

Dell APEX Cloud Platform for Microsoft Azure Tech Book

An infrastructure platform built to optimize Microsoft Azure hybrid cloud deployments

March 2024

H19968

Tech Book

Abstract

This document is a conceptual and architectural review of the Dell APEX Cloud Platform for Microsoft Azure. The tech book first describes how important is innovating in multicloud environments and then focuses on the Dell APEX Cloud Platform for Microsoft Azure technology solution

Copyright

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

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

Copyright © 2024 Dell Inc. or its subsidiaries. Published in the USA March 2024 H19968.H19969

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

Contents

| | |
|--|----|
| Executive summary..... | 4 |
| Innovation in multicloud environments..... | 5 |
| Multicloud by design..... | 6 |
| Dell APEX Cloud Platforms | 7 |
| Dell APEX Cloud Platform for Microsoft Azure introduction..... | 19 |
| Dell APEX Cloud Platform for Microsoft Azure Hardware platforms | 22 |
| Dell APEX Cloud Platform for Microsoft Azure networking topologies | 26 |
| Dell APEX Cloud Platform for Microsoft Azure Operations | 29 |
| Dell APEX Cloud Platform for Microsoft Azure Security | 41 |
| Dell Services..... | 45 |
| The Solution sizing and Configuration | 49 |
| Dell APEX Cloud Platform for Microsoft Azure Use Cases | 53 |
| References..... | 60 |

Executive summary

Overview

This tech book is a conceptual and architectural review of the Dell APEX Cloud Platform for Microsoft Azure. First, the tech book describes how important it is for customers to innovate in multicloud environment and then focuses on the Dell APEX Cloud Platform for Microsoft Azure technology solution.

Audience

This document is intended for Dell Technologies field personnel, partners, and customers involved in designing, acquiring, managing, or operating a Dell APEX Cloud Platform for Microsoft Azure.

Revisions

| Date | Part number/ revision | Description |
|------------|--------------------------|-----------------|
| March 2024 | H19969 | Initial release |

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Inigo Olcoz

Note: For links to other documentation for this topic, see the [Cloud Platforms Info Hub](#).

Innovation in multicloud environments

Application sprawl

There is a new reality today, because more and more companies are embracing modern development tools. While old monolithic applications still exist, and will remain with us for a long time, a high percentage of new applications fall under the cloud native paradigm. New applications proliferate in ways never seen before. According to the IDC (IDC Market Note (Dec 15,2021 update to IDC FutureScape 2020)) 750M apps will be built using modern development tools by 2025. This application demand places an unprecedented burden on IT organizations. This massive number of new applications demands an even higher number of dev/test environments that need to exist to accommodate the applications development lifecycle.

For this specific use case, public cloud providers become an excellent choice. And it is very common to find not one, but various cloud providers servicing this need, which describes a typical multicloud environment.

This reality is well described in ESG's Study- Infrastructure for Cloud-Native Development Enablement, showing how many organizations that are using modern application tooling are also using public clouds as their development and test arena. Some of the key benefits companies are expecting to find in these cloud providers is rapid access to that modern tooling and the ability to have access to extra IT resources for a short period of time.



Figure 1. New applications impact IT operations

Barriers to innovation

The multicloud paradigm is often imposed by the reality that IT departments face in their daily activities. It is not a planned adoption, but more a multicloud by default. That “by default” adoption brings in several challenges which customer must face and solve.

The first one is **Management complexity**. Having vital data for the business on-premises, and in different cloud platforms depicts a complex scenario as migrating applications from one provider to another is far from simple. According to Forrester's Research (Forrester:

Multicloud Answers the Enterprise Call for Scalable, Secure Infrastructure, January 2022), data management across multiple clouds is the leading obstacle of innovation.

Working on a multicloud environment, where each cloud provider has its own management interfaces, creates at least a **limited multicloud visibility** challenge. Managing multiple cloud environments requires IT departments to have a holistic view of where data lives, and what is happening with that data. This can lead to compliance and security issues.

Cloud providers offer a detailed description of their individual cost. However, the challenge is how difficult it is to predict how those initial costs will evolve over time. When new services are added, existing ones get additional features (such as backup or data replication), or the environment deployed grows (in terms of resource utilization), cost starts to differ from what was previously planned. Asset customizations as well as consultancy services from cloud providers also add to the cost deviation. This **cost unpredictability** is yet another challenge of the multicloud approach. IT organizations have faced budget contractions for more than a decade now.

This multicloud reality requires a well-rounded staff of skilled professionals. Multicloud management obliges involves operating a complex on-premises infrastructure and a number of different cloud providers interfaces. This **skills gap** creates an additional and unnecessary burden on the IT organizations' already heavy payload. CIO Insights found that 79% of IT Departments incurred significant unplanned cloud costs (<https://www.cioinsight.com/cloud-virtualization/how-unexpected-costs-create-a-cloud-hangover/>).

Security has become one of the main concerns for companies worldwide. We can learn of important security breaches every week. Compliance with external and internal regulations and policies is also intrinsic to any organization today. These two aspects are not easy to implement and maintain and demand proper internal IT resources. When we incorporate different public cloud providers into the equation, the **lack of consistency** in security models and compliance obligations becomes evident.

All these challenges become real barriers holding many organizations back from achieving their innovation priorities and becoming leaders in their market segments.

Multicloud by design

Is this multicloud by default approach, described in the previous section, the best route to success? Probably not, and that is why we introduced the concept of multicloud by design. This is our vision to help companies unlock the benefits of multicloud without being constrained by siloed ecosystems of proprietary tools and services.

APEX is the way we deliver multicloud by design. This involves:

- Taking our industry leading software to the public cloud – a strategy we refer to as **ground-to-cloud**. We did this first with our data protection software which now protects more than 17 EB of data in the public cloud. Building on that experience, we are enriching the public cloud by bringing our storage software to the public cloud with several new APEX offers that give our customers management consistency and enhanced data mobility in multicloud environments.

This multicloud by design approach with APEX also involves:

- Bringing cloud stacks to data center, colocation, and edge environments – what we refer to as **cloud-to-ground**. To this end we have developed, in collaboration with key partners such as Microsoft, Red Hat, and VMware, a new series of turnkey cloud platforms that integrate popular cloud stacks and Kubernetes orchestration software on Dell platforms so customers can extend and optimize the cloud experience of these stacks across public cloud and on premises environments. This gives customers the agility to run their workloads where their business needs dictate in the most efficient way possible.

Finally, our multicloud by design approach also involves bringing simplified cloud experiences to dedicated IT environments.

- APEX accomplishes this through a range of subscription and as-a-Service offers designed help to reduce IT workload, speed technology deployment, control costs and manage end of life assets sustainably. And now we are expanding APEX to include our client devices and servers, giving our customers more solutions to solve a broader set of business challenges.



Figure 2. Multicloud by design: Get simplified cloud experiences with trusted technology

Dell APEX Cloud Platforms

Overview

One key element in Dell's multicloud by design motion is referred to as cloud-to-ground. It is essentially delivering simplified cloud experiences for customers through collaborative innovation with our partners Microsoft, Red Hat, and VMware.

The outcome of these collaborations is modern cloud platforms and consumption models that enable customers to implement a cloud operating model for their entire IT infrastructure across the company, application, and data space.

These new platforms that help customers break the barriers for a multicloud successful adoption are built on Dell's successful story in the hyperconverged space. Whether we

look at the VMware hyperconverged infrastructure (HCI) or the Microsoft Azure Stack HCI markets we find the common fact that Dell offerings are market leaders. If we analyze in detail the key aspect that is common in both offerings is that they deliver real business outcomes through automated operations and lifecycle management. In the Microsoft space, this is mainly done via the Open Manage Integration with Windows Admin Center (OMIMSWAC), and in the VMware space through VxRail HCI System Software.

These offering had different building blocks on all the domains: Compute, Storage, and Network.

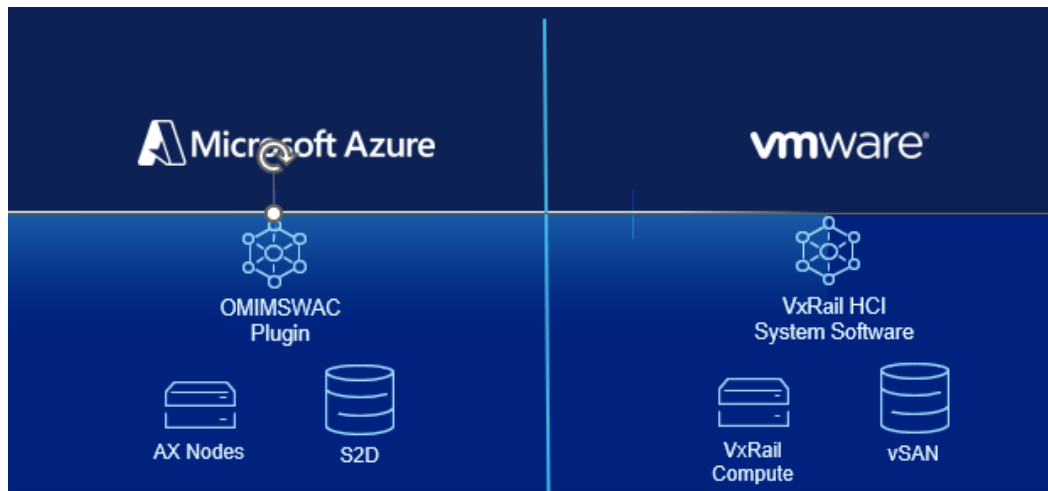


Figure 3. Building blocks for Dell’s Microsoft and VMware HCI offerings: Azure Stack HCI and VxRail

As newly designed solutions, Dell APEX Cloud Platforms have been conceived from a common infrastructure building block approach.

The compute layer is built on Dell MC (Multicloud) nodes. At this stage there are two types of nodes, MC-660 (based on Dell’s PowerEdge R-660) and MC-760 (based on Dell’s PowerEdge R-760).

The common storage layer is built on Dell’s Enterprise Software Defined Storage (SDS). This option is not available at first release for the Microsoft offering but will be ready later. Dell’s Enterprise SDS inherits the technology core to PowerFlex.

For infrastructure Management and Operations, all APEX Cloud Platform leverage the APEX Cloud Platform Foundation Software that plays a key role in:

- Automating initial deployments
- Streamlining on-going operational tasks
- Delivering operational consistency across on-premises and cloud deployments




Figure 4. Dell APEX Cloud Platforms common building blocks

Dell APEX Cloud Platform Foundation Software, a key building block for Dell APEX Cloud Platforms, inherits the technologies and innovation that resulted in the VxRail HCI System Software and its evolution over the years. The Cloud Platform Foundation Software consists of multiple, integrated software elements that extend cloud ecosystem native capabilities to deliver a seamless, automated, operational experience, keeping the infrastructure in continuously validated states to ensure workloads are consistently up and running.

Dell APEX Cloud Platform Foundation Software consists of a collection of microservices running in a Linux-based virtual appliance. This Virtual Machine (VM) is the Cloud Platform Manager VM. It communicates with the underlying infrastructure and integrates automation workflows into the Dell APEX Cloud Platform extension in Microsoft Windows Admin Center. This VM is instantiated on each platform during initial cluster deployment, and it remains running throughout the life of the platform.

Cloud Platform Foundation Software integrates with native partner user interfaces (Windows Admin Center and Azure Portal for Microsoft, OpenShift Web Console for Red Hat, and vCenter for VMware), enabling administrators to have a consistent user experience through familiar tools, managing on-premises resources as they would in the public cloud. This software provides general benefits on automating platforms deployment and operations, as described in previous sections. But it also provides benefits that are specific to each cloud platform.

Although providing a fully automated end-to-end Day 1 deployment and cluster creation experience is a common feature for all APEX Cloud Platforms, the way it is delivered in the Azure space has its unique implications.

At its initial launch, in September 2023, Dell APEX Cloud Platform for Microsoft Azure is not customer installable. Dell offers a white-glove deployment experience via Dell ProDeploy Services to make the process smooth and successful. (This deployment experience will be made available soon for customers to drive.)

The APEX Cloud Platform Foundation Software provides specific benefits for Azure environments:

- Accelerates the creation of real Azure hybrid environments improving platform time to value by fully deploying Azure hybrid cloud in 88% fewer steps than a manual approach.
 - Dell APEX Cloud Platform Foundation Software provides an automated node discovery and election process using IPv6 multicast. If there are restrictions on IPv6 multicast in the environment, MC nodes can be manually added for initial cluster creation.

- The Azure Stack HCI cluster created is registered with your Azure subscription during this automated process. Azure Stack HCI clusters must send usage and billing data to your Azure subscription at least once every 30 days. The cluster nodes are also onboarded to Azure Arc via the installation of the Connected Machine Agent.
- Through the inclusion of Microsoft network ATC in the automated Azure cluster deployment process, host network configuration gets extremely simplified, reducing not only deployment time but also complexity and potential errors. It helps ensuring network configuration consistency across the clusters and eliminate configuration drift.
- Reduces operational expense by allowing right-sizing the Azure MC nodes-based infrastructure to meet current business requirements and scale-out rapidly using a fully automated approach via the automated cluster expansion workflow in the APEX Cloud Platform extension in Windows Admin Center to meet increasing business demands.
- Ensures the stability and security of the Azure platform using full stack lifecycle management - with all updates ready to install within 4 hours of initial release by Microsoft.

A broad set of publicly available RESTful APIs are provided to customers to deliver greater cloud and IT automation extensibility.

Customers always have access to Dell Secure Remote Support (SRS) for all included hardware and software throughout the entire lifecycle of the infrastructure.

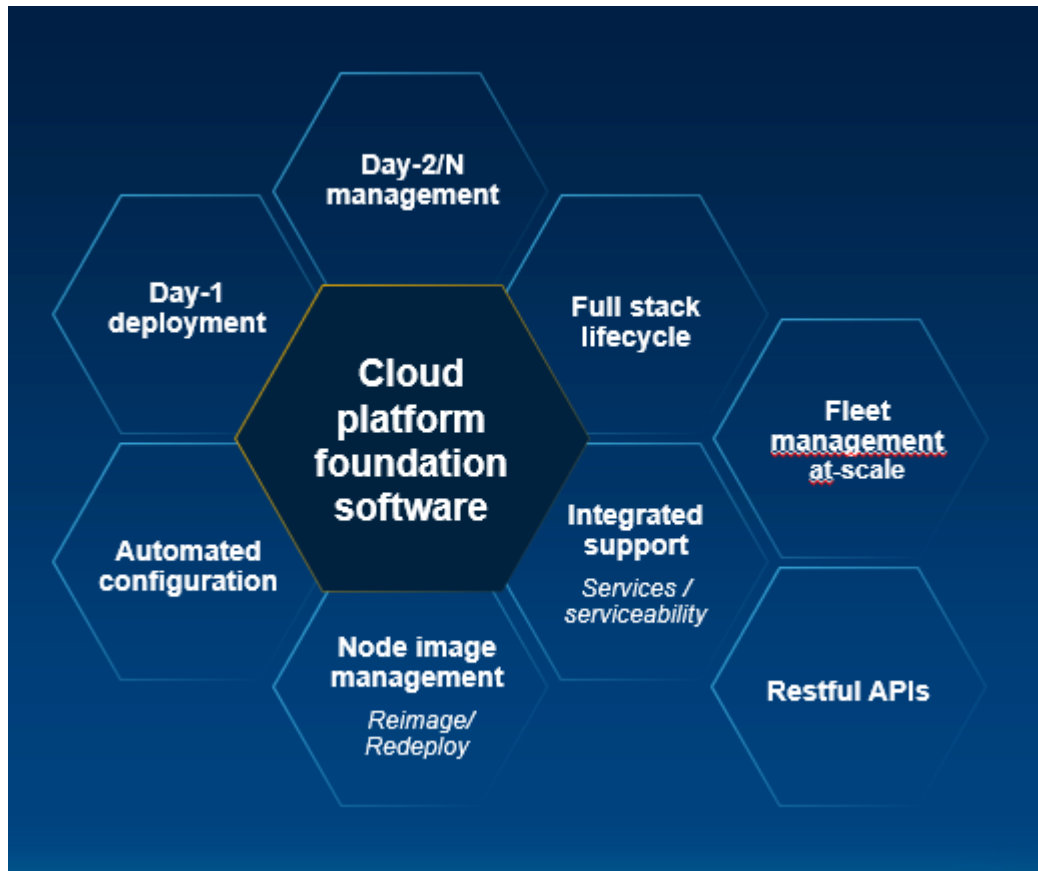


Figure 5. Dell APEX Cloud Platform Foundation Software

Multi-cluster, cloud-based management provides AI driven operations insights to provide detailed health checks, predictive analytics, and further simplifies the upgrade process by pre-staging all required components for individual cluster upgrades making management at scale easier than ever. Multi-cluster management also helps with the ability to rapidly scale up or out by proactively forecasting resource consumption. It learns the existing state of resource capacity, and leverages predictive analytics to forecast when a node might run out of storage, memory, or compute. This insight helps proactively address resource constraints before they happen.

Ecosystem connectors tightly integrate with infrastructure components including Dell Enterprise SDS, physical server components including network interface cards (NICs) and Graphic Processing Units (GPUs), enabling automation and orchestration services across the entire stack.

Automated, intelligent lifecycle management (LCM) functionality automatically updates clusters with pre-validated, pre-tested SW and firmware components, ensuring the infrastructure is in a continuously validated state.

Dell APEX Cloud Platform for Microsoft Azure

Dell APEX Cloud Platform for Microsoft Azure is a on-premises infrastructure platform that is designed to extend and optimize the Azure cloud experience. Dell APEX Cloud Platform for Microsoft Azure is the first offer in the market for Premier Solutions for Microsoft Azure Stack HCI. This platform is the result of extensive engineering

collaboration between Dell and Microsoft, and it is designed to uniquely optimize Azure hybrid cloud operations across data center, edge, and colocation environments.

With over 60+ announced regions, Microsoft’s global reach is second to none, however, with a heritage in the enterprise, just like Dell, Microsoft realizes that meeting customers where they are, and where they need their workloads to run, is incredibly important. Back in 2019, Microsoft announced Azure Arc – a set of technologies that unlock key hybrid scenarios by bringing Azure services and management to any infrastructure. From its single control plane in Azure, organizations can manage, govern, and secure non-Azure resources side by side with native Azure resources, all from the comfort of the Azure portal. Since then, Azure Arc has grown in breadth and depth, providing more comprehensive range of services, along with management, governance, and security capabilities.

Azure Arc is a bridge that extends the Azure platform to help organizations build and run applications and services with the flexibility to run across on-premises datacenters, at the edge, and in multicloud environments and allows developers to build cloud-native applications with consistent development, operations, and security models. Through Azure Arc-enabled infrastructure, organizations can project non-Azure-based resources such as Azure Stack HCI, VMware vSphere, physical and virtual Windows and Linux servers, SQL Servers and Kubernetes clusters, into Azure, for purposes of management, security and governance and more. From there, Azure Arc-enabled services can be pushed out from Azure to run locally on-prem, unlocking a consistent experience to help modernize apps and workloads across your hybrid and multicloud environments.

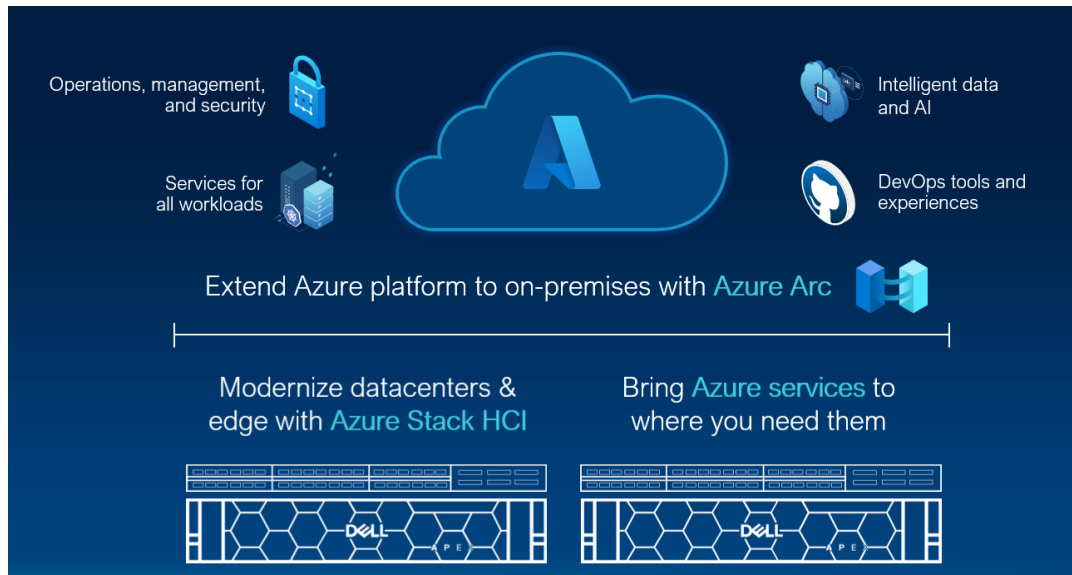


Figure 6. Microsoft innovation with Azure and Azure Arc extensions

What are the specific capabilities that make Dell APEX Cloud Platform for Microsoft Azure unique?

- Dell APEX Cloud Platform Foundation Software takes customers from factory delivered nodes to fully operational Azure hybrid cloud in 88% fewer steps using our automated deployment and cluster creation experience versus taking a manual approach. That means we offer the fastest path to Azure hybrid cloud.

- Our software also provides the fastest path to Azure updates in a turnkey system. Even Microsoft updates appear within our Windows Admin Center extension within 4 hours of being released.
- Our software provides fully automated, end-to-end cluster expansion.
- And in Q1 of 2024, this platform will also be the only offer in the market that supports 3rd party, extended software defined storage for Azure Stack HCI. Dell software defined block storage will deliver cloud-like linear storage scalability independent of Azure Stack HCI compute and Storage Spaces Direct (S2D) storage resources.

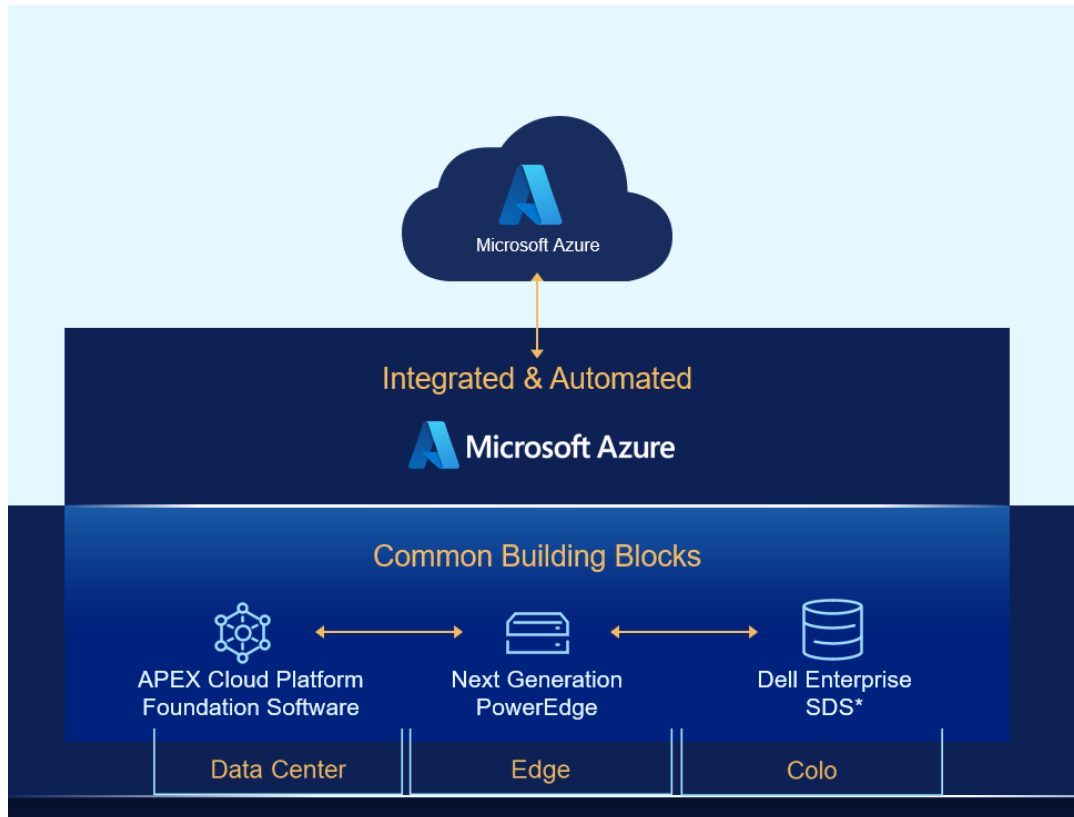


Figure 7. Dell APEX Cloud Platform for Microsoft Azure

We are bringing the simplicity of Azure operations on-prem and to the edge. APEX Cloud Platform for Microsoft Azure empowers organizations to unlock innovation with a consistent and optimized Azure experience across all their IT environments. The platform helps to accelerate productivity by delivering familiar developer experiences across locations and by optimizing delivery of Azure Kubernetes Services and Azure Arc-enabled services.

APEX Cloud Platform for Azure helps mitigate potential security and compliance risks by enforcing consistent Azure management and governance services from cloud to edge. It takes a full-stack approach to security with sophisticated controls and multilayered protection from security threats.

We are running Azure Stack HCI on our intelligently designed MC nodes. You can start with just a single node and scale up to 16 nodes in each platform.

Dell APEX Cloud Platform Foundation Software is one of the foundational stones. It runs in a virtual appliance on every platform and is the brains of the cluster providing management and orchestration at the granular, in-depth cluster-by-cluster level.

The Dell software defined block storage will run on MC storage nodes and will integrate with Azure Stack HCI at the OS level.

IMPORTANT: At initial release, Storage Spaces Direct will be the only storage subsystem we offer. We will release Dell software defined block storage soon after general availability.

As mentioned, Azure Arc is a bridge that extends the Azure platform to help you build cloud-native applications and services with a consistent development, operations, and security model. It also delivers fleet management at-scale with Azure management and governance services.

Dell APEX Cloud Platform for Red Hat OpenShift

For many years, our customers have told us they need to be more flexible, agile, and purposeful about how and where they develop and run their applications and workloads. For most organizations that has meant an evolution to virtualization and more recently to containers to modernize their environments, so they can accelerate development and optimize production workload placement. This has been a rapid evolution, Organizations of all sizes are either currently running containers, actively evaluating containers, or trialing them in non-production environments.

In a multicloud environment, customers will run containers in public clouds and on-prem -- and our goal with the APEX Cloud Platform for Red Hat OpenShift -- is to transform how customers deploy, manage, and run container orchestration on-prem.

There can be multiple challenges to implementing containers and container orchestration platforms, including management inefficiencies, and “marrying” the data plane and the application plane together.

Today, most Kubernetes orchestration platforms run in a virtual environment. This is because it can be extremely complex to deploy Kubernetes on bare metal. However, that can be inefficient from a performance perspective and result in unnecessary expenses - the “virtualization tax”.

However regardless of which deployment method the customers choose, the application plane is generally disparate and uncoordinated from the data plane. Often organizations have to manage these two layers separately, - this could cause application issues if there are storage bottlenecks that go unnoticed and make it difficult to optimize storage services across containers.

Together, our goal with the APEX Cloud Platform for Red Hat OpenShift is to address both of those deployment challenges and help customers transform how they run OpenShift on-premises --- and we do that by delivering a fully integrated stack - including OpenShift, compute, storage resources and automation software.

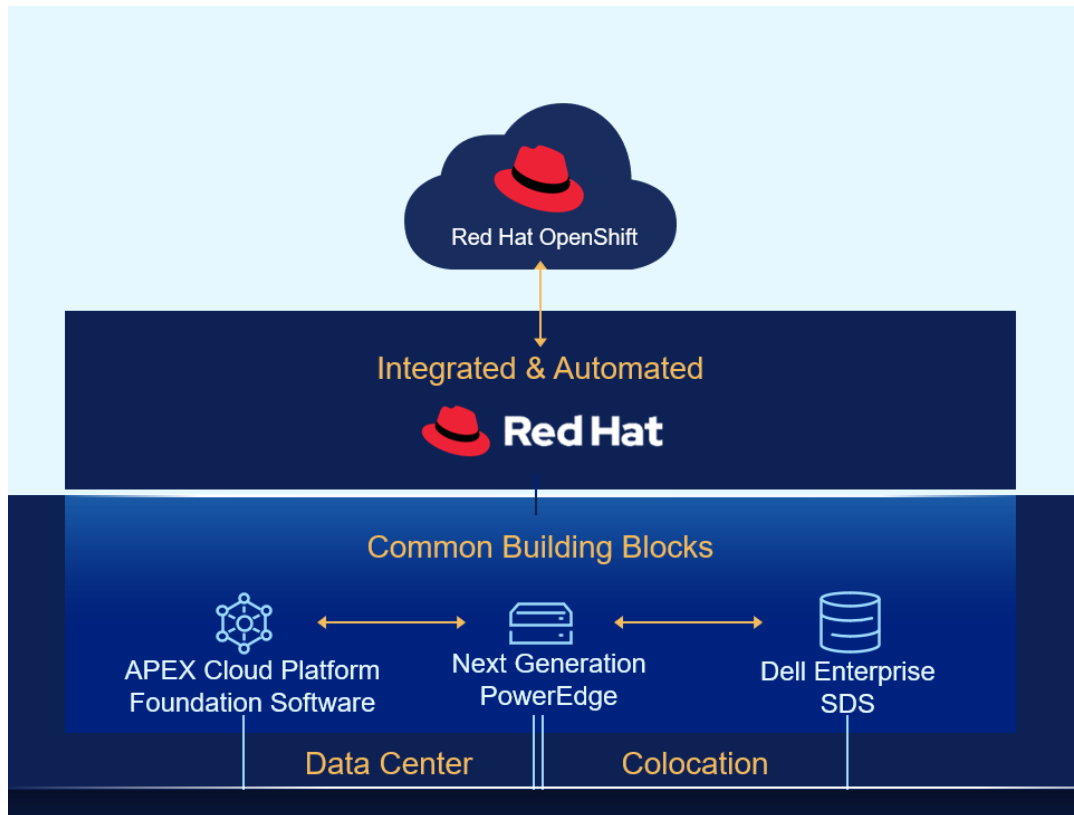


Figure 8. Dell APEX Cloud Platform for Red Hat OpenShift

With this turn-key approach optimized for bare-metal, we help customers overcome complexities and get the best results from their modern applications.

We are uniting the data and application plane with this platform, so customers can operate the whole stack as one unified application delivery platform, simplify how they manage different layers of infrastructure and ensuring optimal application outcomes.

This new, on-premises infrastructure platform has been jointly engineered with Red Hat to provide faster time to value, simplify operations, and ensure a more secure environment, all while reducing licensing costs.

You can think of it as a unified application delivery platform that helps delivering a solution to address our joint customers' top priorities:

- Reducing cost and complexity
- Optimizing workload outcomes
- And enhancing security and governance

Let us take lowering cost and complexity first. Customers can now eliminate the virtualization tax – that means removing the costs associated with licensing and managing a hypervisor in a virtualized environment.

Then there is workload SLAs – there are inherent performance benefits from running your mission-critical applications on bare metal – and that means on-premises.

Finally, when it comes to security and governance, what we are really talking about is reducing the security attack plane. With a smaller attack surface, a bare metal OpenShift deployment mitigates the vulnerability of virtualization.

Customers will benefit from our joint engineering -- from day one -- with deployment. We learned from our own internal testing, depending on a customer's experience and expertise, it can take up to 10 days to deploy and configure OpenShift on bare metal infrastructure. With the APEX Cloud Platform, we have brought that down to just six hours. That means a 90% reduction in deployment time- talk about a faster time to value- customers can start developing and deploying apps almost immediately!

Once they are up and running, the underlying Dell APEX Cloud Platform Foundation Software drastically simplifies the management experience so IT teams can focus more on delivering business value -- through new apps and new features -- and spend less time managing infrastructure.

We all know security is top of mind for customers and the APEX Cloud Platform ensures that all components of the stack – BIOS, firmware, and more – are validated and in compliance when performing upgrades. Through Dell's partnership with Red Hat, we are making critical security patches and upgrades available in real-time.

In terms of deployment, Dell APEX Cloud Platform for Red Hat OpenShift is built as a two-layer architecture that starts with a minimum of eight nodes, four dedicated to compute, and four dedicated to storage. CoreOs is installed on top of the compute cluster, as well as the OpenShift Controller and worker nodes. The rest of the Kubernetes and Management components can be seen in the following figure:

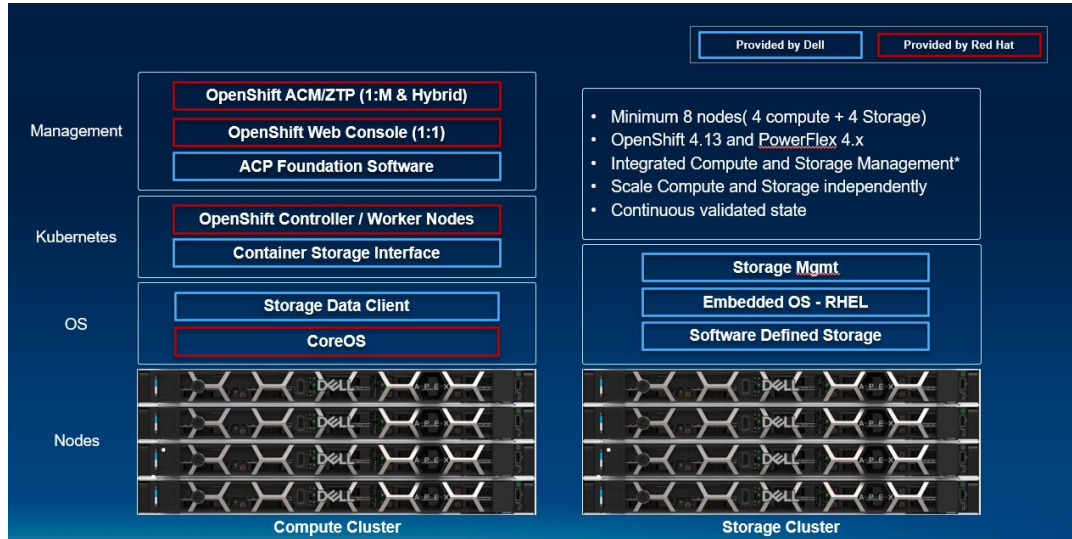


Figure 9. Dell APEX Cloud Platform for Red Hat OpenShift two-layer architecture

The result of these automated Day-1 deployment process is an OpenShift environment, fully operational, including:

- CSI driver
- APEX Cloud Platform Management Software and Web Console
- Configured Operator Hub

- Configured Service Mesh
- Virtualization Operator is not installed by default

The overall picture is shown in the next figure:

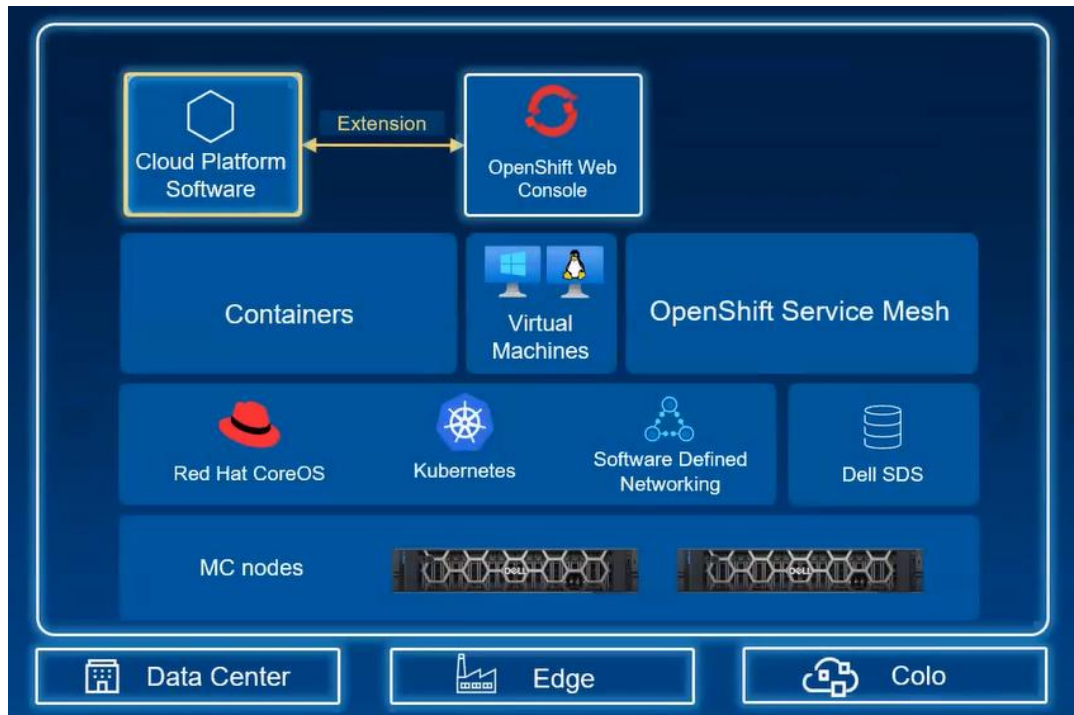


Figure 10. Dell APEX Cloud Platform for Red Hat OpenShift Day1 deployment outcome

Dell APEX Cloud Platform for VMware

Dell Technologies and VMware have a long story partnering on the development of integrated solutions in the HCI and multicloud space. As part of Dell's Multicloud by Design strategy, we are working to release Dell APEX Cloud Platform for VMware, designed with the right balance between simplicity, scale, and flexibility.

Dell APEX Cloud Platform for VMware is a fully integrated turnkey system integrating Dell infrastructure and software and VMware cloud operating stacks delivering consistent multicloud operations by extending vSphere operating models across on-premises and edge environments. Managed through familiar tools like vCenter, combined with integrated automation, and leveraging a common Software Defined Storage (SDS) layer powered by PowerFlex, customers can drive consistency throughout multicloud deployments.

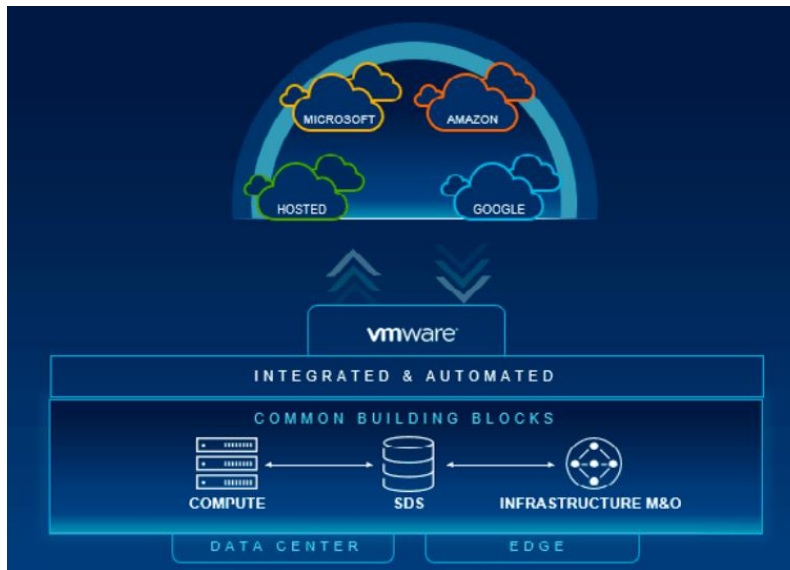


Figure 11. Dell APEX Cloud Platform for VMware

Extend and enhance multicloud for VMware

Simplicity

Dell APEX Cloud Platform for VMware is built on a common sets of building blocks:

- MC nodes based on the latest Dell PowerEdge servers equipped with fourth generation of Intel Xeon Processors.
- The same VMware software and integrated architecture present in most datacenters today.
- Common Software Defined Storage (SDS) to provide data mobility and consistent enterprise SLAs in-cloud and on-premises

Agility

The Dell APEX Cloud Platform for VMware takes advantage of the rich VMware ecosystem with Tanzu and VMware Cloud (VMC) to extend IT and application modernization as well as the ability to manage both VMs and K8s on-premises and in the public cloud. This enables developers and IT teams to develop anywhere and deploy everywhere.

Control

Security is top of mind when managing multicloud infrastructure. Siloed application environments make it difficult to consistently enforce policies to establish a zero-trust approach. The lack of engineered integration between the hardware and ecosystem software increases the likelihood of misconfiguration and discourages updates. Automated full-stack lifecycle management means that upgrades and patches can be applied quickly and non-disruptively to improve security posture. Consistent operating environments across multiple deployments enables teams to apply hardened security and governance across locations. And layers of protection, from the cyber-resilient features embedded in the Dell PowerEdge hardware, to the vSphere security frame works developed by VMware in the application layer, serve to inherently secure deployments.

With this new offering, Dell broadens the range of VMware solutions, providing a choice of deployment for HCI and multicloud, Dell APEX Private Cloud, Dell APEX Hybrid Cloud, HCI and now Dell APEX Cloud Platform for VMware.

Dell APEX Cloud Platform for Microsoft Azure introduction

An evolution for Azure Stack HCI

Dell APEX Cloud Platform for Microsoft Azure is the first offer in the Premier Solutions category for Microsoft Azure Stack HCI.

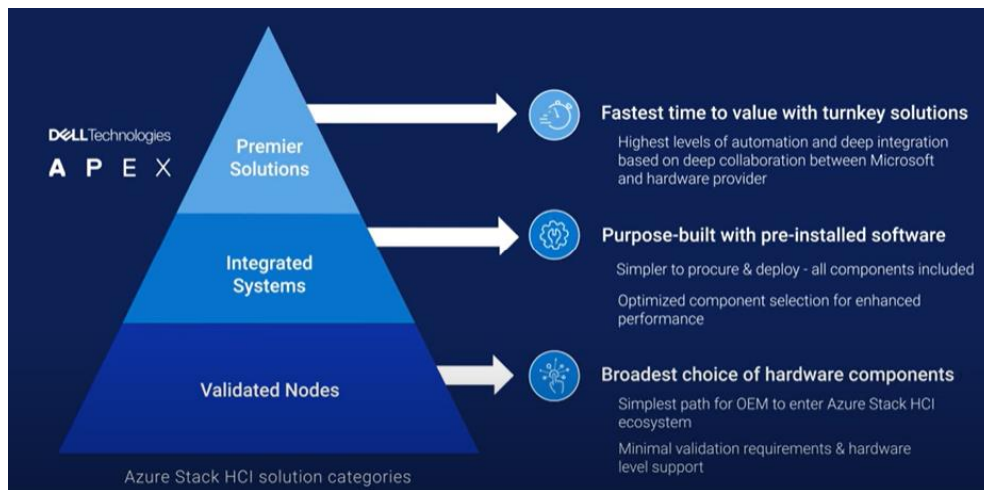


Figure 12. Microsoft Azure Stack HCI categories, including the new Premier one

This new platform comes to the market as the result of a long story of collaboration between Dell and Microsoft. During this long journey we have learned what customers needed to accelerate application development and optimizing IT operations, and worked to provide the best solutions that combined the strength and technologies of two market leaders.

To this innovative platform, Dell brings its consolidated experience creating engineering-validated infrastructure platforms, management automation and orchestration solutions, with the best services experience and portfolio.

Microsoft brings its cloud services leadership, with the ability to simplify customers' management, governance and applications deployment and operations through a centralized Azure control plane.

To this successful collaboration, Dell has added one of its most valuable and proven Intellectual Property pieces, which comes from the leading HCI platform, VxRail.

Importing VxRail Manager technology to build the new Dell APEX Cloud Foundation software, a keystone of this new Dell APEX Cloud Platforms, enables them to inherit the automated operations, orchestration, lifecycle management and improved serviceability that VxRail has had for flagship for years.

Dell APEX Cloud Platform for Microsoft Azure is empowered by the latest generation of Dell PowerEdge servers, incarnated in these platforms in the MultiCloud (MC) nodes. Equipped with the 4th generation of Intel Xeon Processors, capable of offering a wide

range of networking, storage and GPU capabilities, keeping a low footprint and power consumption profile.

On the storage domain, the platform is capable of offering both Microsoft's and Dell's Software define storage technologies (SDS). At initial release, only Microsoft storage spaces direct has been made available, but soon Dell SDS will be available for customer's choice. This Dell's software defined block storage will run on top of MC storage nodes integrating with Azure Stack HCI at OS level.

Core values

With these foundational characteristics, customers can benefit from the core value this next generation platform brings to the Azure marketplace:

- Ability to extend and optimize Azure public cloud experience to on-premises environments.
- Help customers benefit from a rapid time to value through a fully automated and integrated Azure hybrid platform.
- Accelerate developers' productivity by providing a consistent experience on-premises and in the Azure's public cloud.
- Optimize workload placement with flexible and seamless portability to meet application performance and governance requirements.
- Enforce consistent governance and compliance across on-premises and Azure public cloud environments.

Dell APEX Cloud Platform Foundation Software

With this key piece of software, Dell is leveraging and evolving its highly differentiated management and orchestration capabilities from its VxRail system software. We have taken the M&O Intellectual Property that distinguishes VxRail and we have customized it for Azure environments with this software. The software plays a key role in:

1. Automating initial deployments
2. Streamlining on-going operational tasks
3. Delivering operational consistency across your dispersed Azure on-prem and Cloud deployments with centralized management

Dell APEX Cloud Platform Foundation Software consists of multiple, integrated software elements that extend cloud ecosystem native capabilities to deliver a seamless, automated, operational experience, keeping the infrastructure in continuously validated states to ensure workloads are consistently up and running.

Del APEX Cloud Platform Foundation Software integrates with Windows Admin Center and Azure Portal, enabling administrators to have a consistent user experience through familiar tools, managing on-premises resources as they would do in Azure's public cloud

Azure Arc integration

As a hybrid designed solution, Azure Arc plays a key role, extending the Azure public platform to on-premises, helping building cloud-native applications and services with a consistent development, operations, and security model. Through Azure management and governance services, it delivers fleet management at-scale.

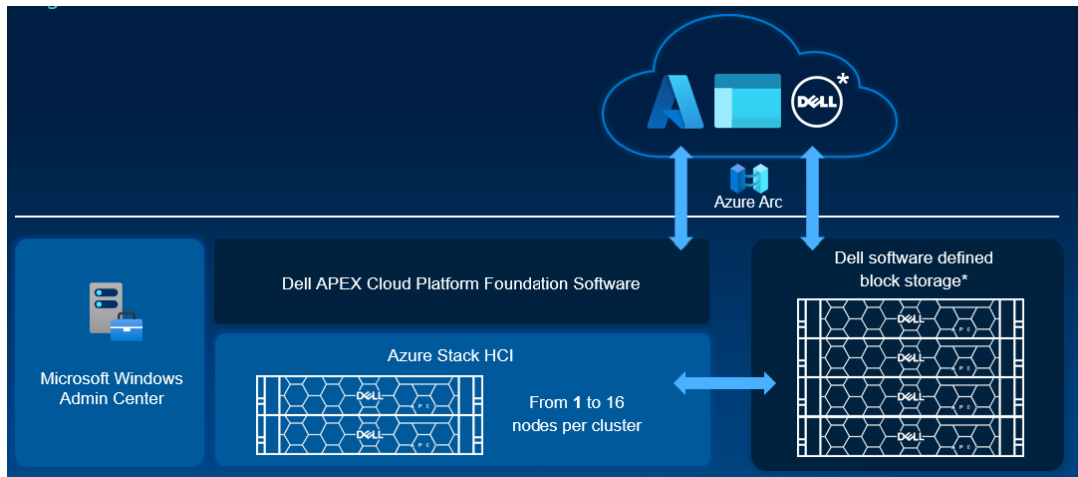


Figure 13. High-level solution architecture. * These will be delivered after initial Dell APEX Cloud Platform for Microsoft Azure launch

Simplify application modernization

Dell APEX Cloud Platform for Microsoft Azure is all about providing the optimal infrastructure foundation to deploy your modern and legacy applications, no matter where you want to deploy them and what specific use cases you are trying to address. This cloud platform will offer a broad choice of infrastructure nodes, some optimized for core datacenter and others for Edge and Remote Office's or Branch Office's (ROBO) use cases. Almost any type of workload deployment and operations gets simplified by the automated platform that has the rightsized infrastructure to host compute, storage, IO or throughput intensive workloads.

In this ample definition we can place generic virtualized workloads, transactional databases, AI/ML analytics, virtual desktop workloads, or containerized cloud-native applications. Edge and ROBO computing will have specific solutions as part of the portfolio, soon after the initial release.

Security and governance hardening

Dell 16th generation PowerEdge servers feature an enhanced Cyber Resilient Architecture that provides a hardened server design to protect, detect, and recover from cyberattacks. Since MC nodes in the APEX Cloud Platforms are built on 16th generation PowerEdge servers, the solution inherits the benefits of the Cyber Resilient Architecture.

Delivering the Cyber Resilient Architecture requires security awareness and discipline at each stage of development. This process is called the Security Development Lifecycle (SDL) model, in which security is not an afterthought but is rather an integral part of the overall server design process. This design process encompasses a view of security that spans the entire server lifecycle.

We are performing vulnerability scanning Black Duck® software composition analysis (SCA), penetration testing, threat modeling, code reviews, and more for every release of our APEX Cloud Platform for Microsoft Azure.

APEX Cloud Platform for Azure helps mitigate potential security and compliance risks by enforcing consistent Azure management and governance services from cloud to edge. It takes a full-stack approach to security with sophisticated controls and multilayered protection from security threats.

Dell APEX Cloud Platform for Microsoft Azure Hardware platforms

At initial release, Dell APEX Cloud Platform for Microsoft Azure hardware platforms are based on Dell's PowerEdge R660 and R760 rackmount servers. Soon we will expand the platform family with the addition of a MultiCloud (MC) node based on the edge-optimized XR4000 series PowerEdge servers.

The MC-660 and MC-760 provide a flexible range of configuration options, best suited to adapt to the customer's performance needs, and workload type.

In the following figure, we can visualize at a glance, the main options available at initial launch:

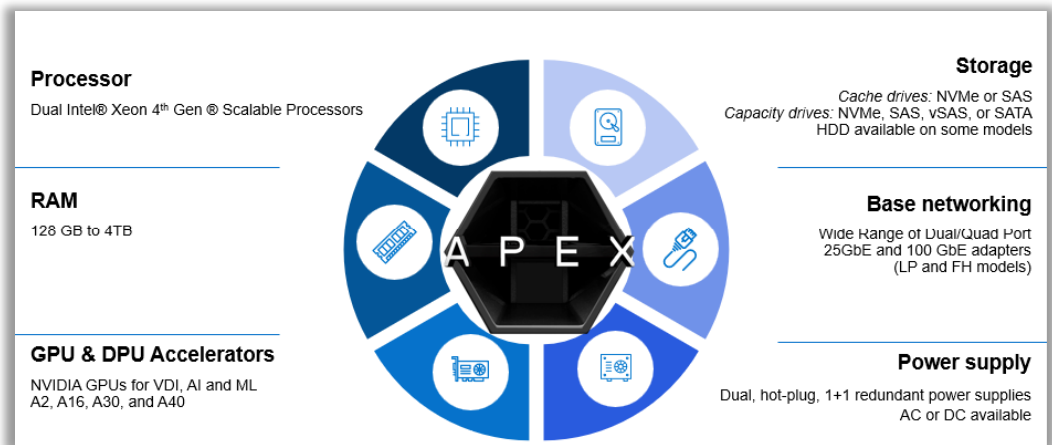


Figure 14. Dell APEX Cloud Platform multicloud nodes configuration options

The MC-660 represent the 1U form factor option, condensing up to 10 SSD or 10 NVMe drives, 4 TB of DDR5 RAM and two Intel Sapphire Rapids processors in a small frame.

For the MC-660 there are two configuration types: those based in SSD drives, named **All Flash**, and those based in NVMe drives, named **All NVMe**. These are summarized in the figure below:

| | All Flash | All NVMe |
|------------------------|--|--|
| Model | MC-660 – All Flash | MC-660 All NVMe |
| Chassis Configuration | 10 x 2.5" disk slots (up to 2 universal) C05-03 | 10 x 2.5" NVMe C05-09 |
| Processor | 2x Intel Sapphire Rapids Processors | 2x Intel Sapphire Rapids Processors |
| Memory | Up to 32 x DDR5 (128 GB – 4 TB) | Up to 32 x DDR5 (128 GB – 4 TB) |
| Storage controller | HBA355i | NA |
| Storage configurations | Up to 76.8 TB SSD | Up to 153.6 TB NVMe |
| OS Boot | BOSS N1 – 2 x 960 GB M.2 (RAID 1) | BOSS N1 – 2 x 960 GB M.2 (RAID 1) |
| NIC | PCIe– DP 25/100 GbE Optional OCP 3.0 | PCIe– DP 25/100 GbE Optional OCP 3.0 |
| GPU | Option for 2 SW | Option for 2 SW |
| PCIe slots | RC2 (R2P + R3P) = 3x LP (2x LP Gen5 + 1x LP Gen4) RC3 (R1P + R4P) = 2x 16 FH (Gen5) | RC2 (R2P + R3P) = 3x LP (2x LP Gen5 + 1x LP Gen4) RC3 (R1P + R4P) = 2x 16 FH (Gen5) |
| Power | 1100/1400/1800 W | 1100/1400/1800 W |

Figure 15. Dell APEX Cloud Platform for Microsoft Azure configurations based on MC-660 nodes

Two riser configurations have been chosen for Dell APEX Cloud Platform for Microsoft Azure:

- Config 2: based on three Low Profile (LP) PCIe slots, two Gen5 and one Gen 4)




Figure 16. First Riser configuration for MC-660

- Config 3: based in two Full Height (FH) PCIe slots (Gen5)



Figure 17. Second Riser configuration for MC-660


The second option for Dell APEX Cloud Platform for Microsoft Azure is based on the MC-760, a 2U form factor server with different storage option combining SSD, NVMe and SAS HDD drives. It also incorporates two Intel Sapphire Rapids processors and up to 32 DDR5 RAM DIMMs.

For the MC-760 4 different configuration types have been created:

- **All Flash:** based in a storage configuration with up to 24+4 (rear) SAS SSD drives
- **GPU Dense:** based in a storage configuration with up to 24 SAS SSD drives, supporting up to four Single Wide (SW) NVIDIA GPUs or 2 Double Wide (DW) NVIDIA GPUs
- **All NVMe:** based in a storage configuration with up to 24 NVMe drives, supporting up to 4 Single Wide NVIDIA GPUs or 2 Double Wide NVIDIA GPUs

- **Hybrid:** based on hybrid storage configurations mixing SAS HDD with NVMe or SAS SSD in a two-tier architecture

Details for these configuration options are summarized in the figure below:



| | All Flash | GPU Dense | All NVMe | Hybrid |
|------------------------|--|---|---|--|
| Model | MC-760 – All Flash (SAS SSD) | MC-760 – All Flash (SAS SSD) | MC-760 All Flash (NVMe SSD) | MC-760 Hybrid |
| Chassis Configuration | 24x2.5" (SAS4) with 4x Universal + 4x2.5" SAS/SATA rear C09-19 | 24x2.5" (SAS4) with 8x Universal C09-10 | 24x2.5" (NVMe) switch C13-01 | 12x3.5" (SAS HDD) + 4x2.5" (NVMe) 12x3.5" (SAS HDD) + 4x2.5" (SAS) C02-05/C02-06 |
| Processor | 2x Intel Sapphire Rapids Processors | 2x Intel Sapphire Rapids Processors | 2x Intel Sapphire Rapids Processors | 2x Intel Sapphire Rapids Processors |
| Memory | Up to 32 x DDR5 (128 GB – 4 TB) | Up to 32 x DDR5 (128 GB – 4 TB) | Up to 32 x DDR5 (128 GB – 4 TB) | Up to 32 x DDR5 (128 GB – 4 TB) |
| Storage controller | HBA355i | HBA355i | NA | HBA355i |
| Storage configurations | Up to 215 TB SSD | Up to 215 TB SSD | Up to 368.64 TB NVMe | Up to 240+ TB HDD |
| OS Boot | BOSS N1 – 2 x 960 GB M.2 (RAID 1) | BOSS N1 – 2 x 960 GB M.2 (RAID 1) | BOSS N1 – 2 x 960 GB M.2 (RAID 1) | BOSS N1 – 2 x 960 GB M.2 (RAID 1) |
| NIC | PCIe – DP 25/100 GbE Optional OCP 3.0 | PCIe – DP 25/100 GbE Optional OCP 3.0 | PCIe – DP 25/100 GbE Optional OCP 3.0 | PCIe – DP 25/100GbE Optional OCP 3.0 |
| GPU | NA | Up to 4 SW GPU or Up to 2 DW GPU | Up to 4 SW GPU or Up to 2 DW GPU | N/A |
| PCIe slots | 2 - x16 LP slots and 2 - x8 FH slots- RC8 | 2 - x16 LP slots and 2 - x8 FH slots- RC8 for Storage Optimized Config 4 x8 FH (Gen5) + 2x 8 FH + 2 x16 LP - RC2 2 x16 LP + 2 x8 FH + 2 x16 DW (Gen5)– RC3-2 for DW GPU Optimized Configuration | 2 - x16 LP slots and 2 - x8 FH slots- RC8 for Storage Optimized Config 4 x8 FH (Gen5) + 2x 8 FH + 2 x16 LP - RC2 2 x16 LP + 2 x8 FH + 2 x16 DW (Gen5)– RC3-2 for DW GPU Optimized Configuration | 2 - x16 LP slots and 2 - x8 FH slots- RC8 |
| Power | 1100/1400/2400/2800 W | 1100/1400/2400/2800 W | 1100/1400/2400/2800 W | 1100/1400/2400/2800 W |

Figure 18. Dell APEX Cloud Platform for Microsoft Azure configurations based on MC-760 nodes

Four riser configurations have been chosen for Dell APEX Cloud Platform for Microsoft Azure:

- Config 2: based on four FH (Gen5), two FH (Gen4), and two LP (Gen4) PCIe slots



Figure 19. First riser configuration option for MC-760

- Config 3-2: based on two LP (Gen4), two FH (Gen4), and two Double Wide (DW) (Gen5) PCIe slots



Figure 20. Second riser configuration option for MC-760

- Config 5-1: based on two LP (Gen4), two FH (Gen4), and two FH (Gen5) PCIe slots




Figure 21. Third riser configuration option for MC-760

- Config 6: based on two LP (Gen4), two FH (Gen5) PCIe slots




Figure 22. Fourth riser configuration option for MC-760

Every MC node requires one or more add in cards that can be installed in either the OCP slot or PCIe riser slots. The available network cards in MC nodes are one that meet Microsoft functional and certification requirements for Azure Stack HCI. The on-board LOM ports are not used in a Dell APEX Cloud Platform for Microsoft Azure deployment.

In this [Microsoft link](#), we can check the industry standards that are mandatory for the specific roles of network switches validated for Azure Stack HCI.

To choose the right network adapters, qualified by network traffic type (management, compute, storage) we can visit the Windows Server Catalog, [Network section](#) for certified models.

These network adapters need to support this set of capabilities:

- Dynamic Virtual Machine Multi-Queue (Dynamic VMMQ or d.VMMQ)
- Remote Direct Memory Access (RDMA)
- Guest RDMA
- Switch Embedded Teaming (SET)

To expand on these technologies, visit the Microsoft [Host network requirements for Azure Stack HCI site](#).

For more details on these Dell MC nodes please check:

[Dell PowerEdge R660 Technical Guide](#)

[Dell PowerEdge R760 Technical Guide](#)

Dell APEX Cloud Platform for Microsoft Azure networking topologies

Overview

Dell APEX Cloud Platform for Microsoft Azure is by design, built to scale from the smallest deployment to the biggest data center. That flexible scalability is manifested through the different network deployment models we have at our disposal.

Single node

We can start as small as a single node Azure Stack HCI cluster, with a cost-effective and easy to maintain server. In which you can run the same workloads you would run in a multi-node Azure Stack HCI cluster.

This deployment model does not need to deploy additional networking elements, such as Top of the Rack (ToR) switches, and although at initial release, does not support scaling beyond one node, that cluster expansion will be enabled to be done manually, with the help of Dell Services and Support, soon after.

This approach is best suited for edge, or branch offices projects, generally very cost sensitive and that can tolerate the resiliency provided by a single server.

Deployment operations, as well as lifecycle managing these environments, is done via Dell APEX Cloud Foundation software.



Figure 23. Single-node Azure Stack HCI cluster

In this scenarios, four NIC ports are required, even though two of the four ports are not used. Two NIC ports are used for forward Management and Compute traffic and the two unused NIC ports are reserved for possible future cluster expansions and should not be connected until additional nodes are added.

Switchless storage network

Another deployment model that may be attractive for edge use cases is the switchless storage network topology. Switchless storage network topologies use direct connections between the PCI adapters dedicated for storage traffic between cluster nodes. This improves Total Cost of Ownership by not requiring high-speed switches (we recommend 25GbE minimum) for storage interconnects. The connectivity to the Top of Rack switches are only for management and VM traffic. Cluster expansion is not supported when using switchless storage topology. You would need to redeploy the cluster if redeployment became necessary.




Figure 24. Switchless storage network architecture

In this case we always recommend the use of dual-link full mesh configurations, as seen the figure above, simplifying storage networking by using direct connections between cluster nodes. For two-node switchless deployments, two network intents are created. One for management and compute network traffic, the second for storage traffic.

Northbound, we recommend the use of two ToR switches.

There are some advantages associated to this topology:

- The Network interface card (NIC) ports used for storage traffic can operate at a higher speed that the TOR switches can support.
- The ToR switches do not need to support Data Center Bridging (DCB) configuration to support Remote Direct Memory Access (RDMA) storage traffic.
- The absence of switches for storage traffic reduces the capital expenditure cost offsetting up this type of architecture. This fact also impacts in the reduction of potential configuration steps required. This added value would be reduced if we increased the number of nodes.

Scalable

The scalable network architecture should be the option of choice when we want to build an environment that, although can start as small as previous network options, two nodes, is designed to enable cluster expansions in a fully automated way using Dell APEX Cloud Platform Foundation Software, to the maximum of 16 nodes.

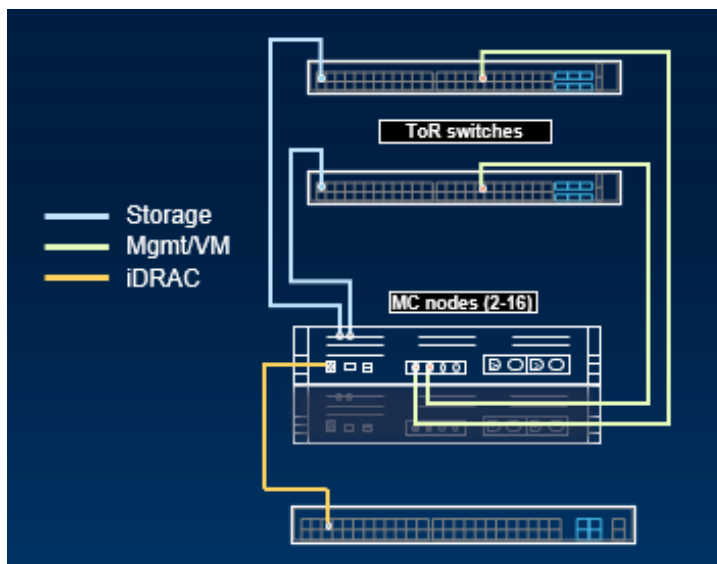


Figure 25. Scalable network topology architecture, detailing all network traffic types

This architecture enables two ways of configuring the networking fabric:

- Fully converged: Management, storage, and VM traffic flow through the same Ethernet ports configured with Switch Embedded Teaming (SET)
- Non-converged: Management and VM traffic use the same Ethernet ports configured as SET, and storage traffic uses dedicated physical adapters

Scalable Network topology supports one or more Azure Stack HCI clusters and allows support of the maximum number of cluster nodes without adding cabling complexity when adding nodes to the initial cluster.

At initial release, node scaling is done one at a time, up to the maximum cluster size of 16. This will be a Dell Services and Support led effort.

The ToR switches require DCB configuration to support SMB Direct traffic over iWARP (computer networking protocol that implements remote direct memory access for efficient data transfer over IP networks) or RDMA over Converged Ethernet (RoCEv2).

Switch ports with a minimum of 25GbE links are recommended for cluster node connectivity.

This topology requires more sophisticated ToR switches, higher bandwidth switch ports and more complex switch configurations.

To choose the best switches for your configuration, we can check Microsoft's physical network requirements URL found [here](#). However, we highly recommend using the Dell PowerSwitch network switches since we have fully validated them and provided sample switch configs on our [InfoHub](#) site.

We can leverage Network ATC for host networking configuration with all the described topologies (Switchless, Non-Converged, and Fully-Converged), obtaining the following benefits in all cases:

- Reduce overall network configuration deployment time, complexity, and reduce human errors due to erroneous inputs.
- Deploy the latest Microsoft validated and supported network configuration best practices.
- Ensure configuration consistency across the nodes in the cluster.
- Eliminate configuration drift with periodic checks every 15 minutes.

Cluster expansion scenarios under these three topologies

- All cluster expansion operations described here will be done with Dell Support and Services assistance.
- Single-node: Expanding to 2-node cluster will be supported and validated by Dell starting in the Dell APEX Cloud Platform for Microsoft Azure March 2024 release.
- Switchless 2-node cluster expansion to 3-node cluster: Microsoft and Dell Technologies do not support this scenario. Beginning with our March release, Dell Support can submit an internal process to have Dell Services convert the 2-node cluster from a switchless storage network topology to a switched storage network

topology. Then, Dell Services can use the Add Node feature in APEX Cloud Platform Foundation Software to expand the cluster to three nodes.

- Switched storage network: Running a two or more nodes cluster and expanding one node at a time: This has always been validated and supported by Dell Technologies and can be accomplished using the Add Node feature in the Dell APEX Cloud Platform Foundation Software. However, today, this must be performed by Dell Services because of the requirement to use the NIM tool on the new node delivered from the factory.

Dell APEX Cloud Platform for Microsoft Azure Operations

Day 0 Operations Node re-image

At initial launch, node imaging is an operation done by Dell Services team. It can be used to cover a number of different scenarios:

- Prepare nodes for deployment with latest validated / supported firmware and Azure Stack HCI OS version
- Re-image / redeploy a faulted or replaced node
- Image and update firmware on a new node for node expansion

These operations are done through the introduction of a new tool, the Node Image Management (NIM) tool. At this point in time, the use of this tool is restricted to Dell support teams.

The NIM tool works through an automated workflow that interacts with the servers IDRAC, Dell's System Update (DSU) and Dell's Repository Manager (DRM) to update firmware, drivers, and finally, install the Azure Stack HCI version.

Day 1 Operations Cluster Deployment

Although providing a fully automated end-to-end Day 1 deployment and cluster creation experience is a common feature for all APEX Cloud Platform, the way it is delivered in the Azure space has its unique implications.

At its initial launch, in September 2023, Dell APEX Cloud Platform for Microsoft Azure is not customer installable. Dell offers a white-glove deployment experience via Dell ProDeploy Services to make the process smooth and successful.

This deployment experience will be made available soon for customers to drive.

The whole deployment process is initiated when the Cloud Platform Manager VM is automatically spawned in one of the MultiCloud (MC) nodes that conform to the Cloud Platform initial cluster. It gets a default IP assigned, so we can browse to that IP and see the Day 1 deployment wizard appear.

Accepting the End User License Agreement is the first requirement to complete. Once done we can start the process by fulfilling deployment prerequisites:

- Create a Service Principal in the Azure Active Directory tenant with the proper permissions in the Azure subscription. This Service Principal was used to register the Azure Stack HCI cluster and onboard the cluster nodes to Azure Arc.

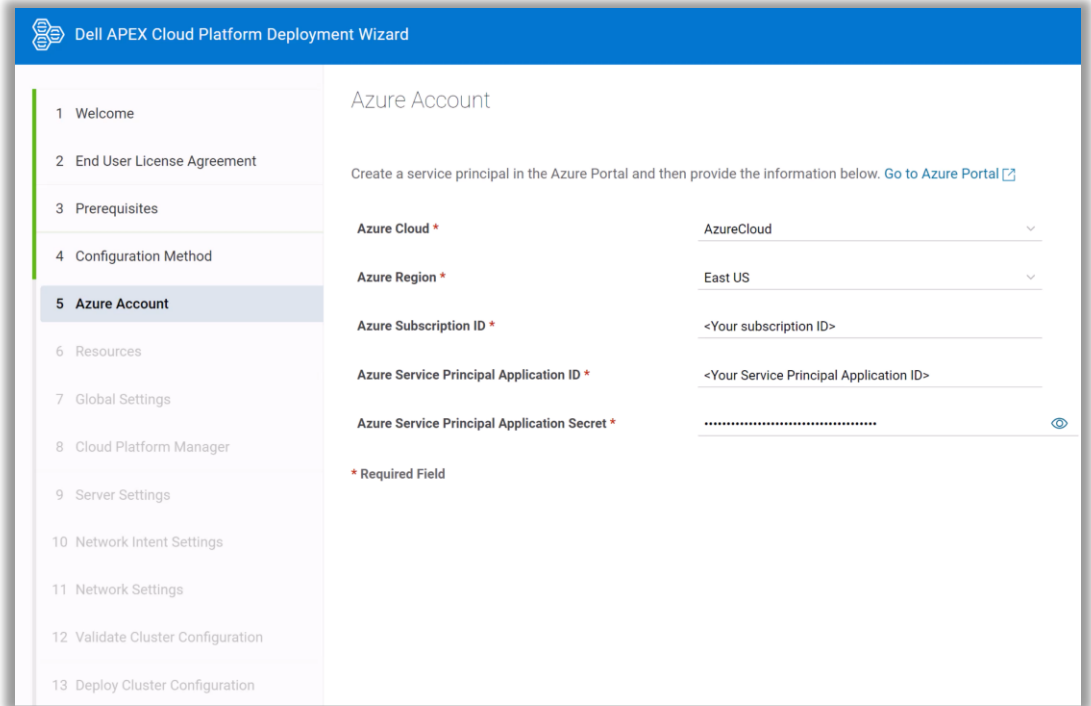



Figure 26. Creating an Azure account Service Principal as a prerequisite for the automated Day 1 deployment

- When prerequisites are completed, available nodes will be automatically discovered, and a primary node will be elected and identified by this house icon . A compatibility check will run to ensure these were valid Dell MC nodes.

We can select the number of servers we want to use to build this first cluster, but our engineering tests have proved that initially creating a cluster with one to six nodes provides the most expedient and consistent deployment experience.

The automated cluster expansion workflow in our APEX Cloud Platform extension in Windows Admin Center can be used to scale out to a maximum cluster size of 16 nodes once the cluster is up and running.

The Wizard driven deployment will inquire for typical infrastructure Global Settings such as cluster name, DNS, LDAPs, Active Directory or Dell’s Windows Admin Center Extension configuration and Azure Arc registration.

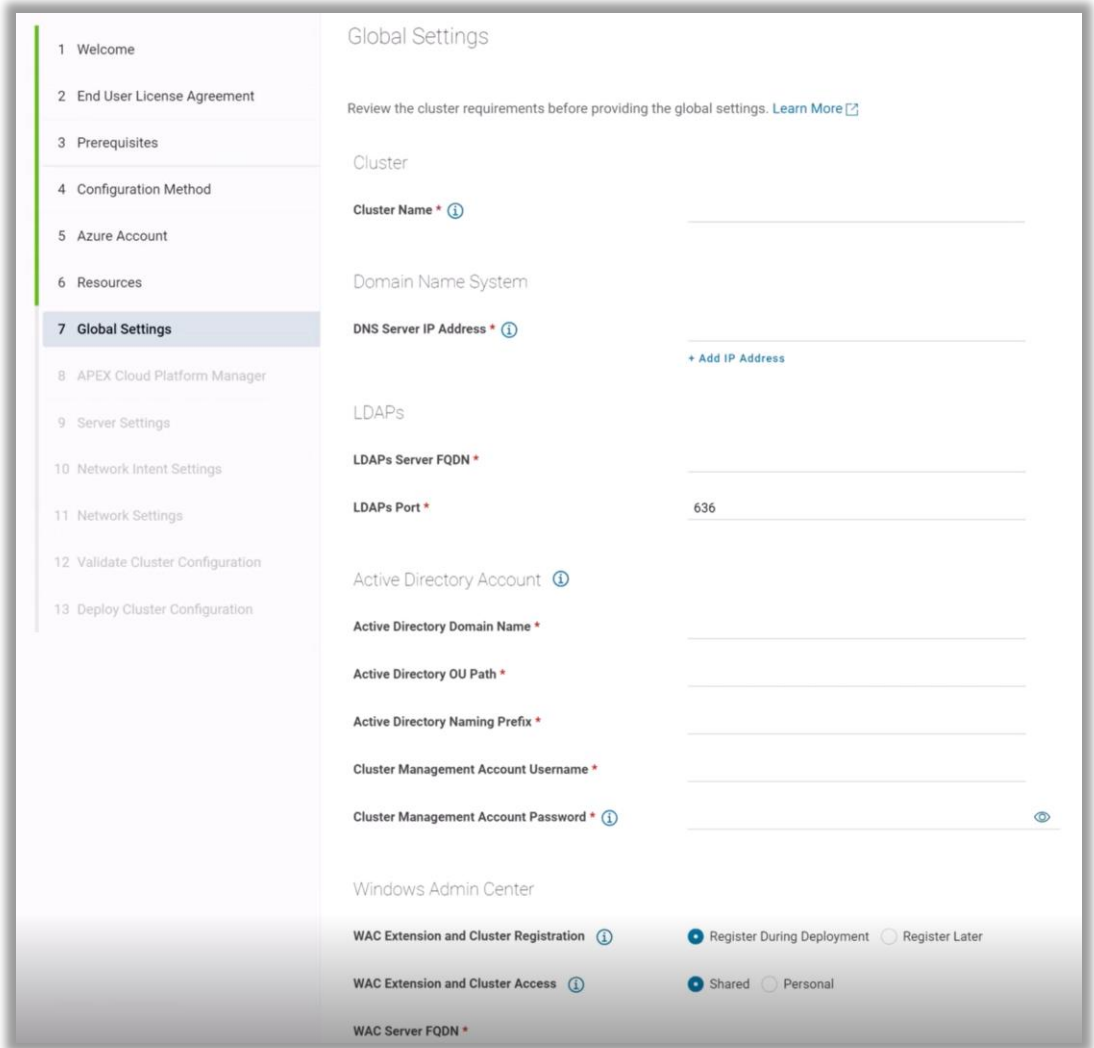


Figure 27. Setting the platform global settings through a wizard driven UI

If Dell's extension did not exist in Windows Admin Center, it can be selected the option for its automated install here:

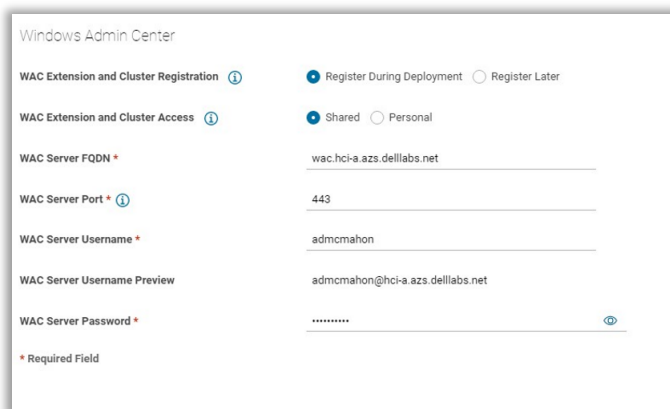


Figure 28. Dell APEX Cloud Platform for Microsoft Azure extension for Windows Admin Center

The automated process continues configuring the Cloud Platform Manager VM credentials, then the servers name and IP settings for the APEX Cloud Platform initial cluster.

Microsoft Network ATC greatly simplifies host network configuration for Azure Stack HCI clusters using an intent-based approach.

In this sample shown, we setup a non-converged network topology creating two intents. Using the integrated Ethernet adapter ports for management and compute traffic and an additional PCI adapter for storage traffic.

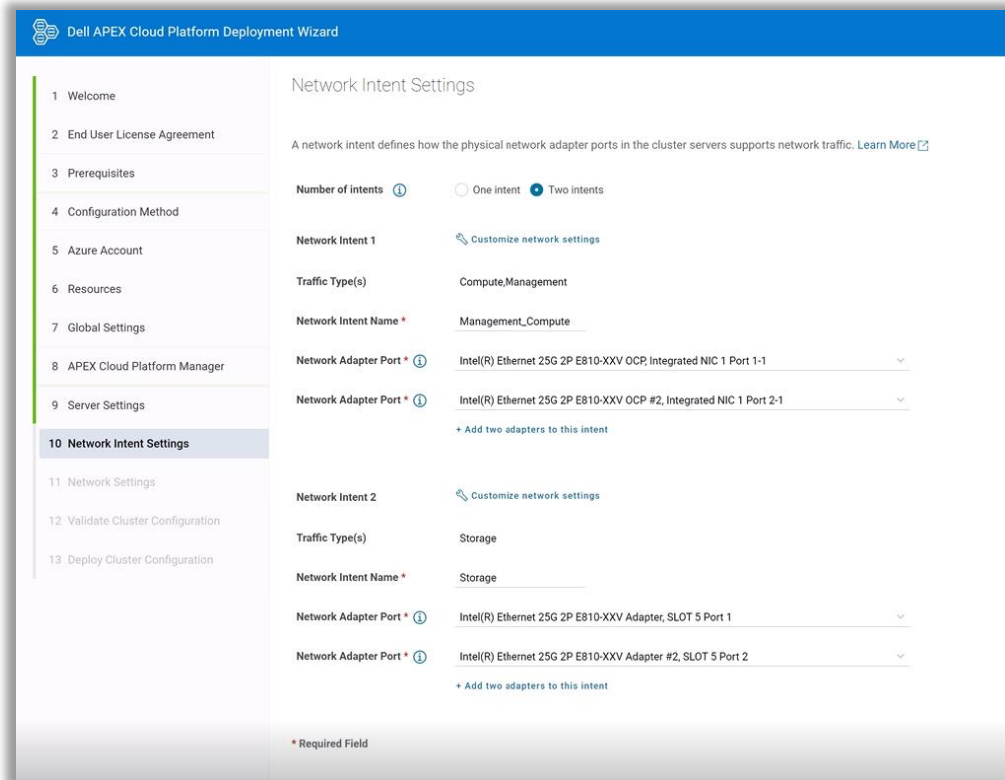


Figure 29. Cluster network configuration using an Intent-based approach.

The wizard driven process allows the user to configure other important network settings such as SR-IOV, jumbo frame size, and RDMA implementation. The last network configuration step will be to provide VLAN IDs for the storage subnets and configure the cluster IP addressing information.

We can perform a final live validation through the **Validate Configuration** option to confirm the data provided is aligned with the rules and best practices proven by our engineering organization to ensure a smooth deployment and cluster creation experience.

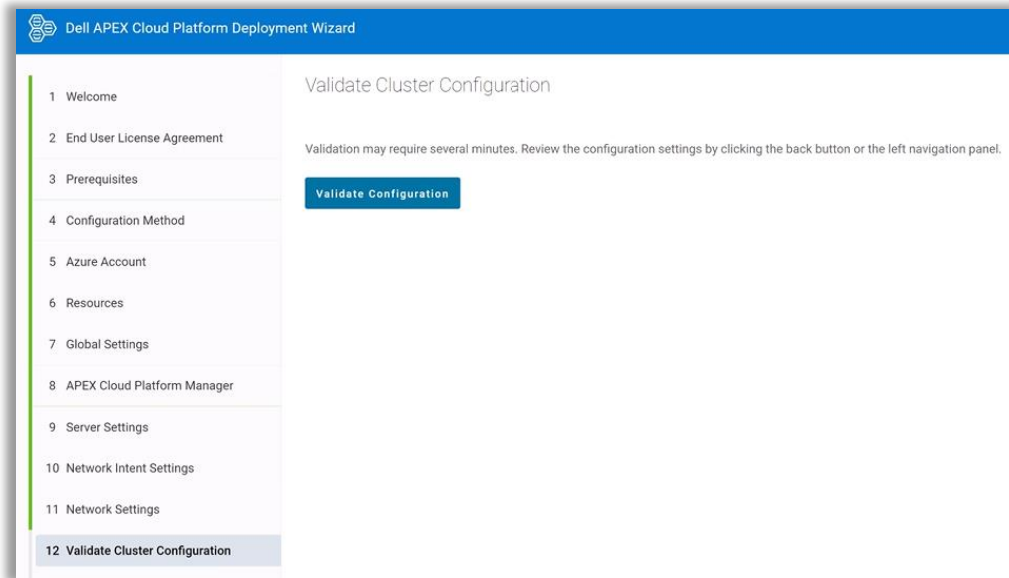


Figure 30. Final validation for initial cluster deployment

Any errors identified at this point can be corrected going back step-by-step to where they were found.

The deployment can take a few hours to run depending on the initial cluster size. During the deployment, we can observe all steps a manual installation would require, and that are automated through this process. There are over 50 steps streaming by in the View Details window. That translates into considerable time savings.

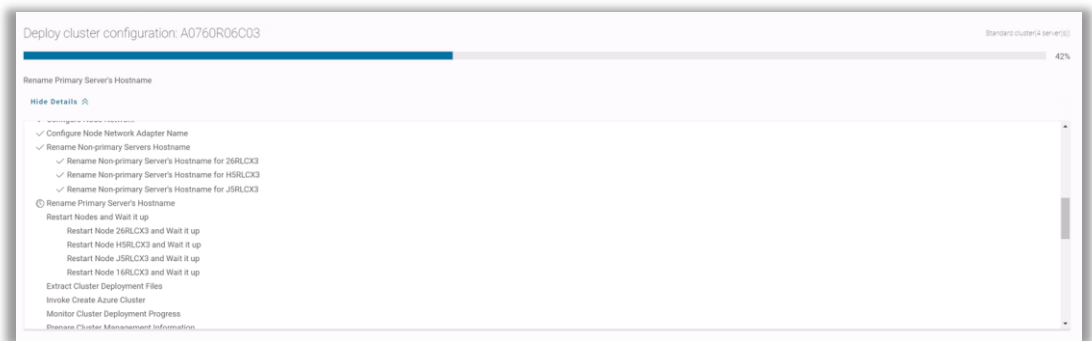


Figure 31. Cluster deployment steps can be observed through the View Details Pane

When the deployment process has successfully completed, we can connect to Windows Admin Center. From there we can perform granular, cluster-level Day 2 operational tasks to administer Dell APEX Cloud Platforms for Microsoft Azure.

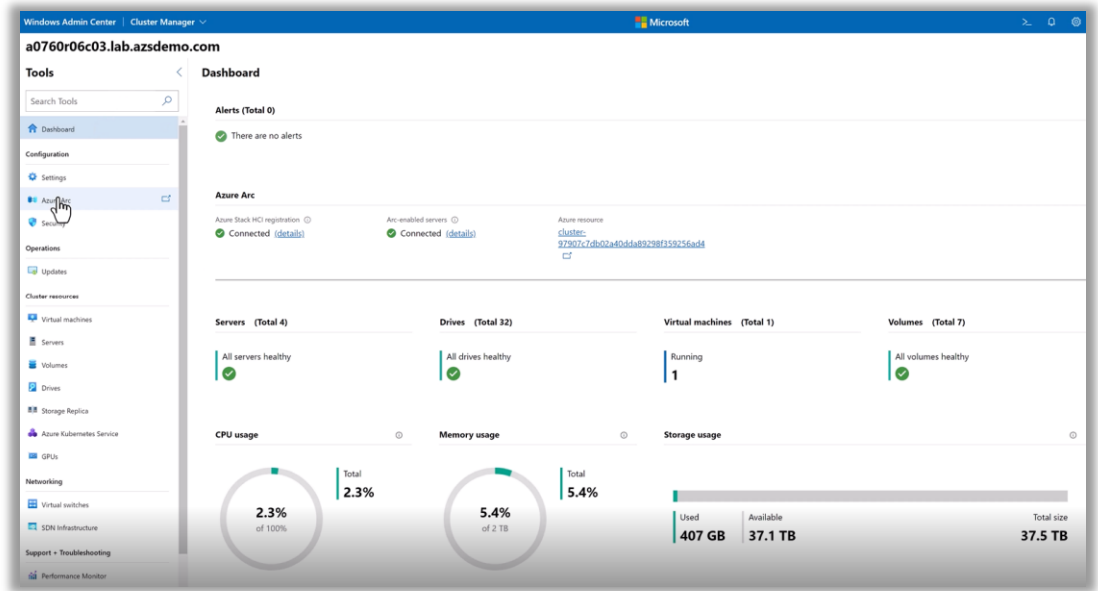


Figure 32. Windows Admin Center dashboard, showing a recently registered Azure Stack HCI node

The automation registered the new Azure Stack HCI cluster and onboarded the cluster-nodes to Azure Arc-enabled servers.




Figure 33. The newly created Azure Stack HCI cluster, a true hybrid cloud platform, registered and onboarded into Azure Arc

Day 2 Operations One of the key goals Dell has set with these cloud platforms is to provide a consistent hybrid management experience. In the case of the Azure platform this is done by leveraging two widely utilized management tools:

- Windows Admin Center extension with automation integrated from Dell APEX Cloud Platform Foundation Software for granular, cluster-level management.
- Azure management and governance services via Arc-enabled infrastructure for fleet management at-scale.

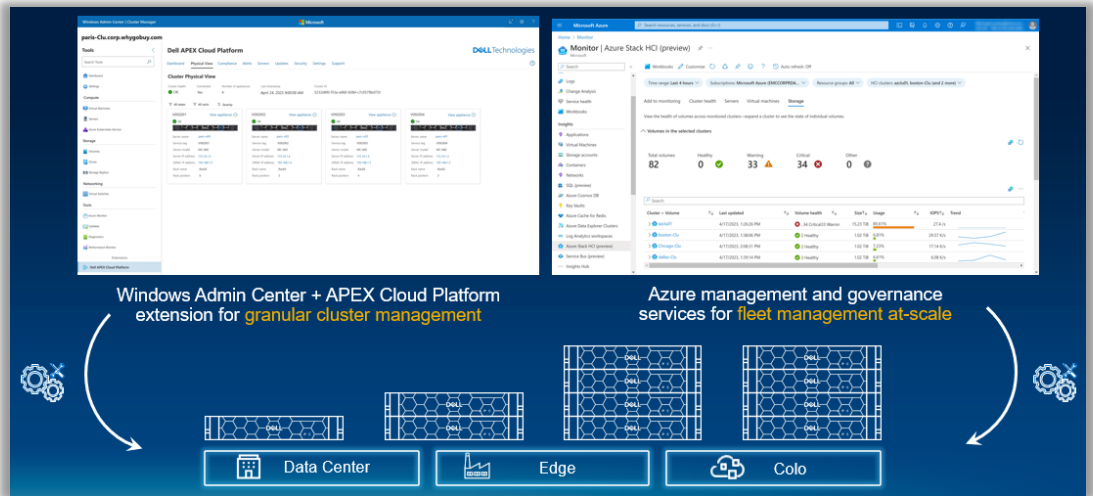


Figure 34. Dell APEX Cloud Platform for Microsoft Azure management experience

Lifecycle Management

Once we have setup our Azure Stack HCI cluster, we need to provide simple, automated ways to keep all the software, drivers, and firmware in a validated stated, certified by Dell and Microsoft, updated to the right version and compatible with the rest of the pieces that conform the solution.

To reach this goal, the platform periodically queries both Dell and Microsoft to discover applicable updates. These will appear in the Updates tab of Dell’s APEX Cloud Platform extension in Windows Admin Center within 4 hours of being released.

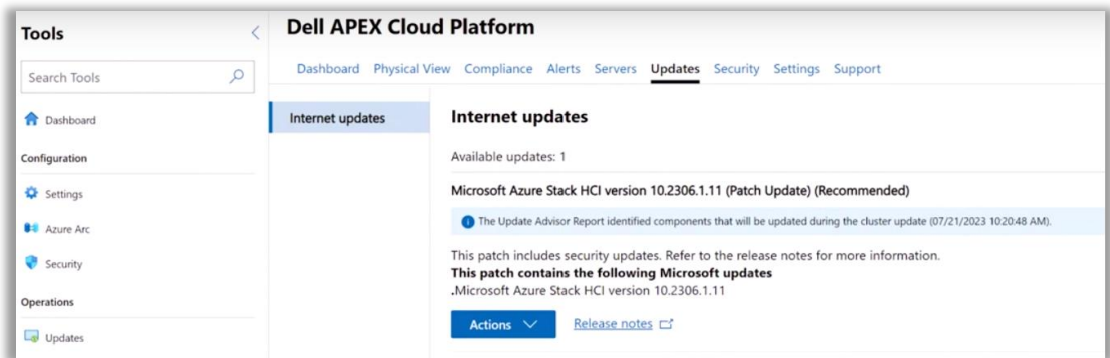


Figure 35. Updates available for Dell APEX Cloud Platform for Microsoft Azure

From there we can apply the updates through an orchestrated process. All the updates published have been jointly tested by Microsoft and Dell in our labs.

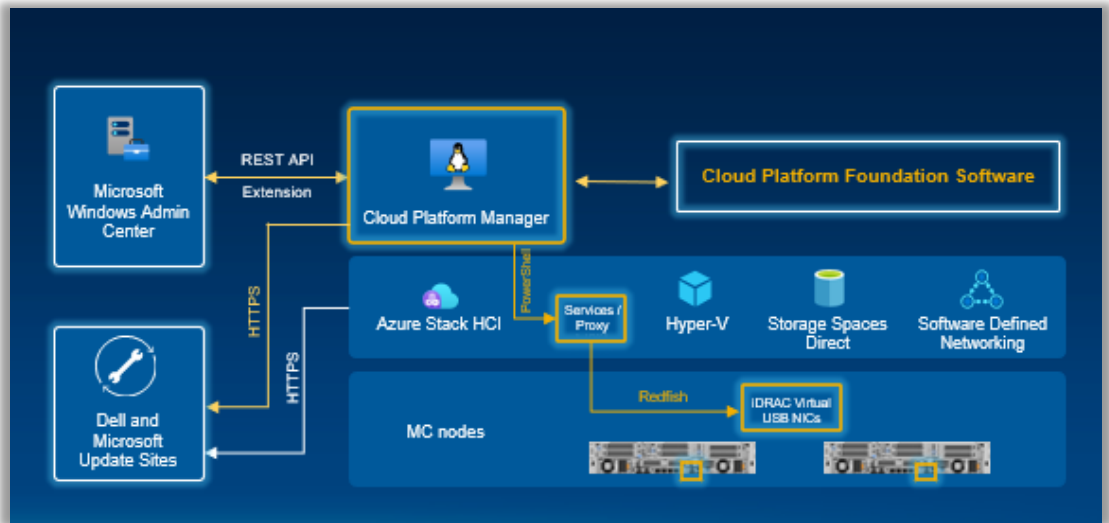


Figure 36. Dell APEX Cloud Platform for Microsoft Azure showing lifecycle management components

Through this automated process we can update, not only the Azure Stack HCI solution, but the BIOS, iDRAC, firmware, and drivers of the MC nodes.

To ensure all component versions are compliant with the platform’s current known good state and that the cluster MC nodes are powered on and do not have any pending or running jobs in their iDRAC job queries, we ran a series of prechecks.

Note: Dell Infrastructure Lock gets automatically disabled during the update process and re-enabled after the updates are completed. We will expand on this and other security features in the Security section.

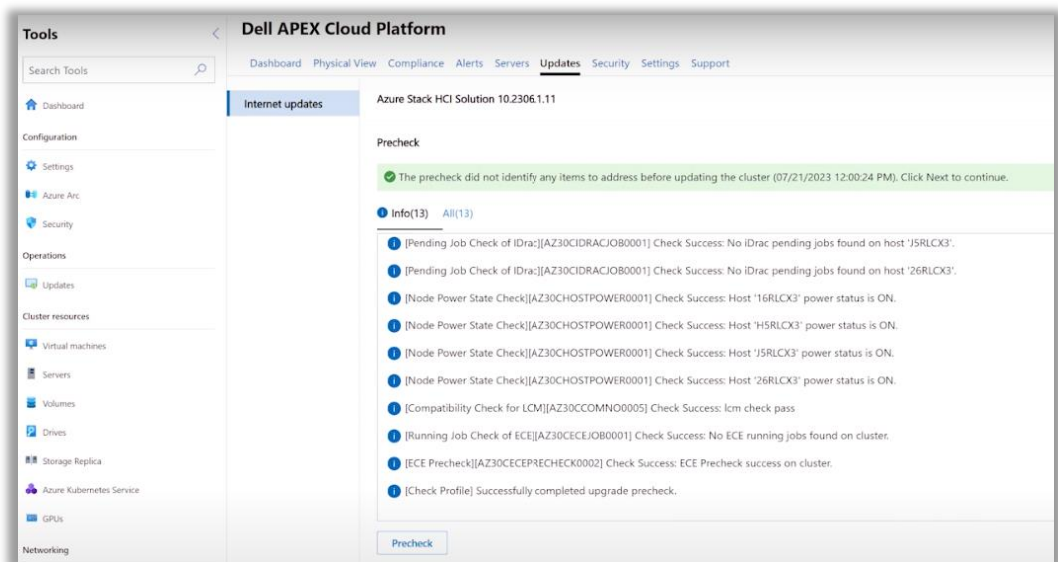


Figure 37. Prechecks completed to validate the launch of the update process

All these updates are applied to the Azure platform in a non-disruptive manner, placing nodes into maintenance mode and orchestrating reboots, one at a time, to avoid disrupting workloads running in the Azure Stack HCI cluster.

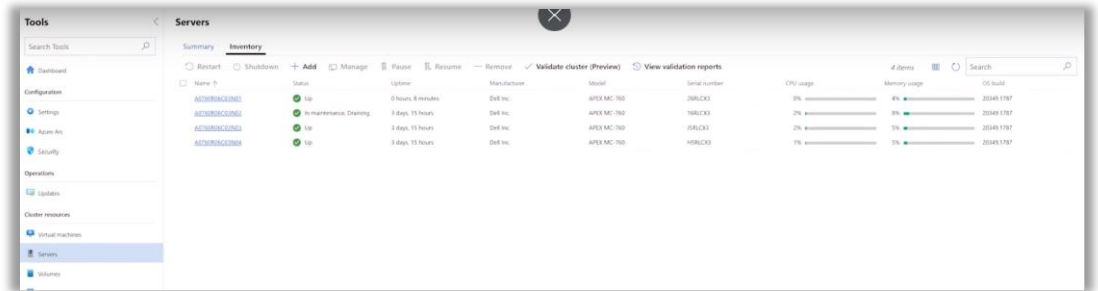


Figure 38. Phased update to a four node Azure Stack HCI cluster

The update process described takes around 90 steps to complete. All these steps have been automated by this Dell APEX Cloud Platform Foundation Software. That represents huge time savings, and much lower chances for a human error.

Through the compliance task we can confirm the update has been completed successfully and the Azure Stack HCI cluster is now fully compliant, as see in the following capture:

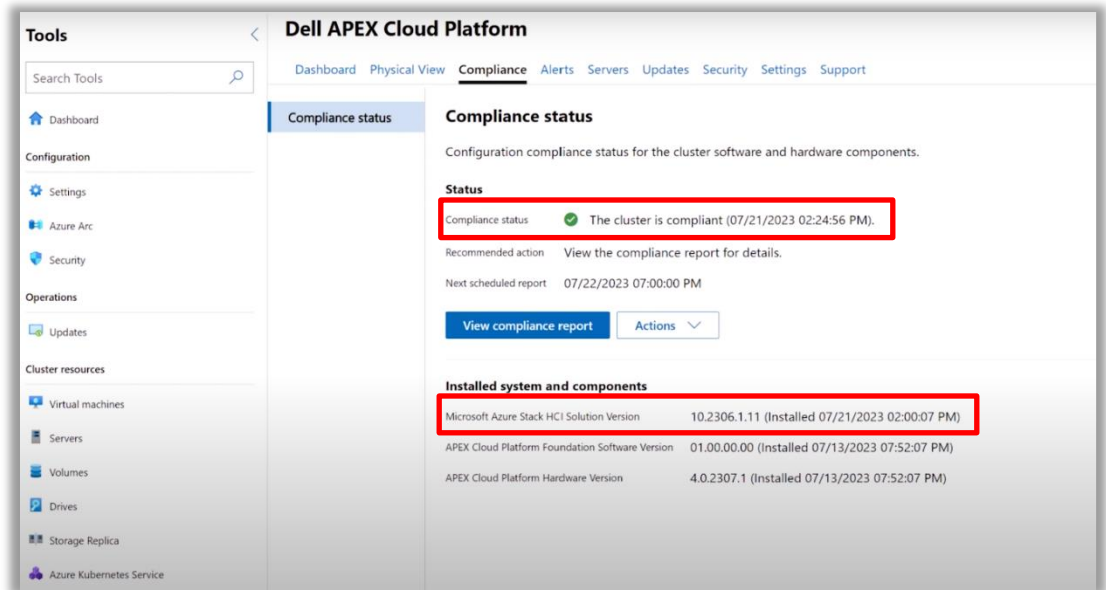


Figure 39. Automated updated process completed, showing a fully compliant Azure Stack HCI cluster

Cluster Expansion

At initial release, the Dell APEX Cloud Platform for Microsoft Azure supports only expanding a cluster running a minimum of two nodes, and it must be configured to use the scalable network topology. With this case, you can add one node at a time to a maximum cluster size of 16 nodes.

There are plans to improve this stance in upcoming platform releases.

As seen with lifecycle management, cluster expansion is a wizard driven, heavily automated operation.

When the nodes for the expansion have been physically installed and properly connected to their network switch ports, we can launch the expansion process. This will launch a node discovery mode that uses IPv6 Multicast Listener Discovery (MLD) snooping to find available servers.

When discovery has been completed, four compatibility checks are run against the new hardware:

- Symmetry
- Disk commonality
- Hardware
- Software

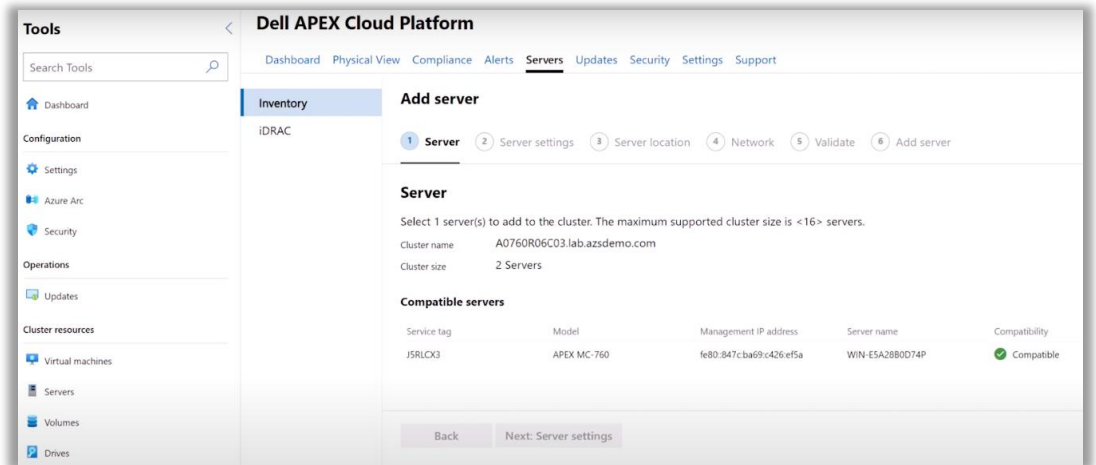


Figure 40. Compatibility checks are run by the platform to validate the addition of new nodes

Had any of these failed, we have been provided the details to correct the failures found.

The workflow also runs additional validation checks designed to ensure consistency and standardization are maintained across the platform after the cluster expansion takes place, and that resiliency has not been compromised.

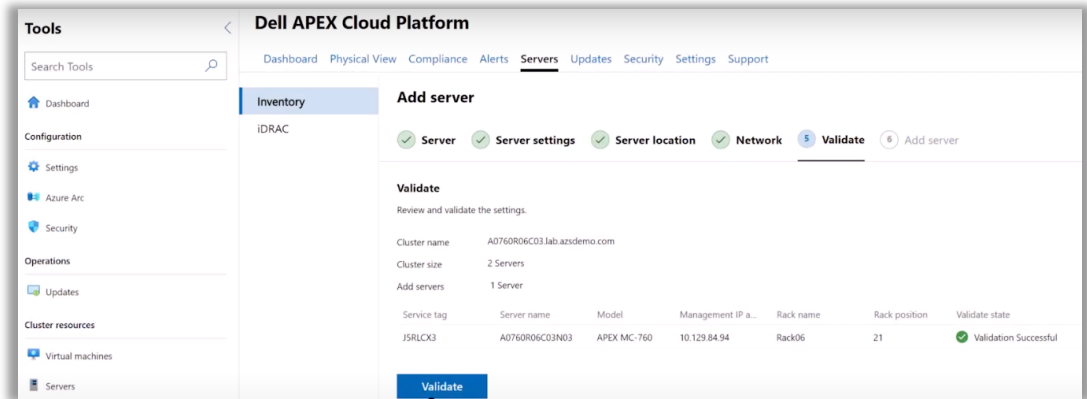


Figure 41. Platform validation checks has been successfully completed

The process to add a node can take up to 2 hours, and the node is added with hostname, networking, and local administrator credentials configured. The process also rebalances the storage pool through a low priority operation not to impact the running workloads. Additional validations are run at this point to ensure normal cluster operation.

In case any failure happens during these configuration steps, the workflow would stop. For this circumstance a retry mechanism is included to rapidly resume the process from the failed step.

Serviceability integrations

Once our Dell APEX Cloud Platform for Microsoft Azure is completely deployed, one of the first actions we want to undertake is establish the necessary protocols to ensure Dell support is fully operational. For that, the integration between the APEX Cloud Platform and Dell's ProSupport backend needs to be setup.

Connection to Dell ProSupport can be done directly, or through a locally installed Dell Secure Connect Gateway (SCG). This connection will enable the automatic upload of logs files and telemetry data, the phone home feature, automatic support case creation, and remote access to the platform by Dell ProSupport for rapid issue resolution.

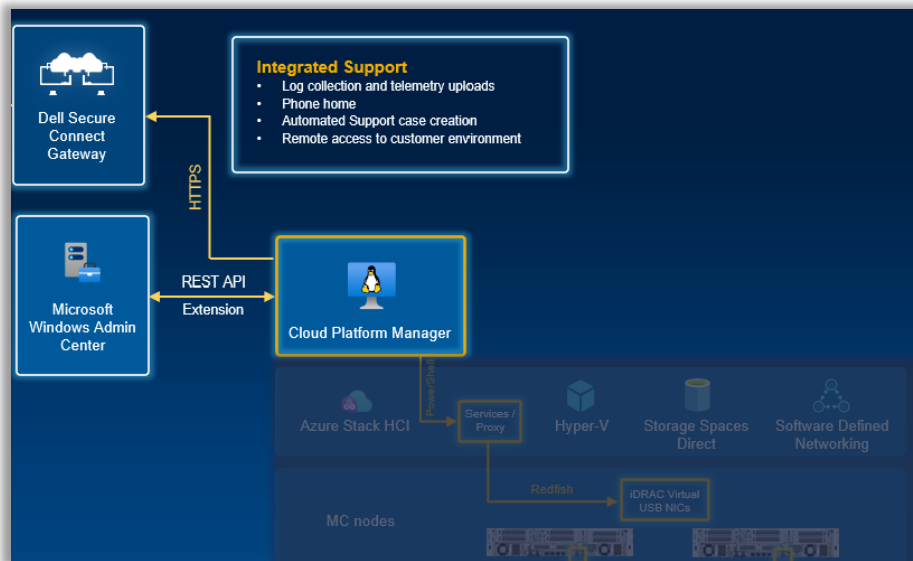


Figure 42. Platform architecture to provide integrated support features

Through the platform Network settings, we can establish a connection option to Dell support and telemetry settings.

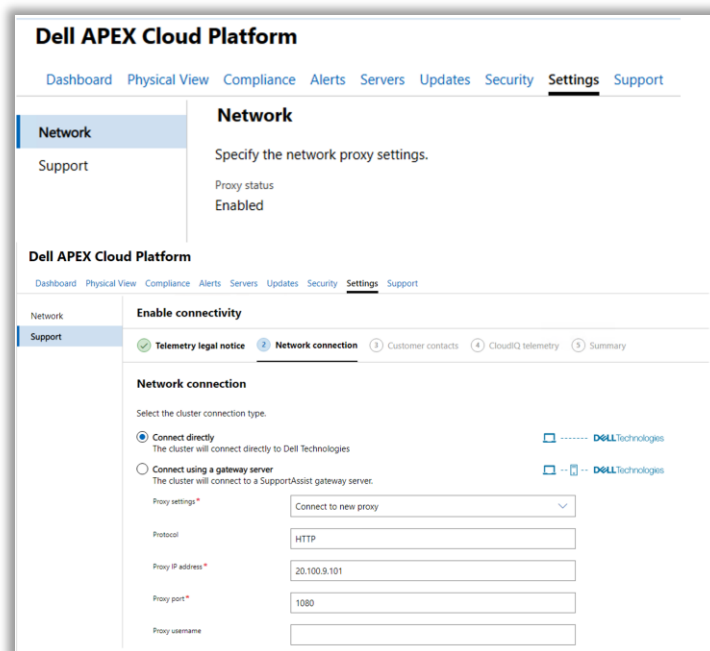


Figure 43. Setting up the integration of Dell APEX Cloud Platform for Microsoft Azure and Dell ProSupport. This capture shows a direct connection

Dell CloudIQ provides a set of key features like proactive health monitoring, predictive analytics, and incident detection and reporting.

We can enable sending telemetry data to Dell Support. At initial release Dell APEX Cloud Platform for Microsoft Azure will not show up in CloudIQ's management console, but that feature will be available very soon.

The Support tab is the central pane for all the platform serviceability features. From there we can:

- Create service requests (SR). Through a template we can input all the necessary data Dell Support needs to solve a potential issue
- We can also search for knowledge base articles and posts on Dell community forums that are relevant to the APEX Cloud Platform from the extension in Windows Admin Center
- We can navigate through the latest cluster configuration data sent to Dell Support or download a configuration report in JSON format
- We can create, download, and manage, log bundles that can help troubleshooting operational issues in the platform. We can choose to create the log bundle from three categories:
 - Cloud Platform Manager
 - iDRAC
 - Azure Stack HCI

As we have setup the connectivity with Dell Support, a Dell Support engineer can establish a secure remote connection into the Cloud Platform Manager VM to troubleshoot and resolve issues.

Another serviceability feature available is the automated case creation. When an alert surfaces in Dell APEX Cloud Platform integration in Windows Admin Center, an event will pop-up in another Dell Support tool (Dell Support CLM Reporting). If the generated alert needs intervention from Dell Support engineers, they will proactively contact the customer to begin working on the case.

Dell APEX Cloud Platform for Microsoft Azure Security

Cyber resilient architecture

Dell's 16th generation PowerEdge servers feature an enhanced Cyber Resilient Architecture that provides a hardened server design to protect, detect, and recover from cyberattacks. Since MC nodes in the APEX Cloud Platforms are built on 16th generation PowerEdge servers, the solution inherits the benefits of the Cyber Resilient Architecture.

Delivering the Cyber Resilient Architecture requires security awareness and discipline at each stage of development. This process is called the Security Development Lifecycle (SDL) model, in which security is not an afterthought but is rather an integral part of the overall server design process. This design process encompasses a view of security that spans the entire server lifecycle.

We are performing vulnerability scanning, Black Duck® software composition analysis (SCA), penetration testing, threat modeling, code reviews, and more for every release of our platform, as an on-going testing and evaluation of new potential vulnerabilities.

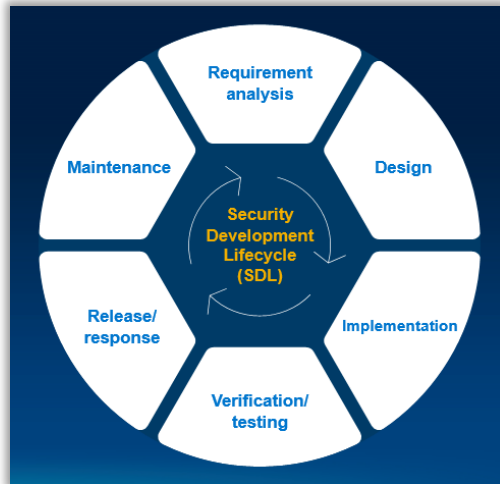


Figure 44. Security present in every phase of a product lifecycle

Through this process, product features are conceived, designed, prototyped, implemented, set into production, deployed, and maintained, with security as a key priority. Server firmware is designed to obstruct, oppose, and counter injection of malicious code during all phases.

Full stack lifecycle management

Full stack lifecycle management is a critical ingredient to maintaining a strong security posture throughout the life of your APEX Cloud Platforms. You cannot benefit from the rigorous SDL release efforts unless you are consistently applying Dell and Microsoft updates. Lifecycle management is also important to keep your platforms operating at peak performance and resiliency. We have made it easy to maintain a Continuously Validated State with our automation.

As shown in previous sections, you do not have to leave your Dell APEX Cloud Platform extension in Windows Admin Center to perform system updates. Dell and Microsoft updates sites are periodically queried for available updates that will show in the Updates tab within 4 hours of being released.

Table 1. Update categories, typical cadence, and description

| Update Type | Typical Cadence | Description |
|------------------|-----------------|---|
| Baseline updates | Quarterly | Baseline updates include new features and improvements. They typically require host system reboots and might take longer. |
| Patch Updates | Monthly | Patch updates primarily contain quality and reliability improvements. Some patches require host system reboots, while others do not. To fix critical bugs or security vulnerabilities, patches might be released sooner than a monthly. |
| Hotfixes | On-demand | Hotfixes address blocking issues that could prevent regular patch or baseline updates. |

All these update types, as a Premier Azure Stack HCI solution, are validated by both Dell and Microsoft engineering teams in our respective labs using mature CI/CD practices.

This ensures updates are stable and applied consistently every time to optimize the performance, stability, and security of the platform.

Full stack lifecycle management introduces significant benefits, such as:

- Update management simplification by consolidating update workflows for various components into a single experience. You never have to leave Dell's interface in Admin Center to apply updates.
- Keeping the platform in a well-tested and optimal configuration state.
- Updates are ready to apply within 4 hours after releasing to the public.
- Helps avoid disrupting running workloads with comprehensive health checks before and during an update. Also, the Cluster-Aware Updating features is leveraged to perform rolling updates one node at a time.
- Improve platform reliability with automatic retry and the remediation of known issues.

Intrinsic infrastructure security features

Dell APEX Cloud Platform for Microsoft Azure comes bundled with intrinsic infrastructure security management features. Combined with the security features that come out-of-the-box with Azure Stack HCI, your platform is fully protected against current and future threats.

The two security management features leveraged in Dell's APEX Cloud Platform for Microsoft Azure are:

- Infrastructure lock
- Secured-core server

Infrastructure lock

Infrastructure Lock is a feature of Dell APEX Cloud Platform Foundation Software. The Dell APEX Cloud Platform extension in Windows Admin Center provides a toggle switch that enables and disables the System Lockdown feature of Dell iDRAC on a cluster-wide basis.

Enabling Infrastructure Lock on all the MC nodes in a cluster:

- Protects against unintentional or malicious changes to critical configuration settings in the BIOS or iDRAC
- Prevents unapproved updates of BIOS, iDRAC, firmware, and drivers

All the MC nodes in a cluster will have System Lockdown disabled after initial cluster deployment. However, it is strongly advised to enable it after the platform is fully operational.

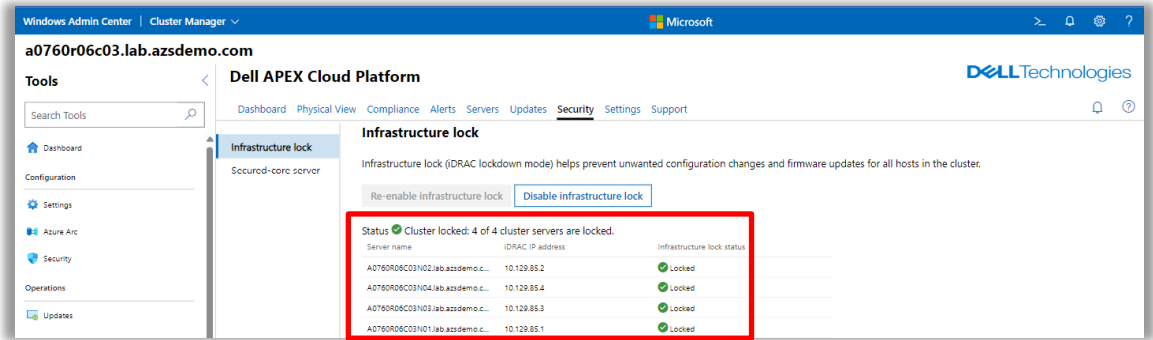


Figure 45. Infrastructure lock can be easily enabled in a cluster wide basis, protecting MC nodes against malicious changes

When infrastructure lock is enabled, the iDRAC UI will alert about the impossibility of performing configuration changes in that server as expected.

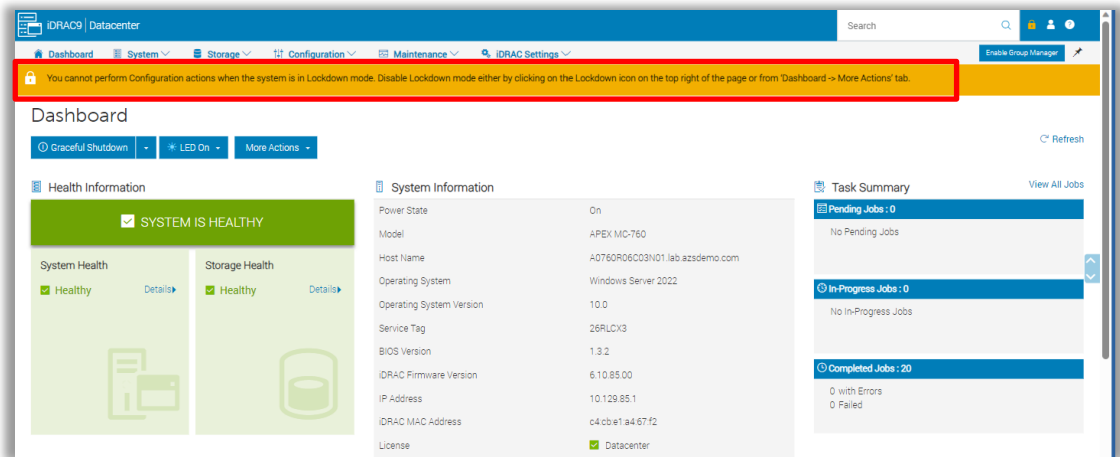


Figure 46. MC-760 iDRAC console showing infrastructure lock enabled, thus impeding configuration changes

Secured-core server

Secured-core server helps proactively defend against and disrupt many of the paths attackers might use to exploit a system by establishing a hardware root-of-trust, protecting firmware, and introducing virtualization-based security.

To fully enable secured-core server, settings must be configured in the BIOS and operating system. In the case of Dell APEX Cloud Platform for Microsoft Azure, a secured-core server is fully enabled during initial deployment.

If any required settings get changed, you can return them to fully secured-core server compliant using the Windows Admin Center extension with no disruption to running workloads.

Secure-core server policies can only be modified in the nodes of the Azure Stack HCI cluster if infrastructure lock is disabled. If this happens in any node of a cluster, it is very simple to bring the node/nodes back to full security by navigating to the Security tab in the Windows Admin Center extension and **Re-enabling** secured-core server.

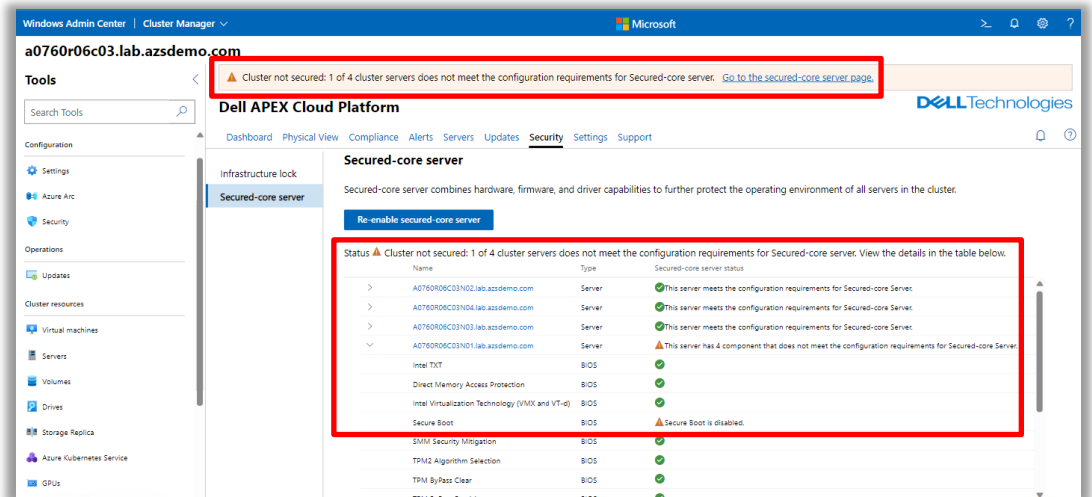


Figure 47. A node in the Azure Stack HCI cluster is not meeting the Secure-core server requirements, as the Secure Boot policy has been disabled

To correct these configuration settings requires restarting the cluster node. This will be done in a non-disruptive manner by placing the affected node in maintenance mode and Live Migrating the running VMs to another node.

By leveraging these two security features, **Infrastructure Lock** and **secured-core server**, we can maintain a strong security posture on these infrastructure platforms.

Dell Services

Companies that want to deliver cloud services to employees and customers through a multicloud approach reap many benefits. However, the complexity of a multicloud approach brings many challenges, some of them already introduced in the first chapter of this book.

Some key challenges, from a Services perspective, include a siloed cloud structure and overprovisioning, application incompatibility across cloud platforms, inconsistent governance practices and more.

Dell services are uniquely positioned to help you achieve a multicloud by design architecture. If you are new to multicloud, we will get you started on the right path, and if you need to get control of siloed, unregulated cloud implementations, we can help you bring them together within a strategic framework, as follows:

- We can help you develop a strategy to optimize your cloud investments, avoiding overprovisioning and using your budgets efficiently.
- Our services ensure you have a sound rationale for application and workload placement based on factors like performance and security requirements.
- When you let our experts deploy, support, and assist with your multicloud infrastructure, you free your internal resources to focus on business and growth priorities.

- Improve efficiency of DevOps and simplify container deployment and management across cloud platforms.
- Your data is the life of your business. Rely on Dell services to help you keep it secure with rational workload and application placements and the implementation of data protection and cyber security solutions.
- Ensure your staff is up to date with the latest technical expertise from Dell education services

Infrastructure Services

Dell Infrastructure Services deliver global best practices to deploy, support, and assist in your IT needs specific to Dell APEX Cloud Platforms.

Dell ProDeploy

Dell ProDeploy for Infrastructure ensures you are set up for multicloud success from day one, including versatile delivery options, like remote or onsite, to fit every budget and operating model.

The ProDeploy Infrastructure Suite provides choice for customers. The minimum requirement for Dell APEX Cloud Platform for Microsoft Azure is ProDeploy. This includes onsite hardware installation with all configurations done remotely. With ProDeploy Plus, customers get 24X7 onsite hardware installation and configuration, additional deployment hours for customer directed tasks and training credits for their staff.

Dell's expertise in the deployment of this solution benefits customers in several ways, some of which came to light in a 2023 study conducted by [Principled Technologies](#). Customers who chose ProDeploy Plus were able to get their systems deployed three times faster and with 67% less project planning time versus their in-house administrators. Also, an internal analysis showed that customers placed 58% fewer support calls after deployment when hiring Dell for the job.

Dell ProSupport

With Dell ProSupport for Infrastructure, customers benefit from a single source of solution-level support for both hardware and software. This includes support for Microsoft Azure Stack HCI operating system, regardless of where the customer purchases their software. At times when collaboration is required with the Microsoft support team, Dell handles the interaction, allowing our customers to handle support with just one vendor.

Dell offers three support options for customers of APEX Cloud Platform for Microsoft Azure:

1. ProSupport Next Business Day, which is the minimum required, and it provides comprehensive, 24X7 proactive and predictive support with next day response.
2. ProSupport is also available with a four-hour response objective.
3. ProSupport Plus delivers mission critical support with a four-hour response objective and a six-hour repair objective for hardware related issues for example, we will be at our customers site within four hours, and we will repair hardware issues within six hours). ProSupport Plus now also includes:
 - On-demand onsite diagnosis from Dell should our customers staff not be able to perform onsite troubleshooting
 - Rapid dispatch of a senior field engineer in parallel with phone-based troubleshooting

- Priority production in the event of a critical situation caused by natural disaster or any other event not normally included under warranty. In many cases, this includes the expedited production of a new Dell system.

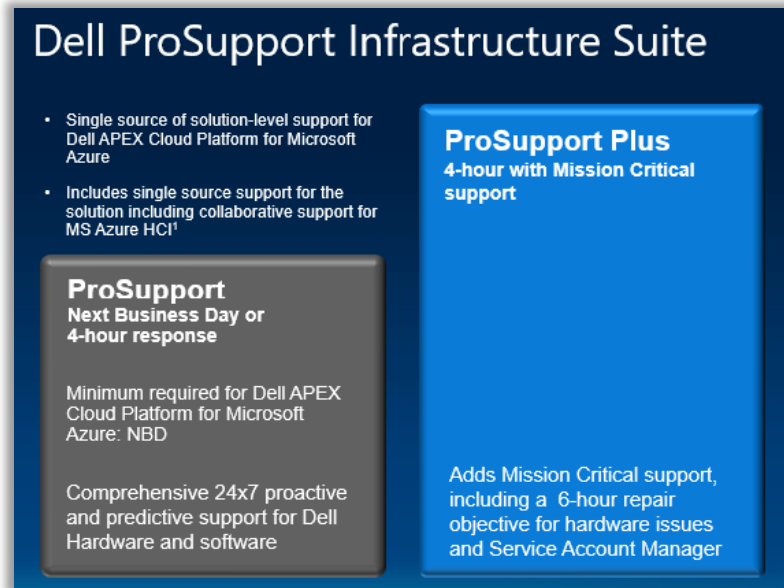


Figure 48. Dell ProSupport Infrastructure suite options

So the key things to remember regarding our ProSupport offers for Dell APEX Cloud Platform for Microsoft Azure:

- Customers benefit from a single source of system-level support, which includes support for both hardware and software, including Microsoft Azure Stack HCI operating system, regardless of where the customer purchases their software. At times when collaboration is required with the Microsoft support team, Dell handles the interaction, allowing our customers to handle support with just one vendor.

IMPORTANT NOTES:

Customers must also have a support contract with Microsoft. However, we encourage customers to make Dell their initial point of contact for any issues, including Microsoft Azure Stack HCI software.

For APEX Cloud Platform for Microsoft Azure, we offer collaborative support for Microsoft HCI software on both ProSupport and ProSupport Plus. This is different to most products, where limited third-party support is only available with ProSupport Plus.

Dell Consulting Services

Dell consulting services can help you navigate through the complexities of a multicloud implementation.

- Dell consulting services help companies develop a multicloud strategy that identifies your organizations' needs and the mix of applications to support.
- Consulting services help you implement a cloud operating model and deliver developer, resiliency, and Azure Kubernetes services.

- And to fully adopt a multicloud strategy, Dell can help you accelerate and de-risk migration of applications, workloads, and data or to modify existing applications to become cloud native.
- When it comes to cloud scale, we help you transform data into actionable insights and implement processes to measure key success criteria for business stakeholders.

The following chart shows the extensive list of services and outcomes that Dell services can deliver to help you achieve a sound multicloud by design. Items boxed in red are mandatory for Dell APEX Cloud Platform for Microsoft Azure:

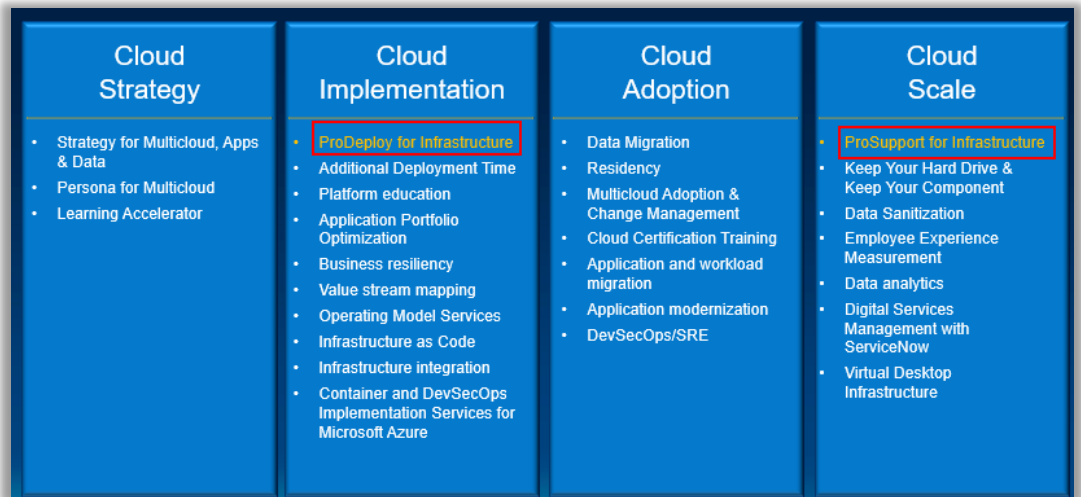


Figure 49. Dell Services for Dell APEX Cloud Platform for Microsoft Azure

Dell Professional Services

Dell also provides optional support services that are often required while managing multicloud services. These include:

- Residency
- Onsite Diagnosis
- Data Sanitization
- Data Destruction
- Keep Your Hard Drive
- Keep your Component

Dell also provides a variety of **consumption options** that give customers the flexibility they need to support their financial strategies.

For customers who prefer to purchase and own their infrastructure, **Dell Financial** services provides different business-class payment solutions to fit varied customer needs.

Dell APEX Custom is ideal for customers who prefer a Pay-Per-Use offer that can be logged as an Operating Expense.

Flex-on-Demand delivers consumption-based infrastructure with the elastic capacity required for peak periods and is ideal for a large number of customers.

Data Center Utility provides a customer the opportunity to build the unique pay-for-use solution that fits their organization and drives the IT business and financial outcomes they require. Customers can highly customize their solution by utilizing single or multiple billing metrics that align to their data, workloads, and use cases.

For more information about Dell multicloud capabilities visit this [Dell Services site](#).

The Solution sizing and Configuration

With APEX Cloud Platforms, customers do not have to spend valuable time with architecture and engineering planning to run workloads. The common infrastructure building blocks are already tested and validated with a broad range of configurations ready to meet the requirements of even the most demanding applications. However, to determine the right-sized system that achieves the best return on investment, it is essential to understand the performance and capacity characteristics of your existing environment.

Dell Technologies follows a consultative and systematic approach to help customers modernize their data estate. Dell Live Optics is a free, online software tool used to collect, visualize, and share data about existing IT infrastructures and workloads.

Optical Prime is an agent that gets installed on-premises that collects metrics like peak CPU utilization, IOPS, used storage, and average daily writes from various operating systems. Optical Prime also comes bundled with workload collectors to gather critical information directly from the application layer.

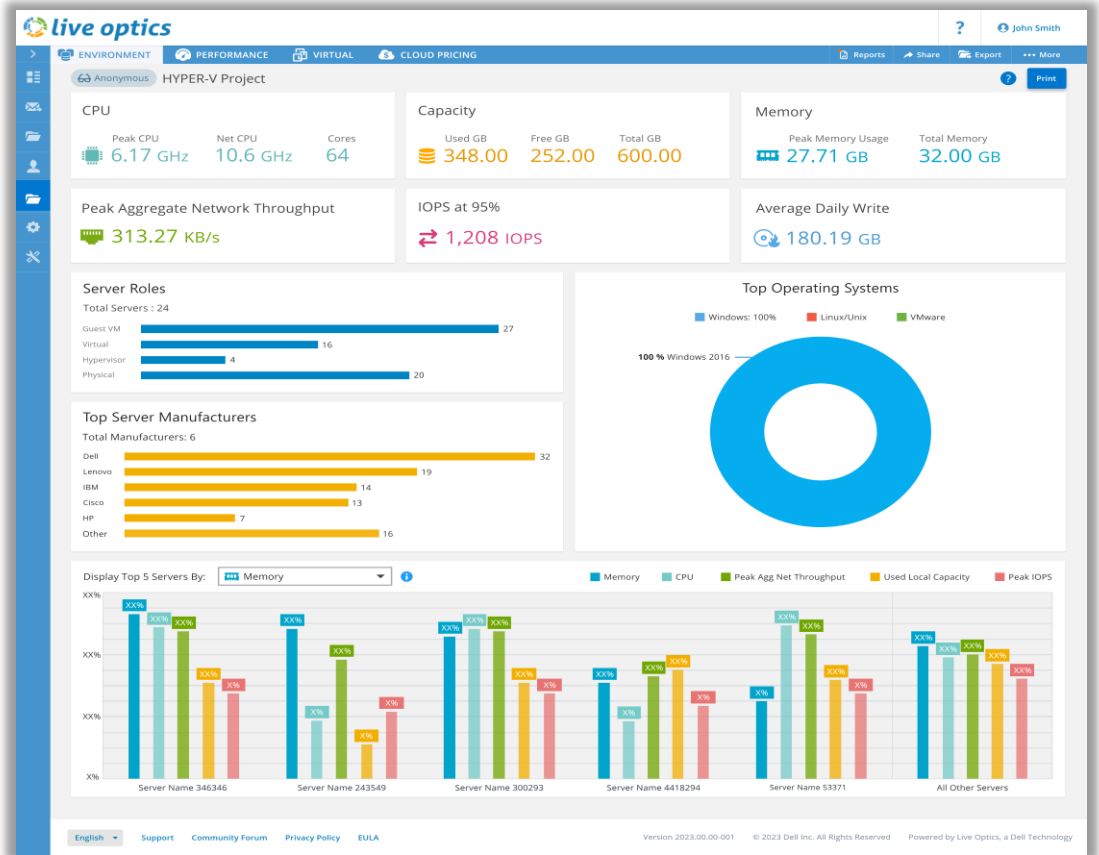


Figure 50. Live Optics free tool gathers critical information about the running workloads to help right sizing the new Dell APEX Cloud Platform for Microsoft Azure

As explained in previous chapters, Dell APEX Cloud Platforms inherit key Dell technologies that have proven successful and effective. That is also the case of the sizing tool, [Dell PowerSizer](#). This tool, internal to Dell and its partners, is able to generate best performing and cost optimal configurations for Dell APEX Cloud Platforms based in MC-660 and 760 nodes. It can create full NVMe, All-Flash, and hybrid configurations of any sizing scenario allowing to save multiple revisions of them.

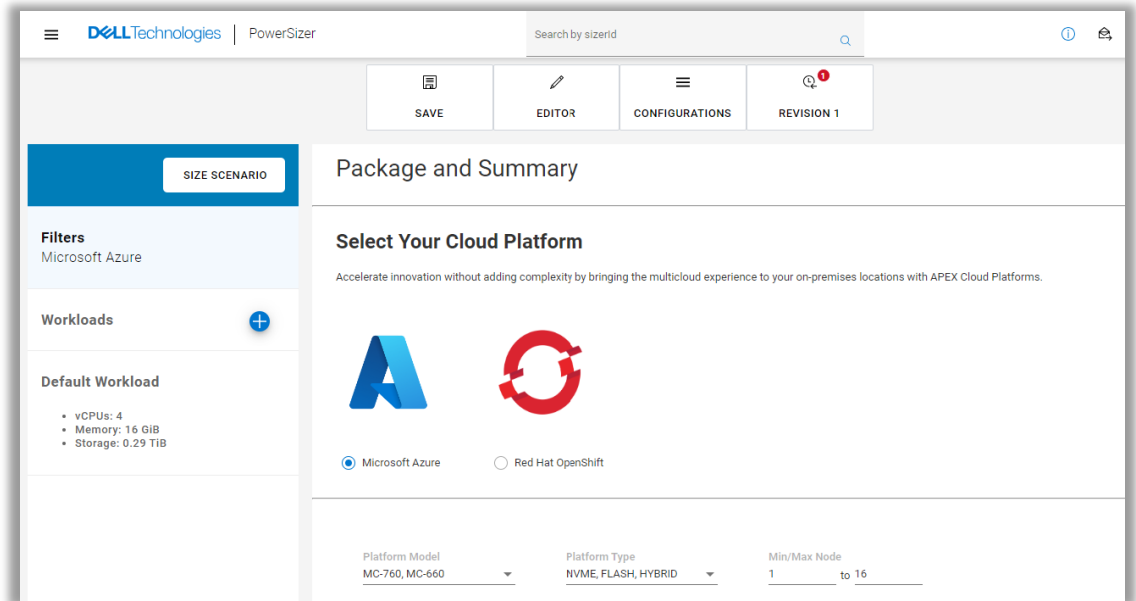


Figure 51. Dell internal (and partners) sizing tool for Dell APEX Cloud Platforms

For these APEX Cloud Platforms, this sizing tool inherits years of sizing experiences and Intellectual Property of Dell leading HCI platform, VxRail.

The sizing workflow follows a series of steps that end up with an optimal quote for that meets all the performance and capacity requirements, with the lower cost:

- Define solution requirements (ideally influenced by the data gathered by Live Optics)
- Configure systems that conform the solution in terms of CPU, memory, storage, network and GPUs.
- Evaluate the systems proposed and workloads characterized and customize the solution as needed.
- Create the customer quote.




Figure 52. Dell APEX Cloud Platform sizing workflow overview

The output from the tool ensures the configurations generated are balanced and viable. This output includes a minimum and recommended cluster configuration that meets the workloads' requirements in terms of performance, capacity, resiliency, and future growth within the scope of the project's budget.

The screenshot shows the Dell PowerSizer interface. At the top, there's a search bar and navigation tabs for 'SAVE', 'EDITOR', 'CONFIGURATIONS', and 'REVISION 1'. The main content area is titled 'OVERVIEW' and 'Configuration Option'. Below this, there's a table with columns for 'Item', 'Specifications', 'Node', and 'Cluster'. The table lists various hardware specifications for a cluster configuration.

| Item | Specifications | Node | Cluster |
|---------|----------------------|--------------------------|--|
| CPU | 6442Y (24C, 2.6 GHz) | 2 (48 Cores - 124.8 GHz) | 22 (528 Cores - 1372.8 GHz) |
| Memory | 32 GiB RDIMMs | 16 (512 GiB) | 176 (5632 GiB) |
| Storage | 0.96 TB NVME | 4 (3.84 TiB) | 44 (Usable: 11.53 TiB, Raw: 38.42 TiB) |
| IOPS | Max IOPS | 70160 | 771760 |

Below the table, the 'Sizer ID' is listed as 'azp_1003'. The interface also shows a 'Cheapest Option' dropdown menu.

Figure 53. Sample proposed configuration for Dell APEX Cloud Platform for Microsoft Azure, based in 11 MC-660 All NVMe nodes, with Three-way mirror, and 512 GB of RAM per node.

Dell APEX Cloud Platform for Microsoft Azure Use Cases

Virtualized Applications

Well into the 21st century, for a platform such as Dell APEX Cloud Platform for Microsoft Azure, is still very relevant the ability to provide Infrastructure as a Service in an optimal manner. Whether it is about provision VMs or containers, application still need a basic infrastructure support to run upon.

Dell APEX Cloud Platform for Microsoft Azure's architecture allows users to provision VMs in a modern, predictable manner.

Azure Arc resource bridge (in preview) is a VM that runs on your APEX Cloud Platform and hosts a management Kubernetes cluster. It supports VM self-service provisioning and management onto on-premises Azure Stack HCI-based platforms through the Azure portal and programmatically using ARM templates and Bicep files.

Resource bridge enables you to deploy Windows and Linux VMs from the extensive Azure Marketplace. There are many VM images supported in the marketplace with a wide variety of applications pre-installed and ready to go. One such application might be SQL Server.

These VMs that are provisioned in this manner get onboarded to Azure Arc through the installation of the Connected Machine Agent. If you have provisioning SQL Server VMs from the marketplace, they will receive those extra benefits available to SQL Servers.

For those VMs that are still running on existing infrastructure, Azure Migrate has a functionality (in preview) to allow migrating VMs from aging Hyper-V environments to Azure Stack HCI. The process is very agile, as no VM data is transferred to Azure, only VM metadata, and there is no need for guest agent for the operation to complete. The migration process involves downtime, and you have the ability to test your migrations before proceeding.

The architectural diagram is shown in the following figure:

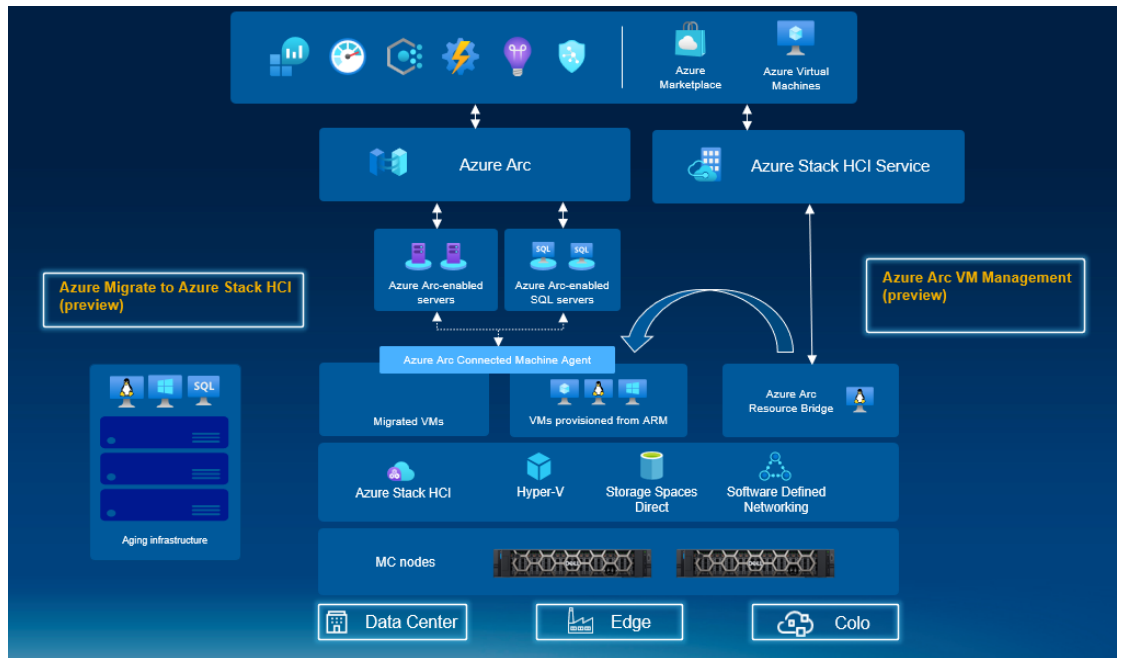


Figure 54. VM self-service provisioning in Dell APEX Cloud Platform for Microsoft Azure

DBaaS use-case

Of particular benefit result deploying and operating SQL server databases and SQL managed instances, and for many customers, databases constitute one of the most valuable applications in their business portfolio. Today, when many customers implement a hybrid approach to their cloud operating model, they must also consider a similar hybrid approach for their database deployments. Data can reside on-premises or in the cloud depending on specific dataset needs. This rich feature usually represents an important obstacle for IT administrators, now obliged to operate and maintain a complex and advanced architecture that can surpass their capabilities.

The challenges found in these hybrid environments are of a varied nature. They include keeping a consistent control version on database engines and their associated security updates or operating a diverse database toolkit between on-premises and cloud instances. If we add infrastructure inflexibility and variable infrastructure configuration needs to cope with database performance requirements, a heavy burden for many IT admins is the result.

In this already complex scenario, business demands can make it even harder. IT might be required to deploy the latest version of a database engine to get the most from its new features and performance, but simultaneously keep running legacy databases as they are key in the company’s strategy.

A smart way to respond to these challenges might be to adapt a Database-as-a-Service (DBaaS) platform. A DBaaS platform enables creating modern, cloud-native applications, minimizing the operational burden described previously, and optimizing cross-functional productivity for DBAs, IT admins, and software developers. Microsoft and Dell Technologies have created a modular, integrated solution that:

- Simplifies overall life cycle management
- Automates common administrative tasks

- Increases platform resiliency
- Ensures performance to meet defined SLAs

Dell APEX Cloud Platform for Microsoft Azure provides a consistent platform to implement a DBaaS architecture. By providing the same Azure experience to data centers and edge locations, simplifying operations, and alleviating the cost of application development, Dell APEX Cloud Platform for Microsoft Azure enables operations automation, life cycle management for the full system stack, as well as flexible consumption models with enterprise-level support and services.

Hybrid DBaaS with Dell APEX Cloud Platform for Microsoft Azure helps IT and database admins, and software developers by providing:

- Self-service database provisioning
- Elastic resource scalability
- Reduction in operational and management costs
- Flexible consumption models
- Automation of common administrative tasks
- Platform resiliency and performance to meet defined SLAs

The backend infrastructure supporting Dell APEX Cloud Platform for Microsoft Azure is designed to provide node-based linear scalability, and the transactional and throughput performance required by DBaaS deployments.

Cloud Native with Azure Kubernetes Service

Today, most companies have a varied application portfolio. A mix of legacy virtualized applications combined with cloud native developments. To extract the maximum value of a platform such as Dell APEX Cloud Platform for Microsoft Azure, unleashing full application value, we need to leverage Azure Kubernetes Service (AKS). This way, developers get a consistent experience with the hyperscale optimized and fully managed Azure Kubernetes services.

There are two ways to create AKS clusters:

- Using PowerShell
- Through Windows Admin Center

When deployed, you can choose default control plane nodes and Kubernetes cluster configuration options. Other advanced settings (monitoring, Microsoft Entra ID, and others) can be configured during and after the deployment process.

When AKS is deployed, you can connect your Kubernetes clusters to Azure. Once connected to Azure Arc-enabled Kubernetes, you can access your on-premises Kubernetes clusters through the Azure portal and deploy management services such as GitOps, Microsoft Defender for Kubernetes, Azure Monitor Container Insights, and Azure Policy for Kubernetes.

Once connected to Azure Arc, you can use Azure Monitor for:

- Health monitoring of your Kubernetes clusters and workloads.

- providing performance visibility by collecting memory and processor metrics from controllers, nodes, and containers. Metrics and container logs are automatically collected and to the metrics database in Azure Monitor.

AKS can host Linux and Windows-based containers. Upon Kubernetes cluster creation on top of the Azure Stack HCI infrastructure, we can choose to create node pools to run Linux, Windows, or both type of containers.

As a hosted Kubernetes service, AKS handles critical day-to-day management, such as upgrades and automatic certificate rotations, so you can focus on running and developing containerized workloads.

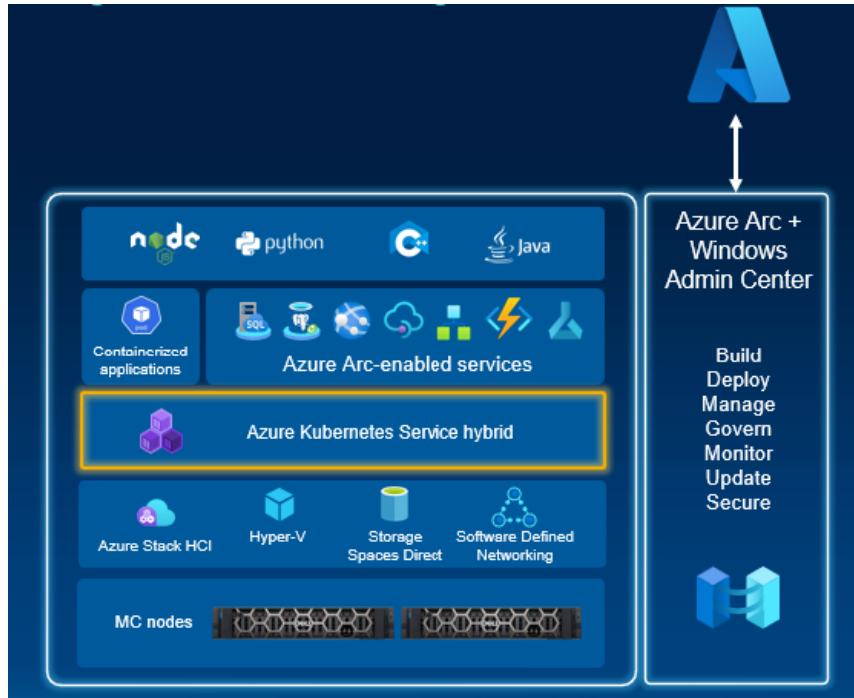


Figure 55. AKS hybrid architecture with Dell APEX Cloud Platform for Microsoft Azure

The management cluster, also known as the service host, gets deployed as a single Linux-based VM running the Azure Linux distribution. The management cluster is responsible for deploying and managing multiple workload clusters and it includes the following components:

- **API Server.** Interacts with the management tools like WAC and Azure Arc
- **Load balancer.** Manages load-balancing rules for the API server of the management cluster.

Then we can create workload clusters. This creates the highly available control plane components and worker node components all in the form of Windows or Linux-based VMs depending on what kind of containers you intend to build. Containerized applications run on a workload cluster. To achieve application isolation, you can deploy up to eight workload clusters. The workload cluster consists of the following components:

- **Control plane.** Runs on a Linux distribution and contains API server components for interaction with Kubernetes API and a distributed key-value store for storing all the configuration and data of the cluster.

- **Load balancer.** Runs on a Linux VM and provides load-balanced services for the workload cluster.
- **Worker nodes.** Run on a Windows or Linux operating system that hosts containerized applications.
- **Kubernetes resources.** Pods represent a single instance of your application, that usually have a one to one mapping with a container, but certain pods can contain multiple containers. Deployments represent one or more identical pods. Pods and deployments are logically grouped into a namespace that controls access to management of the resources.

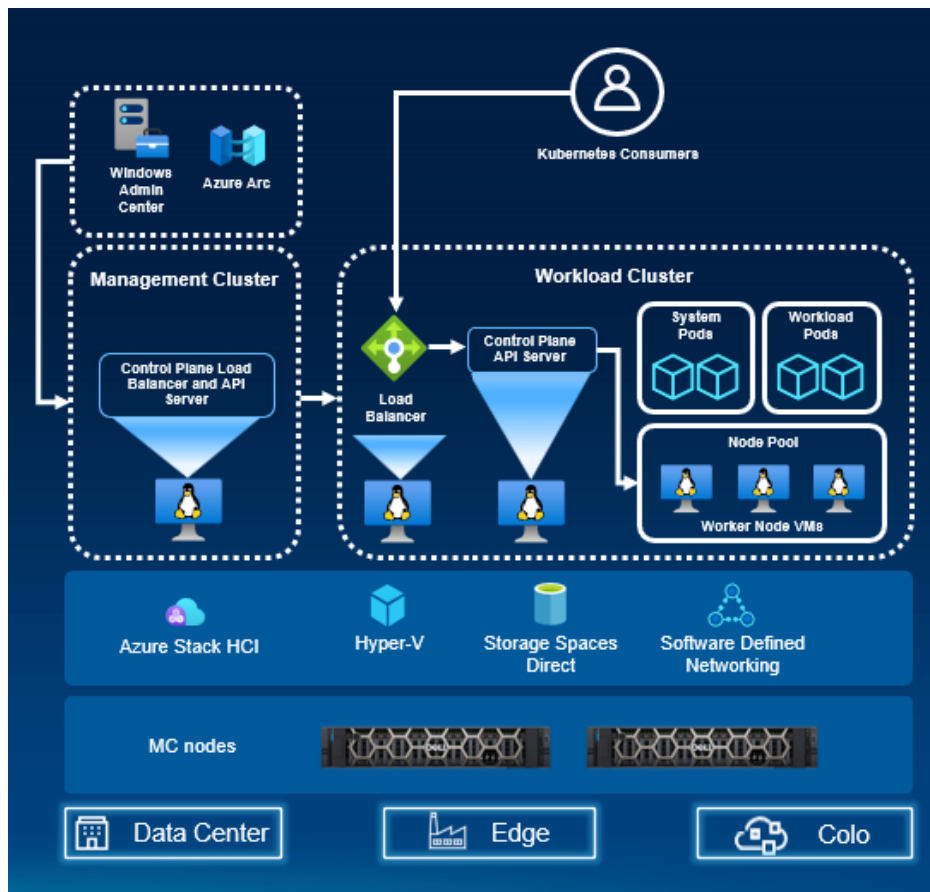


Figure 56. Default AKS hybrid components

For more details on the Kubernetes clusters architecture and components, visit this [Microsoft page](#).

Azure Virtual Desktop (AVD)

Modern organizations are striving to deal with an increasingly complex environment – employees and students must be able to work and learn from anywhere. There is an ever-increasing organizational imperative to do more with less, all of which is happening in a technological environment where cybersecurity vulnerabilities are being exposed at a rapidly increasing pace, resulting in significant financial and reputational risk to businesses. In this environment, desktop and application virtualization is attracting a lot of interest. It is highly secure (because it keeps user data in the data center) and is a true

“do anything from anywhere” powerhouse, allowing users to connect to their desktops and applications using endpoint devices with multiple form factors and running many operating systems.

Azure Virtual Desktop (AVD) is a desktop and application virtualization solution that runs in Azure public cloud and on Azure Stack HCI OS-based private cloud environments. It provides IT with strong, granular controls and a straightforward management experience, while also providing a rich experience to Windows and Microsoft 365 users. As an Azure Stack HCI OS-based platform, Dell APEX Cloud Platform for Microsoft Azure is an ideal platform for AVD-based virtual desktop environments. It allows organizations to benefit from the improved performance and compliance with data sovereignty requirements that are delivered by an on-premises environment while also providing a streamlined, cloud-consistent experience through the same rich comprehensive management plane in Azure.

Organizations that want to maximize the value of their hybrid AVD deployments need to choose an on-premises environment that complements and enhances AVD. Dell APEX Cloud Platform for Microsoft Azure has many capabilities and benefits that are of significant benefit to AVD deployments:

- Ability to move Azure workloads between public and private clouds as needed, allowing virtual desktop users to be deployed on Dell APEX Cloud Platform for Microsoft Azure or on an Azure hybrid cloud depending on organization-specific requirements
- Intrinsic security throughout the entire integrated hardware and software stack, reinforcing the inherent virtual desktop security provided by AVD
- Full stack life cycle management across the entire stack, ensuring that organizations can quickly and seamlessly ensure that the on-premises platform hosting their AVD environment has the latest security, functionality, and performance enhancements
- Broad choice of 1U (MC-660) and 2U (MC-760) rackmount servers with various CPU, memory, storage, and GPU configurations, providing true linear scalability, ensuring that organizations can granularly scale their Dell APEX Cloud Platform for Microsoft Azure environment to deliver the appropriate number of desktops with the precise level of performance

Virtual desktop solutions such as AVD can deliver transformational benefits for organizations. However, these benefits can only be fully captured when the virtual desktops run on a platform that can enable the full potential of these virtual desktop technologies. The unique combination of capabilities and benefits delivered by Dell APEX Cloud Platform for Microsoft Azure positions it perfectly as that virtual desktop platform of choice.

Terminology

The following table provides definitions for some of the terms that are used in this document.

Table 2. Terminology

| Term | Definition |
|-----------------------------|---|
| Multicloud | Companies' strategy to operate business workloads in different cloud providers |
| APEX | Dell APEX is our end-to-end solutions portfolio offered through an as-a-Service and subscription model, where customers pay only for the services consumed. |
| Monolithic applications | Monolithic means Composed all in one piece. Monolithic applications are single tiered, which means multiple components are combined into one large application. |
| Cloud native | Cloud-native architecture and technologies are an approach to designing, constructing, and operating workloads that are built in the cloud and take full advantage of the cloud computing model. |
| Hyperconverged | Hyper-converged infrastructure (HCI) is a software-defined IT infrastructure that virtualizes all of the elements of conventional "hardware-defined" systems. HCI includes, at a minimum, virtualized computing (a hypervisor), software-defined storage, and virtualized networking (software-defined networking). |
| PowerFlex | Software-defined storage solution from Dell Technologies that creates a server-based storage area network from local server storage using x86 servers. |
| RESTful APIs | Software architectural style created to guide the design and development of the architecture of the World Wide Web. It emphasizes uniform interfaces, independent deployment of components, and the scalability of interactions between them. REST API calls can be used from almost any modern programming language through HTTP methods to communicate with the web server (GET, PUT, POST, DELETE, and PATCH). |
| GPU | Specialized processor originally designed to accelerate graphics rendering |
| Azure Arc | Centralized Microsoft Azure console to manage and operate cloud and on-premises Azure assets |
| | |
| Kubernetes | Open-source container orchestration system for automating software deployment, scaling, and management. |
| IPv6 | Most recent version of the Internet Protocol designed to address the limitations of its predecessor, IPv4 |
| Fleet management | Service that allows an easy way of managing and operating multiple virtual machines and or containers |
| Storage Spaces Direct (S2D) | Azure Stack HCI feature that enables cluster servers' internal storage as a shared software define storage |

References

Dell Technologies documentation

The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [*Main product page with spec sheets, solution briefs, infographics, and other great collateral*](#)
- [*Infohub for Dell APEX Cloud Platforms*](#)
- [*VDI Design Guide—Azure Virtual Desktop on Dell Integrated System for Microsoft Azure Stack HCI*](#)
- [*Using Dell ProDeploy Plus Infrastructure can improve deployment times for Dell technology*](#)
- [*Dell multicloud services*](#)
- [*YouTube playlist with educational and demo videos – Product Launch*](#)
- [*YouTube playlist with March's update – Part I*](#)
- [*Dell APEX Cloud Platform for Microsoft Azure interactive demo, ITD-0235*](#)

Microsoft documentation

The following Microsoft documentation.

- [*Azure Virtual Desktop for Azure Stack HCI \(preview\)*](#)
- [*Azure Stack HCI physical network requirements*](#)
- [*Windows 2022 Server catalog- Network*](#)
- [*Host network requirements for Azure Stack HCI*](#)
- [*Kubernetes cluster architecture*](#)
- [*Microsoft Azure Stack HCI solutions catalog*](#)