, both cases provide a seamless and consistent management experience.

Dell OpenManage Integration with Windows Admin Center can integrate Dell HCP policy definitions into Azure policy. Then, Azure Stack HCI configuration compliance checks occur against these business rules to protect against malicious threats and inadvertent changes to cluster-level, BIOS, iDRAC, and component settings at-scale across many integrated systems.

Running Check Compliance automatically compares the recommended rules packaged together in the Dell HCP policy definitions with the settings on the running integrated system.

It's critical to re-align non-compliant policies with the best practices validated by Dell Technologies engineering, and our Dell HCP policy remediation WAC integration (unique at the moment) helps fixing any non-compliant errors by clicking **Fix Compliance**.

Some fixes may require manual intervention and others can be corrected in a fully automated manner using the Cluster-aware updating framework.

You can see Dell HCP Policies for Azure in action here.

# Additional resources for Dell OpenManage Integration with Microsoft Windows Admin Center

- To get a firsthand experience of OMIMSWAC, see the interactive demo: Interactive Demo: OpenManage Integrations with Microsoft Windows Admin Center (delltechnologies.com)
- A collection of short videos hosted on YouTube that will help you learn about Microsoft HCI Solutions from Dell Technologies: Videos | Dell Technologies Info Hub
- Dell Knowledge base article: Support for Dell OpenManage Integration with Microsoft Windows Admin Center
- Operations guide: HCI Operations Guide-Managing and Monitoring the Solution Infrastructure Life Cycle
- Manuals and Documents: https://www.dell.com/support/home/en-us/product-support/product/openmanage-integration-microsoft-windows-admin-center/docs

# Azure Arc extends the Azure control plane

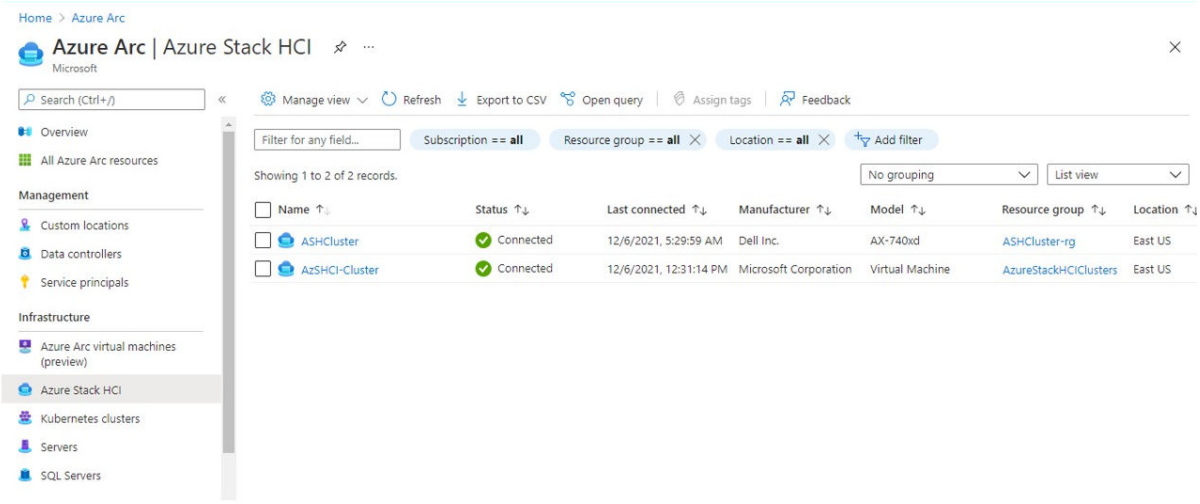


**Figure 58. Azure arc**

Most organizations house traditional infrastructure like rack and stack environments running applications, databases, and server infrastructure that has been growing over the years. These environments still play an important role in the day-to-day business to provide services. There may also be one or more cloud environments like Microsoft Azure and AWS that are becoming more important to provide infrastructure and services that can scale globally across the operational landscape. Edge computing devices are gaining prominence in the connected world, and which must be managed remotely. Today, it can become a complex chore to control and govern these many environments that spread across data centers, multiple clouds, and the edge. More overhead tasks accumulate as IT administrators struggle to learn and operate various management tools.

Azure Arc, one of the breakthrough developments of Microsoft Azure services, allows users to centrally manage multiple environments across on-premises, edge, and multicloud from a single pane of glass. We can think of it like a management plane that can absorb any kind of resource to be managed under one view. Users will be able to extend their Azure management to any infrastructure, adopt cloud practices on premises, and run data and application services anywhere.

Azure Arc enables organizations to:

- Manage entire IT assets with a unified view, by projecting existing non-Azure, on premises, edge devices and other cloud resources into Azure Resource Manager.
- Manage servers, Kubernetes clusters, and databases and make them appear as though they are running in Azure.
- Regardless of the location of resources, use familiar Azure services and management capabilities.
- Along with traditional ITOps, use DevOps practices to support new cloud native patterns.

The following diagram unwraps the management ecosystem supports the Dell Integrated System for Microsoft Azure Stack HCI.

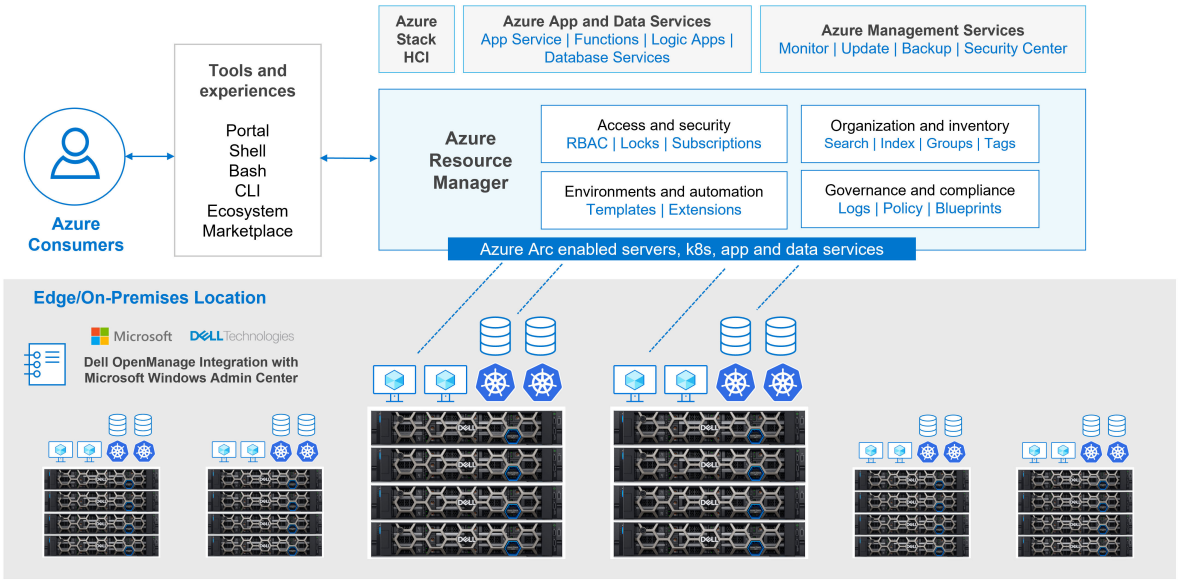


**Figure 59. Managing Azure Stack HCI at scale**

First, we have the Windows Admin Center with Dell OpenManage Integration extension which is the edge/local always available management console. There could be one or more instances of WAC running per physical site in the data center environment intended to manage different fleets of servers. There shall always be IT administrative tasks such as full-stack life cycle management, and more, which can run from WAC for all systems to remain up to date.

As more Azure Stack HCI clusters are deployed, there is a must manage the environment at scale in a more efficient way. Azure Resource Manager, that is, the Azure control plane, is ideally suited to this task. Because Azure Stack HCI is designed to operate as a service in Azure, one can easily integrate it with other Azure management and governance services. In short, the bridge between Azure Resource Manager and infrastructure from other hyper-scaler cloud platforms and on-premises environments is Azure Arc.

# Azure Arc-enabled infrastructure

Azure Arc-enabled infrastructure is simply bringing on-premises and multi-cloud infrastructure to Azure. It consists of three things:

- Servers - Azure Arc-enabled servers
- Databases - Azure Arc-enabled SQL server
- Kubernetes - Azure Arc-enabled Kubernetes

# Azure Arc-enabled servers

Azure Arc-enabled servers enables users to manage Windows and Linux machines running on either physical servers or VMs hosted outside Azure. This management experience is designed to be consistent with how you manage native Azure virtual machines. Each hybrid machine connected to Azure becomes a managed resource in Azure, is assigned a Resource ID, and made a part of the resource group. These hybrid machines can now benefit from standard Azure constructs, such as Azure Policy and applying tags.

# Azure Arc-enabled SQL server

Database instances of SQL server can be now managed from Azure using the Azure Arc-enabled SQL server. It extends Azure services to SQL Server instances hosted outside of Azure, in data centers, on the edge, or in a multi-cloud environment. The running SQL server instances should be registered to Azure Arc to enable Azure services.

# Azure Arc-enabled Kubernetes

With Azure Arc-enabled Kubernetes, Kubernetes clusters located either inside or outside Azure can be attached and configured. When a Kubernetes cluster is connected to Azure Arc, it will:

- Appear in the Azure portal with an Azure Resource Manager ID and a managed identity.
- Be placed in an Azure subscription and resource group.
- Receive tags just like any other Azure resource.

Once Kubernetes clusters managed by Azure Resource Manage are onboarded into Azure Arc, IT administrators can automate creating the same configuration across all Azure Arc-enabled Kubernetes clusters. This helps to standardize the configuration of Kubernetes, at-scale across the enterprise. Azure policy shall then be used to standardize baseline configurations to be applied across the entire inventory of Azure Arc-enabled Kubernetes clusters. Through Azure Role-Based Access Control (RBAC), Kubernetes clusters can be integrated with Azure Active Directory (Azure AD) to provide role-based assignments in Azure to control authorized Connect Kubernetes running outside of Azure for inventory, grouping, and tagging.

Azure Arc-enabled Kubernetes supports the following scenarios for the connected clusters:

- Connect Kubernetes running outside of Azure for inventory, grouping, and tagging.
- Deploy applications and apply configuration using GitOps-based configuration management.
- View and monitor your clusters using Azure Monitor for containers.
- Enforce threat protection using Microsoft Defender for Kubernetes.
- Apply policy definitions using Azure Policy for Kubernetes.
- Use Azure Active Directory for authentication and authorization checks on your cluster.
- Securely access your Kubernetes cluster from anywhere without opening inbound port on firewall using Cluster Connect.
- Deploy Open Service Mesh on top of your cluster for observability and policy enforcement on service-to-service interactions.
- Deploy machine learning workloads using Azure Machine Learning for Kubernetes clusters.
- Create custom locations as target locations for deploying Azure Arc-enabled Data Services (SQL Managed Instances, PostgreSQL Hyperscale.), App Services on Azure Arc (including web, function, and logic apps) and Event Grid on Kubernetes.

For more detailed information on Azure Arc-enabled Kubernetes, see Microsoft Azure Arc-enabled Kubernetes documentation.

# Azure Arc-enabled services

Azure Arc-enabled services are merely running Azure services from anywhere. It consists of two things:

- Azure Arc-enabled data services
- Azure Arc-enabled application services

# Azure Arc-enabled data services

Azure Arc makes it possible to run Azure data services on-premises, at the edge, and in public clouds using Kubernetes and the infrastructure of choice. Currently, the following Azure Arc-enabled data services are available:

- SQL Managed Instance
- PostgreSQL Hyperscale (preview)

**NOTE:** Azure SQL offers several different deployment and management options across IaaS and PaaS for the SQL Server engine hosted on Azure. While Azure Arc-enabled SQL Managed Instance provides the cloud Platform-as-a-Service benefits, SQL Server on Azure Arc-enabled servers provide capabilities such as SQL Server in Azure Virtual Machines.

# Azure Arc-enabled application service

Azure Arc-enabled application service lets App Service, Azure Functions, and Azure Logic Apps run on an Azure Arc-enabled Kubernetes cluster. The Kubernetes cluster can be on-premises or hosted in a third-party cloud. This approach lets app developers take advantage of the features of App Service and lets their IT administrators maintain corporate compliance by hosting the App Service apps on internal infrastructure. It also lets other IT operators safeguard their prior investments in other cloud providers by running App Service on existing Kubernetes clusters.

For more detailed information about Azure Arc, see the Microsoft Azure Arc documentation.

# Dell OpenManage Integration for System Center (OMIMSSC)

Dell provides a family of integrations for full-stack management in Microsoft environments. When installed and configured, OpenManage Integrations for Microsoft System Center enables native management of Dell systems and solutions. These integrations provide visibility and control of hardware infrastructure, operating systems, and virtual machines.

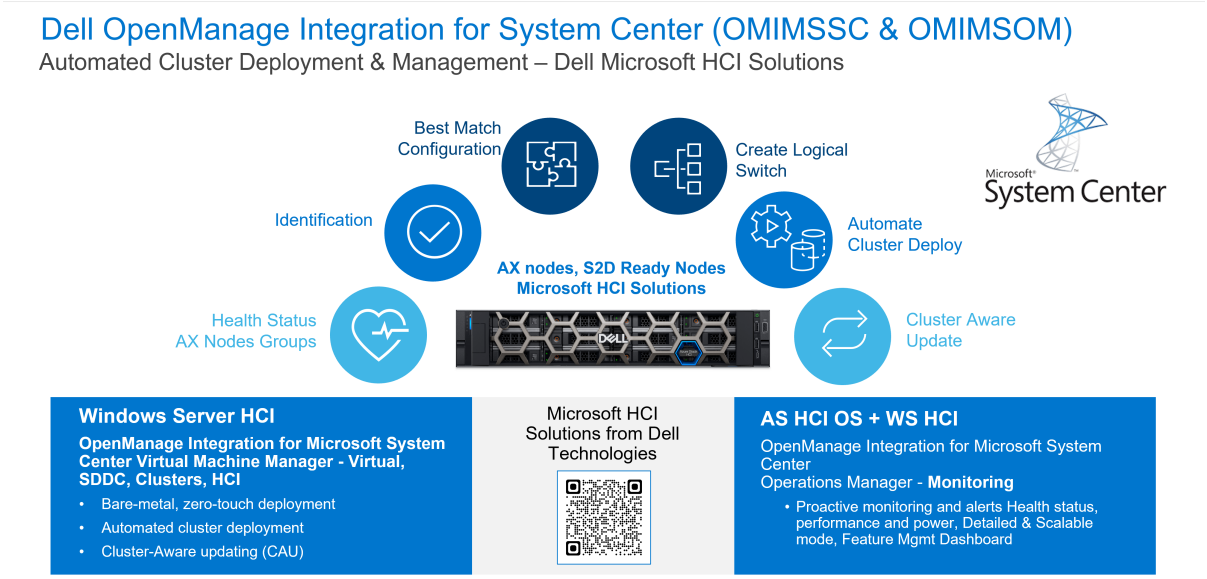


**Figure 60. Automated cluster deployment and management – Dell Microsoft HCI solutions**

OpenManage Integrations for System Center enable users to deploy, configure, update, and monitor within Operations Manager (OM) and Virtual Machine Manager (VMM). The key benefits of OMIMSSC are:

- Optimizes management of Dell hardware in environments where customers have invested in Microsoft operating systems and System Center solutions.
- Provides comprehensive hardware monitoring and management of Dell Servers, AX nodes, Storage, and Ready Nodes through a customer's existing Microsoft System Center console.
- Dell AX nodes, S2D Ready Nodes Lifecycle Management simplified automation of Microsoft HCI solutions cluster on bare-metal Ready Nodes, Infrastructure Cluster-aware updates including driver updates.
- Unifies management of Dell rack and modular servers in multi-vendor operating system and hypervisor environments while protecting customers' existing investments in Microsoft System Center solutions.
- Automates basic IT management tasks resulting in reduced risk of operational errors, increased productivity, and preservation of business continuity with agent-free management using iDRAC.
- Simplifies implementation as a "single, virtual software appliance" delivering end-to-end server management with a consistent look and feel, including enhanced features for converged infrastructure administration.

For more information, see Dell OpenManage Integrations for Microsoft System Center page.

# Automation simplifies life cycle management of Microsoft HCI solutions

Dell Technologies labs conducted a thorough test on updating the component firmware and Windows Server 2019 (EOL June 30, 2023) operating system drivers on a four-node cluster with AX nodes. Two different approaches for updates were chosen:

- A fully automated, cluster-aware approach enabled by Dell OpenManage Integration with Windows Admin Center
- A manual approach enabled by the iDRAC

The results revealed the great benefits that automation can provide for the Life Cycle Management (LCM) of AX nodes. The following are some ways in which an organization can benefit from automation provided by the Dell OpenManage Integration with Microsoft Windows Admin Center [7].

Based directly on testing a four-node cluster in the lab, the automated approach:

- Required 82 percent fewer steps
- Required 90 percent less focused attention from an IT operator
- Reduced by 40 percent the total time required and the resulting maintenance window
- Minimized risk caused by IT operator data entry mistakes or installation option guesswork

For more information about the Life Cycle Management approach, see the white paper Dell HCI Solutions for Microsoft Windows Server: Life Cycle Management Approach Comparison.

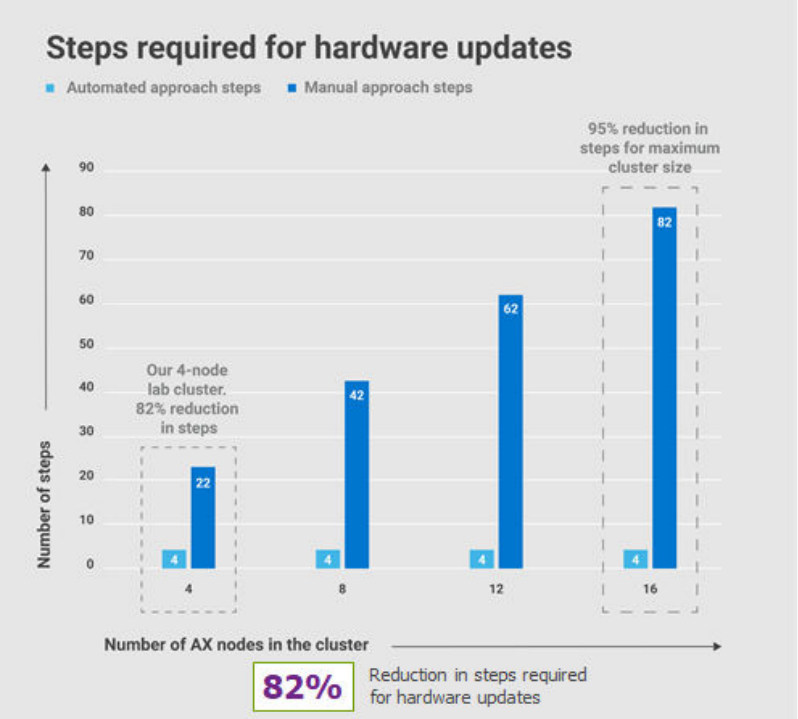


7 Based on Dell internal testing https://infohub.delltechnologies.com/t/dell-emc-solutions-for-microsoft-azure-stack-hci-life-cycle-management-approach-comparison-1/.

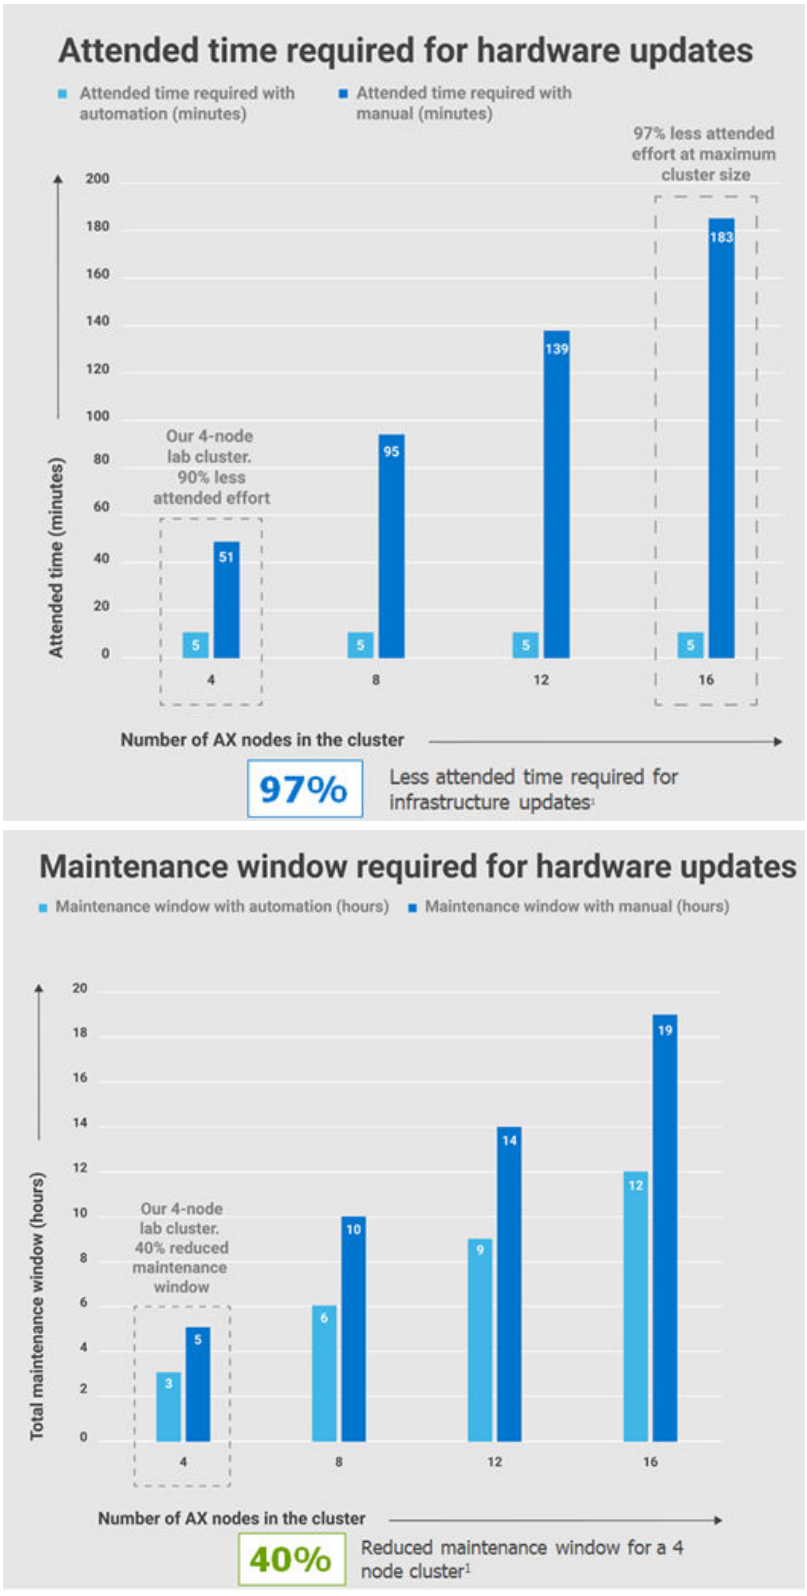


**Figure 61. Dell OpenManage Integration with Windows Admin Center saves time and reduces risk**

The Sizing and Configuration Best Practices chapter describes several Dell tools for planning, designing, and deploying your data centers, using a methodical approach. It also showcases best practices and example configurations to help you choose the right solution, based on the needs of your organization.

10

# Sizing and Configuration Best Practices

**Topics:**



## Overview

KEY TAKEAWAY

As we get through all the aspects of the Dell Integrated System for Microsoft Azure Stack HCI, we shall take a holistic approach to size, design , procure and deploy the solution in the datacenter. Dell Technologies has made it easy to analyze and estimate the hardware equipment required for implementing the solution. This is done by using Dell tools like Live Optics , Sizer and Fabric Design Center (FDC). The first step is to run the Live Optics tools to collect the information about the environment followed by determining the type of system required .The next step is to size the solution based on the information gathered which is done by the Sizer tool. With that in place , the FDC shall help simplify the network design by running either of two Microsoft HCI solutions.

Dell Technologies provides a proven approach and methodology that incorporates modern standards and practices for customers to select and size an optimized cluster configuration which is best suited for their business needs. Dell Integrated System for Microsoft Azure Stack HCI offers a broad range of validated AX nodes and network topologies to meet the requirements of nearly any use case and workload. Example scenarios and recommendations are included the following tables.

**Table 3. Use cases and recommendations**

| Use case | Recommendations |
|---|---|

**Table 3. Use cases and recommendations (continued)**

| | |
|---|---|
| Remote Branch Office (ROBO) and Edge workloads | We may recommend running a two-node cluster with lower core count processors, modest amount of RAM, all flash storage, and a switchless storage network topology.<br>**NOTE:** Best suited if there is no further cluster expansion. |
| Business critical applications with significant storage capacity and expansion requirements. | We may suggest a two-tier, hybrid storage configuration, and scalable networking topology. |
| 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. |

We have intelligently designed the solution portfolio of AX nodes which offers a wide range of factory ready configurations based on the storage and compute requirements. The below table depicts the component combinations included in our AX node models designed for specific business use cases.

| | CPU-dense nodes in small footprint | | | Balance of CPU and storage resources for demanding workloads | | |
|---|---|---|---|---|---|---|
| | AX-640<br>End of Life | AX-650<br>EOL CE Countries | AX-6515 | AX-740xd<br>End of Life | AX-750 | AX-7525 |
| Platform | Intel 14G density optimized node for applications needing high-performance storage and compute balance | Intel 15G density optimized node for applications needing high-performance storage and compute balance | AMD based 15G single socket node optimized for Value | Intel 14G capacity and performance optimized node for applications needing compute and storage balance | Intel 15G capacity and performance optimized node for applications needing compute and storage balance | AMD based 15G two-socket node highly scalable for demanding workloads |
| Targeted Workloads and Use Cases | • Virtualization: dense, powerful compute node<br>• Database with very high-performance storage<br>• Service Providers: application tier | • Virtualization: dense, powerful compute node<br>• Database with very high-performance storage<br>• Service Providers: application tier | • Virtualization<br>• Enterprise LOB<br>• Databases<br>• Retail<br>• ROBO | • Software Defined Storage<br>• Big Data, Unstructured data, Analytics<br>• Service providers: data tier | • Software Defined Storage<br>• Big Data, Unstructured data, Analytics<br>• Service providers: data tier | • Dense virtualization<br>• Database with very high-performance storage<br>• Data analytics<br>• Compute intensive workloads requiring high core count |
| CPU Cores | 16 – 56 | 16 – 80 | 8 – 64 | 16 – 56 | 16 – 80 | 16 - 128 |
| Memory | 96GB – 1.5TB | 128GB – 4TB | 64GB – 1TB | 96GB – 1.5TB | 128GB – 4TB | 128GB – 2TB |
| Max storage | 92TB | 154TB | 60TB | 192TB | 368TB | 368TB |
| GPUs | - | A2 | - | - | A30, A16, A2, A40, T4, A10 | A30, A16, A2, A40, T4, A10 |

**Figure 62. AX node solution portfolio**

# Dell Integrated System for Microsoft Azure Stack HCI - procurement and deployment

## Live optics

Dell Technologies follows a consultative and systematic approach to help IT professionals modernize their infrastructure. The first stage in setting up the new Azure Stack HCI is to start by sizing the new system. To do that we must first collect and analyze the information from the currently running workloads. Dell Live Optics is a free online tool that can gain insights into your environment by collecting configuration and performance data from Microsoft Windows, VMware vCenter, Linux/Unix based servers, and storage. 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.

**NOTE:** Live optics collector captures only configuration related data of the customer's environment and does not include Personal Identifiable Information (PII) data.

To begin with:

- Customers must engage with their Dell Technologies account team to have an account created on the Live Optics portal.
- The next is to 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 various rich reports available in the portal, such as the examples depicted below.


**Figure 63. Demo environment report from live optics portal dashboard**


**Figure 64. Demo performance report from live optics portal dashboard**

# Azure Stack HCI sizer tool

The account team and customers use the data collected by Live Optics to influence selections in our Azure Stack HCI Sizer Tool. This tool integrates all the Azure Stack HCI design constraints and best practices from extensive engineering validation efforts. It creates confidence 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 and to select the one that meets current demand and future growth projections.

https://www.dell.com/sizer/ashci/

**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.

# Dell Fabric Design Center

The Dell Fabric Design Center (FDC) is a cloud-based application that automates the planning, design, and deployment of network fabrics that power Dell compute, storage, and hyperconverged infrastructure solutions. The FDC is ideal for turnkey solutions and automation based on validated deployment guides.

FDC allows design customization and flexibility to go beyond validated deployment guides. For additional information, go to the Dell Fabric Design Center.

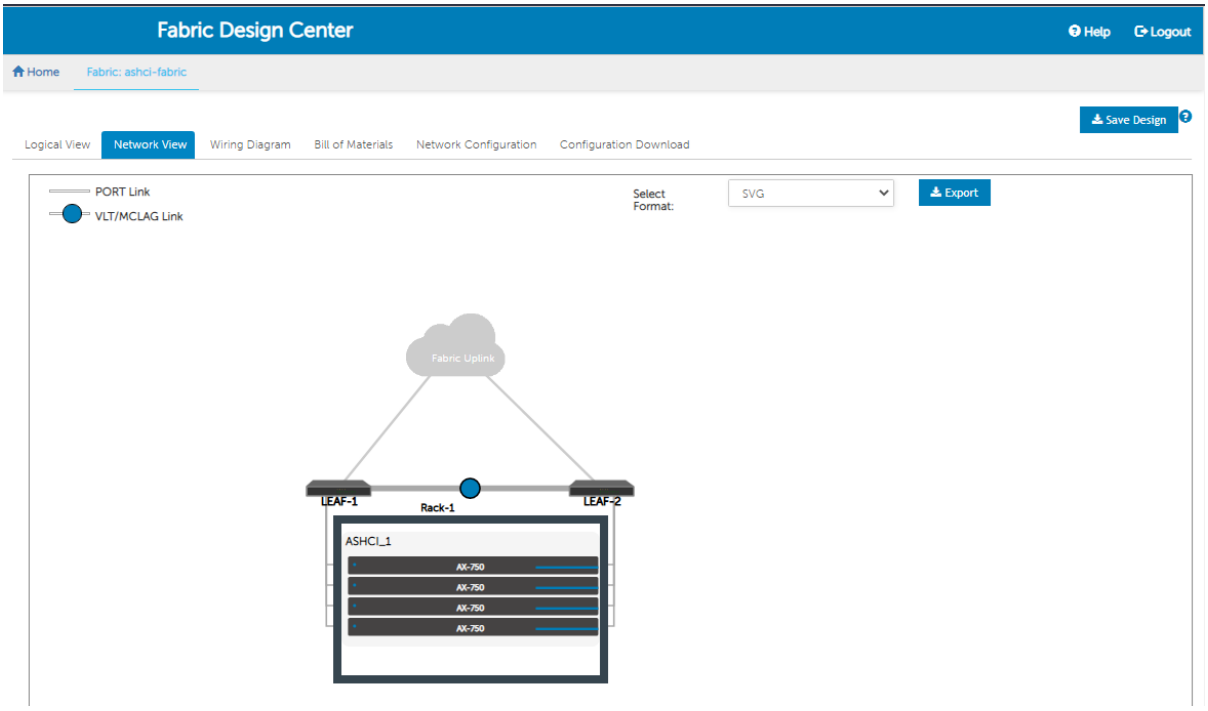


**Figure 65. Fabric design center**

# Best practices and example configurations

Summarizing a few best practices outlined below, enables you to quickly find the right hardware configurations for Azure Stack HCI that we have validated for various needs.

**NOTE:** All the configurations depicted for the following Azure Stack HCI technical use cases are meant for a thorough discussion and refinement of the customer's requirements. The configurations presented here are meant for showcasing sample setups solving a specific business case. From offering expert advice to solving complex problems, we have got you covered. Contact Dell Sales team for any questions.

# Example of ROBO

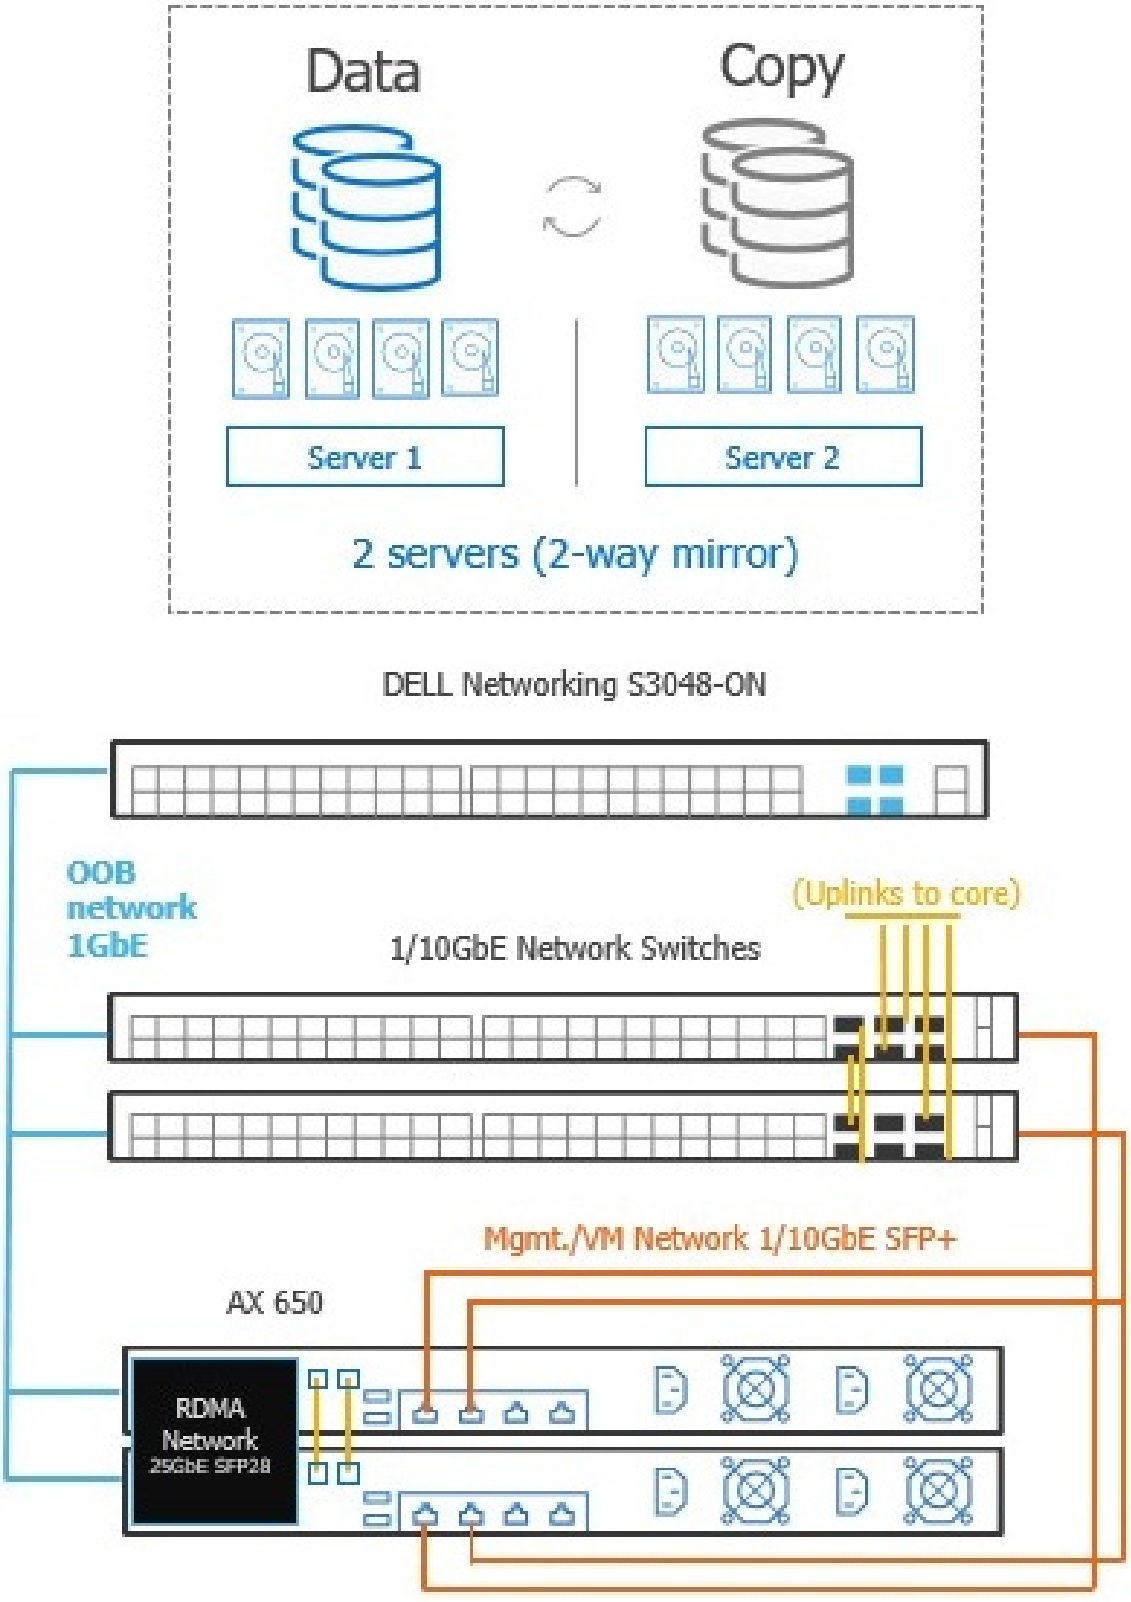


Figure 66. Topology diagram of ROBO

Azure Stack HCI offers excellent performance with affordable price for small deployments with its small footprint but powerful, it can be an appropriate, cost effective solution with limited space and workload requirements. Cluster expansion is not supported therefore this solution is best targeted for specific use cases like remote office branch office (ROBO), retail stores, field sites, and edge nodes.

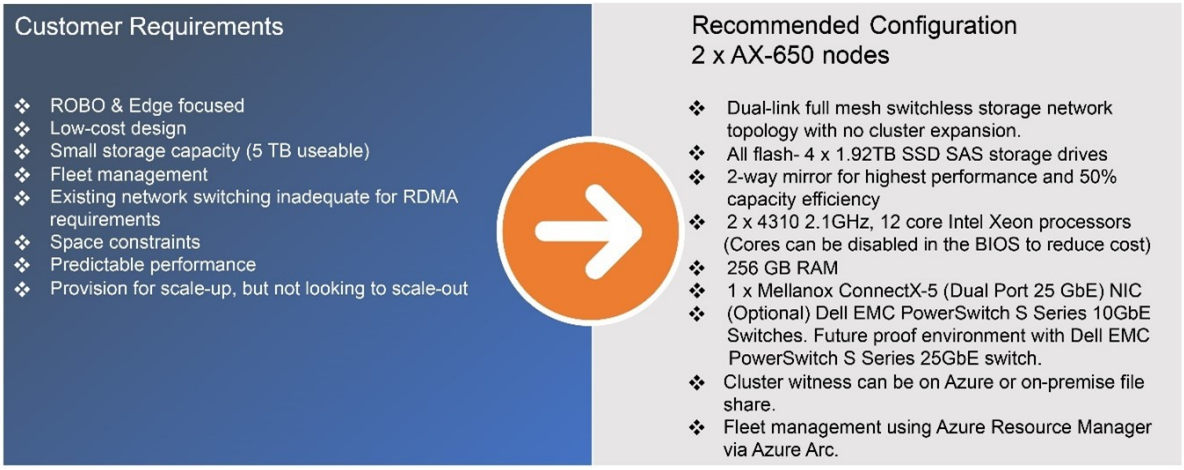


**Figure 67. Recommended configurations for ROBO**

For more information about lab setup and testing results for a ROBO use case, see Value Optimized AX-6515 for ROBO Use Cases.

# Example of Virtual Desktop Infrastructure (VDI)

Azure Virtual Desktop (AVD), formerly Windows Virtual Desktop, is a desktop and app virtualization service that runs in public Azure. It provides IT with comprehensive controls and streamlined management experience, without any infrastructure build-up. It also provides users with a rich experience, running the latest versions of Windows and Microsoft 365 applications.

**NOTE:** AVD is currently in preview and is expected to become generally available in late 2022 or early 2023.

As more organizations must meet specific network and data compliance requirements, Microsoft recently extended the AVD functionality to allow virtual desktops to run on-premises on Azure Stack HCI but be managed from the centralized management plane in Azure. Azure Virtual Desktop is completely hosted and managed by Azure cloud infrastructure which means IT organizations are not required to set up and operate their own complex VDI infrastructure. To comply with data sovereignty and locality requirements, IT administrators are only required to set up and manage the Azure Stack HCI clusters, and the VMs hosted on them, in their data centers.

Besides meeting data locality requirements by keeping user data and apps on-premises, AVD for Azure Stack HCI also improves performance for users in areas with poor connectivity to the public cloud. In those cases, users access their desktops closer to their location with a low latency direct access method instead of a round trip through the cloud. This improves access to on-premises legacy applications which allows IT to provide a streamlined, cloud-consistent experience through the same rich comprehensive management plane in Azure. Best of all, organizations can reduce costs and improve user experience by providing multi-session access concurrently to VMs that are running the newest versions of Windows.

Dell Technologies engineering teams have been working closely with Microsoft to deliver a well-designed and validated solution portfolio that offers optimum performance for hosting VDI workloads. For more information, refer to the VDI Design Guide-Azure Virtual Desktop on Dell Integrated System for Microsoft Azure Stack HCI and its accompanying Implementation Guide.

# AVD Architecture

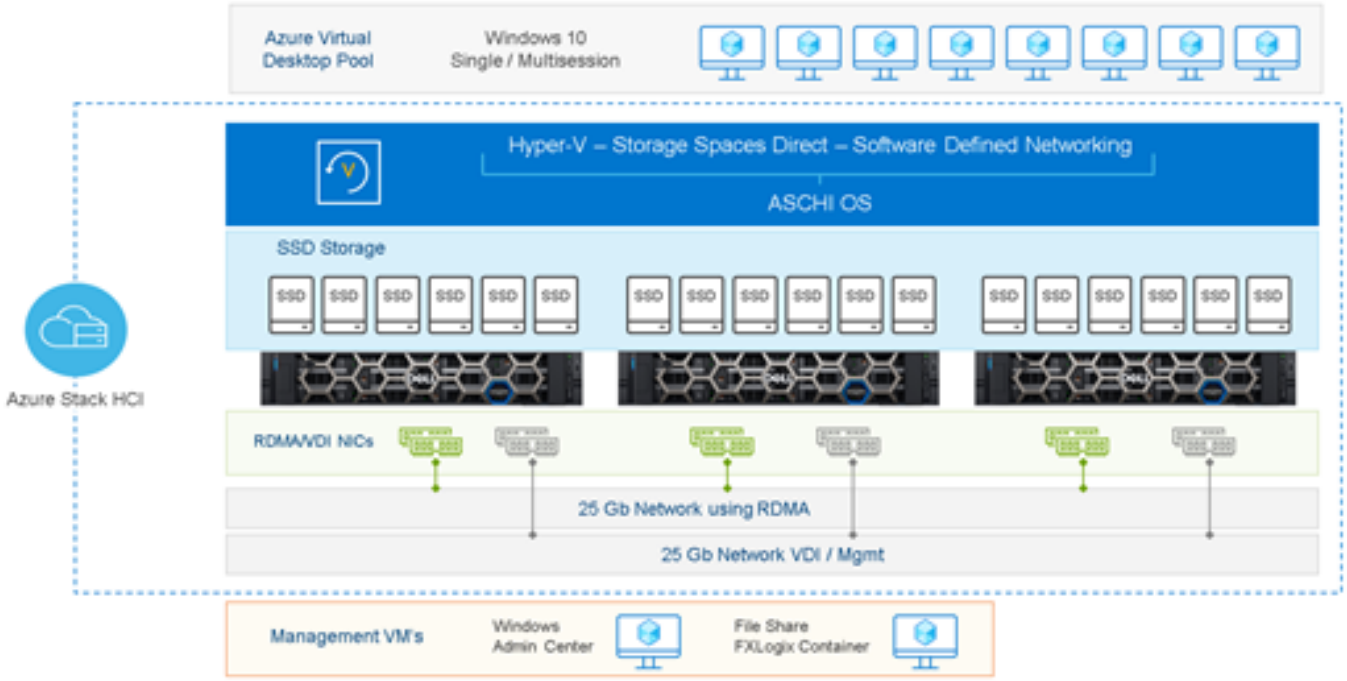

# Example of high-performance Microsoft SQL server

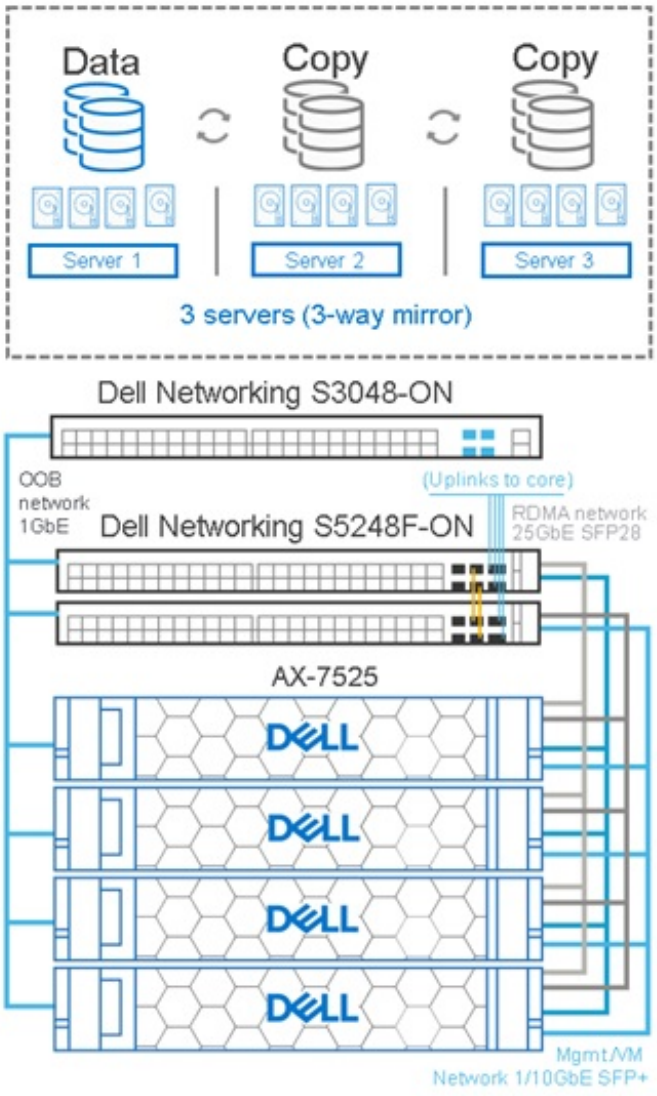


**Figure 68. Topology diagram for hosting SQL server workloads**

Excellent performance is necessary and a crucial requirement for SQL server deployments because of its high reliance in the business world. Azure Stack HCI offers a great platform for implementing a high performing, scalable and manageable SQL server deployment. It also offers flexibility to host Windows VMs and Linux VMs, giving customers a choice to run Windows or Linux version of SQL Server depending on their needs.

Integration with Azure adds the capability of business continuity and disaster recovery with Azure Site Recovery.

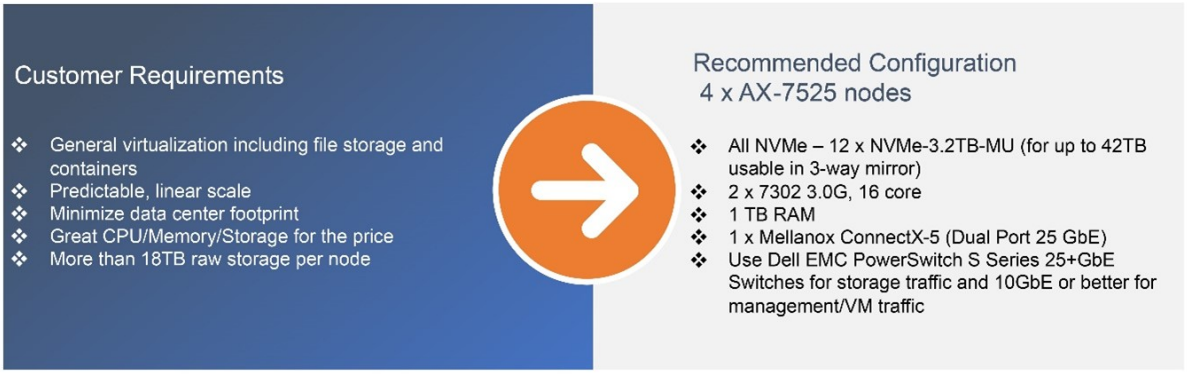


**Figure 69. Topology diagram for hosting SQL server workloads**

For more information about evaluation of Dell Integrated System for Microsoft Azure Stack HCI for database workloads, see Modernize or Perish: Dell Integrated System for Microsoft Azure Stack HCI Throw a Lifeline to Database Workloads.

For more information about a building a Hybrid Database-as-a-Service Platform, see Building a Hybrid Database-as-a-Service Platform with Azure Stack HCI.

# Example of general virtualization

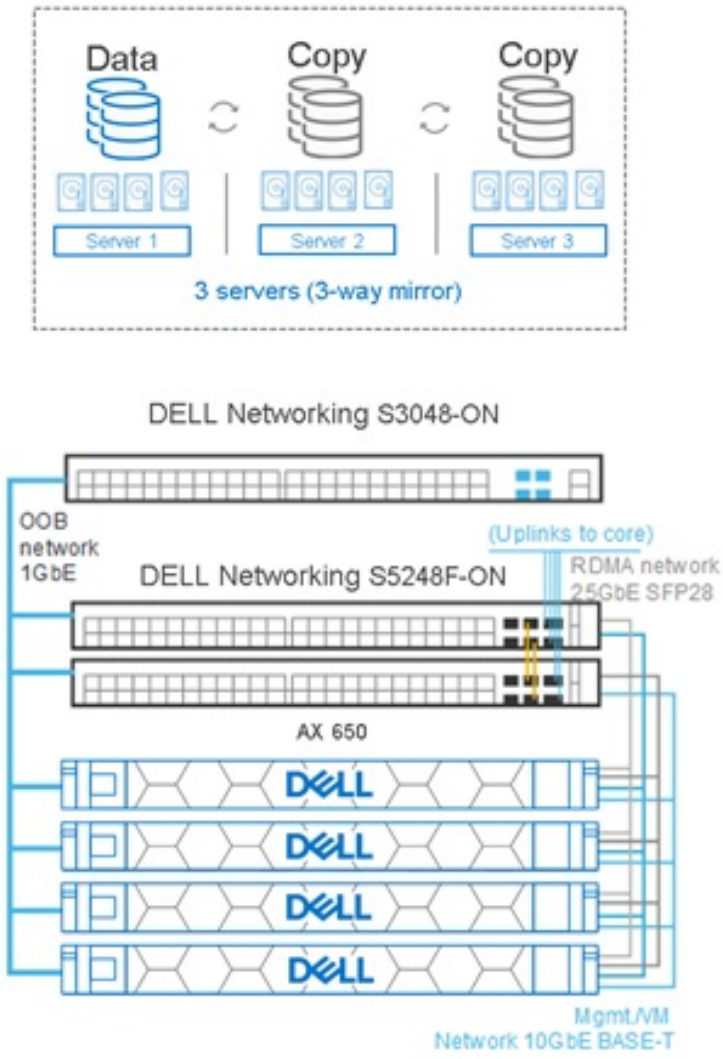


**Figure 70. Topology diagram for trusted enterprise virtualization**

Azure Stack HCI offers a robust environment for serving applications and workloads hosted in VMs with high availability. It provides a trusted enterprise virtualization platform by providing a highly secure infrastructure for workloads through virtualization-based security (VBS). VBS uses the Hyper-V hypervisor to create and isolate a secure region of memory, virtual secure mode (VSM), from the normal operating system.

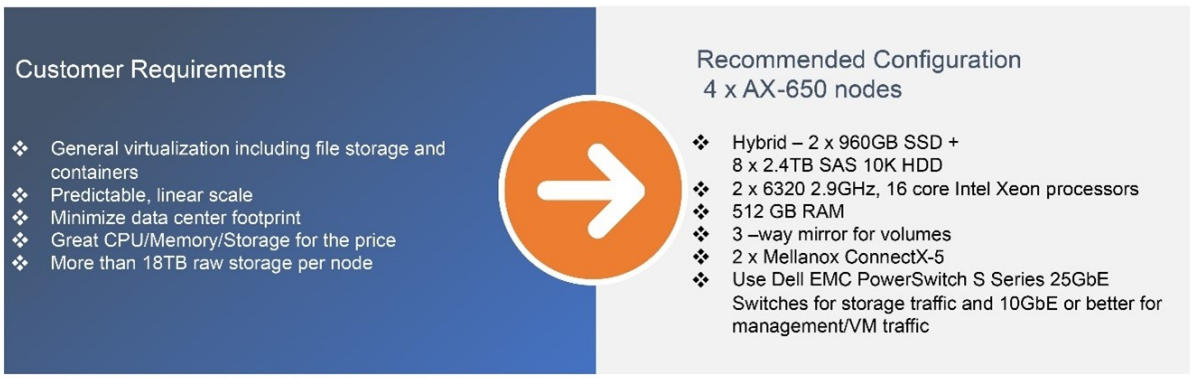


**Figure 71. Recommended configurations for general virtualization**

To see our full test environment details and results and to learn more about Dell Integrated System for Microsoft Azure Stack HCI, download the white paper, Crash Through Workload Performance Boundaries with Azure Stack HCI.

# Cluster sets

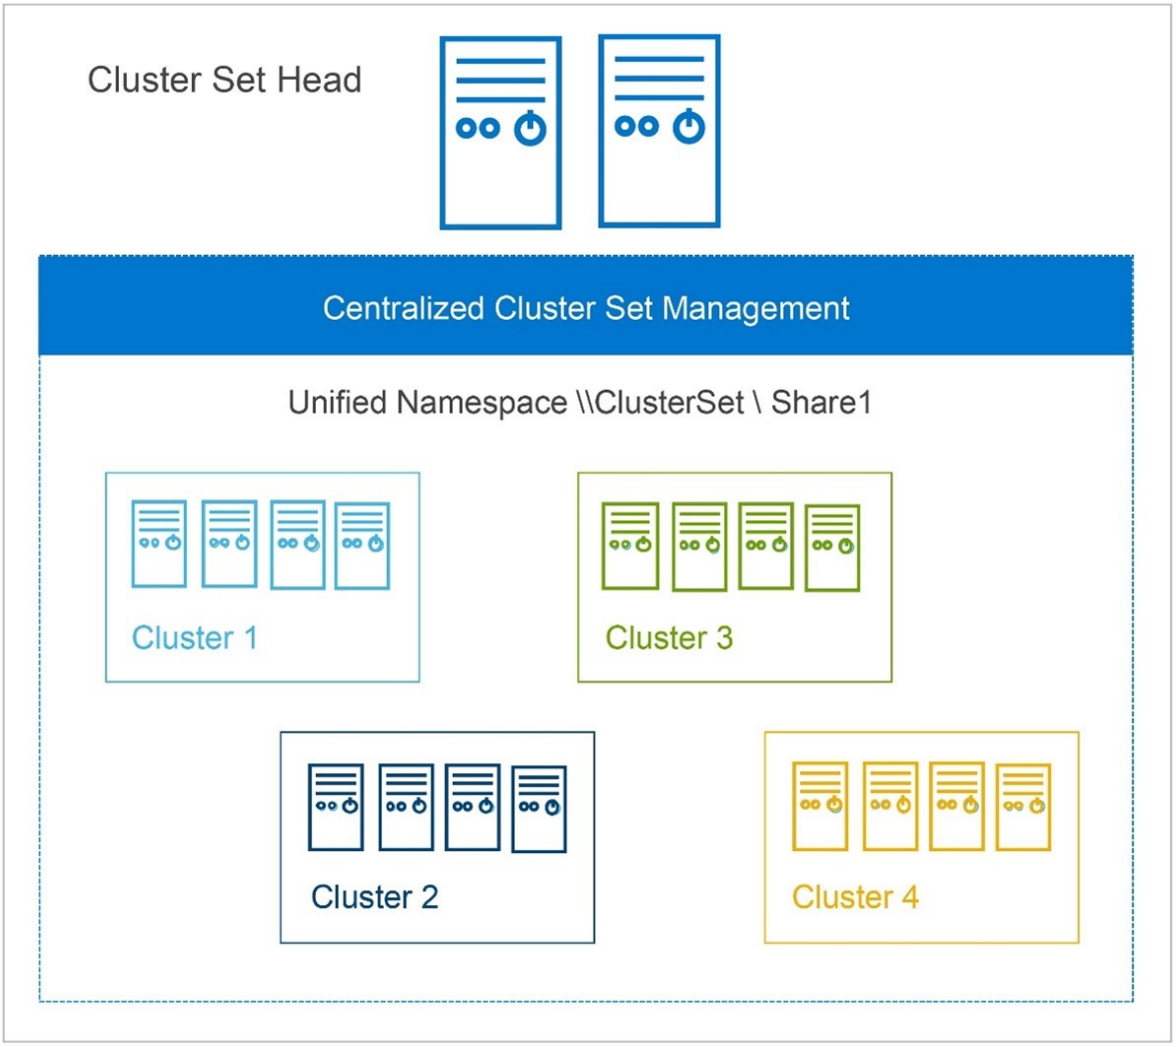


**Figure 72. Block diagram of clusters sets**

Cluster sets are a cloud scale-out technology that provides the benefit of combining multiple clusters without sacrificing resiliency. It encapsulates a cluster within a cluster-set by loosely coupling a group of multiple clusters. The great thing about cluster sets is that a virtual machine (VM) can seamlessly live migrate from one cluster to a host in a different cluster and continue to access its storage.

As mentioned in earlier chapters, we do not recommend expanding a 2-node cluster or expanding a cluster built on switchless networking topology. However, with cluster sets we can get around that limitation. Cluster sets can also be valuable for addressing the homogeneity requirements of HCI clusters. For example, if an Azure Stack HCI cluster is deployed on 14G hardware and we now want to use 15G hardware for expansion, we can do that with the help of cluster sets.

**NOTE:** From a performance perspective, it is important to note that if you want to take advantage of another cluster in a cluster set, it is best to perform not only a Live Migration of VMs but also a Storage Migration so that VM compute and storage reside on the same cluster.

Benefits of cluster sets:

- Combine multiple smaller clusters into single namespace without sacrificing resiliency.
- Manage Failover Cluster life cycle without impacting running VMs.

- Benefit from Azure-like Fault Domains and Availability Sets across clusters.
- Scalability beyond 16 nodes.
- Key for multi-generational clusters in the enterprise, but Cluster Sets must be all Windows Server HCI or Azure Stack HCI. These platforms cannot be mixed in a cluster set.

For more information about Cluster sets, see Microsoft documentation.

The Dell Support and Services chapter describes the best-in-class Dell Support and Services options available to customers.

# 11

# Dell Support and Services

**Topics:**



## Overview

At Dell Technologies, we have invested over 30+ years building best-in-class deployment practices and tools, backed by elite professionals with broad and deep knowledge. Our established global scale drives consistent deployments to help you drive greater business results, around the clock and around the globe. Dell Technologies services will support your ongoing development and innovation in your cloud. We are a one-stop-shop for installation, configuration, and comprehensive source of support.

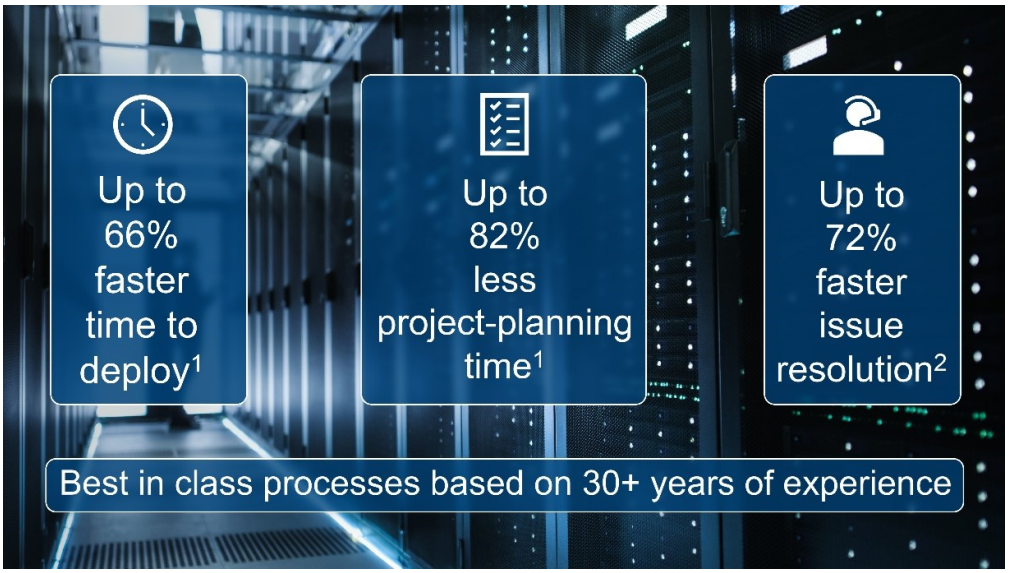


**Figure 73. Dellsupport and deployment services**

Dell Technologies makes implementation simple, flexible and worry free. Certified deployment engineers ensure accuracy and speed, reduce risk and downtime and free IT staff to work on those higher value priorities. The flexible options of deployment services are designed to fit the needs and budgets of an IT department [8], [9]

Our single source of cluster level support covers the hardware, the operating system, hypervisor, and Storage Spaces Direct software, whether the license is purchased from Dell or Microsoft. Dell`s enterprise-class support offers timely and reliable issue resolution delivered from our 24x7 support centers spread across the globe.

8 Based on a November 2017 Principled Technologies Test Report commissioned by Dell comparing in-house deployment vs. Dell ProDeploy for Enterprise deployment service for Dell PowerEdge R730 servers, Dell Storage SC9000 and SC420 and networking components.

9 Based on June 2018 Principled Technologies Report commissioned by Dell, “Save time and IT effort resolving server hardware issues with ProSupport Plus and SupportAssist”

Through Microsoft Cloud Solutions Provider (CSP) program, Dell Technologies can be your single-source of procurement and support for the entire infrastructure stack, including Azure Stack HCI software subscription, reducing time-to-value and simplifying on-going support activities.

# Solution order and deployment path

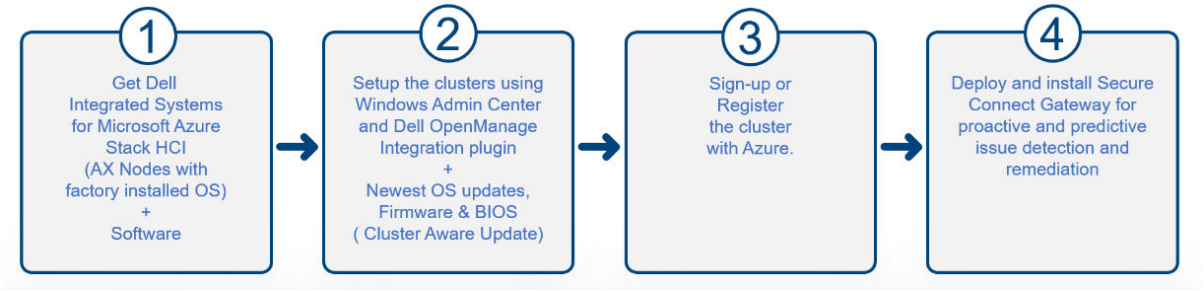


**Figure 74. Azure Stack HCI deployment**

# Factory installation of operating system

Microsoft Azure Stack HCI operating system is factory installed on all the AX nodes and does not include upfront license cost. The factory install also enables the following Windows features:

- Hyper-V
- Failover Clustering
- FS Data Deduplication

**NOTE:** Deduplication is supported but based on your requirement check whether it is necessary or not.

- Datacenter Bridging
- Bit locker

Windows Server 2019 (EOL June 30, 2023) and Windows Server 2022 data centers continue to be available on systems (VMs) with Azure Stack HCI experience as a Secondary operating system. Enough licenses should be included to cover all cores configured in the node.

Azure Stack HCI licensing is billed monthly like other Azure services and only enabled cores are billed, allowing customers to disable cores in the BIOS to reduce costs. This is a great technique for early cost savings with the ability to enable more compute capability as the node utilization grows.

# Do It Yourself (DIY) installation of operating system

Customers can choose to deploy themselves using our standard Deployment Guide or deploy using automated cluster creation in Windows Admin Center. We have a deployment guide which provides an overview of Microsoft HCI Solutions from Dell Technologies, guidance on how to integrate solution components, and instructions for preparing and deploying the solution infrastructure. The guide can be accessed from the below mentioned link.

HCI Deployment Guide—Microsoft HCI Solutions from Dell Technologies

# ProDeploy

For faster, better deployment with less IT effort, we have the Dell ProDeploy deployment services lined up for our customers. It has never been easier maximizing new technology on day one. Dell's ProDeploy offers an expert installation and configuration services at your premises to get the Azure Stack HCI operating system up and running quickly. The ProDeploy suite is designed to complement your business model whether you have an experienced IT staff or no staff at all. We understand the demands on your side and provide the right assistance to keep up with your evolving technology needs. Here are 2 choices for getting started with streamlined deployments.

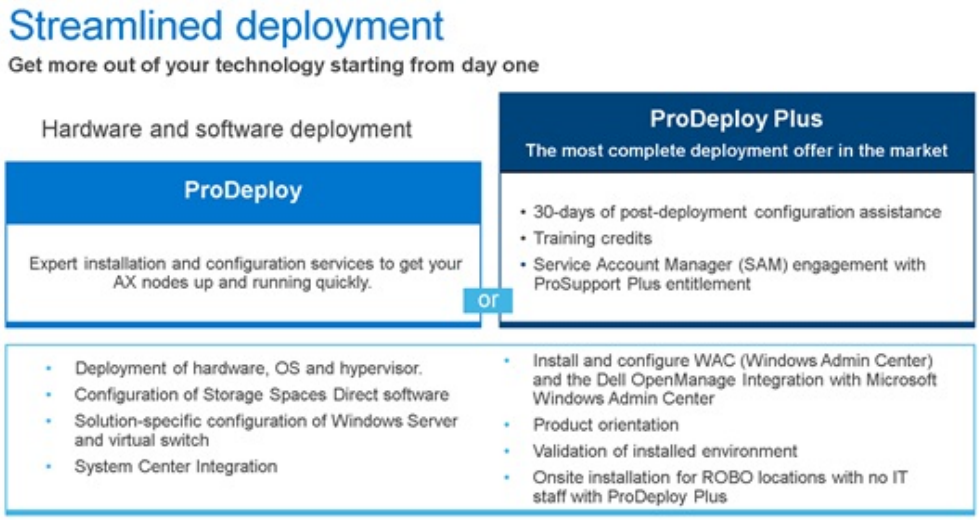


**Figure 75. Dell ProDeploy services**

Both ProDeploy and ProDeploy Plus include installation of Storage Spaces Direct software, and the hypervisor and virtual switches, ultimately creating usable storage in a Storage Spaces Direct cluster. We will also install and configure Windows Admin Center (WAC) and the Dell OpenManage Integration for Microsoft Windows Admin Center (OMIMSWAC). If the customer has an existing System Center Virtual Machine Manager (SCVMM) or System Center Operations Manager (SCOM) environment, they can elect to have their AX node clusters integrated into these environments for the purposes of monitoring and management.

**NOTE:** ProDeploy and ProDeploy Plus do not include installation and configuration of the OpenManage Integration for Microsoft System Center (OMIMSSC). If this is required, submit a custom services request to Dell Sales team.

# ProSupport

Critical workloads and applications running on Dell AX nodes require constant availability for business to run efficiently. The systems supporting them need more than just a break/fix support. Dell ProSupport Plus provides an automated, proactive, and predictive measure to get ahead of problems before they happen. Crafted in response to today's changing data centers and complex environments ProSupport Plus combines expertise, technology, and analytics to improve performance and stability with automated proactive and predictive recommendations.

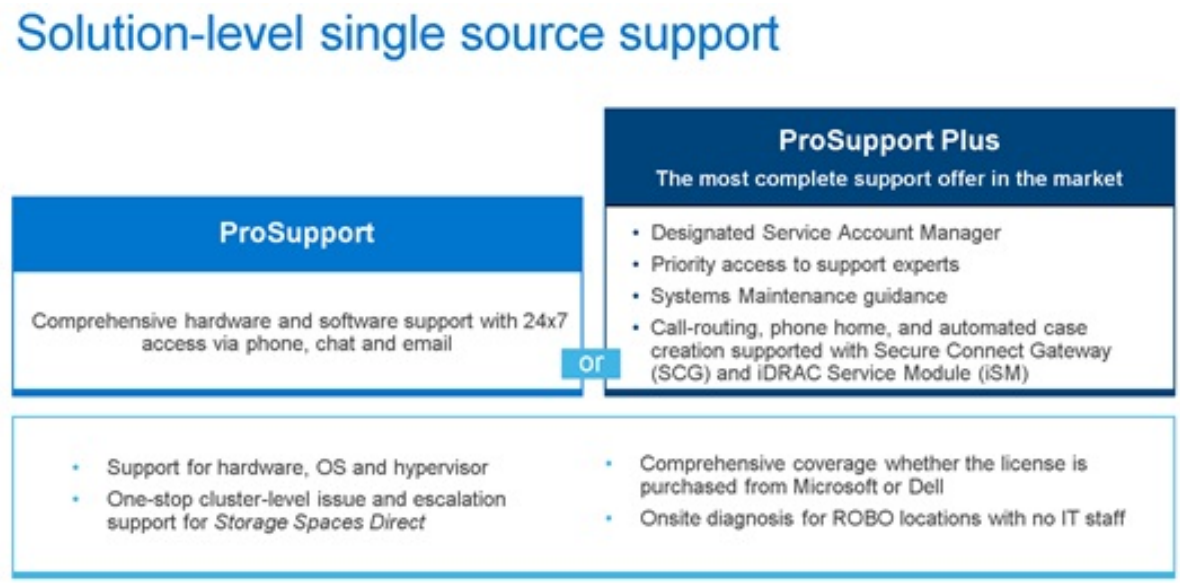


**Figure 76. Dell ProSupport services**

## Secure connect gateway

Dell Technologies provides secure connectivity technology that takes the guesswork out of issue prevention, providing you more time to focus on the projects that matter most. The virtual appliance and application editions provide a secure two-way connection between your environment and Dell Technologies Services ideal for monitoring Dell devices across your data center, including data storage, servers, networking, CI/HCI and data protection, all in one place.

Features of secure connect gateway:

- Proactive, automated issue detection, case creation and notification
- Predictive analysis failure detection for server hard drives and backplanes
- Accelerated issue resolution with remote access and secure, two-way communication between Dell Technologies and the customer's environment
- Enables analytics-based recommendations for support and services

## iDRAC service module

iDRAC Service Module (iSM) which is deployed by Dell personnel, facilitates in collecting additional logs and diagnostics information per AX node for better supportability and serviceability.

# 12

# Conclusion

**Topics:**

- Summary
- Unrivaled performance
- Manageability
- Scalability
- Simplicity
- Resiliency
- Resource efficiency
- Security

## Summary

Dell Integrated System for Microsoft Azure Stack HCI offers an unmatched combination of performance, scalability, convenience, and simplicity. Its hybrid design provides a crucial link between the existing on-premises infrastructure and cloud services. As a result, for digital transformation initiatives, HCI is one of the top areas of strategic IT investments.

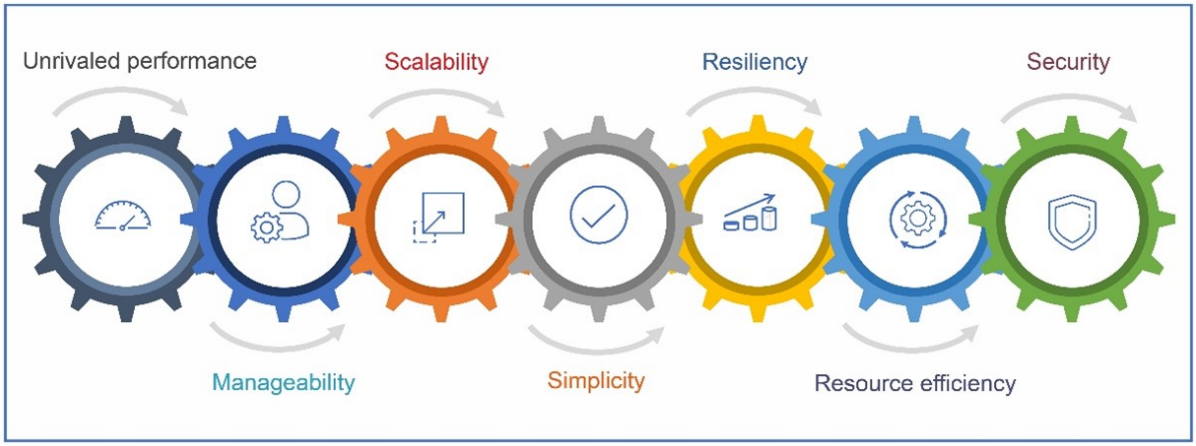


**Figure 77. Dell Integrated System for Microsoft Azure Stack HCI**

## Unrivaled performance

The all-flash drive configuration offers excellent performance with faster rebuilds. Innovative solutions such as RDMA over 100 GbE connections with RoCE frees up CPU resources offering higher throughput and lower latency for compute and storage traffic.

3-way mirror volumes built on all-flash drives offers excellent performance with faster volume rebuilds. Innovative solutions such as RDMA over 100 GbE connections with RoCE frees up CPU resources offering higher throughput and lower latency for compute and storage traffic. iWARP as the RDMA protocol works out-of-box which is easy to setup and troubleshoot with added benefit of no additional network switch configuration required. Though Data Center Bridging (DCB) is not required, we recommend using DCB if NVMe drives are in a converged network topology.

# Manageability

Streamline operational efficiency and flexibility by directly managing your server and cluster environment using the Dell OpenManage Integration with Microsoft Windows Admin Center. Manage at scale with Azure Resource Manager and Azure Arc, and natively integrate with other Azure services.

# Scalability

Easily scale up to 16 nodes in a cluster based on business needs. Two-node and switchless storage cluster scale-out is not supported, but cluster sets can be used.

For scalability, you can start with 3 nodes. For better resiliency and stretch clustering, you can start with 4 nodes.

# Simplicity

Productized and validated AX nodes with factory installed Azure Stack HCI operating systems speeds up deployment time. Create clusters following our deployment guide or use automated cluster creation in Windows Admin Center. Use ProDeploy and ProSupport to create a turnkey solution-level experience.

# Resiliency

Start with the four-node or greater clusters. DR with stretch clustering is available natively with an integrated system eliminates the need for an additional software.

# Resource efficiency

Right size and future proof your IT environments with Dell tools such as Live Optics, HCI Sizing Tool, and Fabric Design Center to optimize price, performance, and capabilities. Use CPU core management in OpenManage Integration to achieve a balance between cost and performance.

# Security

Cyber resilient architecture built into the AX nodes provides a hardened server design to protect, detect, and recover from cyberattacks. The new Secured core feature further safeguards the AX nodes from various threats.

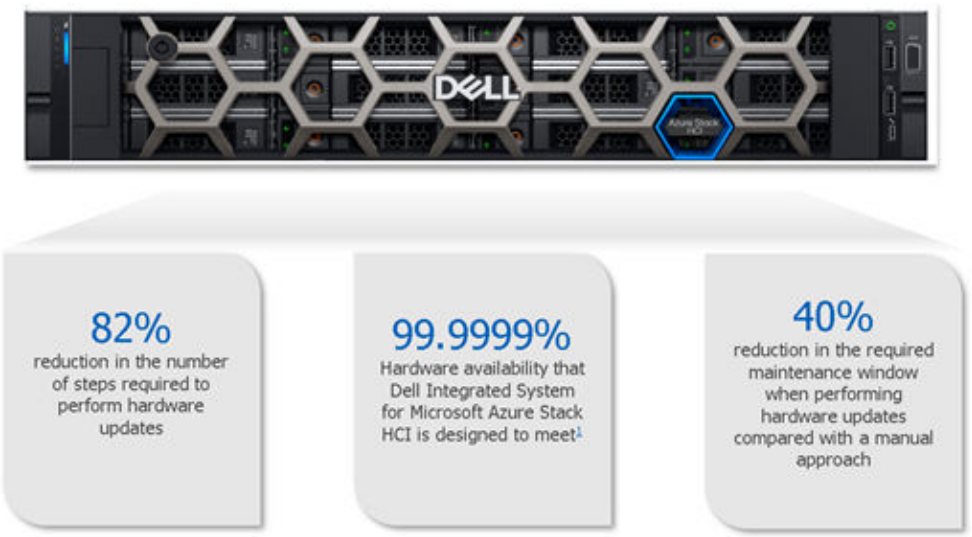


**Figure 78. AX nodes**

Dell Integrated System for Microsoft Azure Stack HCI is an all-in-one solution that offers a tightly integrated HCI operating system with a purpose-built hardware platform.

The Dell OpenManage Integration with Windows Admin Center keeps the cluster up to date with end-to-end, full-stack management that is cluster-aware, enabling operating system, BIOS, firmware, and driver updates a single reboot with no impact to the workloads running on the VMs.

The AX nodes with the Azure Stack HCI operating system deliver a well-designed architecture with optimized and tested configurations to suit various workloads. Dell has also simplified the customer experience starting from designing the data centers, initial configuration, ordering, setup, and deployment with the ProDeploy service offerings. Our best-in-class ProSupport is a single source of support for hardware, software, and operating system issues backed by trained technical support personnel.

We believe that the success made possible by the Dell-Microsoft partnership for the hybrid cloud adoption should be accessible to every business and organization. Choosing Dell Technologies is a great opportunity for companies to adopt a hybrid cloud strategy while still maintaining their traditional workloads and infrastructure architecture.

# 13

# Additional Resources

**Topics:**

- Additional resources

# Additional resources

## Dell resource list

- Product page: Dell Integrated System for Microsoft Azure Stack HCI
- Dell Technologies Info Hub: Microsoft HCI Solutions from Dell Technologies
- Interactive Journey: Azure Stack HCI Interactive Journey

## Microsoft resource list

- Azure Stack HCI operating system releases: What's new in Azure Stack HCI
- About Azure Stack HCI: Azure Stack HCI documentation