

Microsoft HCI Solutions from Dell Technologies:

Managing and Monitoring the Solution Infrastructure
Life Cycle
Operations Guide

Dell Technologies Solutions

Notes, cautions, and warnings

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

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

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

Chapter 1: Introduction.....	4
Document scope.....	4
Audience and assumptions.....	4
Known issues.....	4
Microsoft HCI Solutions from Dell Technologies overview.....	4
Deployment guidance.....	6
 Chapter 2: Day 0 Operations.....	 7
Azure onboarding for Azure Stack HCI OS.....	7
Licensing for Azure Stack HCI for Windows Server 2016 and 2019.....	7
Creating virtual disks.....	7
Managing and Monitoring Azure Stack HCI Cluster using Windows Admin Center.....	8
Installing Windows Admin Center.....	8
Adding the HCI cluster connection.....	9
Accessing the HCI cluster.....	10
Viewing server details.....	10
Viewing drive details.....	11
Managing and monitoring volumes.....	12
Enabling data deduplication on Storage Spaces Direct.....	14
Monitoring and managing VMs.....	14
Managing virtual switches.....	16
Dell EMC OpenManage Integration with Windows Admin Center.....	17
Firmware updates using Dell EMC OpenManage Integration for Microsoft System Center for System Center Virtual Machine Manager.....	26
Firmware and driver updates using the manual method.....	29
Restarting a cluster node or taking a cluster node offline.....	33
Expanding the Azure Stack HCI cluster.....	33
Extending volumes.....	35
Performing AX node recovery.....	36
Operating system recovery.....	38

Introduction

Topics:

- [Document scope](#)
- [Audience and assumptions](#)
- [Known issues](#)
- [Microsoft HCI Solutions from Dell Technologies overview](#)
- [Deployment guidance](#)

Document scope

This operations guide focuses on operational aspects of a hyperconverged infrastructure solution on Azure Stack HCI with Hyper-V and Storage Spaces Direct.

This guide includes an overview of Microsoft HCI Solutions from Dell Technologies, guidance to monitor and manage bare metal, and instructions for performing operations on an Azure Stack HCI cluster and updating the cluster-aware system. This guide is applicable only to infrastructure that is built by using the validated and certified Microsoft HCI Solutions from Dell Technologies.

Microsoft HCI Solutions from Dell Technologies refers to:

- Dell EMC Integrated System for Azure Stack HCI (based on Azure Stack HCI OS v20H2)
- Dell EMC HCI Solutions for Microsoft Windows Server (based on Windows Server 2016/2019 OS)

Instructions in this guide are applicable only to the generally available operating system build of Windows Server 2016, Windows Server 2019, and Azure Stack HCI operating system with the latest applicable updates. These instructions are not validated with Windows Server version 1709. Microsoft HCI Solutions from Dell Technologies do not support the Windows Server Semi-Annual Channel release. Dell Technologies recommends updating the host operating system with the latest cumulative updates from Microsoft before starting the cluster creation and configuration tasks.

Audience and assumptions

The audience for this operations guide includes systems administrators, systems engineers, field consultants, partner engineering team members, and customers with a fair amount of knowledge in deploying hyperconverged infrastructures with Windows Server 2016, Windows Server 2019, the Azure Stack HCI operating system, Hyper-V, and Storage Spaces Direct. We assume that deployment personnel have prerequisite knowledge, including of:

- Microsoft HCI Solutions from Dell Technologies, and the deployment and configuration of BIOS and integrated Dell Remote Access Controller (iDRAC) settings on AX nodes from Dell Technologies
- Deploying and configuring Windows Server core operating system Hyper-V infrastructure

Known issues

Before starting the cluster deployment, see [Dell EMC Solutions for Microsoft Azure Stack HCI - Known Issues](#) for known issues and workarounds.

Microsoft HCI Solutions from Dell Technologies overview

Microsoft HCI Solutions from Dell Technologies encompass various configurations of AX nodes from Dell Technologies to power the primary compute cluster that is deployed as a hyperconverged infrastructure. This hyperconverged infrastructure that is

built by using these AX nodes uses a flexible solution architecture rather than a fixed component design. The following figure illustrates one of the flexible solution architectures. It consists of a compute cluster alongside the redundant top-of-rack (ToR) switches, a separate out-of-band network, and an existing management infrastructure in the data center.

NOTE: Microsoft HCI Solutions from Dell Technologies are available in both hybrid and all-flash configurations. For more information about available configurations, see the [solution overview](#).

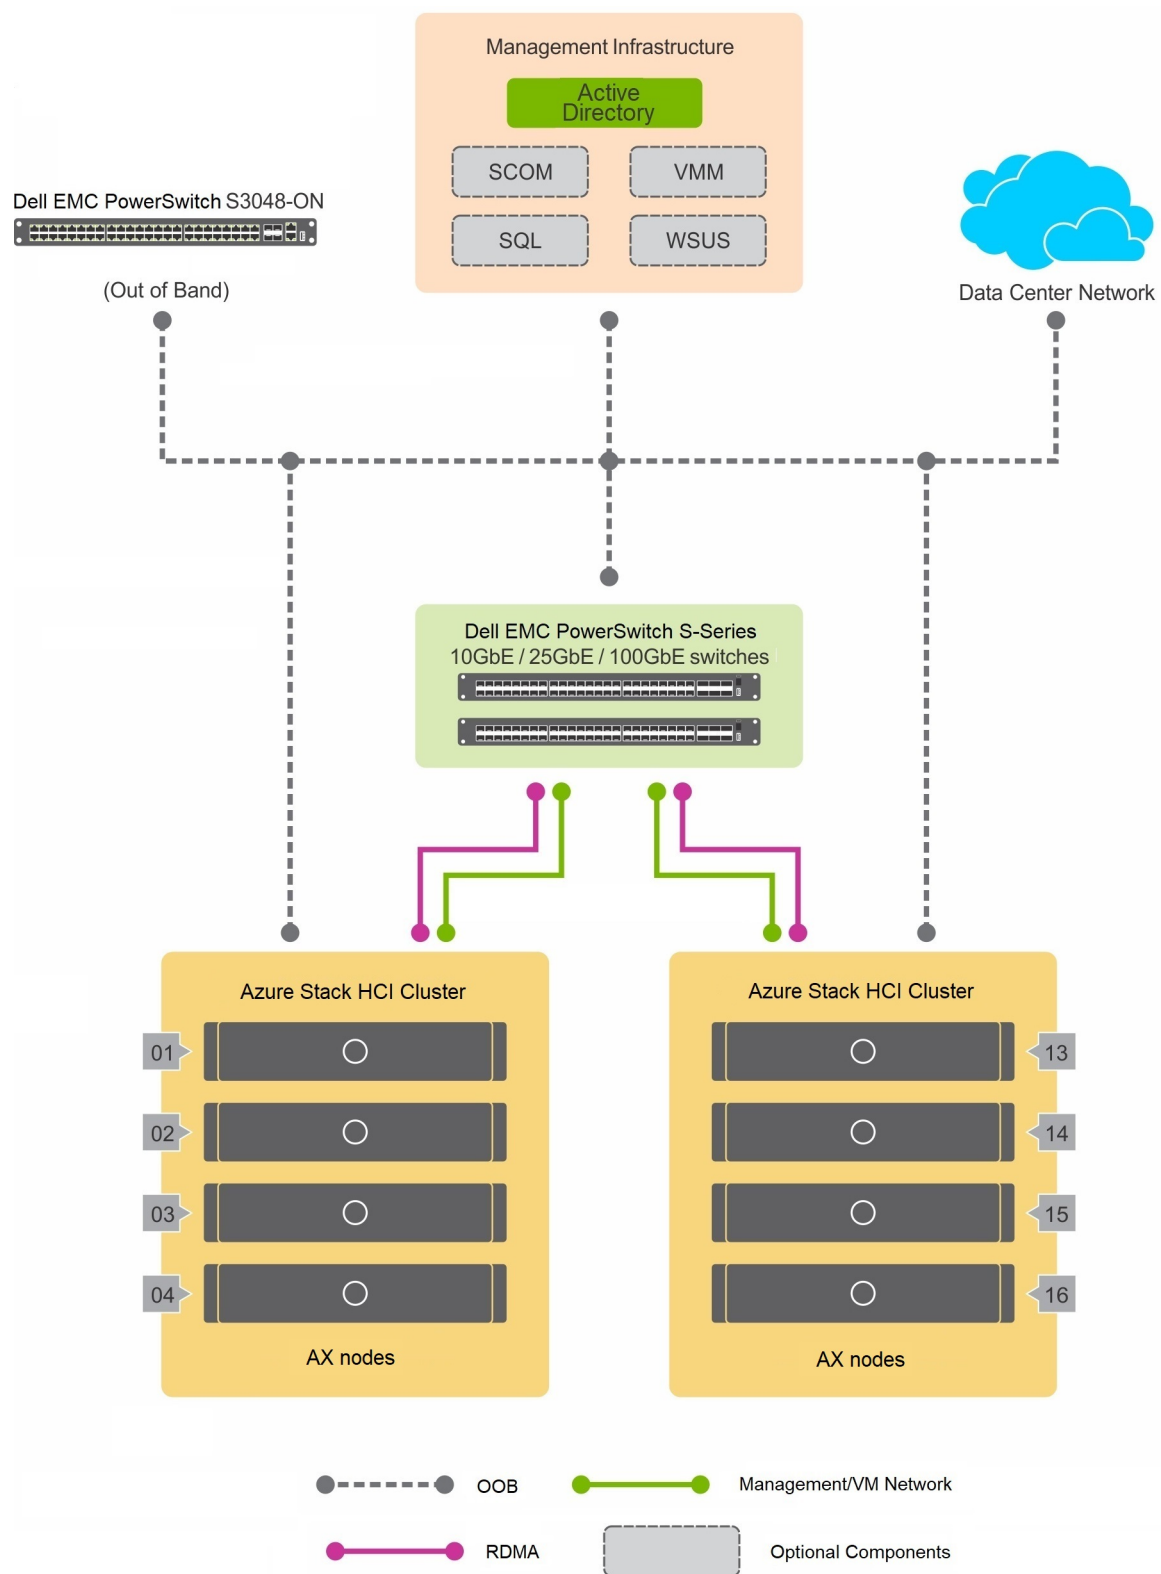


Figure 1. Hyperconverged virtualized solution using AX nodes

Deployment guidance

For deployment guidance and instructions for configuring a cluster using Dell EMC Solutions for Azure Stack HCI, see [Microsoft HCI Solutions from Dell Technologies](#). This operations guidance is applicable only to cluster infrastructure that is built using the instructions provided in the deployment documentation for AX nodes.

Day 0 Operations

After deploying the Azure Stack HCI cluster, complete day 0 operations.

Topics:


- [Azure onboarding for Azure Stack HCI OS](#)
- [Licensing for Azure Stack HCI for Windows Server 2016 and 2019](#)
- [Creating virtual disks](#)
- [Managing and Monitoring Azure Stack HCI Cluster using Windows Admin Center](#)

Azure onboarding for Azure Stack HCI OS

Clusters deployed using Azure Stack HCI OS must be onboarded to Microsoft Azure for full functionality and support. For more information, see [Connect Azure Stack HCI to Azure](#).

Licensing for Azure Stack HCI for Windows Server 2016 and 2019

When the server operating system is installed using the retail or volume licensing media, the operating system license must be activated. On the Server Core operating system, activate the license by using either the Server Configuration tool (`sconfig` command) or the `slmgr` command.

 **NOTE:** Windows activation is not required if the operating system is factory installed.

To activate the operating system license by using `slmgr`, see [Slmgr.vbs Options for Volume Activation](#).

To activate the operating system license by using the `sconfig` command, see [Configure a Server Core installation of Windows Server 2016 or Windows Server, version 1709, with Sconfig.cmd](#).

For volume activation of the Windows operating system in the data center, see the Microsoft documentation for using the Key Management Service (KMS).


Creating virtual disks

Cluster creation and enabling Storage Spaces Direct configuration on the cluster creates only a storage pool and does not provision any virtual disks in the storage pool. Use the `New-Volume` cmdlet to provision new virtual disks as the cluster shared volumes for the Azure Stack HCI cluster.

When creating volumes in the Azure Stack HCI cluster infrastructure:

- Ensure that you create multiple volumes—a multiple of the number of servers in the cluster. For optimal performance, each cluster node should own at least one virtual disk volume. Virtual machines (VMs) on each volume will perform optimally when running on the volume owner node.
- Limit the number of volumes in the cluster to 32 on Windows Server 2016 and 64 on Windows Server 2019 and the Azure Stack HCI operating system.
- Ensure that the storage pool has enough reserve capacity for any in-place volume repairs arising out of failed disk replacement. The reserved capacity should be at least equivalent to the size of one capacity drive per server and up to four drives.

For general guidance about planning volume creation, see [Planning volumes in Storage Spaces Direct](#).

 **NOTE:** We recommend that you use the following resiliency settings when you create virtual disks:

- On Windows Server 2016, Windows Server 2019, and the Azure Stack HCI operating system clusters with three or more nodes—Three-way mirror
- On Windows Server 2019 and Azure Stack HCI operating system clusters with four or more nodes—Three-way mirror or mirror-accelerated parity

Managing and Monitoring Azure Stack HCI Cluster using Windows Admin Center

Windows Admin Center is a browser-based management tool developed by Microsoft to monitor and manage Windows servers, failover clusters, and hyperconverged clusters.

The AX nodes for Storage Spaces Direct offer software-defined storage building blocks for creating highly available and highly scalable hyperconverged Infrastructure (HCI). The AX nodes are preconfigured with certified components and validated as a Storage Spaces Direct solution that includes Dell EMC PowerSwitch S-Series switches, with simplified ordering and reduced deployment risks. Dell Technologies offers configuration options within these building blocks to meet different capacity and performance points. With Windows Admin Center, you can seamlessly monitor and manage the HCI clusters that are created on these building blocks.

Installing Windows Admin Center

You can download Windows Admin Center version 2103 from [Microsoft download center](#) and install it on Windows 10, Windows Server 2016, Windows Server 2019, or Windows Server version 1709. You can install Windows Admin Center directly on a managed node to manage itself. You can also install Windows Admin Center on other nodes in the infrastructure or on a separate management station to manage the AX nodes remotely. It is possible to implement high availability for Windows Admin Center by using failover clustering. When Windows Admin Center is deployed on nodes in a failover cluster, it acts as an active/passive cluster providing a highly available Windows Admin Center instance.

The Windows Admin Center installer wizard performs the configuration tasks that are required for Windows Admin Center functionality. These tasks include creating a self-signed certificate and configuring trusted hosts for remote node access. Optionally, you can supply the certificate thumbprint that is already present in the target node local certificate store. By default, Windows Admin Center listens on port 443 (you can change the port during the installation process).

NOTE: The automatically generated self-signed certificate expires in 60 days. Ensure that you use a certificate authority (CA)-provided SSL certificate if you intend to use Windows Admin Center in a production environment.

For complete guidance about installing Windows Admin Center on Windows Server 2016 and Windows Server 2019 with desktop experience or Server Core, see [Install Windows Admin Center](#).

NOTE: This section assumes that you have deployed the Azure Stack HCI cluster from Dell Technologies by using the deployment guidance that is available at: <https://dell.com/azurestackhcimanuals>.

After the installation is complete, you can access Windows Admin Center at `https://managementstationname:<PortNumber>`.




Figure 2. Windows Admin Center start screen

Adding the HCI cluster connection

About this task

For monitoring and management purposes, add the hyperconverged cluster that is based on Dell EMC Solutions for Azure Stack HCI as a connection in Windows Admin Center.

Steps

1. Go to **Windows Admin Center > Cluster Manager**, as shown in the following figure.

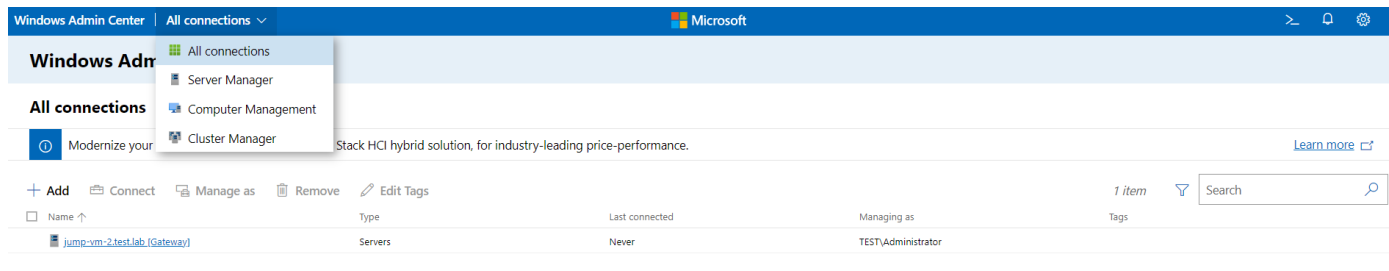


Figure 3. HCI cluster navigation

2. Click **Add**.
The **Add Cluster** window is displayed.
3. Enter the cluster FQDN and select **Also add servers in the cluster**, as shown in the following figure.

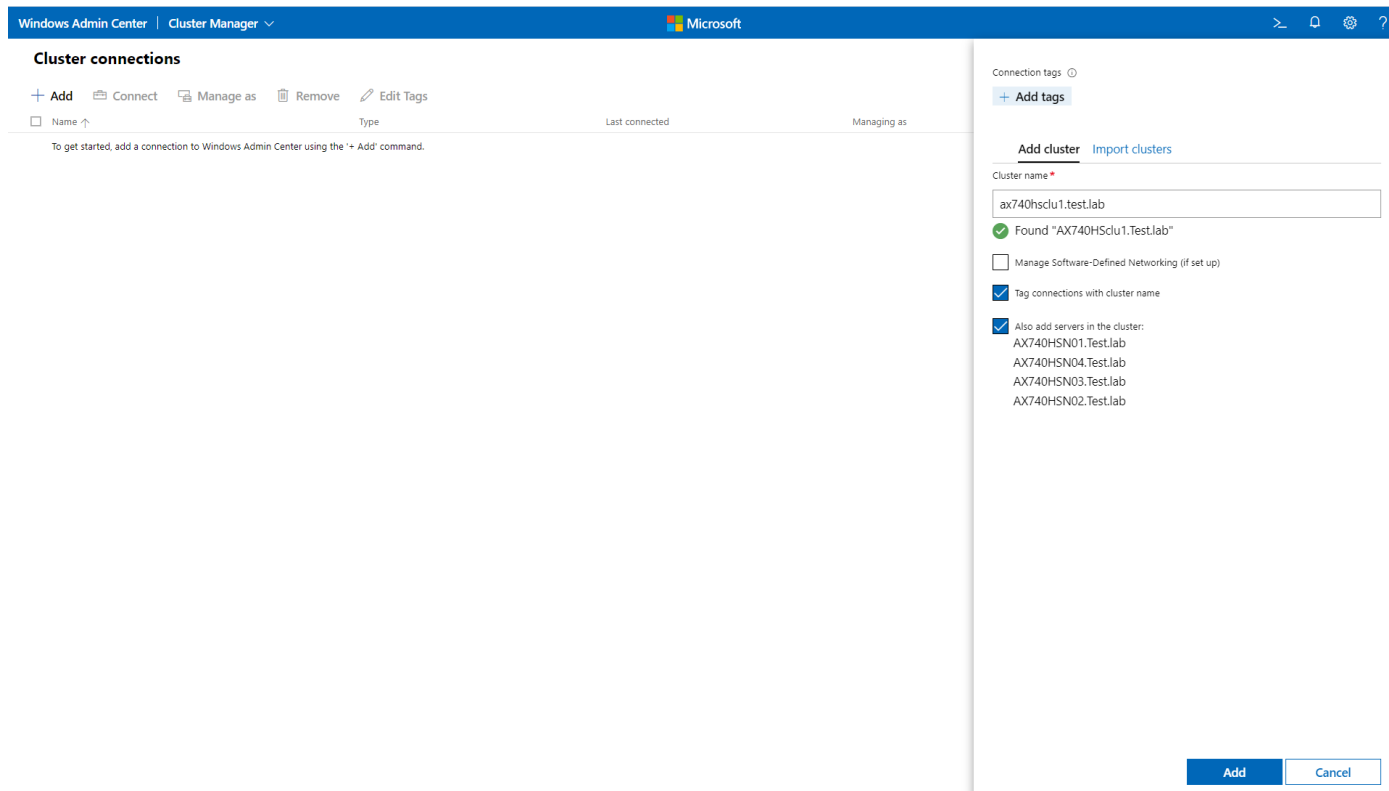


Figure 4. Adding the HCI cluster

Windows Admin Center discovers the cluster and the nodes that are part of the cluster.

4. Click **Add**.
The cluster is added to the connection list and Windows Admin Center is configured to monitor and manage the HCI cluster.

Accessing the HCI cluster

To view the dashboard for the HCI cluster that you have added to Windows Admin Center, in the **Cluster Manager** window, click the cluster name.

This dashboard provides the real-time performance view from the HCI cluster. This view includes total IOPS, average latency values, throughput achieved, average CPU usage, memory usage, and storage usage from all cluster nodes. It also provides a summarized view of the Azure Stack HCI cluster with drives, volumes, and VM health.

You can examine an alert by clicking the alerts tile in the dashboard.

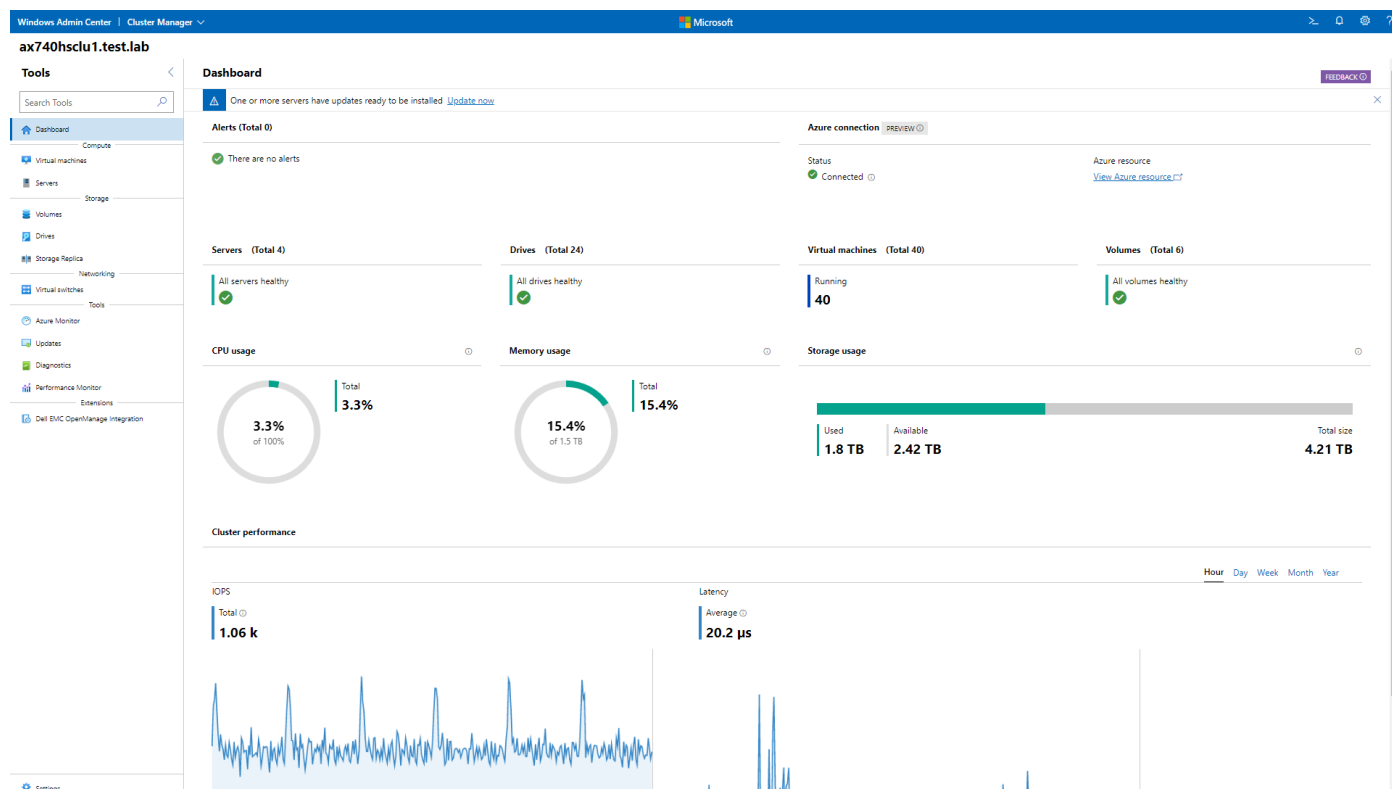


Figure 5. HCI dashboard in Windows Admin Center

Viewing server details

To view the server details, click the tools pane and go to **Servers > Inventory**.

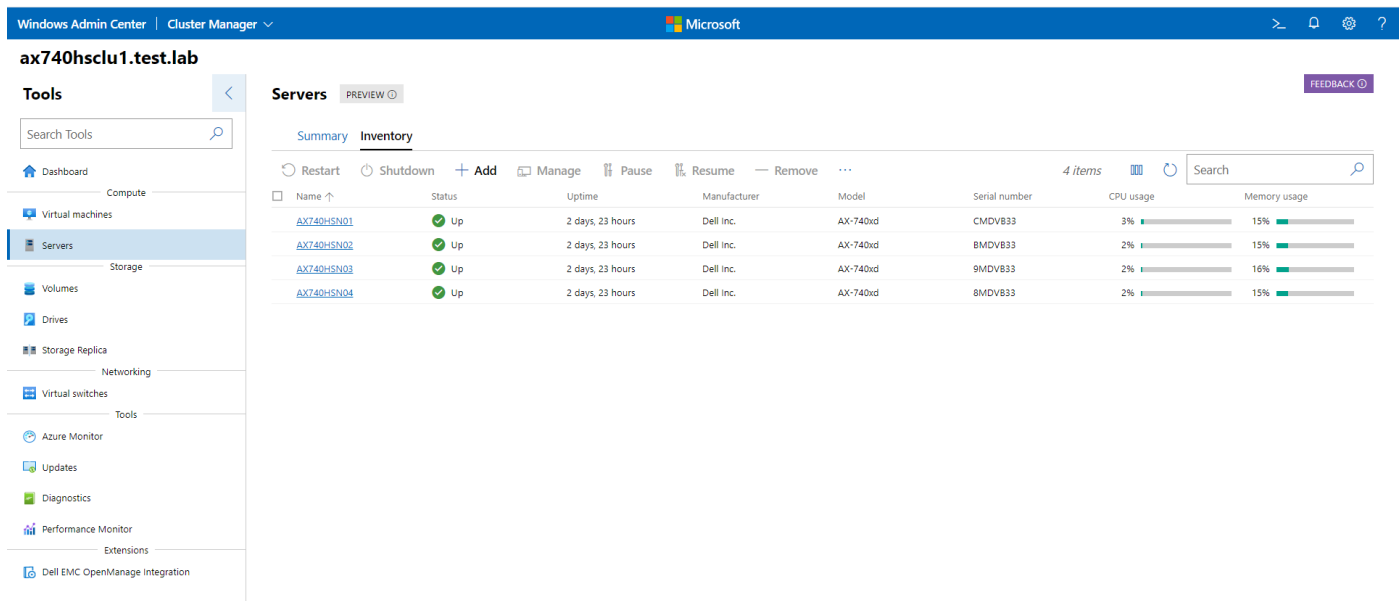


Figure 6. Servers: Inventory tab

NOTE: The metrics in the figure are for a four-node Azure Stack HCI cluster with all-flash drive configuration.

Viewing drive details

About this task

View the total number of drives in the cluster, the health status of the drives, and the used, available, and reserve storage of the cluster as follows.

Steps

1. In the left pane, select **Drives**.
2. Click the **Summary** tab, as shown in the following figure.

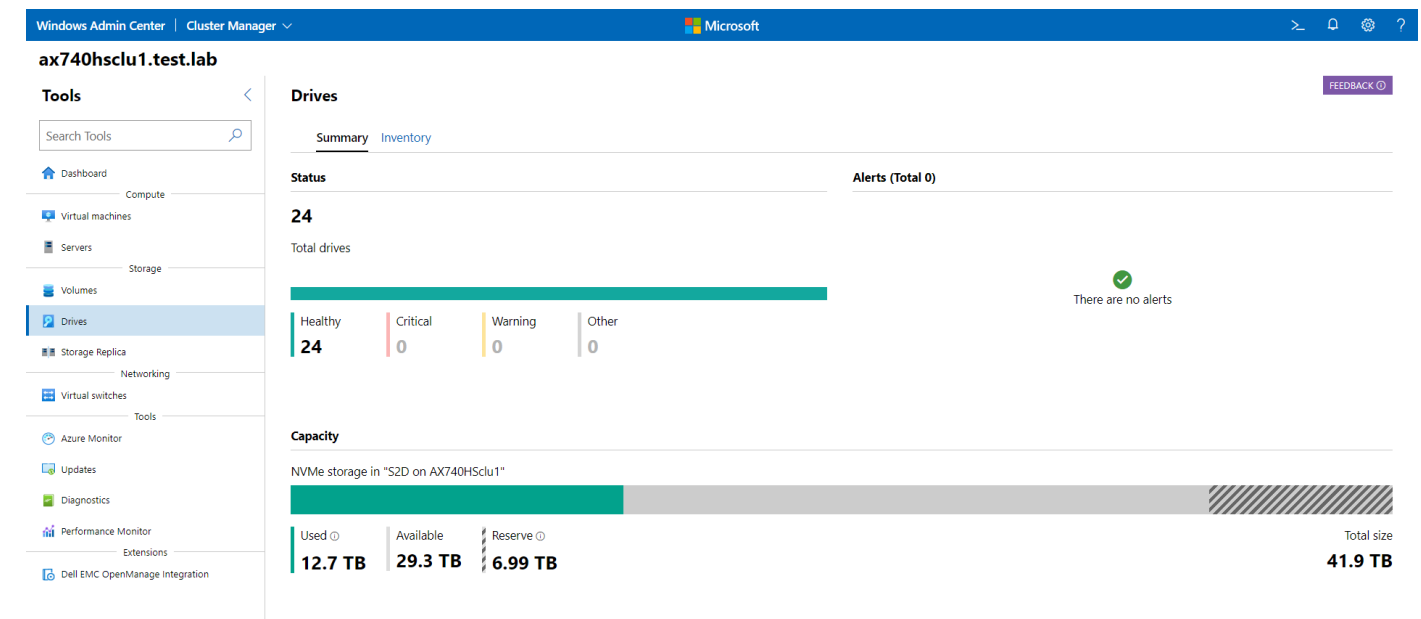


Figure 7. Drives: Summary tab

To view the drive inventory from the cluster nodes, from the left pane, select **Drives**, and then click the **Inventory** tab.

ax740hsc1u1.test.lab

Tools

Search Tools

Dashboard

Compute

Virtual machines

Servers

Storage

Volumes

Drives

Storage Replica

Networking

Virtual switches

Tools

Azure Monitor

Updates

Diagnostics

Performance Monitor

Extensions

Dell EMC OpenManage Integration

Settings

Drives

Summary **Inventory**

Light On Light Off Retire Unretire

24 items

Search

Serial number	Status	Model	Size	Type	Used for	Location	Server	Storage pool	Storage usage
90J0A003TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 21 ...	ax740hscn01	S2D on AX740H5clu1	31%
90J0A004TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 23 ...	ax740hscn01	S2D on AX740H5clu1	30%
90J0A006TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 20 ...	ax740hscn01	S2D on AX740H5clu1	31%
90J0A014TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 22 ...	ax740hscn01	S2D on AX740H5clu1	31%
90J0A016TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 19 ...	ax740hscn01	S2D on AX740H5clu1	31%
90J0A01ATBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 18 ...	ax740hscn01	S2D on AX740H5clu1	31%
90J0A005TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 19 ...	ax740hscn02	S2D on AX740H5clu1	28%
90J0A014TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 22 ...	ax740hscn02	S2D on AX740H5clu1	30%
90J0A012TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 18 ...	ax740hscn02	S2D on AX740H5clu1	29%
90J0A01TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 22 ...	ax740hscn02	S2D on AX740H5clu1	28%
90T0A001TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 23 ...	ax740hscn02	S2D on AX740H5clu1	29%
90T0A002TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 20 ...	ax740hscn02	S2D on AX740H5clu1	29%
90J0A005TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 23 ...	ax740hscn03	S2D on AX740H5clu1	31%
90J0A00TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 18 ...	ax740hscn03	S2D on AX740H5clu1	30%
90J0A015TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 20 ...	ax740hscn03	S2D on AX740H5clu1	31%
90J0A017TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 21 ...	ax740hscn03	S2D on AX740H5clu1	31%
90J0A01CTBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 19 ...	ax740hscn03	S2D on AX740H5clu1	30%
90J0A010TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 22 ...	ax740hscn03	S2D on AX740H5clu1	31%
90J0A00CTBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 23 ...	ax740hscn04	S2D on AX740H5clu1	31%
90J0A004TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 19 ...	ax740hscn04	S2D on AX740H5clu1	30%
90J0A00MTBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 18 ...	ax740hscn04	S2D on AX740H5clu1	31%
90J0A00TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 22 ...	ax740hscn04	S2D on AX740H5clu1	31%
90J0A002TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 20 ...	ax740hscn04	S2D on AX740H5clu1	31%
90J0A018TBC7	OK	Dell Ent NVMe FIPS ...	1.75 TB	NVMe	Capacity	PCIe SSD in Slot 21 ...	ax740hscn04	S2D on AX740H5clu1	31%

Figure 8. Drives: Inventory tab

The HCI cluster is built using four AX-740xd nodes, each with two 1.92 TB NVMe drives.

By clicking the serial number of the drive, you can view the drive information, which includes health status, slot location, size, type, firmware version, IOPS, used or available capacity, and storage pool of the drive.

Also, from the dashboard, you can set the drive options as **Light On** or **Light Off**, or **Retire** or **Unretire** from the storage pool.

Managing and monitoring volumes

You can manage and monitor the Storage Spaces Direct volumes using Windows Admin Center.

The following features are supported in Windows Admin Center:

- Create volume
- Browse volume
- Expand volume
- Delete volume
- Make volume offline or online

To access the volumes on the HCI cluster, select the cluster and, in the left pane, click **Volumes**. In the right pane, the **Summary** and **Inventory** tabs are displayed.

The **Summary** tab shows the number of volumes in the cluster and the health status of the volumes, alerts, total IOPS, latency, and throughput information of the available volumes.

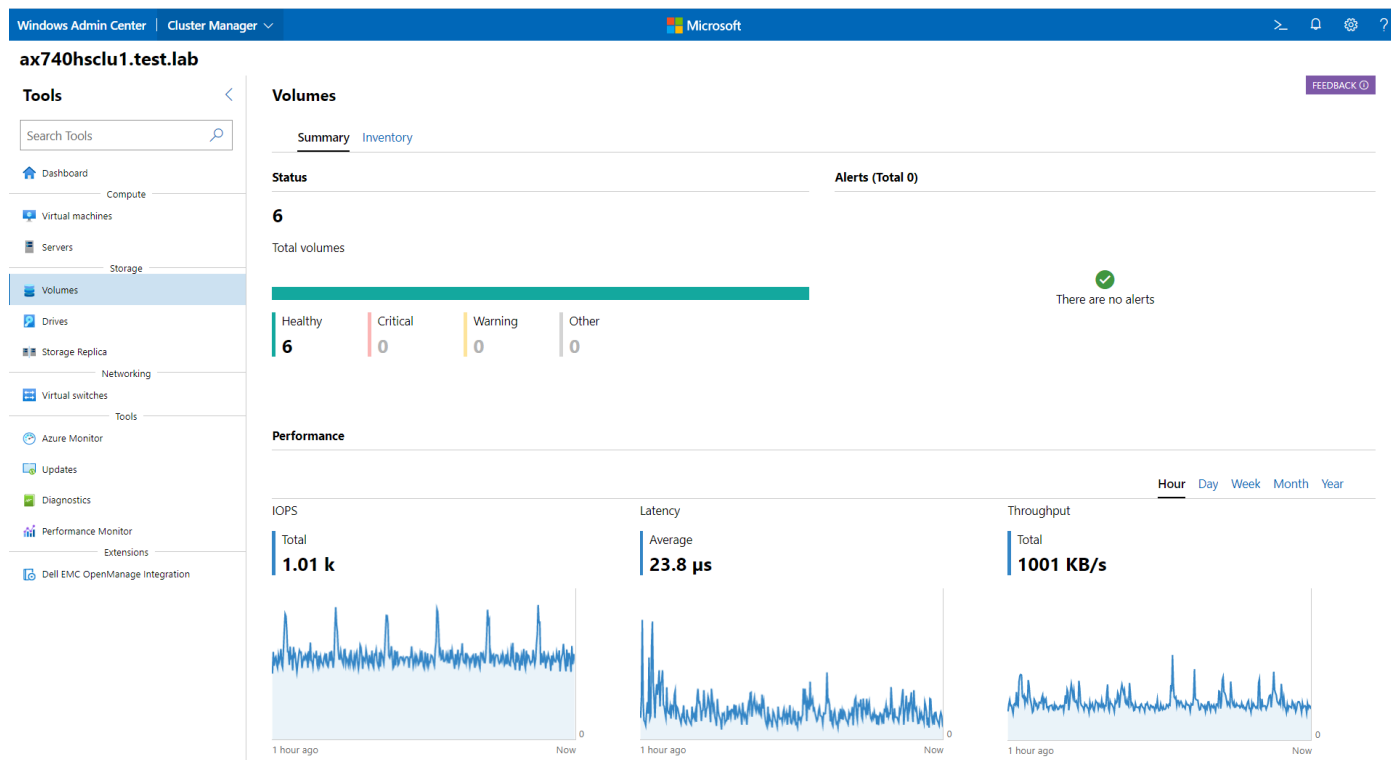


Figure 9. Volumes: Summary tab

The **Inventory** tab provides the volume inventory from the HCI cluster nodes. You can manage and monitor the volumes.

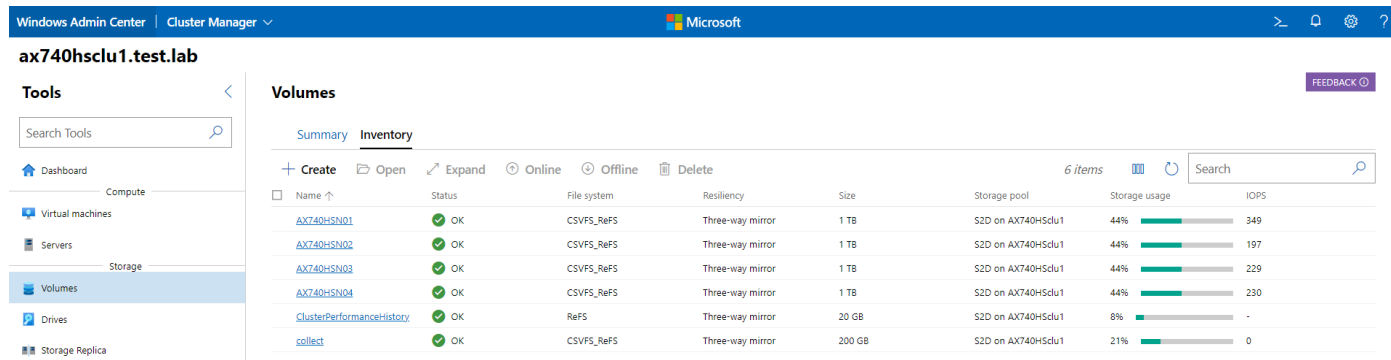


Figure 10. Volumes: Inventory tab

Creating volumes in Storage Spaces Direct

About this task

Create volumes in Storage Spaces Direct in Windows Admin Center as follows.

Steps

1. Go to **Volumes > Inventory**.
2. Click **Create**.
The **Create volume** window is displayed.
3. Enter the volume name, resiliency, and size of the volume, and then click **Create**.
The volume is created.

Managing volumes

About this task

Open, expand, delete, or make a volume offline as follows.


Steps

1. Go to **Volumes > Inventory**.
2. Click the volume name.
3. Click **Open** to open the volume folder.
4. Click **Offline** or **Delete** to make the volume offline, or to delete the volume.
5. Click **Expand** to expand the volume.
The **Expand volume** window is displayed.
6. Enter the additional size of the volume.
7. Select the volume size from the drop-down list and click **Expand**.

Enabling data deduplication on Storage Spaces Direct

About this task

Data deduplication helps to maximize free space on the volume by optimizing duplicated portions on the volume without compromising data fidelity or integrity.

 **NOTE:** To enable data deduplication on an HCI cluster, ensure that the data deduplication feature is enabled on all the cluster nodes. To enable the data deduplication feature, run the following PowerShell command: `Install-WindowsFeature FS-Data-Deduplication`.

Enable data deduplication and compression on a Storage Spaces Direct volume as follows.

Steps

1. Go to **Volumes > Inventory**.
2. Click the volume on which to enable data deduplication.
3. In the optional features, switch the ON button to enable deduplication and compression on that volume.
The **Enable Deduplication** window is displayed.
4. Click **Start** and select **Hyper-V** from the drop-down list.
5. Click **Enable Deduplication**.
Deduplication is enabled and the Storage Spaces Direct volume is compressed.

Monitoring and managing VMs

You can use Windows Admin Center to monitor and manage the VMs that are hosted on the HCI cluster.

To access the VMs that are hosted on the HCI cluster, click the cluster name and, in the left pane, select **Virtual machines**. In the right pane, the **Inventory** tab and the **Summary** tab are displayed.

The **Inventory** tab provides a list of the VMs that are hosted on the HCI cluster and provides access to manage the VMs.

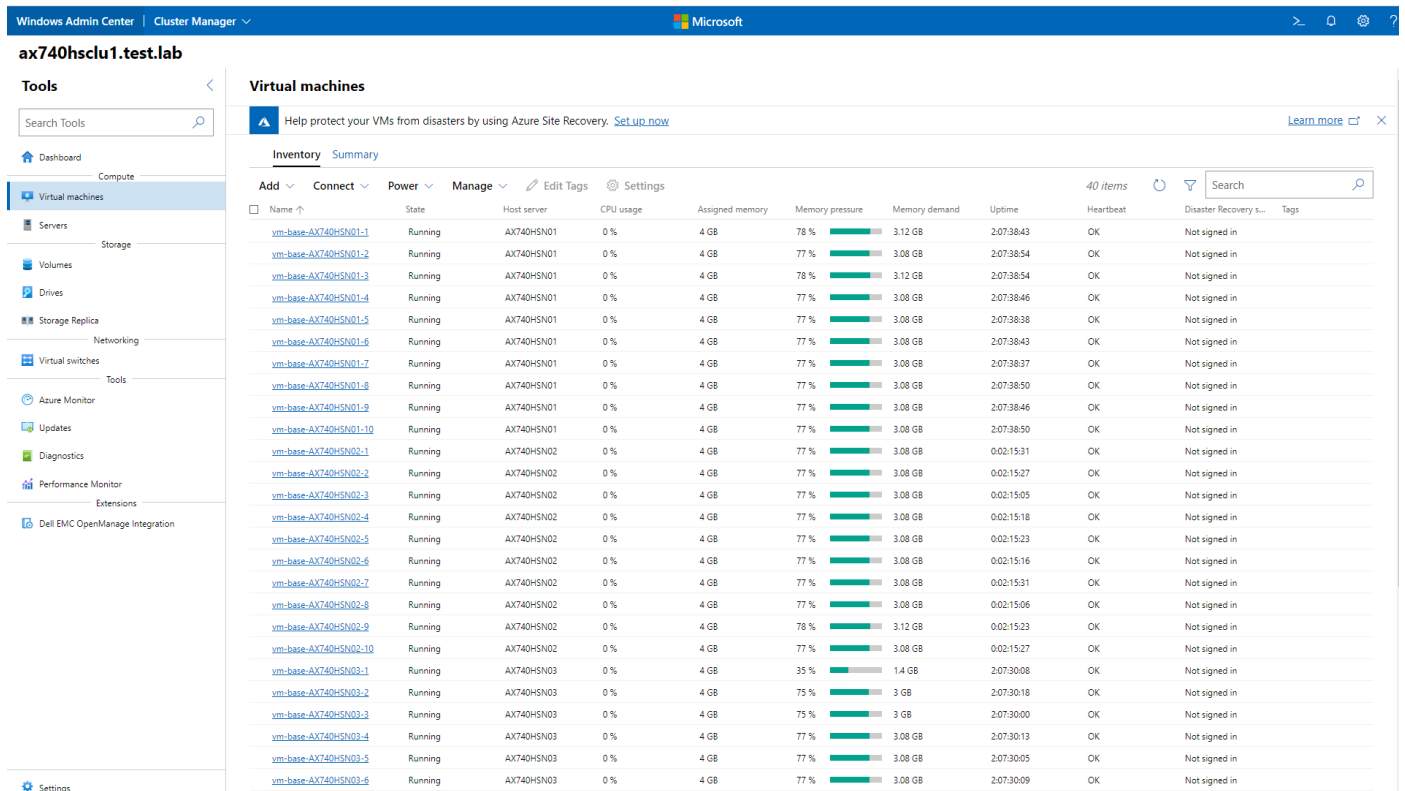


Figure 11. VMs: Inventory tab

The **Summary** tab provides the following information about the VM environment of the HCI cluster:

- Total number of VMs, their state, and alerts
- Host and guest CPU utilization
- Host and guest memory utilization
- VM total IOPS, latency, and throughput information

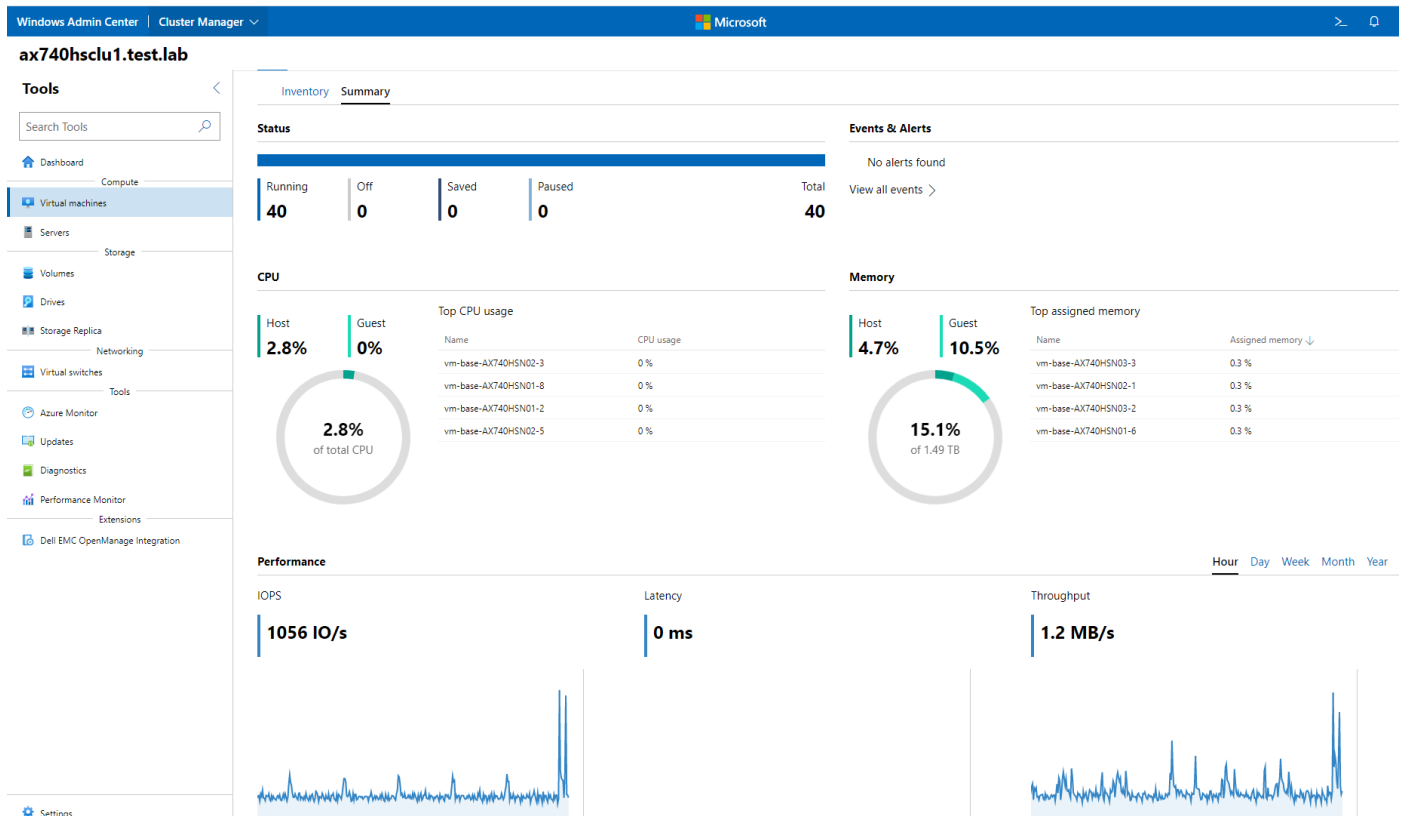


Figure 12. VMs: Summary tab

You can perform the following tasks from the Windows Admin Center console:

- View a list of VMs that are hosted on HCI cluster.
- View individual VM state, host server information, virtual machine uptime, CPU, memory utilization, and so on.
- Create a new VM.
- Modify VM settings.
- Set up VM protection.
- Delete, start, turn off, shut down, save, delete saved state, pause, resume, reset, add new checkpoint, move, rename, and connect VMs.

Managing virtual switches

The virtual switches tool in Windows Admin Center enables you to manage Hyper-V virtual switches of the cluster nodes.

The virtual switches tool supports the following features:

- View existing virtual switches on the server.
- Create a new virtual switch.
- Modify virtual switch properties.
- Delete a virtual switch.

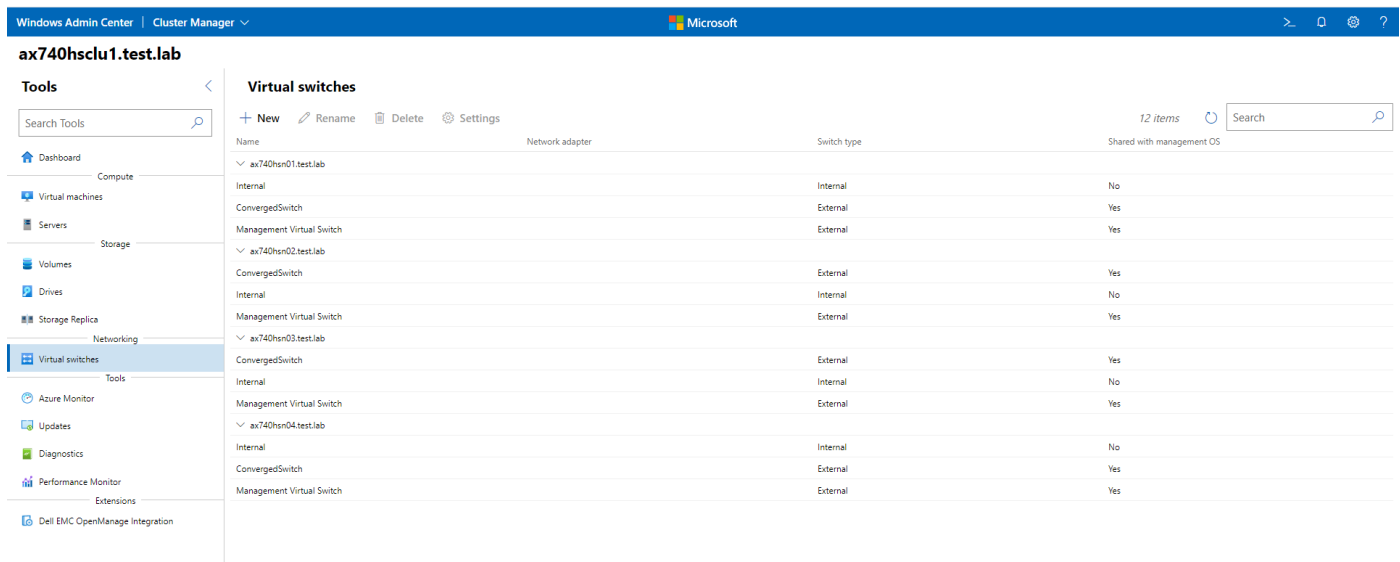


Figure 13. Virtual switches

Dell EMC OpenManage Integration with Windows Admin Center

Dell EMC OpenManage Integration with Windows Admin Center enables IT administrators to manage the hyperconverged infrastructure (HCI) that is created by using Microsoft HCI Solutions from Dell Technologies. OpenManage Integration with Windows Admin Center simplifies the tasks of IT administrators by remotely managing the AX nodes and clusters throughout their life cycle.

For more information about the features, benefits, and installation of OpenManage Integration with Windows Admin Center, see the documentation at <https://www.dell.com/support/home/product-support/product/openmanage-integration-microsoft-windows-admin-center/docs>.

NOTE: For Storage Spaces Direct Ready Node, if you want to use the "Cluster Aware Update" premium feature to update the cluster using the Dell extension you should contact a Dell sales representative to get an Azure Stack HCI license. See the "Firmware and driver updates using the manual method" section.

Prerequisites for managing AX nodes

The prerequisites for managing AX nodes are:

- You have installed the following:
 - Windows Admin Center version 2103 and you are logged in as a gateway administrator.
 - Dell EMC OpenManage Integration with Microsoft Windows Admin Center extension version 2.0.0. For more information about the installation procedure, see the [Dell EMC OpenManage Integration Version 2.0.0 with Microsoft Windows Admin Center Installation Guide](#).
 - Microsoft Cluster Creation Extension version 1.506.0 release or above.
 - Microsoft failover cluster extension version 1.243.0 release or above.
 - An OMIWAC Premium License on each AX node.
- You have added the [HCI cluster connection](#) in Microsoft Windows Admin Center.
- You can access the Windows Admin Center remotely using domain administrator credentials. Otherwise, use local administrator credentials to access the Windows Admin Center locally. For more information, see [What type of installation is right for you?](#)

Installing the Azure Stack HCI license (Ready Nodes only)

AX nodes have a preinstalled Azure Stack HCI license. Storage Spaces Direct Ready Nodes require the installation of an After Point of Sale (APOS) license.

Steps

1. Log in to iDRAC.
2. Select **Configuration > Licenses**.
3. Select **Import**, browse to and select the license, and then click **Upload**.

Managing Azure Stack HCI clusters

Steps

1. In the upper left of Windows Admin Center, select **Cluster Manager** from the menu.
2. In the **Cluster Connections** window, click the cluster name.
3. In the left pane of Windows Admin Center, under **EXTENSIONS**, click **OpenManage Integration**.
4. Review the **Dell EMC Software License Agreement and Customer Notice**, and select the check box to accept the terms of the license agreement.

Health Status

Health Status is the default dashboard that provides details about the Azure Stack HCI cluster nodes.

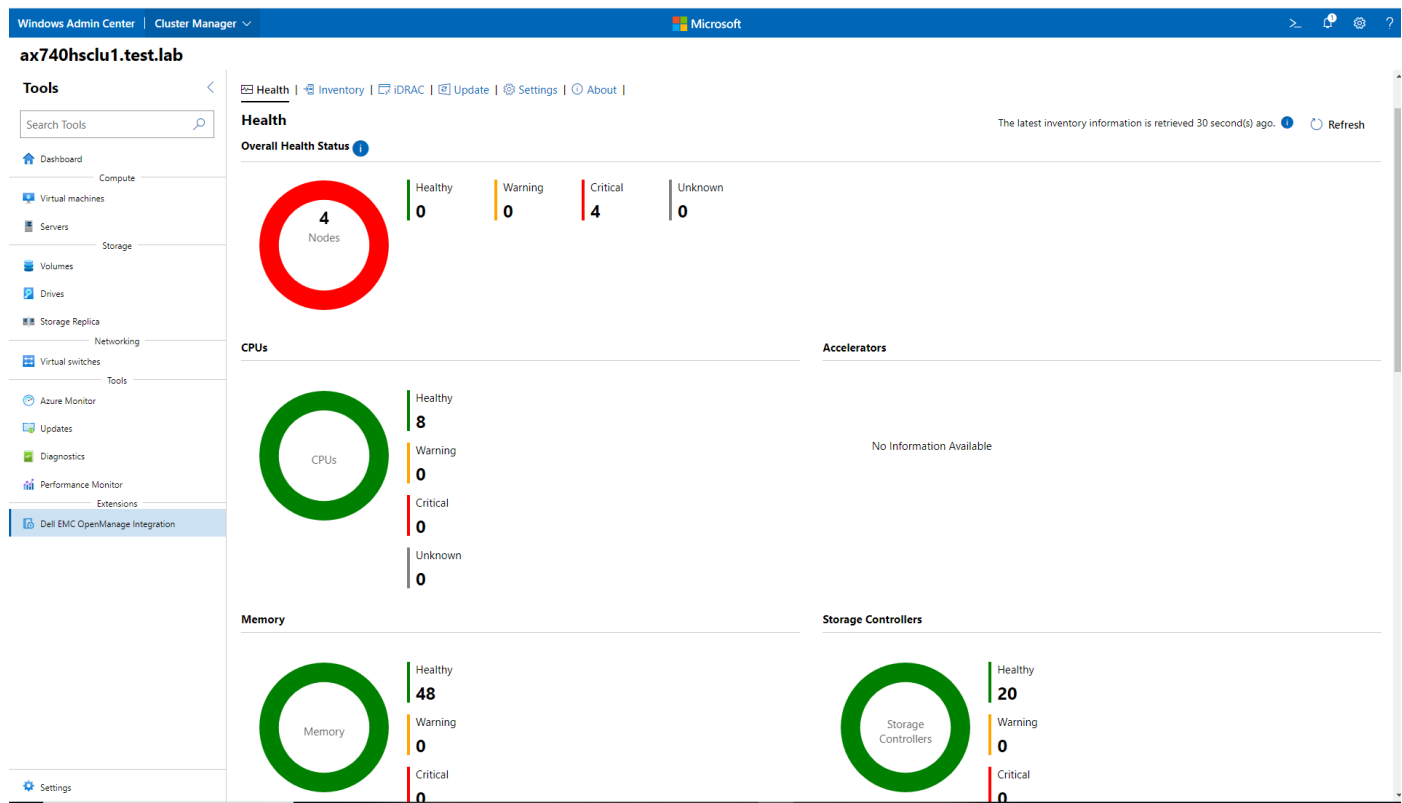


Figure 14. Health Status dashboard

On the **Cluster - Azure Stack HCI** page, click the **Health Status** tab to view the overall health status of the HCI cluster and the health status of the following components of the Azure Stack HCI cluster and nodes:

- Memory
- Power supplies

- CPUs
- Fans
- Storage controllers
- iDRAC
- Storage enclosures
- Physical disks
- Voltages
- Temperatures

Selecting the Critical or Warning section in the overall health status doughnut chart displays the nodes and components that are in the critical or warning state respectively.

Select sections in the doughnut chart to filter the health status of the components. For example, selecting the red section displays only the components with critical health status.

Selecting sections of the chart for individual components shows the respective nodes with the component health status listed. Expand the nodes to view the components.

Inventory

The **Inventory** tab lists the servers that are part of the cluster.

Clicking a server name on the inventory list provides details about the following components:

- System
- Firmware
- CPUs
- Memory
- Storage controllers
- Storage enclosures
- Network devices
- Physical disks
- Power supplies
- Fans

Locating physical disks and viewing their status

The Blink/Unblink feature of Windows Admin Center enables you to locate physical disks or view disk status.

Steps

1. Under the **Inventory** tab, from the **Components** list, select **Physical Disks**.
2. For each physical disk, select **Blink** or **Unblink** to control the disk's LED.

iDRAC

Clicking the **iDRAC** tab displays the **Integrated Dell Remote Access Controller** dashboard. The dashboard lists the servers that are part of the Azure Stack HCI cluster. By selecting each iDRAC, you can view iDRAC details, such as the iDRAC firmware version, the iDRAC IP of the target node, and license information, and can directly launch the iDRAC console.

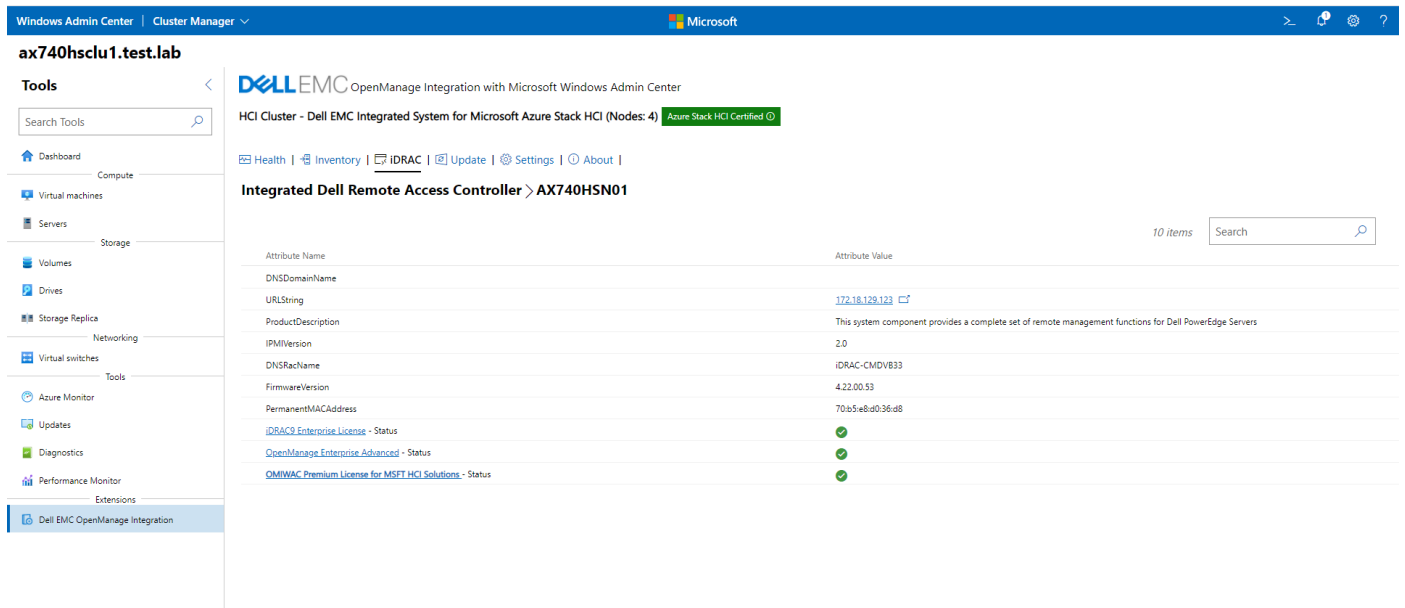


Figure 15. iDRAC dashboard

Settings

Use the **Settings** tab in the Dell EMC OpenManage Integration with Windows Admin Center UI to view the latest update compliance report, update the cluster, and configure proxy settings.

Update tools

To view the latest update compliance report and update the cluster using an offline catalog, OpenManage Integration with Windows Admin Center requires that you configure the settings for the update compliance tools.

In the OpenManage Integration UI, on the **Update Tools** page under the **Settings** tab, specify the download locations for the following update compliance tools:

- Dell EMC Inventory Collector (IC)—Collects the hardware inventory information from Azure Stack HCI cluster nodes
- Dell EMC System Update Utility (DSU)—Updates cluster node firmware and drivers, using Dell Update Packages

Proxy settings

To download the online catalog, update tools, and Dell Update Packages, you must configure the proxy settings unless you can access the Internet without proxy settings.

Configure the proxy settings under the **Settings** tab in the Dell EMC OpenManage Integration UI.

Viewing update compliance and updating the cluster

OpenManage Integration with Windows Admin Center enables you to view the update compliance details (firmware, driver, application, and BIOS). You can update the Azure Stack HCI cluster by using OpenManage Integration.

About this task

Use the **Update** tab of the OpenManage Integration with Windows Admin Center UI to view update compliance and update the cluster.

Steps

1. At **Update source**, select the online catalog or offline catalog.

If you select the online catalog, OpenManage Integration downloads the Azure Stack HCI catalog, system tools, and the required Dell Update Packages from the Internet.

To use an offline catalog, the update tools must be configured under the **Settings** tab, and the catalog file must be exported using the Dell Repository Manager and placed in a shared folder. See [Obtaining the firmware catalog for AX nodes or Ready Nodes using Dell EMC Repository Manager](#).

2. Click **Next: Compliance Details** to generate the update compliance report.

By default, all the upgrades are selected, but you can make alternate selections as needed.

The screenshot shows the Dell EMC OpenManage Integration with Microsoft Windows Admin Center interface. The left sidebar contains navigation options like Dashboard, Virtual machines, Servers, Volumes, Drives, Storage Replica, Networking, Tools, Azure Monitor, Updates, Diagnostics, Performance Monitor, and Extensions. The main area displays the 'Update' section for an HCI Cluster - Dell EMC Integrated System for Microsoft Azure Stack HCI (Nodes: 4). The 'Update' section has four tabs: Update source, Compliance report (selected), Summary, and Cluster aware update. The 'Compliance report' tab shows a 'Component Compliance Summary' with a bar chart indicating 64 Compliant, 0 Urgent, 4 Recommended, and 0 Optional updates, totaling 68. Below this is a 'Compliance Report' table with columns: Component Name, Compliance, Criticality, Current Version, Baseline Version, Type, and Compliance Type. The table lists various components like IDrac, BIOS, Dell HBA330 Adp Driver, Intel Family of Server Adapter, QLogic Family of Server Adapt..., Chipset INF, BOSS, Driver for Marvell Unify Confi..., Intel(R) Gigabit 4P X710/1350 r..., and Intel(R) Ethernet 10G 4P X710/1350 r... with their respective compliance status and versions.

Component Name	Compliance	Criticality	Current Version	Baseline Version	Type	Compliance Type
AX740HSN01 (Licensed)	Compliant	Recommended	4.22.00.53	4.40.00.00	Firmware	Upgradable
AX740HSN02 (Licensed)	Compliant	Recommended	2.10.0	2.10.0	BIOS	Same
AX740HSN03 (Licensed)	Compliant	Recommended	2.51.25.2	2.51.25.02	Driver	Same
AX740HSN04 (Licensed)	Compliant	Recommended	19.5.0	19.5.0	Driver	Same
QLogic Family of Server Adapt...	Compliant	Recommended	35.17.03	35.17.03	Driver	Same
Chipset INF	Compliant	Recommended	10.1.18243.8188	10.1.18243.8188	Driver	Same
BOSS	Compliant	Recommended	2.5.13.3024	2.5.13.3024	Firmware	Same
Driver for Marvell Unify Confi...	Compliant	Recommended	1.2.0.1051	1.2.0.1051	Driver	Same
Intel(R) Gigabit 4P X710/1350 r...	Compliant	Recommended	19.5.12	19.5.12	Firmware	Same
Intel(R) Ethernet 10G 4P X710/1350 r...	Compliant	Recommended	19.5.12	19.5.12	Firmware	Same

Figure 16. Compliance Details

3. Click **Next: Summary** to view the selected component details.

NOTE: Cluster Aware Update is a license feature. Ensure that the Azure Stack HCI license is installed before proceeding.

Windows Admin Center | Cluster Manager

Microsoft

ax740hsc1u1.test.lab

Tools

Search Tools

Dashboard

Virtual machines

Servers

Volumes

Drives

Storage Replica

Networking

Virtual switches

Tools

Azure Monitor

Updates

Diagnostics

Performance Monitor

Extensions

Dell EMC OpenManage Integration

Settings

DELL EMC OpenManage Integration with Microsoft Windows Admin Center

HCI Cluster - Dell EMC Integrated System for Microsoft Azure Stack HCI (Nodes: 4) Azure Stack HCI Certified

Health | Inventory | iDRAC | **Update** | Settings | About

Update

Update source

Compliance report

Summary

Cluster aware update

Update options:

Run now

Schedule update

Following components have been selected for the update. Please click **Cluster aware update** to start the update.

Note: The Update process may take several hours and the target node(s) will be rebooted if needed.

Component Name	Compliance	Criticality ↓	Current Version	Baseline Version	Type
AX740HSN01					
iDRAC	Non-Compliant	Recommended	4.22.00.53	4.40.00.00	Firmware
AX740HSN02					
iDRAC	Non-Compliant	Recommended	4.22.00.53	4.40.00.00	Firmware
AX740HSN03					
iDRAC	Non-Compliant	Recommended	4.22.00.53	4.40.00.00	Firmware
AX740HSN04					
iDRAC	Non-Compliant	Recommended	4.22.00.53	4.40.00.00	Firmware

Back

Next: Cluster aware update

Exit

Figure 17. Update Summary

- To schedule the update for a later time, click **Schedule later**, select **Date/time** and click **Next cluster aware update** to download the required updates.
To use the schedule later feature, download the required downloads and keep them ready to update at the specified time.
- Click **Next: Cluster Aware Update** to begin the update process and click **Yes** at the prompt to enable Credential Security Service Provider (CredSSP) to update the selected components.

22 Day 0 Operations

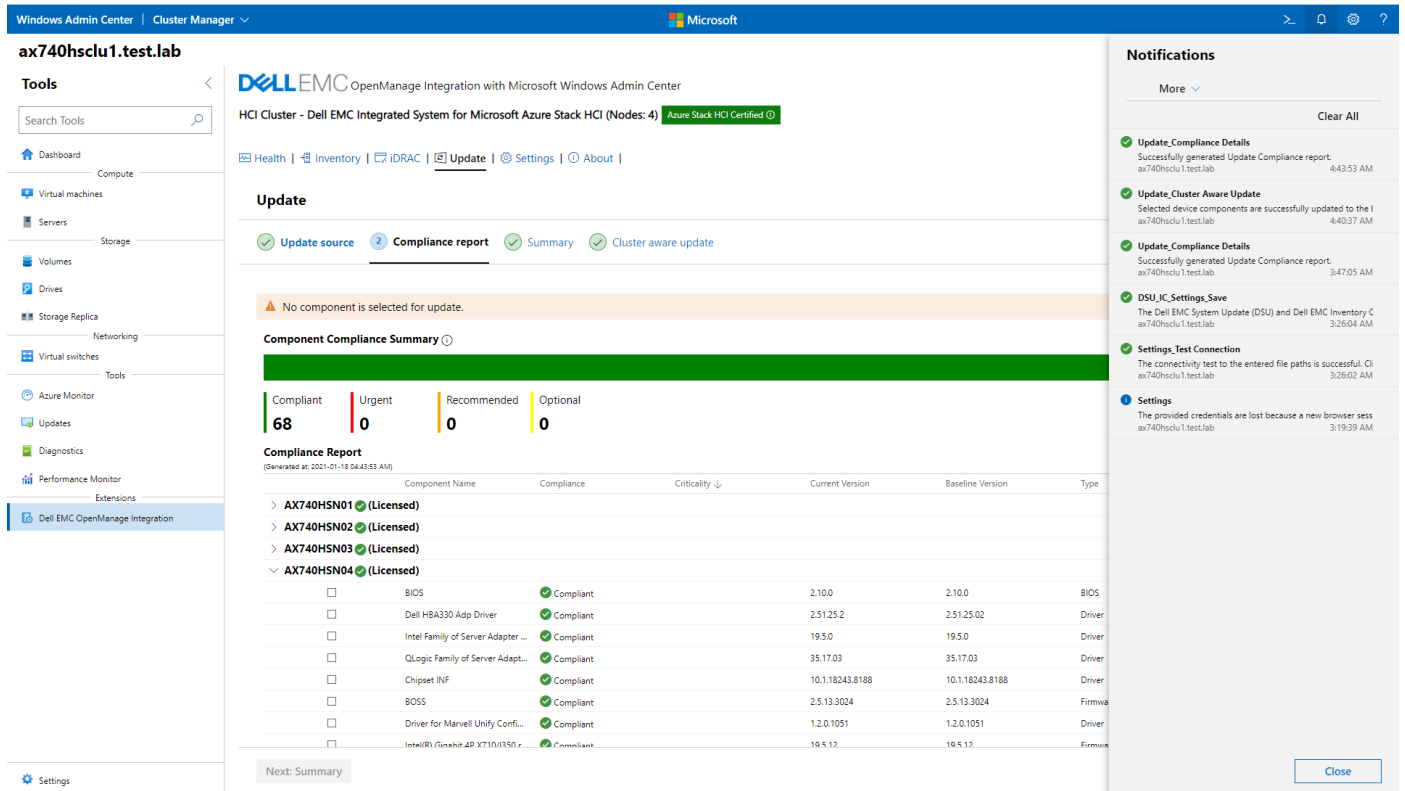


Figure 18. Cluster Aware Update

When the update job is completed, the compliance job is triggered automatically.

Full Stack Cluster-Aware Updating prerequisites for AX-7525 and AX-6515 nodes (offline update)

About this task

If an Internet connection is not available, run Full Stack Cluster-Aware Updating (CAU) in offline mode as follows:

Steps

1. Download the `asHCISolutionSupportMatrix.json` and `asHCISolutionSupportMatrix.json.sign` files from <http://downloads.dell.com/omimswac/supportmatrix/>.
2. Place these files in the `C:\Users\Dell\SymmetryCheck` folder in the gateway system where Windows Admin Center is installed.
3. Run Full Stack CAU.


Results


For more information about CAU, see the [Cluster-Aware Updating Overview](#).

Full Stack Cluster-Aware Updating for Azure Stack HCI clusters using the OpenManage Integration snap-in

About this task


Windows Admin Center with the Dell EMC extension makes it easy to update an Azure Stack HCI cluster using the cluster aware update feature. The feature updates the operating system and Dell EMC-qualified firmware and drivers. When an update is selected, all the updates are installed on the cluster nodes. A single reboot is required to install operating system, firmware, and driver updates per server.


 **NOTE:** Full Stack Cluster-Aware Updating (CAU) is only available on Azure Stack HCI clusters built using the Azure Stack HCI operating system. For more information about CAU, see the [Cluster-Aware Updating Overview](#).

 **NOTE:** Full Stack CAU is a licensed feature. Ensure that the Azure Stack HCI license is installed before proceeding.

To perform both operating system updates and hardware upgrades on Azure Stack HCI cluster nodes, carry out the following steps:

Steps

1. In Windows Admin Center, select **Updates** from the **Tools** menu.
You must enable CredSSP and provide explicit credentials. When asked if CredSSP should be enabled, click **Yes**.
The **Updates** page is displayed.
2. For an operating system update, see [Microsoft's Azure Stack HCI documentation](#).
3. On the **Install updates** page, review the operating system updates and select **Next: Hardware updates**.
4. If the Dell EMC OpenManage Integration extension is not installed, click **Install** to accept the license terms and install the extension. If you have already installed the OpenManage Integration extension version 2.0, click **Get updates** to move to the **Hardware updates** page.
5. On the **Hardware updates** page, review the prerequisites listed to ensure that all nodes are ready for hardware updates and then click **Next: Update Source**. Click **Re-Run** to run the prerequisites again.
You must meet all the prerequisites listed on the **Prerequisites** tab, otherwise you cannot proceed to the next step.
6. To generate a compliance report against the validated Azure Stack HCI catalog, follow these steps on the **Update source** page:
 - a. Select one of these methods to download catalog files:
 - Select **Online (HTTPs) - Update Catalog for Microsoft HCI Solutions** to download the catalog automatically from dell.com. The online catalog option is selected by default. Online catalog support requires direct Internet connectivity from the Windows Admin Center gateway. The overall download time of a catalog depends on the network bandwidth and the number of components being updated.
 **NOTE:** Accessing to the Internet using proxy settings is not supported.
 - Select **Offline - Dell EMC Repository Manager Catalog** to use the DRM catalog configured in a CIFS location. OMIMSWAC with or without Internet access allows you to select Offline - Dell EMC Repository Manager Catalog to generate a compliance report. You can use this option when the Internet is not available. For more information, see [Obtaining the firmware catalog for AX nodes or Ready Nodes using Dell EMC Repository Manager](#).
 - To use the offline catalog, select DRM Settings to ensure that the CIFS share path is configured with the DRM catalog.
 - b. To use the Dell EMC System Update (DSU) and Inventory Collector (IC) tools, select **Advance setting** and then do the following:
 - Select **Manually configure DSU and IC** and then select **Settings** to manually download and configure DSU and IC tools in a shared location. We recommend using this option when OMIMSWAC is not connected to the Internet. DSU and IC settings that are configured using **Update Tool** settings in the OpenManage Integration extension are also available under **Advanced settings** in the OpenManage Integration snap-in.

OMIMSWAC downloads the catalog, collects the DSU and IC tools that are configured in the **Settings** tab, and generates a compliance report. If DSU and IC tools are not configured in the **Settings** tab, then OMIMSWAC downloads them from <https://downloads.dell.com> to generate the compliance report.
7. On the **Compliance report** tab, view the compliance report. When finished, click **Next: Summary**.
The 'upgradable' components that are 'non-compliant' are selected by default for updating. You can clear the check box for the selected components or select the 'non-compliant,' 'downgradable' components. However, if you want to change any of the default selections, ensure that the dependencies between the corresponding component firmware and drivers are met.
8. On the **Summary** tab, review the components to be updated and then click **Next: Download updates**. The download task continues in the background whether the UI session is live or not. If the UI session is live, the node level progress status is displayed. OMIMSWAC creates a notification when the download task is finished.
 **NOTE:** While the download is in progress, it is recommended that you do not exit or close the browser. If you do, the download update operation may fail.
9. If the download operation fails, check the log files stored at the following paths for troubleshooting purposes:
 - Gateway system—<Windows
Directory>\ServiceProfiles\NetworkService\AppData\Local\Temp\generated\logs

- Windows 10 gateway system—<Windows installed drive>\Users\<user_name>\AppData\Local\Temp\generated\logs
- After the cluster update is over, DSU logs for individual nodes can be found in the <Windows Directory>\Temp\OMIMSWAC folder on the respective nodes.

To run the compliance report again, click **Re-run Compliance** and repeat steps 4 to 7.

10. After the updates are downloaded, follow the instructions in the Windows Admin Center to install both operating system and hardware updates. If the UI session is live, the node level progress status is displayed. Windows Admin Center creates a notification once the update is completed.

Updating a stand-alone node before adding it to the cluster

Before creating a cluster, ensure that each node is updated with the latest versions of firmware and drivers.

Steps

1. In Windows Admin Center, in the left pane, click **Add**.
2. In the **Windows Server** tile, click **Add**.
3. Enter the node name and click **Add**.
4. Under **All connections**, select the server and click **Manage as**.
5. Select **use another account for this connection**, and then provide the credentials in the domain\username or hostname\username format.
6. Click **Confirm**.
7. In the **Connections** window, click the server name.
8. In the left pane of Windows Admin Center, under **EXTENSIONS**, click **OpenManage Integration**.
9. Review the **Dell EMC Software License Agreement and Customer Notice**, and select the check box to accept the terms of the license agreement.
10. Go to the **Update** tab and, select either the online catalog or offline catalog.
11. Click **Next: Compliance Details**, select all the updates, and click **Next: Summary** and then **Next: Update**.

Known issues

The following table lists known issues and workarounds related to OpenManage Integration with Microsoft Windows Admin Center with Microsoft HCI Solutions from Dell Technologies clusters.

NOTE: For details about troubleshooting steps and known issues, see the *Dell EMC OpenManage Integration Version 2.0.0 with Microsoft Windows Admin Center User's Guide* at <https://www.dell.com/support/home/product-support/product/openmanage-integration-microsoft-windows-admin-center/docs>.

Table 1. Known issues

Issue	Resolution/workaround
Running <code>Test-Cluster</code> fails with network communication errors. With USB NIC enabled in iDRAC, if you run the <code>Test-Cluster</code> command to verify the cluster creation readiness or cluster health, the validation report includes an error indicating that the IPv4 addresses assigned to the host operating system USB NIC cannot be used to communicate with the other cluster networks.	This error can be safely ignored. To avoid the error, temporarily disable the USB NIC (labeled as Ethernet, by default) before running the <code>Test-Cluster</code> command.
The USB NIC network appears as a partitioned cluster network. When the USB NIC is enabled in iDRAC, cluster networks in the failover cluster manager show the networks associated with the USB NIC as partitioned. This issue occurs because the cluster communication is enabled by default on all network	Remove the USB NIC from any cluster communication by using the following script: <pre>\$rndisAdapter = Get-NetAdapter - InterfaceDescription 'Remote NDIS Compatible Device' -ErrorAction SilentlyContinue</pre>

Table 1. Known issues (continued)

Issue	Resolution/workaround
adapters, and USB NIC IPv4 addresses cannot be used to communicate externally, which, therefore, breaks cluster communication on those NICs.	<pre> if (\$rdisAdapter) { Write-Log -Message 'Remote NDIS found on the system. Cluster communication will be disabled on this adapter.' # Get the network adapter and associated cluster network \$adapterId = [Regex]::Matches(\$rdisAdapter.InstanceID, '(?<={})(.*?)(?=})').Value \$usbNICInterface = (Get- ClusterNetworkInterface).Where({\$_ .adapter Id -eq \$adapterId}) \$usbNICClusterNetwork = \$usbNICInterface.Network # Disable Cluster communication on the identified cluster network (Get-ClusterNetwork -Name \$usbNICClusterNetwork.ToString()).Role = 0 } </pre>
While triggering full stack updates, the Tests Summary page might appear.	<p>As a workaround, verify whether the pre-update or post-update scripts are part of the cluster role. If they are present, remove the scripts from the cluster node by running the following command in PowerShell:</p> <pre> Set-CauClusterRole -PreUpdateScript \$null -PostUpdateScript \$null </pre> <p>For more information about the prerequisites required for a cluster update, see Update Azure Stack HCI clusters.</p>
The update status takes a long time to refresh.	<p>During full stack cluster updates, the update status shown in the Updates page might take a long time to refresh. If this issue occurs, it is recommended that you stay on the Updates page and wait for the update to complete. The update status will automatically be displayed once the update is complete.</p>
When using CredSSP authentication to run scripts on a remote machine, the update job might fail with an error. This failure occurs because CredSSP has been disabled in the gateway machine.	<p>To resolve the issue, follow these steps:</p> <ol style="list-style-type: none"> 1. From the PowerShell window, run <code>gpedit</code>. 2. In the Group Policy Editor window, browse to Computer Configurations > Administrative Templates > System > Credentials Delegation. 3. Select Allow delegating fresh credentials with NTLM-only server authentication and enable it. 4. Run <code>gpupdate /force</code> in PowerShell.


Firmware updates using Dell EMC OpenManage Integration for Microsoft System Center for System Center Virtual Machine Manager

Dell EMC OpenManage Integration for Microsoft System Center is an appliance-based integration with the System Center suite of products.

OpenManage Integration for Microsoft System Center enables full life-cycle management of Dell EMC PowerEdge servers by using iDRAC with Lifecycle Controller (LC).

OpenManage Integration for Microsoft System Center provides operating system deployment, Azure Stack HCI cluster creation, hardware and firmware updating, and maintenance of servers and modular systems. Integrate OpenManage Integration for Microsoft System Center with Microsoft System Center Virtual Machine Manager (SCVMM) to manage your PowerEdge servers in virtual and cloud environments.

Checking compliance and updating firmware

 **NOTE:** This method is applicable only for Storage Spaces Direct Ready Nodes.

Perform compliance checks, bare-metal firmware updates, and firmware updates using the cluster-aware update feature. To perform these tasks, within SCVMM, first discover the Storage Spaces Direct Ready Nodes and create or edit an update source.

Before performing these tasks, ensure that:

- SCVMM and the OpenManage Integration for Microsoft System Center appliance have been deployed and configured.
For more information, see the installation guide at <https://www.dell.com/support/home/us/en/04/product-support/product/omimssc-sccm-scvmm-v7.2/docs>.
- The Azure Stack HCI cluster has been deployed by using OpenManage Integration for Microsoft System Center.
For more information about deploying an Azure Stack HCI cluster, see the user guide at <https://www.dell.com/support/home/us/en/04/product-support/product/omimssc-sccm-scvmm-v7.2/docs>.

Discovering the Storage Spaces Direct Ready Nodes

To perform compliance checks and firmware updates, first discover the Storage Spaces Direct Ready Nodes.

Steps

1. Launch SCVMM.
2. In the left pane, click **Fabric**, and then, under **Servers**, select **All Hosts**.
3. On the top banner, click **DELL EMC OMIMSSC**.
4. Expand **Configuration and Deployment** and select **Server View**.
5. On the **Server View** page, click **Discover**.
6. In the **Discover** window, select **Discover using an IP Address** or **Discover using an IP Range**.
7. Click **Create new** to create a credential profile for iDRAC.
 - a. On the **Credential Profile** page:
 - At **Credential Type**, select **Device Credential Profile**.
 - At **Profile Name**, enter the profile name.
 - At **Profile Description**, enter a profile description (optional).
 - At **User name**, enter the iDRAC username.
 - At **Password**, enter the iDRAC password.
 - At **Default Profile for**, select **IDRAC**.
 - b. Click **Finish**.
8. For **Apply this Credential Profile**, select the credential profile that you created.
9. For **iDRAC IP Address**, enter the IP address details of the Storage Spaces Direct Ready Nodes.
10. For **Job Name**, enter a job name for discovery.

Creating or editing an update source

After discovering the Storage Spaces Direct Ready Nodes, create or edit an update source before performing compliance checks and firmware updates within SCVMM.

Steps

1. Launch SCVMM.
2. In the left pane, click **Fabric**, and then, under **Servers**, select **All Hosts**.
3. On the top banner, click **DELL EMC OMIMSSC**.

4. In the left pane, select **Maintenance Center**, and then, at the top of the window, select **Maintenance Settings**.
5. Update the systems by using the online catalog or offline catalog.
Using the online catalog:
 - a. Select **DELL ONLINE HTTPS CATALOG (default)**, and then click **edit**.
 - b. On the **Firmware Update Source** page, keep the default values, create a proxy credentials profile, and select the proxy credentials to connect to the Internet.
 - c. Click **Test Connection** to test the Internet connection to the catalog path.
 - d. Click **Save**.

Using the offline (Dell Repository Manager) catalog:

- a. Click **Create**.
- b. On the **Firmware Update** page:
 - For **Firmware Update Source Name**, enter a friendly source name.
 - For **Description**, enter a description (optional).
 - For **Source Type**, select Dell Repository manager sources.
 - For **Location**, enter the shared path location: \\<servername>\folder\filename.xml>.
 - For **Credentials**, create a credentials profile or use an existing profile to connect to the shared path.
- c. Click **Test Connection** to test the connection to the shared path.
- d. Click **Save**.

Updating the firmware on a bare-metal server

With OpenManage Integration for Microsoft System Center on a bare-metal server, you can update firmware or schedule firmware updates.

Steps

1. Launch SCVMM.
2. In the left pane, click **Fabric**, and then, under **Servers**, select **All Hosts**.
3. On the top banner, click **DELL EMC OMIMSSC**.
4. In the left pane, select **Maintenance Center**.
5. In the **Device Group/Servers** list, select **Default Unassigned Servers Update Group**, which is where the discovered servers are listed.
6. Under **Select Update Source**, select **DELL ONLINE HTTPS S2D CATALOG**.
Use the **Filter Updates** menu to filter the compliance report based on the nature of the update, component type, or server model.

The compliance report of the servers in the selected group is displayed.

7. Click **Run Update**.
8. In the **Update Details** window:
 - At **Firmware Update Job Name**, enter the job name.
 - At **Firmware Update Job Description**, enter the job description.
 - If you want to downgrade the firmware to the catalog version (not recommended), select **Allow Downgrade**.
 - At **Schedule Update**, select **Run Now** or schedule an update for a later time.
 - At **Update Method**, select an update method.
 - **Agent Free Update**—Updates are applied, and the system restarts immediately.
 - **Agent-Free Staged Update**—Updates that do not require a system restart are applied immediately. Updates that require a restart are applied when the system restarts.
9. Click **Finish**.

Updating the firmware with the cluster-aware feature

With OpenManage Integration for Microsoft System Center, you can update firmware or schedule firmware updates using the cluster-aware feature.

Steps

1. Launch SCVMM.
2. In the left pane, click **Fabric**, and then, under **Servers**, select **All Hosts**.
3. On the top banner, click **DELL EMC OMIMSSC**.
4. In the left pane, select **Maintenance Center**.
5. In the **Device Group/Servers** list, select **Cluster_<ClusterName>_Group**, which is where the discovered servers are listed.
6. Under **Select Update Source**, select **DELL ONLINE HTTPS S2D CATALOG**.
Use the **Filter Updates** menu to filter the compliance report based on the nature of the update, component type, or server model.

The compliance report of the servers in the selected group is displayed.

7. Click **Run Update**.
8. In the **Update Details** window:
 - At **Firmware Update Job Name**, enter the job name.
 - At **Firmware Update Job Description**, enter the job description.
 - If you want to downgrade the firmware to the catalog version (not recommended), select **Allow Downgrade**.
 - At **Schedule Update**, select **Run Now** or schedule an update for a later time.
 - At **Update Method**, select an update method.
 - **Agent Free Update**—Updates are applied, and the system restarts immediately.
 - **Agent-Free Staged Update**—Updates that do not require a system restart are applied immediately. Updates that require a restart are applied when the system restarts.
9. Click **Finish**.

Firmware and driver updates using the manual method

These procedures describe how to prepare and update firmware and drivers on an Azure Stack HCI cluster manually.

Preparing for maintenance operations

About this task

Use the following PowerShell commands to ensure that all the requirements are met before proceeding with the maintenance operation of an AX node in an Azure Stack HCI cluster. These steps ensure that all the requirements are met and that no faults exist before placing an AX node into maintenance mode.

Steps

1. Verify that all the nodes in the cluster are available by running the `Get-clusternode` command.
2. Verify that all the cluster networks are available by running the `Get-ClusterNetwork` command.
3. Verify that the cluster status is healthy by running the following commands:
 - `Get-ClusterS2D`
 - `Get-StoragePool`
 - `Get-StorageSubSystem -FriendlyName *Cluster* | Get-StorageHealthReport`
4. Verify that all the physical and virtual drives are healthy by running the following commands:
 - `Get-physicaldisk`
 - `Get-virtualdisks`
5. Run the `Get-storagejob` command to verify that no back-end repair jobs are running.

Placing an AX node in maintenance mode

About this task

After ensuring that the prerequisites are met and before performing the platform updates, place the AX node in maintenance mode (pause and drain). You can move roles or VMs and gracefully flush and commit data in the AX node.

Steps

1. Run the following command to put the node in maintenance mode (pause and drain). Verify that all the roles and virtual drives are drained properly and operational in other nodes after they are moved:

```
Suspend-ClusterNode -name "Hostname" -Drain
```

2. Place the target node in maintenance mode:

```
Get-StorageFaultDomain -type StorageScaleUnit | Where-Object {$_.FriendlyName -eq "<Hostname>"} | Enable-StorageMaintenanceMode
```

3. Run the `Get-Physical Disk` command, and ensure that the Operational Status value is in maintenance mode for the drives that belong to that server.

You can also run the following command and verify that the drives all belong to the paused node:

```
Get-Storagepool -IsPrimordial 0 |Get-PhysicalDisk | ? operationalstatus -eq 'In Maintenance Mode' |Get-StorageNode -PhysicallyConnected
```

4. Turn off the System Lockdown mode.

Obtaining the firmware catalog for AX nodes or Ready Nodes using Dell EMC Repository Manager

About this task

For a qualified set of firmware and drivers for AX nodes or Ready Nodes, we recommend that you use an Azure Stack HCI catalog.

You can generate the firmware catalog along with the firmware and drivers by using Dell EMC Repository Manager (DRM) and copy it to a shared path.

Steps

1. Install DRM version 3.0.1.423 or later.
2. On the DRM home page, click the **Dell EMC Repository Manager** drop-down list.
3. In the **Manage** section, click **Application Preferences**.
The **Preferences** window is displayed.
4. Click **Plug-ins**.
5. Select all the plug-ins and click **Update**.
A message is displayed about the successful completion of the update.
6. Click **Catalogs**.
7. Select all the catalogs and click **Update**.
8. Click **Close** to close the **Preferences** window.
9. On the home page, click **Add Repository**.
The **Add Repository** window is displayed.
10. Enter the **Repository name** and **Description**.
11. Select **Index Catalog- <version>** from the **Base Catalog** drop-down menu.
12. Select **Update Catalog for Microsoft HCI Solutions** from the **Catalog Group**.
13. Select the latest catalog from the **Catalogs** section.
14. Click **Save**.

The **Update Catalog for Microsoft HCI Solutions** is populated in the **Base Catalog** section.

15. In the **Manual** Repository Type, click **All systems in base catalog** and then click **Add**.
The repository is displayed on the repository dashboard available in the home page.
16. Select the repository and click **Export**.
The **Export Deployment Tools** window is displayed.
17. Select the location to export files and click **Export**.
The files are exported to the specified location.

Updating the AX node by using iDRAC out of band

About this task

AX nodes offer device firmware updates remotely through iDRAC. For Azure Stack HCI clusters, the recommended option is to use an Azure Stack HCI catalog for a qualified set of firmware and BIOS. Generate the latest [Dell EMC Azure Stack HCI catalog file](#) through Dell EMC Repository Manager (DRM) and copy the file to a network location before proceeding with the update process.

Steps

1. Log in to the iDRAC web interface.
2. Click **Maintenance > System Update**.
The **Firmware Update** page is displayed.
3. On the **Update** tab, select **Network Share** as the file location.
4. Provide the details of the network share, as shown in the following figure:

Integrated Dell Remote Access Controller 9 | Enterprise

Dashboard | System | Storage | Configuration | Maintenance | iDRAC Settings | Enable Group Manager

Manual Update

Location Type: Network Share

Catalog Location (optional):

Catalog Name(optional): catalog.xml

Network Settings

Protocol: CIFS

IP Address*: 172.18.45.9

Share Name*: q1\Firmware

Domain Name:

User Name*: test\administrator

Password*:

Check for Update

Test network connection

Activate Windows
Go to Settings to activate Windows.

Figure 19. Check for updates

5. Click **Check for updates**.
A list of available updates is displayed, as shown in the following figure.

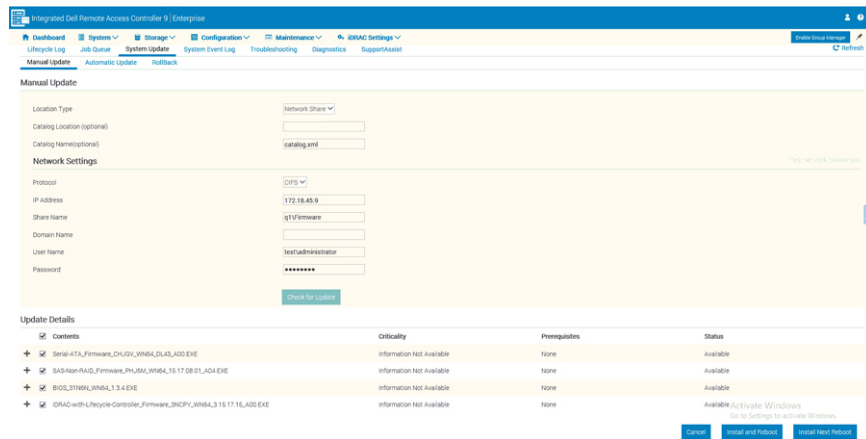


Figure 20. Select updates

6. Select the updates and click **Install Next Reboot** to install and reboot the system.

Updating the out-of-box drivers

For certain system components, you might need to update the drivers to the latest Dell supported versions, which are listed in the Supported Firmware and Software Matrix.

Run the following PowerShell command to retrieve the list of all driver versions that are installed on the local system:

```
Get-PnpDevice | Select-Object Name, @[{l='DriverVersion';e={ (Get-PnpDeviceProperty -
InstanceId $_.InstanceId -KeyName 'DEVPKEY_Device_DriverVersion').Data}} -Unique | Where-
Object {($_.Name -like "*HBA*") -or ($_.Name -like "*mellanox*") -or ($_.Name -like
"*Qlogic*") -or ($_.Name -like "*X710*") -or ($_.Name -like "*intel*") -or ($_.Name
-like "*Broadcom*")}]
```

Run the following PowerShell command to check the chipset driver installation status. If there is an error, install the chipset driver.

```
Get-PnpDevice -PresentOnly | Where-Object {($_.Status -ne 'OK') -and ($_.Problem -ne
'CM_PROB_NONE' -and $_.Problem -ne 'CM_PROB_DISABLED')}
```

After you identify the required driver version, including for the chipset and the HBA, download the driver installers from <https://www.dell.com/support> or by using the Dell EMC Repository Manager (DRM) as described in [Obtaining the firmware catalog for AX nodes or Ready Nodes using the Dell EMC Repository Manager](#).

After the drivers are downloaded, copy the identified drivers to AX nodes from where you can manually run the driver DUP files to install the drivers and restart the node.

Alternatively, to install the drivers silently, go to the folder and run the following command: `DriverUpdate.EXE /s /f`

Exiting the AX node from maintenance mode

After updating the AX node, exit the storage maintenance mode and node maintenance mode by running the following commands:

```
Get-StorageFaultDomain -type StorageScaleUnit | Where-Object {$_ .FriendlyName -eq
"<Hostname>"} | Disable-StorageMaintenanceMode
```

```
Resume-ClusterNode -Name "Hostname" -Failback Immediate
```

These commands initiate the operation of rebuilding and rebalancing the data to ensure load balancing.

For the remaining cluster nodes, repeat the preceding procedures for conducting maintenance operations.

Restarting a cluster node or taking a cluster node offline

About this task

Use the following procedure to restart a cluster node or to take a cluster node offline for maintenance:

Steps

1. Verify the health status of your cluster and volumes:
 - `Get-StorageSubSystem -FriendlyName *Cluster* | Get-StorageHealthReport`
 - `Get-physicaldisk`
 - `Get-virtualdisks`
2. Suspend the cluster node:
 - `Suspend-ClusterNode -name "Hostname" -Drain`
3. Enable storage maintenance mode:
 - `Get-StorageFaultDomain -type StorageScaleUnit | Where-Object {$_.FriendlyName -eq "<Hostname>"} | Enable-StorageMaintenanceMode`
4. Restart the server or shut it down for maintenance.
5. Disable storage maintenance mode.
 - `Get-StorageFaultDomain -type StorageScaleUnit | Where-Object {$_.FriendlyName -eq "<Hostname>"} | Disable-StorageMaintenanceMode`
6. Resume the cluster node:
 - `Resume-ClusterNode -Name "Hostname" -Failback Immediate`

Results

For more information, see [Taking a Storage Spaces Direct server offline for maintenance](#).

Expanding the Azure Stack HCI cluster

Expanding cluster compute or storage capacity are tasks performed during cluster operations. This section provides instructions for performing these tasks.

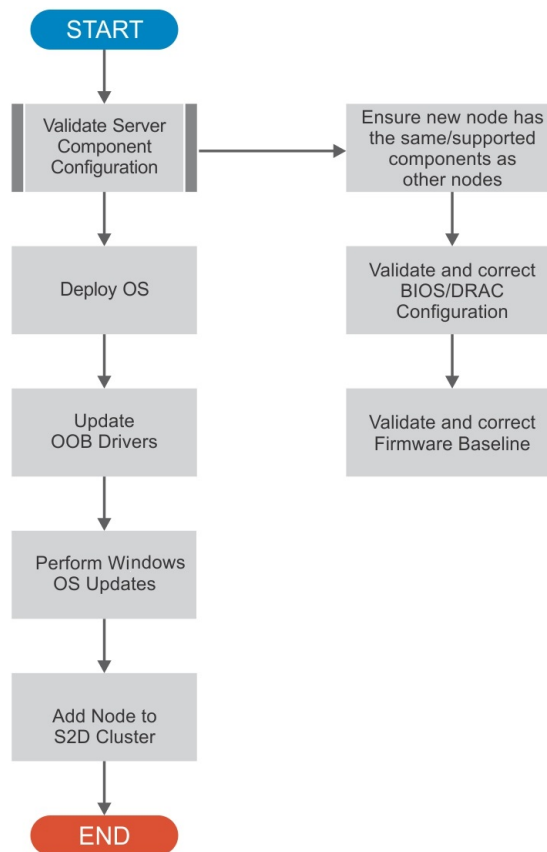


Figure 21. Expanding the Azure Stack HCI cluster

Azure Stack HCI node expansion

In an HCI cluster, adding server nodes increases the storage capacity, improves the overall storage performance of the cluster, and provides more compute resources to add VMs. Before adding new server nodes to an HCI cluster, complete the following requirements:

- Verify that the processor model, HBA, and NICs are of the same configuration as the current nodes on the cluster and PCIe slots.
- Ensure that all disk types and the amount in each node are the same as the node in use. Do not combine two different disk types in the same cluster or node. For example, you cannot combine SATA and SAS HDD/SSD drives in the same node or cluster. The following table shows the supported options for expanding storage capacity of the cluster.


Table 2. Options to expand storage capacity of the cluster

Option 1 conditions	Option 2 conditions
<ul style="list-style-type: none"> ○ Drive is listed in the Support Matrix ○ Same drive manufacturer ○ Same capacity and endurance ○ Latest model ○ Latest firmware 	<ul style="list-style-type: none"> ○ Drive is listed in the Support Matrix ○ Different drive manufacturer ○ Same capacity and endurance ○ Different model ○ Different firmware

- Ensure that the BIOS, drivers, firmware, and chipset are as listed in the support matrix.
- Apply the BIOS configuration to the node and configure iDRAC. For more information about configuring the node, see the [Dell EMC HCI Solutions for Microsoft Windows Server Deployment Guide](#). Do not run the PowerShell commands in the following sections of the deployment guide again because the cluster is already created, Storage Spaces Direct is already enabled, and the management network is already excluded:
 - Creating the host cluster
 - Enabling Storage Spaces Direct
 - Configuring the host management network as a lower priority network for live migration
- Ensure that the following tasks are completed:

1. Pass cluster validation and SES device compliance tests.
2. Verify that the nodes are compliant with the firmware baseline.
3. Update the hardware timeout configuration for the Spaces port.
4. After the node configuration, update Microsoft Windows to bring the node to the same level as the cluster.

Adding server nodes manually

 **NOTE:** The procedure is applicable only if the cluster and Storage Spaces Direct configuration is done manually.

To manually add server nodes to the cluster, see <https://technet.microsoft.com/windows-server-docs/storage/storage-spaces/add-nodes>.

Storage Spaces Direct storage expansion

In an HCI cluster, expanding storage by adding drives on the available slots on the cluster nodes adds storage capacity to the cluster and improves storage performance. Before the storage expansion, ensure that all disk types and the amount in each node are the same and are equal to that of the node in use. Do not combine two different disk types in the same cluster or node. For example, you cannot combine SATA and SAS HDD/SSD drives in the same node or cluster.

The following options for expanding the storage capacity of the cluster are supported:

- Option 1: Expand the storage with the same drive manufacturer, capacity, endurance, latest model, and latest firmware. Determine if it is available on the AX node support matrix.
- Option 2: Expand the storage with a different drive manufacturer, model, firmware, and the same capacity and endurance. Determine if it is available on the AX node support matrix.

When new disks are added to extend the overall storage capacity per node, the Azure Stack HCI cluster starts claiming the physical disks into an existing storage pool.

After the drives are added, they are shown as available for pooling (CanPool set to True) in the output of the `Get-PhysicalDisk` command.

Within a few minutes, the newly added disks are claimed in the existing pool and Storage Spaces Direct starts the rebalance job. Run the following command to verify that the new disks are a part of the existing pool:

```
PS C:\> Get-StorageSubSystem -FriendlyName *Cluster* | Get-StorageHealthReport
CPUUsageAverage           :    2.66 %
CapacityPhysicalPooledAvailable :    8.01 TB
CapacityPhysicalPooledTotal   :   69.86 TB
CapacityPhysicalTotal        :   69.86 TB
CapacityPhysicalUnpooled     :         0 B
CapacityVolumesAvailable     :   15.09 TB
CapacityVolumesTotal         :   16.88 TB
IOLatencyAverage            :   908.13 us
IOLatencyRead               :         0 ns
IOLatencyWrite              :   908.13 us
IOPSRead                    :         0 /S
IOPSTotal                   :         1 /S
IOPSWrite                   :         1 /S
IOThroughputRead            :         0 B/S
IOThroughputTotal           :   11.98 KB/S
IOThroughputWrite           :   11.98 KB/S
MemoryAvailable             :   472.87 GB
MemoryTotal                 :    768 GB
```

After all available disks are claimed in the storage pool, the `CapacityPhysicalUnpooled` is 0 B.

The storage rebalance job might take a few minutes. You can monitor the process by using the `Get-StorageJob` cmdlet.

Extending volumes

You can resize volumes that are created in Spaces Direct storage pools by using the `Resize-VirtualDisk` cmdlet. For more information, see <https://technet.microsoft.com/windows-server-docs/storage/storage-spaces/resize-volumes>.

Performing AX node recovery

If a cluster node fails, perform node operating system recovery in a systematic manner to ensure that the node is brought up with the configuration that is consistent with other cluster nodes.

The following sections provide details about operating system recovery and post-recovery configuration that is required to bring the node into an existing Azure Stack HCI cluster.

NOTE: To perform node recovery, ensure that the operating system is reinstalled.

Configuring RAID for operating system drives

Prerequisites

The Dell EMC PowerEdge servers offer the Boot Optimized Storage Solution (BOSS) controller as an efficient and economical way to separate the operating system and data on the internal storage of the server. The BOSS solution in the latest generation of PowerEdge servers uses one or two BOSS M.2 SATA devices to provide RAID 1 capability for the operating system drive.

NOTE: All Dell EMC Solutions for Azure Stack HCI are configured with hardware RAID 1 for the operating system drives on BOSS M.2 SATA SSD devices. The steps in this section are required only when recovering a failed cluster node. Before creating a new RAID, the existing or failed RAID must be deleted.

About this task

This procedure describes the process of creating operating system volumes.

Steps

1. Log in to the iDRAC web interface.
2. Go to **Storage > Controllers**.

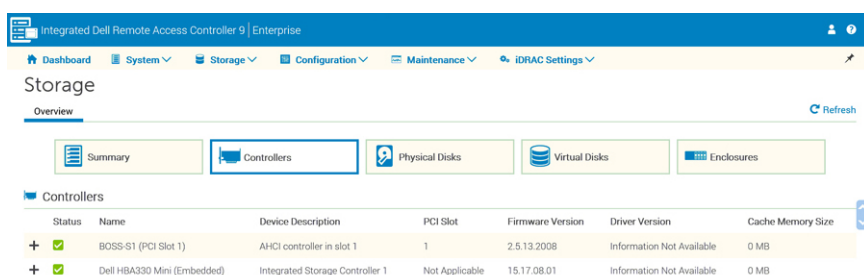


Figure 22. View controllers

3. Go to **Configuration > Storage Configuration > Virtual Disk Configuration**, and then click **Create Virtual Disk**.

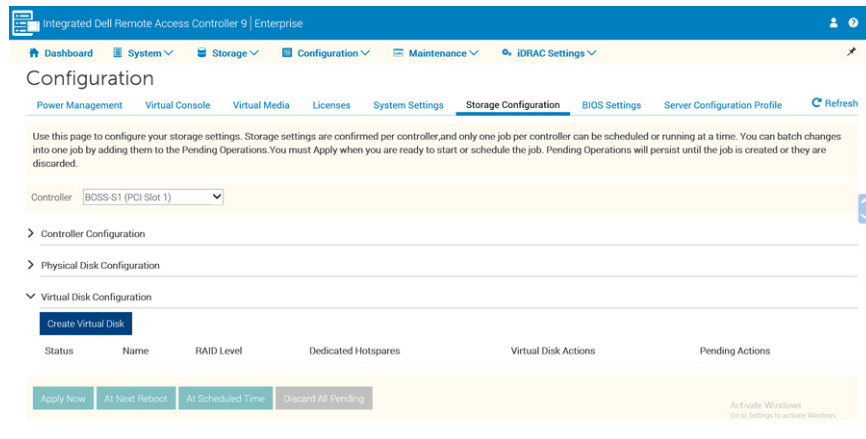


Figure 23. Create a virtual disk

4. Provide a virtual disk name and select BOSS M.2 devices in the physical disks.

Create Virtual Disk

Name

OS_VD

×

Layout

RAID-1

▼

Media Type

SSD

▼

Stripe Element Size

64 KB

▼

Read Policy

No Read Ahead

▼

Write Policy

Write Through

▼

Disk Cache Policy

Default

▼

Span Count

1

▼

A system reboot may be required to create a Virtual Disk.

The available Virtual Disk settings are limited to match the system configuration.

Cancel

Add to Pending Operations

Figure 24. Provide virtual disk name

Create Virtual Disk

A system reboot may be required to create a Virtual Disk.

The available Virtual Disk settings are limited to match the system configuration.

RAID 10, 50 & 60 Make sure to select specific sets of physical disks for the required configurations.

Number of Physical Disks - Minimum : [2] Maximum : [2] Current Selection : [2]

Virtual Disk Size - Minimum : [223.57 GB] Maximum : [223.57 GB] Specified Size : [223.57 GB]

The **Span Count** can only be adjusted for **RAID 50 & 60** after selecting the physical disks.

A Red Diamond icon highlights the physical disks that are already supporting one or more virtual disks.

Select Physical Disks

Status	Name	Capacity	Media Type
<input checked="" type="checkbox"/>	SSD 0	223.57 GB	SSD
<input checked="" type="checkbox"/>	SSD 1	223.57 GB	SSD

Cancel

Add to Pending Operations

Figure 25. Set Physical Disks

- Click **Add Pending Operations**.
- Go to **Configuration > Storage Configuration > Virtual Disk Configuration**.

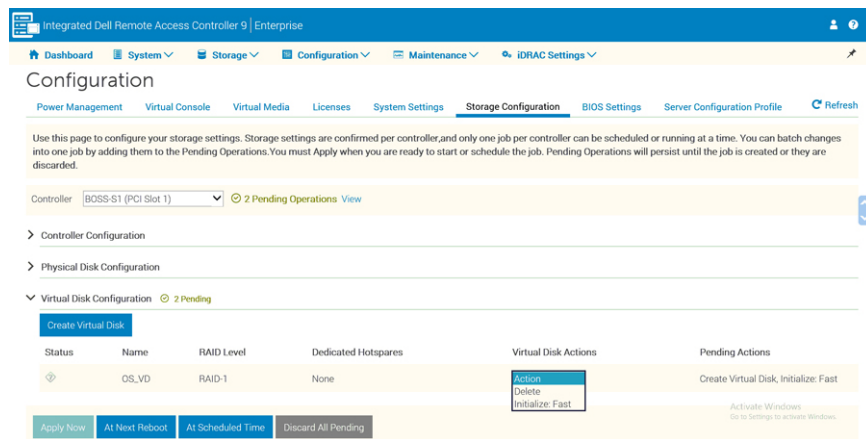


Figure 26. Initialize configuration

- Select the virtual disk, and then select **Initialize: Fast in Virtual Disk Actions**.
- Reboot the server.
- NOTE:** The virtual disk creation process might take several minutes to complete.
- After the initialization is completed successfully, the virtual disk health status is displayed.

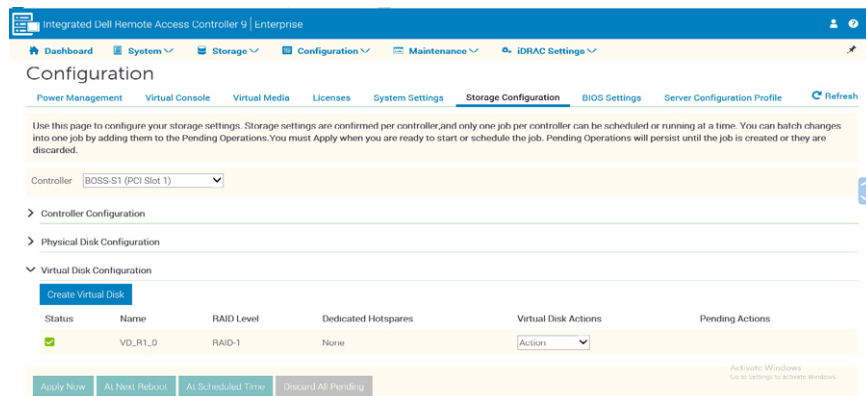


Figure 27. Virtual disk health status

Operating system recovery

This section provides an overview of steps involved in operating system recovery on the Dell EMC Solutions for Azure Stack HCI.

- NOTE:** Ensure that the RAID 1 VD created on the BOSS M.2 drives is reinitialized.
- NOTE:** Do not reinitialize or clear the data on the disks that were a part of the Storage Spaces Direct storage pool. This helps to reduce repair times when the node is added back to the same cluster after recovery.

Manual operating system recovery

For manually deployed nodes, you can recover the operating system on the node by using any of the methods that were used for operating system deployment.

Factory operating system recovery

For the factory-installed OEM license of the operating system, Dell Technologies recommends that you use the operating system recovery media that shipped with the PowerEdge server. Using this media for operating system recovery ensures that the operating system stays activated after the recovery. Using any other operating system media triggers the need for activation after operating system deployment. Operating system deployment using the recovery media is the same as either retail or other operating system media-based installation.

After completing the operating system deployment using the recovery media, perform the following steps to bring the node into an existing Azure Stack HCI cluster:

1. Update CPU chipset, network, and storage drivers.
2. Configure host networking.
3. Change the hostname.
4. Perform AD Domain Join.
5. Configure the QoS policy (for RoCE for RDMA only).
6. Configure RDMA.
7. Configure the firewall.
8. Perform Day 0 operating system updates.
9. Add server nodes to the cluster.

For instructions on steps 1 through 7, see the [Dell EMC HCI Solutions for Microsoft Windows Server Deployment Guide](#).

SupportAssist Enterprise

Dell SupportAssist Enterprise (SAE) version (2.0.70.62) automates support by proactively identifying hardware and software issues. When an issue is detected, SupportAssist notifies you about the issue and automatically creates a Support Request with Dell EMC. It will also attach server hardware logs required for troubleshooting and efficient issue resolution.


For more information, see [Dell EMC SupportAssist Enterprise 2.x – Guide and Download](#).

iDRAC Service Module (iSM) for AX nodes and Storage Spaces Direct Ready Nodes

The iDRAC Service Module (iSM) is a lightweight software module that you can install on AX nodes to complement the iDRAC interfaces—the user interface (UI), RACADM CLI, Redfish, and Web Services-Management (WS-Man)—with additional monitoring data.

About this task

Follow these steps to install iSM on the operating system.

 **NOTE:** The ISM application package will be installed as part of firmware and driver updates using the ASHCI catalog.

Steps

1. Log in to iDRAC-->iDRAC Settings-->Settings-->iDRAC Service Module Setup
2. Start the virtual console.
3. Log in to the host operating system as an administrator.
4. From the device list, select the mounted volume that is identified as SMINST and then click the ISM_win.bat script to start the installation.

Results

After the installation is completed, iDRAC indicates that the iSM is installed and specifies the latest installation date.

FullPowerCycle

FullPowerCycle is a calling interface function that provides a way to reset the server auxiliary power. An increasing amount of server hardware runs on server auxiliary power. Troubleshooting some server issues requires you to physically unplug the server power cable to reset the hardware running on auxiliary power.

The FullPowerCycle feature enables the administrator to connect or disconnect auxiliary power remotely without visiting the data center. This feature is supported on AX nodes and Storage Spaces Direct Ready Nodes.

These are the relevant commands to run in the PowerShell console:

- To request FullPowerCycle on your system: `Invoke-FullPowerCycle -status request`
- To get the status of FullPowerCycle on your system: `Invoke-FullPowerCycle -status Get`
- To cancel FullPowerCycle on your system: `Invoke-FullPowerCycle -status cancel`