

# Microsoft HCI Solutions from Dell Technologies

## Deployment Guide

### Abstract

This guide focuses on deploying a scalable hyperconverged infrastructure that is built by using the validated and certified Dell Storage Spaces Direct Ready Nodes, Windows Server 2016, Windows Server 2019, Windows Server 2022, and Azure Stack HCI operating system versions 20H2 and 21H2.

## 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.

<b>Chapter 1: Introduction.....</b>	<b>5</b>
Document overview.....	5
Audience and scope.....	5
Known issues.....	6
<b>Chapter 2: Solution Overview.....</b>	<b>7</b>
Solution introduction.....	7
Deployment models.....	7
Switchless storage networking.....	7
Scalable infrastructure.....	8
Stretched cluster infrastructure.....	10
Single node cluster.....	10
Solution integration and network connectivity.....	10
Secured-Core.....	11
<b>Chapter 3: Solution Deployment.....</b>	<b>12</b>
Solution deployment introduction.....	12
Deployment prerequisites.....	12
Predeployment configuration.....	13
Configuring network switches.....	13
Configuring iDRAC and BIOS.....	13
Configure QLogic NICs.....	14
Operating system deployment.....	14
Manual operating system deployment.....	15
Factory-installed operating system deployment.....	15
Upgrade using sConfig (Stand-alone systems).....	16
Install roles and features.....	16
Verifying firmware and software compliance with the support matrix.....	16
Update out-of-box drivers.....	17
Change the hostname.....	17
Configuring host networking.....	17
Registry settings.....	18
Joining cluster nodes to an Active Directory domain.....	19
Deploying and configuring a host cluster.....	19
Creating the host cluster.....	19
Enabling Storage Spaces Direct.....	20
Configuring the host management network as a lower-priority network for live migration.....	20
Updating the page file settings.....	20
Configuring a cluster witness.....	21
Azure onboarding for Azure Stack HCI operating system.....	21
Best practices and recommendations.....	21
Disable SMB Signing.....	21
Update the hardware timeout for the Spaces port.....	21
Enable jumbo frames.....	22

Recommended next steps.....	22
Deployment services.....	22
<b>Chapter 4: References.....</b>	<b>23</b>
Dell Technologies documentation.....	23
Microsoft documentation.....	23
<b>Chapter 5: GPU Integration.....</b>	<b>24</b>
Installing and configuring GPUs on AX-750/7525.....	24
AX-750 GPU configuration.....	24
AX-7525 GPU configuration.....	24
GPU physical installation.....	25
Deploy GPU devices using Discrete Device Assignment (DDA).....	25
<b>Appendix A: Persistent Memory for Windows Server HCI.....</b>	<b>33</b>
Configuring persistent memory for Windows Server HCI.....	33
Persistent memory requirements.....	33
Configuring persistent memory BIOS settings.....	34
Configuring Windows Server HCI persistent memory hosts.....	35
Managing persistent memory using Windows PowerShell.....	36

# Introduction

## Topics:

- [Document overview](#)
- [Audience and scope](#)
- [Known issues](#)

## Document overview

This deployment guide provides an overview of Microsoft HCI Solutions from Dell Technologies, guidance on how to integrate solution components, and instructions for preparing and deploying the solution infrastructure. Microsoft HCI Solutions from Dell Technologies includes:

- Dell EMC Integrated System for Microsoft Azure Stack HCI (based on Azure Stack HCI OS v20H2 or v21H2)
- Dell EMC HCI Solutions for Microsoft Windows Server (based on Windows Server 2016/2019/2022 OS)

For end-to-end deployment steps, use the information in this guide with the information in [Reference Guide-Network Integration and Host Network Configuration Options](#).

This guide applies only to infrastructure that is built using the validated and certified AX nodes from Dell Technologies and Windows Server 2016, Windows Server 2019, Windows Server 2022, and the Azure Stack HCI operating system.

## Audience and scope

The audience for this document includes systems engineers, field consultants, partner engineering team members, and customers with knowledge about deploying hyperconverged infrastructures (HCIs) with Windows Server 2016, Windows Server 2019, Windows Server 2022, the Azure Stack HCI operating system, Hyper-V, and Storage Spaces Direct. Customers who do not have Volume License agreements with Microsoft can order AX nodes from Dell Technologies with a factory-installed operating system and OEM license or as a bare-metal installation.

An Azure Stack HCI cluster can be deployed in the following ways:

- Dell Technologies Services-led—Certified deployment engineers ensure accuracy, speed, and reduced risk and downtime.
- Customer led—Customers who have the qualified level of technical expertise follow the instructions in this guide.
- Using Windows Admin Center—Deployment services engineers or customers perform the Azure Stack HCI solution deployment using Windows Admin Center.

**i NOTE:** Instructions in this deployment guide are applicable only to:

- The generally available (GA) operating system build of Windows Server 2016 with the latest applicable updates
- The Windows Server 2019 GA build with the latest operating system updates
- The Windows Server 2022 GA build with the latest operating system updates
- The GA operating system builds of the Azure Stack HCI operating system with the latest cumulative update

These instructions are not validated with Windows Server version 1709. AX nodes from Dell Technologies do not support the Windows Server Semi-Annual Channel release. Dell Technologies recommends that you update the host operating system with the latest cumulative updates from Microsoft before starting the Azure Stack HCI cluster creation and configuration tasks.

## Assumptions

This guide assumes that deployment personnel understand the following technologies and tasks:

- AX nodes from Dell Technologies

- Deploying and configuring BIOS and integrated Dell Remote Access Controller (iDRAC) settings
- 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](#).

# Solution Overview

## Topics:

- [Solution introduction](#)
- [Deployment models](#)
- [Single node cluster](#)
- [Solution integration and network connectivity](#)
- [Secured-Core](#)

## Solution introduction

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

For information about supported AX nodes and operating system support for each of the AX nodes, see the [Support Matrix for Microsoft HCI Solutions](#).

The solutions are available in both hybrid and all-flash configurations. For more information about available configurations, see the [AX nodes specification sheet](#).

## Deployment models

Microsoft HCI Solutions from Dell Technologies offer the following types of cluster infrastructure deployments:

- Switchless storage networking
- Scalable infrastructure
- Stretched cluster infrastructure

**NOTE:** This guide does not provide deployment instructions for stretched cluster infrastructure. For information about this infrastructure, see the [Dell EMC Integrated System for Microsoft Azure Stack HCI: Stretched Cluster Deployment Reference Architecture Guide](#).

## Switchless storage networking

This Microsoft HCI Solutions from Dell Technologies infrastructure type offers two to four nodes in a switchless configuration for storage traffic. This infrastructure can be implemented using any of the validated and supported AX nodes. However, the number of nodes in a cluster varies between the AX node models and the number of network adapters that each model supports.

Switchless storage networking offers two full-mesh configurations:

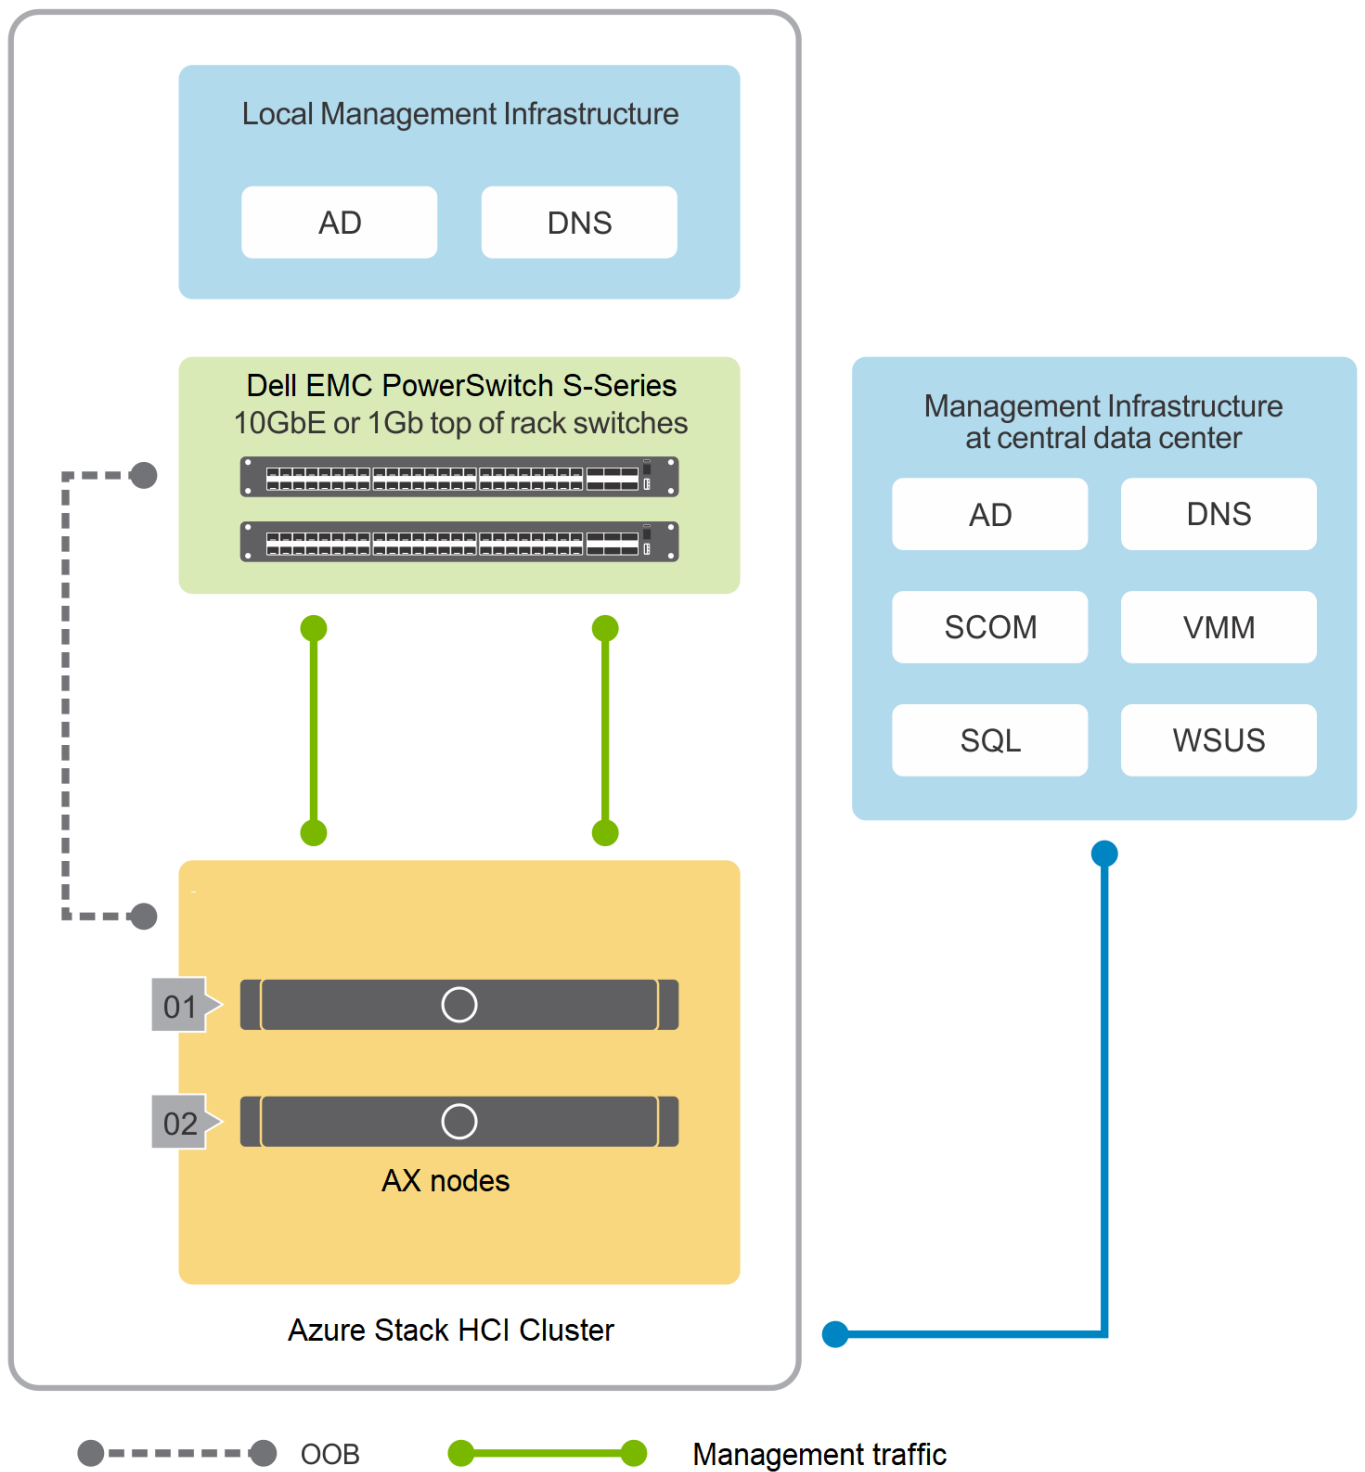
- Single-link
- Dual-link

For more information about these configurations, see [Network Integration and Host Configuration Options](#).

For the two-node cluster deployment, configure a cluster witness. For details, see [Configuring a cluster witness](#).

The switchless storage configuration has been validated with back-to-back connections for storage connectivity.

The following figure shows the switchless storage networking infrastructure:



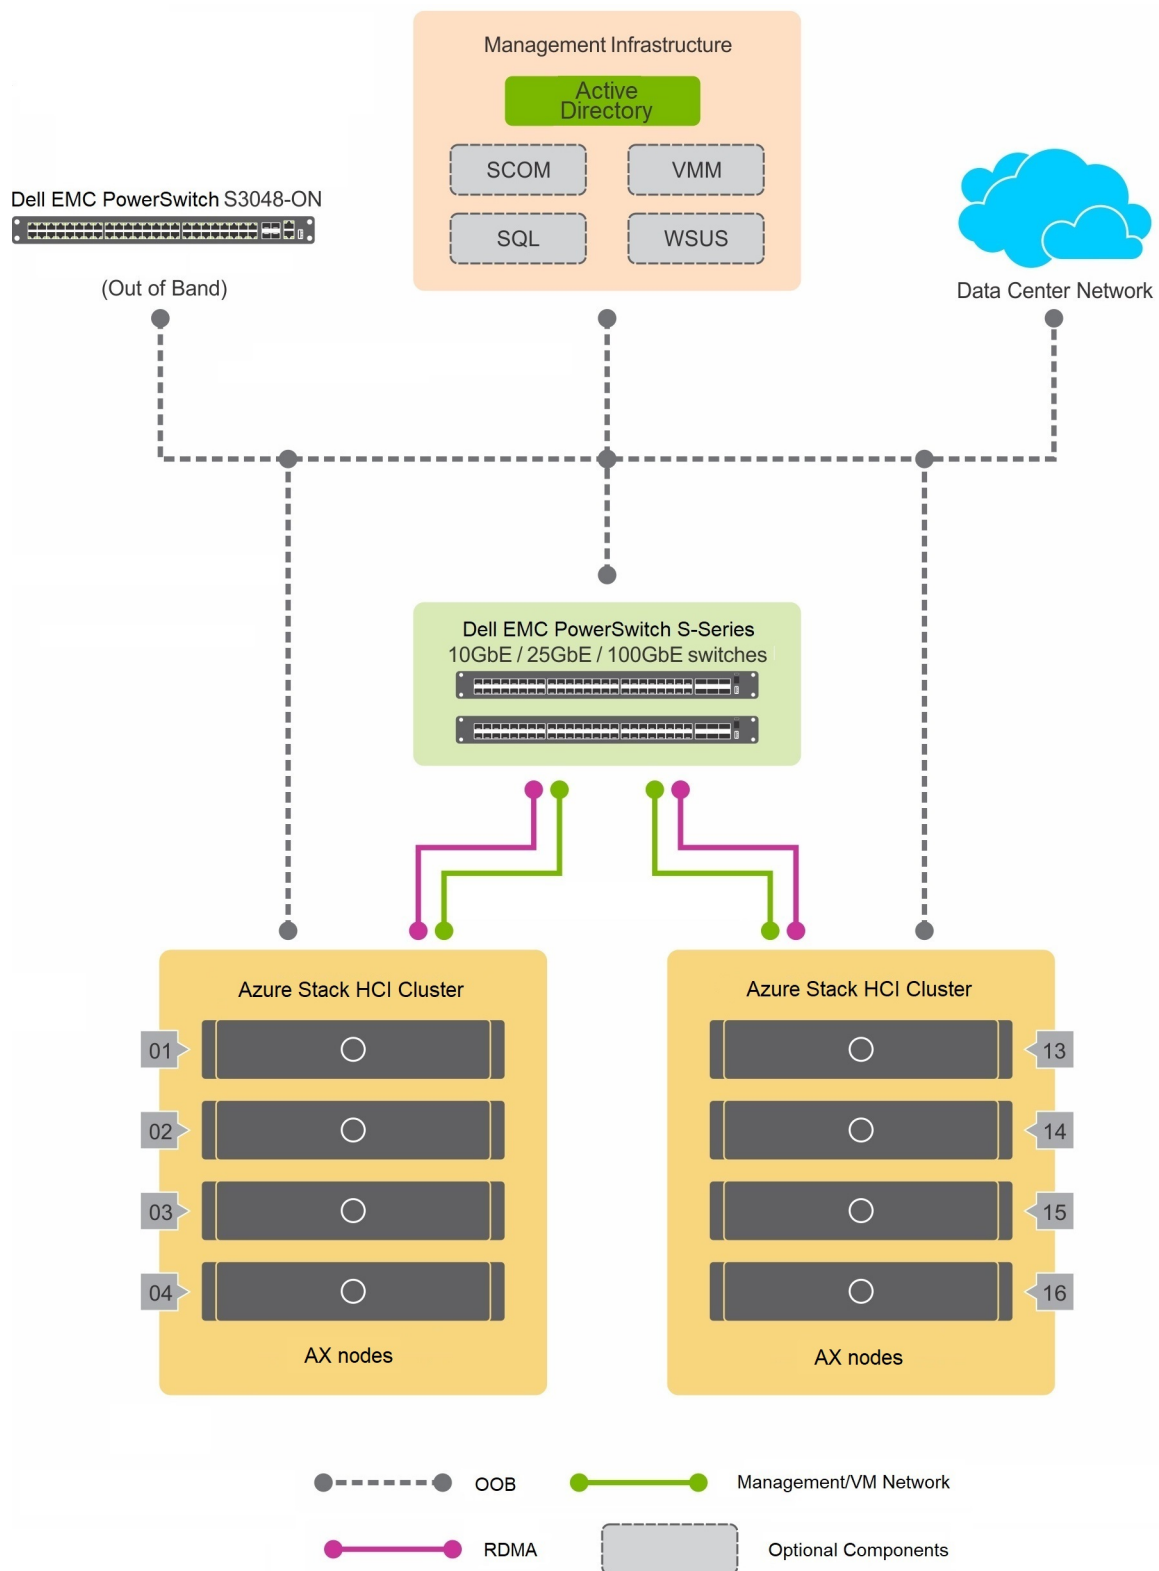
**Figure 1. Switchless storage networking**

## Scalable infrastructure

The scalable offering within Microsoft HCI Solutions from Dell Technologies encompasses various AX node configurations. In this Windows Server HCI solution, as many as 16 AX nodes power the primary compute cluster.

The following figure illustrates one of the flexible solution architectures. It includes the Azure Stack HCI cluster, redundant top-of-rack (ToR) switches, a separate out-of-band (OOB) network, and an existing management infrastructure in the data center.





**Figure 2. Scalable solution architecture**

Microsoft HCI Solutions from Dell Technologies does not include:

- Management infrastructure components such as a cluster for hosting management VMs
- Services such as Microsoft Active Directory, Domain Name System (DNS), and Windows Server Update Services (WSUS)
- Microsoft System Center components such as Operations Manager (SCOM)

The instructions in this guide do not include deployment of any of these services and components, and they assume that at least an Active Directory domain controller is available in the existing management infrastructure. In a remote office scenario, Dell Technologies recommends that you deploy either an Active Directory replica or read-only domain controller (RODC) at the remote office. If you are using an RODC at the remote site, connectivity to the central management infrastructure with a writeable domain controller is mandatory during deployment of the Azure Stack HCI cluster.

**NOTE:** Dell Technologies does not support expansion of a two-node cluster to a larger cluster size. A three-node cluster provides fault-tolerance only for simultaneous failure of a single node and a single drive. If the deployment requires future expansion and better fault tolerance, consider starting with a four-node cluster at a minimum.

**NOTE:** For recommended server and network switch placement in the racks, port mapping on the top-of-rack (ToR) and OOB switches, and details about configuring the ToR and OOB switches, see [Reference Guide-Network Integration and Host Network Configuration Options](#).

This deployment guide provides instructions and PowerShell commands to manually deploy an Azure Stack HCI cluster. For information about configuring host networking and creating an Azure Stack HCI cluster by using System Center Virtual Machine Manager (VMM), see [Preparing and Using SCVMM for Azure Stack HCI Network and Cluster Configuration](#).

## Stretched cluster infrastructure

The Azure Stack HCI operating system supports disaster recovery between two sites using Azure Stack HCI clusters. With Storage Replica as its foundation, stretched clusters support both synchronous and asynchronous replication of data between two sites. The replication direction (unidirectional or bi-directional) can be configured for either an active/passive or active/active stretched cluster configuration.

**NOTE:** Stretched clustering infrastructure is supported only with the Azure Stack HCI operating system. For more information, see the [Dell EMC Integrated System for Microsoft Azure Stack HCI: Stretched Cluster Deployment Reference Architecture Guide](#).

## Single node cluster

Microsoft added a feature to support single node clusters on the Azure Stack HCI operating system. Single node clusters are similar to stand-alone Storage Spaces nodes but are delivered as an Azure service. Single node clusters support all the Azure services that a multi-node Azure Stack HCI cluster supports. On a single node cluster, there is no automatic provision of failover to another node. A physical disk is the fault domain in a single node cluster and only a single-tier configuration (All-Flash or All-NVMe) is supported.

While node-level High Availability cannot be supported on a single node cluster, you can choose application-level or VM-level replication to maintain High Availability in your infrastructure.

**NOTE:** Steps to set up an application-level or VM-level replication is beyond the scope of this document.

## Solution integration and network connectivity

Each of the variants in Microsoft HCI Solutions from Dell Technologies supports a specific type of network connectivity. The type of network connectivity determines the solution integration requirements.

- For information about all possible topologies within both fully converged and nonconverged solution integration, including with switchless storage networking and host operating system network configuration, see [Network Integration and Host Network Configuration Options](#).
- For switchless storage networking, install the server cabling according to the instructions detailed in [Cabling Instructions](#).
- For sample switch configurations for these network connectivity options, see [Sample Network Switch Configuration Files](#).

### Fully converged network connectivity

In the fully converged network configuration, both storage and management/VM traffic use the same set of network adapters. These adapters are configured with Switch-Embedded Teaming (SET). When using RoCE in a fully converged network configuration, you must configure Data Center Bridging (DCB).

The following table describes when to configure DCB based on the chosen network card and switch topology:

**Table 1. DCB configuration based on network card and switch topology**

Network card on node	Fully converged switch topology	Nonconverged switch topology	Switchless topology
Mellanox (RoCE)	DCB (required)	DCB (required) for storage adapters only	No DCB/QoS required
QLogic (iWARP)	DCB (required for All-NVMe configurations only)	No DCB	No DCB/QoS required

**NOTE:** Enable DCB only on storage (RDMA) adapters connected to the ToR switches that require it.

**NOTE:** Manually disable the DCB on the management adapters using the `Disable-NetAdapterQos <nicName>` command.

## Nonconverged network connectivity

In the nonconverged network configuration, storage traffic uses a dedicated set of network adapters either in a SET configuration or as physical adapters. A separate set of network adapters is used for management, VM, and other traffic classes. In this connectivity method, DCB configuration is optional for QLogic (iWARP), but mandatory for Mellanox (RoCE) adapters.

The switchless storage networking deployment model also implements nonconverged network connectivity without the need for network switches for storage traffic.

## Network connectivity for single node clusters

A single node cluster has adapters configured for only management and VM traffic. However, Azure Stack HCI Engineering still recommends configuring a virtual network interface for Live Migration in case you intend to use Shared Nothing Live Migration to move workloads to other clusters at a later time. You can also use the adapter to configure application or VM replication. For guidance on network configurations/topologies, see [Network Integration and Host Network Configuration Options](#).

## Secured-Core

Secured-Core is a Windows Server security feature available in Windows Server 2022 and Azure Stack HCI OS 21H2. Enabling Secured-Core involves modifying BIOS and operating system level settings.

The AX-7525 (with AMD Epic 7xx3 Milan CPUs), AX-650, and AX-750 platforms support the required BIOS settings to enable the Secured-Core feature. Also, each of these platforms ship with the required TPM 2.0 v3 hardware installed from the factory.

You can enable operating system-level settings using Windows Admin Center.

To enable Secured-Core, see [Managing and Monitoring the Solution Infrastructure Life Cycle Operations Guide](#)

# Solution Deployment

## Topics:

- [Solution deployment introduction](#)
- [Deployment prerequisites](#)
- [Predeployment configuration](#)
- [Operating system deployment](#)
- [Install roles and features](#)
- [Verifying firmware and software compliance with the support matrix](#)
- [Update out-of-box drivers](#)
- [Change the hostname](#)
- [Configuring host networking](#)
- [Registry settings](#)
- [Joining cluster nodes to an Active Directory domain](#)
- [Deploying and configuring a host cluster](#)
- [Best practices and recommendations](#)
- [Recommended next steps](#)
- [Deployment services](#)

## Solution deployment introduction

Microsoft HCI Solutions from Dell Technologies can be deployed in the following ways:

- Manual operating system deployment—Begin by manually installing the operating system on AX nodes.
- Factory-installed operating system deployment—Begin with AX nodes that have factory-installed Windows Server 2019 with the Desktop Experience feature or the Azure Stack HCI operating system. The core edition is not available as a factory-installed operating system for Windows Server 2019.

Each deployment method has its own prerequisites, including configuring network switches, as described in this guide.

**NOTE:** Instructions in this deployment guide are applicable to only the following:

- Windows Server 2016 generally available (GA) build with the latest applicable updates
- Windows Server 2019 GA build
- Azure Stack HCI operating system, version 20H2 and 21H2
- Windows Server 2022

AX nodes from Dell Technologies do not support the Windows Server Semi-Annual Channel release. Dell Technologies recommends that you update the host operating system with the latest cumulative updates from Microsoft before starting the cluster creation and configuration tasks.

**NOTE:** Each task in this deployment guide requires running one or more PowerShell commands. Dell Technologies recommends using these commands to complete the deployment tasks because the UI might not work as expected in some scenarios. For example, the cluster validation UI wizard within the Failover Cluster Manager fails intermittently due to a known issue in the Microsoft code.

## Deployment prerequisites

Dell Technologies assumes that the management services that are required for the operating system deployment and cluster configuration are in the existing infrastructure where the Azure Stack HCI cluster is being deployed.

The following table describes the management services:

**Table 2. Management services**

Management service	Purpose	Required/optional
Active Directory	User authentication	Required
Domain Name System	Name resolution	Required
Windows Software Update Service (WSUS)	Local source for Windows updates	Optional
SQL Server	Database back-end for System Center VMM and System Center Operations Manager (SCOM)	Optional

## Predeployment configuration

Before deploying AX nodes, complete the required predeployment configuration tasks.

Complete the following predeployment configuration before deploying the Azure Stack HCI solution.

**NOTE:** If the cluster has persistent memory devices, preoperating-system deployment configuration, as well as postoperating-system deployment configuration is required. For more information, see [Appendix: Persistent Memory for Azure Stack HCI](#).

## Configuring network switches

Based on the selected network topology, configure the ToR network switches to enable storage and VM/management traffic.

**NOTE:** The switchless storage networking deployment model requires configuration of switches that are deployed for host management and OOB traffic only. Storage traffic uses full-mesh connectivity between the nodes.

**NOTE:** Management network redundancy is a combination of either iDRAC or operating system DNS/IP resolution. Dell Technologies recommends that you deploy a network topology that supports a dual control plane while sharing a single data plane. Virtual Link Trunking (VLT) is Dell Technologies proprietary technology that provides network resiliency for data I/O.

Configuring switch VLT redundancy with Virtual Router Redundancy Protocol (VRRP) provides a virtual floating IP address that any node can reference as a gateway. If a switch fails, the virtual IP address is transferred to a peer switch.

VRRP is an active/standby, first-hop redundancy protocol (FHRP). When used with VLT peers, VRRP becomes an active/active protocol. The VRRP virtual MAC address is the local destination address in the forwarding information base (FIB) table of both VLT peers. Using this method, the backup VRRP router forwards intercepted frames that have a destination MAC address that matches the VRRP virtual MAC address.

A standard Storage Spaces Direct deployment requires three basic types of networks: OOB management, host management, and storage. The number of network ports (two or four) that are used within the storage configuration determines whether you have two or four fault domains.

For sample switch configurations, see [Microsoft Azure Stack HCI Solutions from Dell Technologies: Switch Configurations – RoCE Only \(Mellanox Cards\)](#) and [Microsoft Azure Stack HCI Solutions from Dell Technologies: Switch Configurations – iWARP Only \(Qlogic Cards\)](#).

For configuration choices and instructions about different network topologies and host network configuration, see [Network Integration and Host Network Configuration Options](#).

## Configuring iDRAC and BIOS

The AX nodes are factory-configured with optimized system BIOS and iDRAC settings. This preconfiguration eliminates the need for you to manually configure the settings to a recommended baseline.

The iDRAC in AX nodes can be configured to obtain an IP address from DHCP or can be assigned a static IP address. When the OOB network in the environment does not provide DHCP IP addresses, you must manually set a static IPv4 address on each iDRAC network interface. You can access the physical server console to set the addresses by using KVM or other means.

## Configure BIOS settings including the IPv4 address for iDRAC

Perform these steps to configure the IPv4 address for iDRAC. You can also perform these steps to configure any additional BIOS settings.

### Steps

1. During the system boot, press F12.
2. At **System Setup Main Menu**, select **iDRAC Settings**.
3. Under **iDRAC Settings**, select **Network**.
4. Under **IPV4 SETTINGS**, at **Enable IPv4**, select **Enabled**.
5. Enter the static IPv4 address details.
6. Click **Back**, and then click **Finish**.

## Configure QLogic NICs

The QLogic FastLinQ 41262 network adapter supports both iWARP and RoCE.

When used with the QLogic network adapters, the AX nodes are validated only with iWARP. Manually configure the adapter to enable iWARP based on the chosen network configuration.

## Configure the QLogic NIC


Configure the QLogic network adapter for each port.

### Steps

1. During system startup, press F2 to enter System Setup.
2. Click **System BIOS** and select **Device Settings**.
3. Select the QLogic network adapter from the list of adapters.
4. Click **Device Level Configuration** and ensure that **Virtualization Mode** is set to **None**.
5. Click **Back**, and then click **NIC Configuration**.
6. On the **NIC Configuration** page, select the following options:
  - Link Speed: SmartAN
  - NIC + RDMA Mode: Enabled
  - RDMA Operational Mode: iWARP
  - Boot Protocol: None
  - Virtual LAN Mode: Disabled
7. Click **Back**, and then click **Data Center Bridging (DCB) Settings**.
8. On the **Data Center Bridging (DCB) Settings** page, set **DCBX Protocol** to **Disabled**.
9. Click **Back**, click **Finish**, and then click **Yes** to save the settings.
10. Click **Yes** to return to the **Device Settings** page.
11. Select the second port of the QLogic adapter and repeat the preceding steps.
12. Click **Finish** to return to the **System Setup** page.
13. Click **Finish** to reboot the system.

## Operating system deployment

These instructions are for manual deployment of the Windows Server 2016, Windows Server 2019, Windows Server 2022, or Azure Stack HCI operating system version 20H2 or 21H2 on AX nodes from Dell Technologies. Unless otherwise specified, perform the steps on each physical node in the infrastructure that will be a part of Azure Stack HCI.

 **NOTE:** The steps in the subsequent sections are applicable to either the full operating system or Server Core.

**NOTE:** The command output that is shown in the subsequent sections might show only Mellanox ConnectX-4 LX adapters as physical adapters. The output is shown only as an example.

## Manual operating system deployment

Dell Lifecycle Controller and iDRAC provide operating system deployment options. Options include manual installation or unattended installation by using virtual media and the operating system deployment feature in Lifecycle Controller for Windows Server 2016, Windows Server 2019, Windows Server 2022, and Azure Stack HCI operating system version 20H2 and 21H2.

A step-by-step procedure for deploying the operating system is not within the scope of this guide.

The remainder of this guide assumes that:

- Windows Server 2016, Windows Server 2019, Windows Server 2022, or Azure Stack HCI operating system version 20H2 or 21H2 installation on the physical server is complete.
- You have access to the iDRAC virtual console of the physical server.

**NOTE:** For information about installing the operating system using the iDRAC virtual media feature, see the ["Using the Virtual Media function on iDRAC 6, 7, 8 and 9"](#) Knowledge Base article.

**NOTE:** The Azure Stack HCI operating system is based on Server Core and does not have the full user interface. For more information about using the Server Configuration tool (Sconfig), see [Deploy the Azure Stack HCI operating system](#).

## Factory-installed operating system deployment

If the cluster nodes are shipped from Dell Technologies with a preinstalled operating system, complete the out-of-box experience (OOBE):

- Select language and locale settings.
- Accept the Microsoft and OEM EULAs.
- Set up a password for the local administrator account.
- Update the operating system partition size, and shrink it as needed.

**NOTE:** Partition size adjustment is not available with the factory-installed Azure Stack HCI operating system.

The OEM operating system is preactivated and the Hyper-V role is predeployed for Windows Server 2019/2022. After completing the OOBE steps, perform the steps in [Install roles and features](#) to complete the cluster deployment and Storage Spaces Direct configuration. For the Azure Stack HCI operating system, roles and features are preinstalled.

The Azure Stack HCI operating system factory image has multilingual support for these languages: English, German, French, Spanish, Korean, Japanese, Polish, and Italian. To change language:

- Run the following PowerShell commands. <LANGUAGE> can be en-US, fr-FR, ja-JP, ko-KR, de-DE, pl-PL, it-IT, or es-ES.
  - `Set-WinUserLanguageList <LANGUAGE>`
  - `Set-WinSystemLocale -systemlocale <LANGUAGE>`
- Reboot after running the commands.

**NOTE:** For all of the languages except Polish, the main screen changes to the appropriate font. For Polish, the English font is used.

If you purchased a license for a secondary operation system to run your virtual machines, the VHD file is located in the `C:\Dell_OEM\VM` folder. Copy this VHD file to a virtual disk in your Azure Stack HCI cluster and create virtual machines using this VHD.

**NOTE:** Do not run your virtual machine with the VHD file residing on the BOSS device (for example, `c:\`). It should always be on `c:\clusterstorage\<VD1>`

## Upgrade using sConfig (Stand-alone systems)

Using a **sconfig** menu (Server Configuration from command prompt), you can update the servers one at a time. Customers who receive HCI OS 20H2 may update to 21H2 using this method so that a cluster can be created on 21H2 without having to upgrade after cluster creation.

### Steps

1. On the sconfig menu, select option 6 and update all quality updates (at least to October 19, 2021 CU (KB5006741). This step is to ensure that you are at a supported KB for 21H2 update.
2. Once all quality updates are completed, go to **Feature Updates** on the sconfig menu and perform an OS upgrade from 20H2 to 21H2. After completing the OS upgrade, perform step 1 to install all the quality updates for 21H2. You may have to run this multiple times to get to the latest cumulative update.
3. Use Windows Admin Center to update each node to the latest hardware support matrix. See [Deployment Guide-Creating an Azure Stack HCI cluster using Windows Admin Center](#).
4. Once the operating system on all nodes is updated to the latest CU of 21H2, you may go to creating the cluster using PowerShell or Windows Admin Center.

## Install roles and features


Deployment and configuration of a Windows Server 2016, Windows Server 2019, Windows Server 2022, or Azure Stack HCI operating system version 20H2 or 21H2 cluster requires enabling specific operating system roles and features.


Enable the following roles and features:

- Hyper-V service (not required if the operating system is factory-installed)
- Failover clustering
- Data center bridging (DCB) (required only when implementing fully converged network topology with RoCE and when implementing DCB for the fully converged topology with iWARP)
- BitLocker (optional)
- File Server (optional)
- FS-Data-Deduplication module (optional)
- RSAT-AD-PowerShell module (optional)

Enable these features by running the `Install-WindowsFeature` PowerShell cmdlet:

```
Install-WindowsFeature -Name Hyper-V, Failover-Clustering, Data-Center-Bridging, BitLocker, FS-FileServer, RSAT-Clustering-PowerShell, FS-Data-Deduplication -IncludeAllSubFeature -IncludeManagementTools -verbose
```

 **NOTE:** Install the `storage-replica` feature if Azure Stack HCI operating system is being deployed for a stretched cluster.

 **NOTE:** Hyper-V and the optional roles installation require a system restart. Because subsequent procedures also require a restart, the required restarts are combined into one (see the Note in the ["Changing the hostname"](#) section).

## Verifying firmware and software compliance with the support matrix

Microsoft HCI Solutions from Dell Technologies are validated and certified with certain firmware versions that are related to the solution infrastructure components.

Use the validated firmware and software, as specified in the [support matrix](#), to ensure that the solution infrastructure remains supported and delivers optimal performance.


You can verify compliance and update the nodes with an Azure Stack HCI online or offline catalog by using Dell EMC OpenManage Integration with Windows Admin Center. To verify compliance and to update firmware and drivers on a stand-alone node, see the [Dell EMC HCI Solutions for Microsoft Windows Server—Managing and Monitoring the Solution Infrastructure Life Cycle](#).



# Update out-of-box drivers

For certain system components, you might have to update the driver to the latest Dell Technologies supported version.

## About this task

 **NOTE:** This section is optional if you are using OpenManage Integration with Windows Admin Center to update the nodes.

## Steps

1. Depending on the platform, many devices may not automatically be recognized with in-box drivers. Install the proper Intel or AMD chipset drivers.
2. Run the following PowerShell command to retrieve a 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 "*Broadcom*") -or
($_.Name -like "*marvell*") }
```


3. Update the out-of-box drivers to the required versions, if necessary.  
For the latest Dell Technologies supported versions of system components, see the [Support Matrix for Microsoft HCI Solutions](#).  
Download the driver installers from <https://www.dell.com/support> or by using the [Dell EMC Azure Stack HCI Solution Catalog](#).
4. Attach a folder containing the driver DUP files to the system as a virtual media image:
  - a. In the iDRAC virtual console menu, click **Virtual Media**.
  - b. In the **Create Image from Folder** window, click **Create Image**.
  - c. Click **Browse**, select the folder where the driver DUP files are stored, and, if required, change the name of the image.
  - d. Click **Create Image**.
  - e. Click **Finish**.
  - f. From the **Virtual Media** menu, select **Connect Virtual Media**.
  - g. Select **Map Removable Disk**, click **Browse**, and select the image that you created.
  - h. Click **Map Device**.After the image is mapped, it appears as a drive in the host operating system.
5. Go to the driver DUP files and run them to install the updated out-of-box drivers.

# Change the hostname

By default, the operating system deployment assigns a random name as the host computer name. For easier identification and uniform configuration, Dell Technologies recommends that you change the hostname to something that is relevant and easily identifiable.

Change the hostname by using the `Rename-Computer` cmdlet:

```
Rename-Computer -NewName S2DNode01 -Restart
```

 **NOTE:** This command induces an automatic restart at the end of rename operation.

# Configuring host networking

Configure Microsoft HCI Solutions from Dell Technologies to implement a fully converged or nonconverged network for storage and management connectivity.

Complete the following steps to configure host networking:

1. Create VM switches (based on topology).

2. Create VM adapters, and configure VLANs and IP addresses.
3. Map the VM storage adapters (based on topology).
4. Enable RDMA for the storage adapters.
5. Change RDMA settings on the QLogic NICs.
6. Configure the QoS policy.
7. Disable the DCBX Willing state in the operating system.

**NOTE:** Dell Technologies recommends implementing a nonconverged network and using physical network adapters for the storage traffic rather than using SET. However, in a nonconverged configuration, if virtual machine adapters must have RDMA capability, SET configuration is necessary for storage adapters.

For more information about each of the preceding steps and all possible topologies within both fully converged and nonconverged solution integration (including switchless storage networking solution integration) and host operating system network configuration, see [Network Integration and Host Network Configuration Options](#).

**NOTE:** The host operating system network configuration must be complete before you join cluster nodes to the Active Directory domain.

**NOTE:** Steps on the configuration of RDMA adapters do not apply to single node clusters.

## Registry settings

**NOTE:** This registry key procedure only applies if the operating system is Azure Stack HCI OS. It does not apply for Windows Server OS (WS2019 or WS2022).

- Azure Stack HCI OS is available in multiple languages. You can download the ISO file of Azure Stack HCI OS (in any language) from [Microsoft | Azure Stack HCI software download](#).
- After installing Azure Stack HCI OS on an AX node, you must perform the following registry configuration before Azure onboarding to ensure the node gets identified as one sold by Dell EMC.
- PowerShell commands:

```
New-ItemProperty "HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\OEMInformation"
-name SupportProvider -value DellEMC
```

To verify "DellEMC" is successfully entered into the registry, run the following command:

```
(get-itemproperty "HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\OEMInformation").SupportProvider
```

The screenshot shows a Windows PowerShell window titled "Administrator: Windows PowerShell". The command prompt shows the following commands and output:

```
PS C:\> New-ItemProperty "HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\OEMInformation" -name SupportProvider -value DellEMC
```

The output displays the properties of the newly created registry item:

```
SupportProvider : DellEMC
PSPath          : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\OEMInformation
PSParentPath    : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion
PSChildName     : OEMInformation
PSDrive         : HKLM
PSProvider      : Microsoft.PowerShell.Core\Registry
```

Below this, the verification command is run:

```
PS C:\> (Get-ItemProperty "HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\OEMInformation").SupportProvider
```

The output of the verification command is:

```
DellEMC
PS C:\>
```

Figure 3. Registry settings

# Joining cluster nodes to an Active Directory domain

Before you can create a cluster, the cluster nodes must be a part of an Active Directory domain.

**NOTE:** Connecting to Active Directory Domain Services by using the host management network might require routing to the Active Directory network. Ensure that this routing is in place before joining cluster nodes to the domain.

You can perform the domain join task by running the `Add-Computer` cmdlet on each host that will be a part of the Azure Stack HCI cluster.

**NOTE:** Optionally, you can add all newly created computer objects from the cluster deployment to a different Organizational Unit (OU) in Active Directory Domain Services. In this case, you can use the `-OUPath` parameter along with the `Add-Computer` cmdlet.

```
$credential = Get-Credential
Add-Computer -DomainName S2dlab.local -Credential $credential -Restart
```

**NOTE:** This command induces an automatic restart at the end of the domain join operation.

## Deploying and configuring a host cluster

After joining the cluster nodes to an Active Directory domain, you can create a host cluster and configure it for Storage Spaces Direct.

### Creating the host cluster

Verify that the nodes are ready for cluster creation, and then create the host cluster.

#### Steps

1. Run the `Test-Cluster` cmdlet:

```
Test-Cluster -Node S2Dnode01, S2DNode02, S2dNode03, S2dNode04 -Include 'Storage Spaces Direct', 'Inventory', 'Network', 'System Configuration'
```

The `Test-Cluster` cmdlet generates an HTML report of all performed validations and includes a summary of the validations. Review this report before creating a cluster.

2. Run the `Get-PhysicalDisk` command on all cluster nodes.

Verify the output to ensure that all disks are in the healthy state and that the nodes have an equal number of disks. Verify that the nodes have homogenous hardware configuration.

3. Run the `New-Cluster` cmdlet to create the host cluster.

**NOTE:** For the `-IgnoreNetwork` parameter, specify all storage network subnets as arguments. Switchless configuration requires that all storage network subnets are provided as arguments to the `-IgnoreNetwork` parameter.

```
New-Cluster -Name S2DSystem -Node S2Dnode01, S2DNode02, S2dNode03, S2dNode04
-StaticAddress 172.16.102.55 -NoStorage -IgnoreNetwork 172.16.103.0/27, 172.16.104.0/27
-Verbose
```

In this command, the `StaticAddress` parameter is used to specify an IP address for the cluster in the same IP subnet as the host management network. The `NoStorage` switch parameter specifies that the cluster is to be created without any shared storage.

The `New-Cluster` cmdlet generates an HTML report of all performed configurations and includes a summary of the configurations. Review the report before enabling Storage Spaces Direct.

## Enabling Storage Spaces Direct

After you create the cluster, run the `Enable-ClusterS2D` cmdlet to configure Storage Spaces Direct on the cluster. Do not run the cmdlet in a remote session; instead, use the local console session.

Run the `Enable-ClusterS2D` cmdlet as follows:

```
Enable-ClusterS2D -Verbose
```

The `Enable-ClusterS2D` cmdlet generates an HTML report of all configurations and includes a validation summary. Review this report, which is typically stored in the local temporary folder on the node where the cmdlet was run. The verbose output of the command shows the path to the cluster report. At the end of the operation, the cmdlet discovers and claims all the available disks into an auto-created storage pool. Verify the cluster creation by running any of the following commands:

```
Get-ClusterS2D
```

```
Get-StoragePool
```

```
Get-StorageSubSystem -FriendlyName *Cluster* | Get-StorageHealthReport
```

## Configuring the host management network as a lower-priority network for live migration

After you create the cluster, live migration is configured by default to use all available networks.

During normal operations, using the host management network for live migration traffic might impede the overall cluster role functionality and availability. Rather than disabling live migration traffic on the host management network, configure the host management network as a lower-priority network in the live migration network order:

```
$clusterResourceType = Get-ClusterResourceType -Name 'Virtual Machine'
$hostNetworkID = Get-ClusterNetwork | Where-Object { $_.Address -eq '172.16.102.0' } |
Select-Object -ExpandProperty ID
$otherNetworkID = (Get-ClusterNetwork).Where({$_ .ID -ne $hostNetworkID}).ID
$newMigrationOrder = ($otherNetworkID + $hostNetworkID) -join ';'
Set-ClusterParameter -InputObject $clusterResourceType -Name MigrationNetworkOrder -Value
$newMigrationOrder
```

## Updating the page file settings

To help ensure that the active memory dump is captured if a fatal system error occurs, allocate sufficient space for the page file. Dell Technologies recommends allocating at least 50 GB plus the size of the CSV block cache.

### About this task

1. Determine the cluster CSV block cache size value by running the following command:

```
$blockCacheMB = (Get-Cluster).BlockCacheSize
```

**NOTE:** On Windows Server 2016, the default block cache size is 0. On Windows Server 2019, Windows Server 2022, and Azure Stack HCI operating system version 20H2 and 21H2, the default block cache size is 1 GB.

2. Run the following command to update the page file settings:

```
$blockCacheMB = (Get-Cluster).BlockCacheSize

$pageFilePath = "C:\pagefile.sys"
$initialSize = [Math]::Round(51200 + $blockCacheMB)
$maximumSize = [Math]::Round(51200 + $blockCacheMB)


$system = Get-WmiObject -Class Win32_ComputerSystem -EnableAllPrivileges
if ($system.AutomaticManagedPagefile) {
    $system.AutomaticManagedPagefile = $false
    $system.Put()
}
```

```
$currentPageFile = Get-WmiObject -Class Win32_PageFileSetting
if ($currentPageFile.Name -eq $pageFilePath)
{
    $currentPageFile.InitialSize = $InitialSize
    $currentPageFile.MaximumSize = $MaximumSize
    $currentPageFile.Put()
}
else
{
    $currentPageFile.Delete()
    Set-WmiInstance -Class Win32_PageFileSetting -Arguments @{Name=$pageFilePath;
InitialSize = $initialSize; MaximumSize = $maximumSize}
}
```

## Configuring a cluster witness

A cluster witness must be configured for a two-node cluster. Microsoft recommends configuring a cluster witness for a four-node Azure Stack HCI cluster. Cluster witness configuration helps maintain a cluster or storage quorum when a node or network communication fails and nodes continue to operate but can no longer communicate with one another.

A cluster witness can be either a file share or a cloud-based witness.

 **NOTE:** If you choose to configure a file share witness, ensure that it is outside the two-node cluster.

For information about configuring a cloud-based witness, see [Cloud-based witness](#).

## Azure onboarding for Azure Stack HCI operating system

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

After Microsoft Azure registration, use the `Get-AzureStackHCI` command to confirm the cluster registration and connection status.


## Best practices and recommendations

Dell Technologies recommends that you follow the guidelines that are described in this section.

### Disable SMB Signing

Storage Spaces Direct uses RDMA for SMB (storage) traffic for improved performance. When SMB signing is enabled the network performance of SMB traffic is significantly reduced.

For more information, see [Reduced networking performance after you enable SMB Encryption or SMB Signing in Windows Server 2016](#).

 **NOTE:** By default, SMB Signing is disabled. If SMB Signing is enabled in the environment through a Group Policy Object (GPO), you must disable it from the domain controller.

### Update the hardware timeout for the Spaces port

For performance optimization and reliability, update the hardware timeout configuration for the Spaces port.

The following PowerShell command updates the configuration in the Windows registry and induces a restart of the node at the end of the registry update. Perform this update on all Storage Spaces Direct nodes immediately after initial deployment. Update one node at a time and wait until each node rejoins the cluster.


```
Set-ItemProperty -Path HKLM:\SYSTEM\CurrentControlSet\Services\spaceport\Parameters -Name
HwTimeout -Value 0x00002710 -Verbose
```

`Restart-Computer -Force`

## Enable jumbo frames

Enabling jumbo frames specifically on the interfaces supporting the storage network might help improve the overall read/write performance of the Azure Stack HCI cluster. An end-to-end configuration of jumbo frames is required to take advantage of this technology. However, support for jumbo frame sizes varies among software, NIC, and switch vendors. The lowest value within the data path determines the maximum frame size that is used for that path.

For the storage network adapters in the host operating system, enable jumbo frames by running the `Set-NetworkAdapterAdvancedProperty` cmdlet.

 **NOTE:** Network adapters from different vendors support different jumbo packet sizes. The configured value must be consistent across the host operating system and network switch configuration.

For information about configuring jumbo frames at the switch port level, see [Sample Network Switch Configuration Files](#).

## Recommended next steps

Before you proceed with operational management of the cluster, Dell Technologies recommends that you validate the cluster deployment, verify that the infrastructure is operational, and, if needed, activate the operating system license.

1. Run the `test-Cluster` cmdlet to generate a cluster validation report:

```
Test-Cluster -Node S2DNode01, S2DNode02, S2DNode03, S2DNode04 -Include 'System  
Configuration', 'Inventory', 'Network', 'Storage Spaces Direct'
```

This command generates an HTML report with a list of all the tests that were performed and completed without errors.

2. If the operating system was not factory-installed, activate the operating system license.

By default, the operating system is installed in evaluation mode. For information about activating the operating system license as well as the management and operational aspects for the Azure Stack HCI solution, see the [Operations Guide—Managing and Monitoring the Solution Infrastructure Life Cycle](#).

## Deployment services

Dell Technologies recommends using the company's deployment services to install Microsoft HCI Solutions from Dell Technologies. Issues that arise during do-it-yourself installations and configuration are not covered even if you have purchased Dell Technologies ProSupport or ProSupport Plus. Support for installation and configuration issues is provided under a separate paid services package.

When you call Dell Technologies with an installation and configuration issue, Dell Tech Support routes you to your Sales Account Manager. The Account Manager will then help you to purchase the onsite deployment services package.

# References

**Topics:**

- [Dell Technologies documentation](#)
- [Microsoft documentation](#)

## Dell Technologies documentation

These links provide more information from Dell Technologies:

- [iDRAC documentation](#)
- [Support Matrix for Microsoft HCI Solutions](#)
- [Operations Guide—Managing and Monitoring the Solution Infrastructure Life Cycle](#)

## Microsoft documentation

The following link provides more information about Storage Spaces Direct:

[Storage Spaces Direct overview](#)

# GPU Integration

## Topics:

- Installing and configuring GPUs on AX-750/7525
- Deploy GPU devices using Discrete Device Assignment (DDA)

## Installing and configuring GPUs on AX-750/7525

This section provides the procedure to install and configure the GPUs on AX-750/7525.

**NOTE:** Ensure to order the AX nodes preconfigured for GPU integration, or converted for GPU use including the required components.

### AX-750 GPU configuration

- AX-750 GPU ready server order includes:
  - Microsoft Azure Stack HCI operating system (version 21H2 required).
  - GPU enablement.
  - GPU ready configuration cable install kit R750.
  - Heatsink for two CPUs with GPU configuration.
  - Very high-performance Fan x6 or very high-performance Fan x6 with 30C max ambient temp.
  - Fan foam, HDD 2U.
  - Riser Config 2, full length, 4x16, 2x8 slots, DW GPU capable or Riser Config 2, half length, 4x16, 2x8 slots, SW GPU capable.
  - PSU configuration must be large enough to support additional GPU(s) power requirements
- AX-750 GPU slot priority matrix:
  - Dual width GPUs in slot(s): 7,2




Figure 4. Riser Config 2.

### AX-7525 GPU configuration

- AX-7525 GPU ready server order includes:
  - Microsoft Azure Stack HCI operating system (version 21H2 required) .
  - GPU ready configuration cable install kit R7525.
  - Heatsink for two CPUs and GPU/FPGA/Full Length card configs configuration.
  - High performance Fan x6 or very high performance Fan x6.
  - Fan foam, HDD 2U.
  - Riser Config 3 full length V2 or Riser Config 3, half length, 5 x16 slots.
  - PSU configuration must be large enough to support additional GPU(s) power requirements.
- GPU Slot Priority Matrix:
  - Dual width GPUs in slot(s): 2,5,7





Figure 5. Riser Config 3.

## GPU physical installation

For deployed clusters, place each node into maintenance mode before changing the hardware. For undeployed nodes, install the cards as per the instructions.

- For information on installing a GPU and slot matrix of GPU on **AX750**, see section *Installing a GPU* and Table: *Configuration 2-1: R1A + R2A + R3B + R4A (FL)* in [Dell EMC PowerEdge R750](#).
- For GPU kit, see [Dell EMC PowerEdge R750 Installation and Service Manual](#)
- For information on installing a GPU and slot matrix of GPU on **AX7525**, see section *Installing a GPU* and Table: *Configuration 3-1: R1A+R2A+R3A+R4A (FL)* in [Dell EMC PowerEdge R7525](#).
- For GPU kit, see [Dell EMC PowerEdge R7525 Installation and Service Manual](#)
- [Video](#) Show installation of M10 GPU but is representative of all GPU cards.

**NOTE:** Remove the BOSS card power and data cables from Riser 1 before removing them. The power connection is fragile.

## Deploy GPU devices using Discrete Device Assignment (DDA)

### Prerequisites

- Deploy the HCI cluster, see [Deployment models](#)
- Download and install the latest Nvidia driver package to each GPU equipped cluster node. For free CUDA drivers (see [NVIDIA Driver Downloads](#) for the latest versions)
- Create the cluster from a machine running HCI operating system 21H2 at or greater than the nodes. It is recommended to create the cluster from the node.
- Download and install the latest INF package from [NVIDIA GPU Passthrough Support](#). Place the correct INF for the GPU on the system. Ensure that all the cluster nodes install it. To install, run the PNPUTIL command. For example:

```
pnputil /add-driver .\nvidia_azure_stack_GPUMOEL_base.inf /install
```

### About this task

Dell Technologies supports GPU integration through Discrete Device Assignment (DDA) on Azure Stack HCI. DDA allows the GPU to be assigned to a virtual machine (VM) directly as a hardware component.

**NOTE:** This technology does not support GPU partitioning.

DDA is accomplished either through the Windows Admin Center (WAC) or PowerShell command line on the node.

### Steps

1. Create one VM per physical GPU on the cluster. Install a GPU supported operating system (see supported operating system list on Nvidia website), complete installation, and power off the VMs.
2. Mapping GPU through DDA.
  - a. For Microsoft instruction, see [Discrete Device Assignment](#)
  - b. For Powershell Script, see [Deploy graphics devices using DDA](#)
  - c. General idea:
    - i. Set the VM configuration (MMIO limits).
      - Formula: 2 x (card memory) x (number of cards)

- Example with 16 GB A2's (1 per server) -> 2 x 16 x 1 = 32 GB ~ 32768 MB
  - Example with 24 GB A30's (2 per server) -> 2x24x2=96 GB ~ 98304 MB
- Locate the correct PCI location for the card.
  - Disable the GPU.
  - Dismount the GPU.
  - This removes the PCI device from the host server.
  - Assign PCI resource (GPU) to VM.

 **WARNING: Multiple cards mean multiple device locations. Be careful not mixing them.**

Set the VM configuration (MMIO limits).

- Log in to the node through RDP or iDRAC.
- Run the following script one line at a time.

```
#Configure the VM for a Discrete Device Assignment
$vm = "ddatest1"
#Set automatic stop action to TurnOff
Set-VM -Name $vm -AutomaticStopAction TurnOff
#Enable Write-Combining on the CPU
Set-VM -GuestControlledCacheTypes $true -VMName $vm
#Configure 32 bit MMIO space
Set-VM -LowMemoryMappedIoSpace 3Gb -VMName $vm
#Configure Greater than 32 bit MMIO space
Set-VM -HighMemoryMappedIoSpace 33280Mb -VMName $vm

#Find the Location Path and disable the Device
#Enumerate all PNP Devices on the system
$pnpdevs = Get-PnpDevice -presentOnly
#Select only those devices that are Display devices manufactured by NVIDIA
$gpudevsnames = $pnpdevs | where-object {$_.Class -like "Display" -and $_.Manufacturer -like "NVIDIA"}
#Select the location path of the first device that's available to be dismounted
#by the host.
$locationPath = ($gpudevsnames | Get-PnpDeviceProperty
DEVKEY_Device_LocationPaths).data[0]
#Disable the PNP Device
Disable-PnpDevice -InstanceId $gpudevsnames[0].InstanceId

#Dismount the Device from the Host
Dismount-VMHostAssignableDevice -force -LocationPath $locationPath

#Assign the device to the guest VM.
Add-VMAssignableDevice -LocationPath $locationPath -VMName $vm
```

- It is critical that you iterate the device number when you install the second GPU.
- Mapping a GPU though WAC (recommended).
    - Go to the main settings on the upper right corner and click **Extensions**.


The **Extensions** page is displayed.

- Select **GPUs** and click **install**.

This installs the GPU controls at a cluster level.

## Extensions

Windows Admin Center might restart after installing an extension, temporarily affecting anyone using this instance of Windows Admin Center.

Automatically update extensions  On

Available extensions Installed extensions Feeds

Uninstall Update

Name	Version	Created by	Status
App & Features	2.14.0	Microsoft	Installed
Azure Backup	2.8.0	Microsoft	Installed
Azure File Sync	2.151.0	Microsoft	Installed
Azure hybrid center	2.21.0	Microsoft	Installed
Azure Kubernetes Service	2.31.0	Not Available	Update available (2.56.0)
Certificates	2.8.0	Microsoft	Installed
Cluster Creation	5.0.0	Not Available	Installed
Cluster Manager	2.49.0	Not Available	Update available (2.69.0)
Dell OpenManage Integration	2.3.0	Dell	Installed
Developer Guide	2.237.0	Microsoft	Installed
Devices	2.13.0	Microsoft	Installed
Events	2.31.0	Microsoft	Installed
Fallover cluster tools	2.29.0	Microsoft	Update available (2.47.0)
Files & file sharing	2.49.0	Microsoft	Installed
Firewall	2.8.0	Microsoft	Installed
GPUs	1.1.0	Microsoft	Installed
Local users & groups	2.14.0	Microsoft	Installed
Microsoft Defender for Cloud	3.0.0	Not Available	Update available (3.0.1)
Network Controller tools and SDN Virtual networks	1.28.0	Microsoft	Installed
Networks	2.6.0	Microsoft	Installed
Packet monitoring	2.49.0	Microsoft	Installed
Performance Monitor	2.14.0	Microsoft	Installed
PowerShell	2.18.0	Microsoft	Installed

### Details - GPUs



#### Description

Manage and view all available GPU devices on the host machine

#### ID

msft.sme.gpu

#### Version

1.1.0

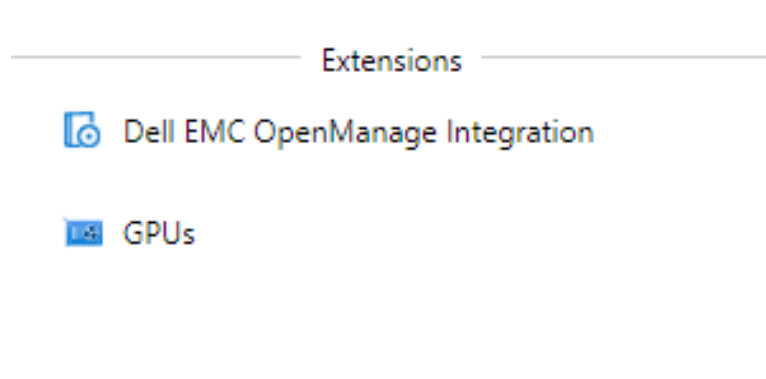
#### Authors

Microsoft

#### License info

<https://licenses.nuget.org/MIT>

- c. Go to the cluster and perform the following:
  - i. Ensure that the VM is turned off.
  - ii. Under Extensions, click the **GPUs** tab.



Under the GPUs tab, all the GPUs are listed per node.

**NOTE:** The Nvidia driver package on installed to each server does not include the '3D Video Controller' INF driver. After the GPU is dismounted from the host, the Nvidia driver package that is installed on the VM will properly recognize the full Nvidia device ID.

**GPU** PREVIEW

Navigate to GPU pools to assign VMs to GPUs

**GPUs** GPU pools

Update driver Enable Disable Mount Dismount

Name	Status	Manufacturer	Location	Driver version	Assignment status
ad7525r10n068 (3)					
Microsoft Basic Display Adapter	OK	(Standard display types)	PCI bus 98, device 0, function 0	10.0.20348.1	Not assignable
3D Video Controller	Error		PCI bus 37, device 0, function 0	-	Ready for assignment
3D Video Controller	Error		PCI bus 226, device 0, function 0	-	Ready for assignment
ad7525r10n069 (3)					
Microsoft Basic Display Adapter	OK	(Standard display types)	PCI bus 98, device 0, function 0	10.0.20348.1	Not assignable
3D Video Controller	Error		PCI bus 37, device 0, function 0	-	Ready for assignment
3D Video Controller	Error		PCI bus 226, device 0, function 0	-	Ready for assignment
ad7525r10n070 (3)					
Microsoft Basic Display Adapter	OK	(Standard display types)	PCI bus 98, device 0, function 0	10.0.20348.1	Not assignable
3D Video Controller	Error		PCI bus 37, device 0, function 0	-	Ready for assignment
3D Video Controller	Error		PCI bus 226, device 0, function 0	-	Ready for assignment
ad7525r10n071 (3)					
Microsoft Basic Display Adapter	OK	(Standard display types)	PCI bus 98, device 0, function 0	10.0.20348.1	Not assignable
3D Video Controller	Error		PCI bus 37, device 0, function 0	-	Ready for assignment
3D Video Controller	Error		PCI bus 226, device 0, function 0	-	Ready for assignment

iii. Click the **GPU Pools** tab to create the GPU pools. Enter the following details:

- Servers – Enter the server details.
- GPU pool name – Enter the name of the GPU pool.
- Select GPUs - Only one GPU per node per pool is allowed.
- Select **Assign without mitigation driver (not recommended)** - This is because the mitigation driver is not available in the current release.

#### New GPU pool

Enter the name and select the GPUs you want associated with this GPU pool below

To ensure high availability, a GPU pool with the same name must be created on each node of the cluster

Servers \*

GPU pool name \*

Select GPUs \*

Name	Status	Manufacturer
ad7525r10n70		
3D Video Controller	Error	
3D Video Controller	Error	
ad7525r10n71		
3D Video Controller	Error	
3D Video Controller	Error	

Advanced >

☒ Assign without mitigation driver (not recommended)

iv. Select the **Assign VM to GPU Pool** to assign the GPU to the VM.

v. Select the **server**, **pool** and **virtual machine**.

vi. Click **advance** and enter **input the memory** requirements.

- Low memory mapped I/O space (in MB)
- High memory mapped I/O space (in MB) – Adjust the maximum memory mapped I/O space to match your particular GPU. The formula is 2 x GPU RAM per attached GPU. So, a VM with two 16 GB A2s will need ((2x16GB) x 2) = 64 GB memory.

Assigning with DDA may affect the AutomaticStopAction, GuestControlledCacheType, and memory-mapped IO spaces values of your VM. Furthermore, you

Advanced >

Low memory mapped IO space (in MB)

3072

High memory mapped IO space (in MB)

33280

The MMIO space values above are reasonable values to set for experimenting with a single GPU. If after starting the VM, the device is reporting an error related to MMIO requirements, see [learn](#) how to precisely calculate MMIO requirements.

☐ Configure offline action to force shutdown

DDA assignment does not support live migration of virtual machines. Clicking this checkbox will change the VM's default offline action to be a force shutdown to successfully migrate.

**NOTE:** This can be calculated by known specifications, or you can run Survey DDA.ps1 provided by Microsoft. For more information, see [Virtualization-Documentation](#) and [Discrete Device Assignment](#)

- MMIO - Set on per VM basis

### Assign VM to GPU pool

GPU pool assignments enable software running in VMs to access the GPU hardware connected to the host. [Read more](#)

Type of assignment \*

Assign whole GPU (DDA)

For security reasons, Microsoft recommends installing a device mitigation driver from your GPUs' vendor. You may assign a GPU without a mitigation driver

Server \*

adc7525r10n068.s2dm.nat

GPU pool \*

testpool

The pool named testpool consists of the following GPUs:

Virtual machine \*

PcsBaseVhd-023352CF-9F34-4CB2-AC34-2ADCE19C9E9C

Assigning with DDA may affect the AutomaticStopAction, GuestControlledCacheType, and memory-mapped IO spaces values of your VM. Furthermore, you

Advanced >

Low memory mapped IO space (in MB)

3072

High memory mapped IO space (in MB)

33280

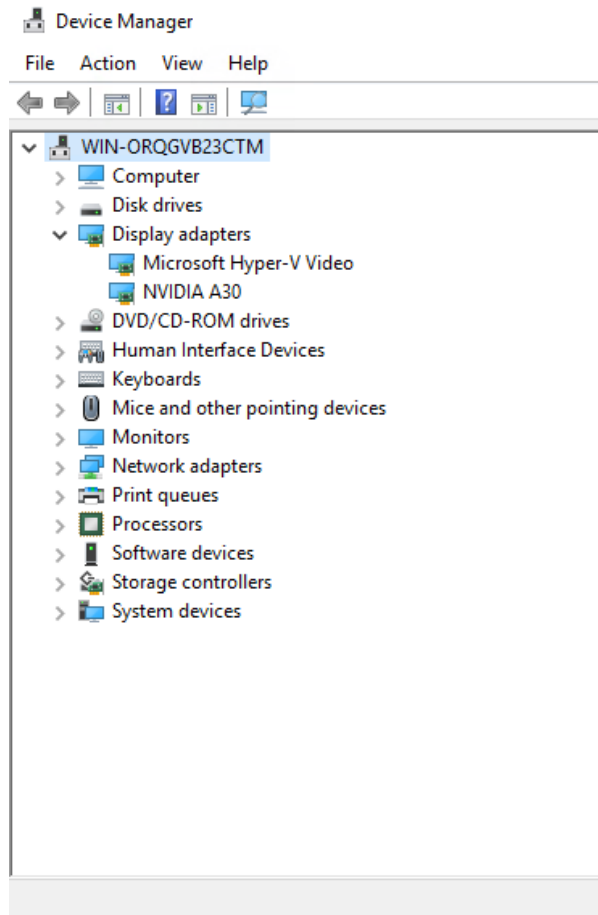
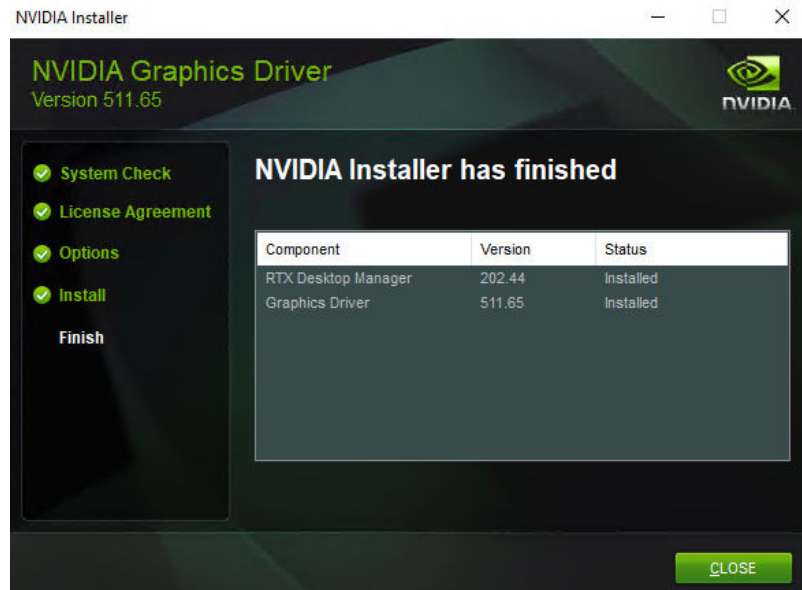
The MMIO space values above are reasonable values to set for experimenting with a single GPU. If after starting the VM, the device is reporting an error related to MMIO requirements, see [learn](#) how to precisely calculate MMIO requirements.

☐ Configure offline action to force shutdown

DDA assignment does not support live migration of virtual machines. Clicking this checkbox will change the VM's default offline action to be a force shutdown to successfully migrate.

- Install the drivers inside the VM.
- Transfer the executable files to the VM and run.





#### 4. Migration of VMs

##### **NOTE:**

- There must be a free GPU on the destination node to migrate a GPU VM otherwise the VM stops.
- This only works with the February '22 Windows Update – earlier versions are not supported.

##### **NOTE:**

Microsoft February 8, 2022 Security update (KB5010354) or later is required for proper VM migration functionality. For more information, see [February 8, 2022 Security update \(KB5010354\)](#)

- Perform the following steps to migrate a VM between nodes when there is a per-node-pool:

Server	GPU	Assignment status	Assigned VMs
gpn1 (1)			
AXESOR17N1	NVIDIA Tesla T4	Assigned	node1_windows
node2 (1)			
AXESOR17N2	NVIDIA Tesla T4	Available (Not assigned)	Ready for assignment
node3-4 (2)			
AXESOR17N3	NVIDIA Tesla T4	Available (Not assigned)	Ready for assignment
AXESOR17N4	NVIDIA Tesla T4	Assigned	node4_windows

- i. Turn off the VM.
  - ii. Remove the GPU pool assignment inside WAC.
  - iii. Quick/Live Migrate VM to the new node.
  - iv. Assign the GPU pool on the destination node to the VM (if none exists, create one).
  - v. Power on the VM.
  - vi. Log in to ensure that the VM is created successfully.
- b. Perform the following steps to migrate a VM if the GPUs per node are all combined into one pool:

Server	GPU	Assignment status	Assigned VMs
gpn1 (1)			
AXESOR17N1	NVIDIA Tesla T4	Assigned	node1_windows
node2 (1)			
AXESOR17N2	NVIDIA Tesla T4	Available (Not assigned)	Ready for assignment
node3-4 (2)			
AXESOR17N3	NVIDIA Tesla T4	Available (Not assigned)	Ready for assignment
AXESOR17N4	NVIDIA Tesla T4	Assigned	node4_windows

- i. Turn off the VM.
  - ii. Remove the GPU pool assignment inside WAC.
  - iii. Quick/Live Migration is possible, and the GPU assignment must change (automatic assignment of GPU on the correct node).
- This triggers a driver failure in the VM.
- iv. Unassign the VM, and then reassign the VM to the correct GPU in the pool.
  - v. Power on the VM.
  - vi. Log in to ensure that the VM is created successfully.
- c. Migrating the VMs with attached GPU is not currently supported – only failover.
5. Linux VMs
- a. Linux VMs behave in the same way both through WAC and DDA.
  - b. Run the command “lspci” to reveal the pci device that is attached to the VM.

There are different driver versions.



# Persistent Memory for Windows Server HCI


## Topics:

- [Configuring persistent memory for Windows Server HCI](#)
- [Configuring Windows Server HCI persistent memory hosts](#)

## Configuring persistent memory for Windows Server HCI

Intel Optane DC persistent memory is designed to improve overall data center system performance and lower storage latencies by placing storage data closer to the processor on nonvolatile media. The persistent memory modules are compatible with DDR4 sockets and can exist on the same platform with conventional DDR4 DRAM DIMMs. In App Direct Mode, the operating system distinguishes between the RAM and persistent storage memory spaces.

Intel Optane DC persistent memory provides an ideal capacity to be used as a cache device (SBL) for Microsoft Storage Spaces Direct. Storage data is interleaved between the Intel Optane DC persistent memory DIMMs within each CPU socket to achieve the best performance. A single region per CPU socket is configured in the system BIOS. Thus, a dual CPU server platform provides Storage Spaces Direct with two persistent memory cache devices. These high-endurance, write-intensive cache devices can be used to enhance the performance of many slower-performing NVMe/SAS/SATA devices that are used for storage capacity, as shown in the following figure:



**Figure 6. Persistent memory**

## Persistent memory requirements

Persistent memory requirements for Microsoft HCI Solutions from Dell Technologies:

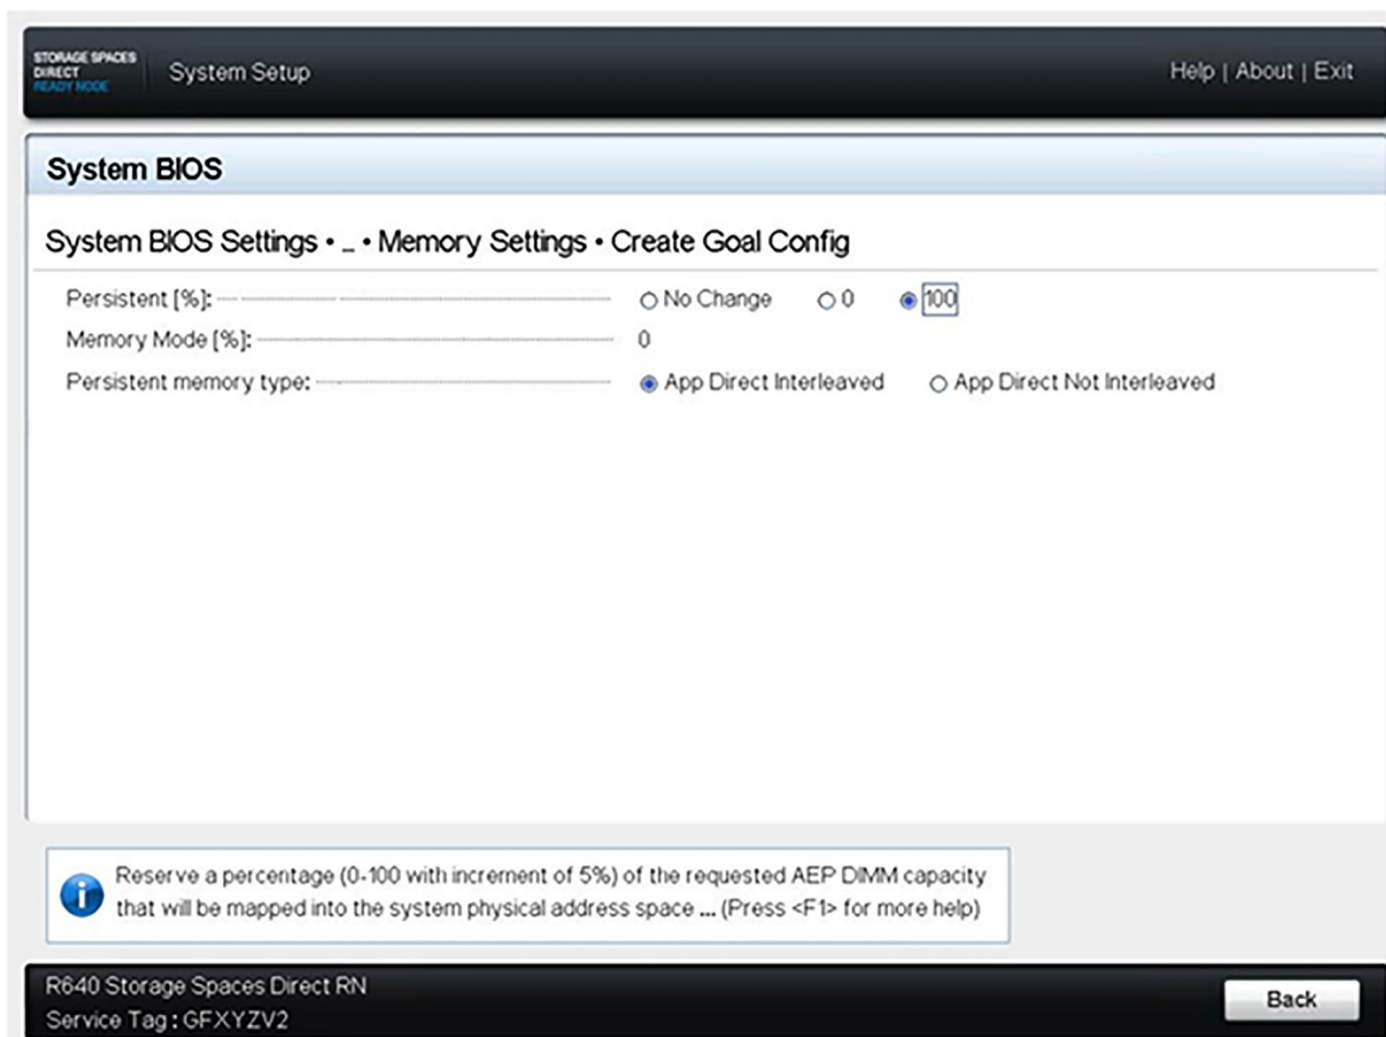
- AX-640 nodes
- 2 x Intel Xeon Cascade Lake-SP Gold or Platinum CPUs (models 52xx, 62xx, or 82xx) per server
- 12 x 32 GB RDIMMs in DIMM slots A1-A6 and B1-B6 (white slots) per server, totaling 384 GB of RAM per server
- 12 x 128 GB Intel Optane DC Persistent DIMMs in DIMM slots A7-A12 and B7-B12 (black slots) per server, totaling 2 x 768 GB cache devices per server
- Windows Server 2019 or Windows Server 2022 Datacenter

## Configuring persistent memory BIOS settings

Configure the BIOS to enable persistent memory.

### Steps

1. During system startup, press F12 to enter **System BIOS**.
2. Select **BIOS Settings > Memory Settings > Persistent Memory**.
3. Verify that **System Memory** is set to **Non-Volatile DIMM**.
4. Select **Intel Persistent Memory**.  
The **Intel Persistent Memory** page provides an overview of the server's Intel Optane DC persistent memory capacity and configuration.
5. Select **Region Configuration**.  
To be used as a Storage Spaces Direct cache device, Intel Persistent Memory must be configured in App Direct Interleaved mode. App Direct Interleaved Mode creates two regions—one region for each CPU socket, as shown in the following figure:



**Figure 7. Persistent memory region configuration**

6. If the configuration is missing or incorrect, select **Create goal config** to reconfigure the persistent memory regions.  
**CAUTION:** Performing these steps erases all previous persistent memory regions.
  - a. In **Create Goal Config**, for **Persistent (%)**, select **100**.
  - b. For **Persistent memory type**, select **App Direct Interleaved**.  
A warning is displayed. All Intel Persistent Memory data is erased when changes are saved to the BIOS configuration.
7. Exit the BIOS and save the configuration.

# Configuring Windows Server HCI persistent memory hosts

Three types of device objects are related to persistent memory on Windows Server 2019 and Windows Server 2022: the NVDIMM root device, physical INVDIMMs, and logical persistent memory disks. In Device Manager, physical INVDIMMs are displayed under **Memory devices**, while logical persistent memory disks are under **Persistent memory disks**. The NVDIMM root device is under **System Devices**. The `scmbus.sys` driver controls the NVDIMM root device.

The `nvdimm.sys` driver controls all NVDIMM devices, while the `pmem.sys` driver controls the logical disks. Both the `nvdimm.sys` and `pmem.sys` drivers are the same for all types of persistent memory, such as NVDIMM-N and Intel Optane DC Persistent Memory (INVDIMM).

The following figure shows the Device Manager on a system with 12 INVDIMMs across dual CPU sockets and two persistent storage disks:

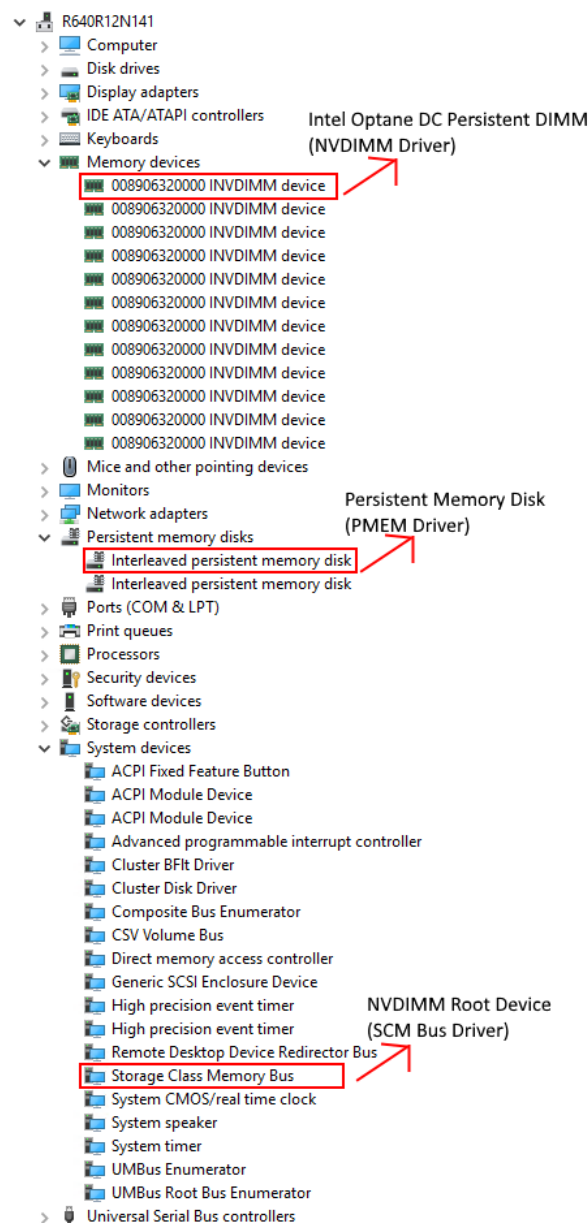


Figure 8. Device Manager example

# Managing persistent memory using Windows PowerShell

Windows Server 2019 and Windows Server 2022 provides a PersistentMemory PowerShell module that enables user management of the persistent storage space.

```
PS C:\> Get-Command -Module PersistentMemory
CommandType      Name                      Version      Source
-----
Cmdlet            Get-PmemDisk              1.0.0.0      PersistentMemory
Cmdlet            Get-PmemPhysicalDevice    1.0.0.0      PersistentMemory
Cmdlet            Get-PmemUnusedRegion      1.0.0.0      PersistentMemory
Cmdlet            Initialize-PmemPhysicalDevice 1.0.0.0      PersistentMemory
Cmdlet            New-PmemDisk              1.0.0.0      PersistentMemory
Cmdlet            Remove-PmemDisk           1.0.0.0      PersistentMemory
```

- **Get-PmemDisk**—Returns one or more logical persistent memory disks that were created by **New-PmemDisk**. The returned object includes information about size, health status, and the underlying physical NVDIMM devices.
- **Get-PmemPhysicalDevice**—Returns one or more physical persistent memory NVDIMM devices. The returned object includes information about size, firmware, physical location, and health status. In App Direct Interleaved mode, each INVDIMM device displays its full capacity as **Persistent memory size**.
- **NOTE:** The Intel Optane DC Persistent Memory firmware might have to be periodically updated. For the supported firmware version, see the [Support Matrix for Microsoft HCI Solutions](#). After identifying the required firmware version, download the Intel Optane DC Persistent Memory Firmware Package from [Dell Technologies Support](#) or use the Microsoft HCI Solutions from Dell Technologies Update Catalog.
- **Get-PmemUnusedRegion**—Returns aggregate persistent memory (Pmem) regions that are available for provisioning a logical device. The returned object has a unique region ID, total size, and list of physical devices that contribute to the unused region.
- **Initialize-PmemPhysicalDevice**—Writes zeroes to the label storage area, writes new label index blocks, and then rebuilds the storage class memory (SCM) stacks to reflect the changes. This cmdlet is intended as a recovery mechanism and is not recommended for normal use.
- **New-PmemDisk**—Creates a disk out of a given unused region. This cmdlet writes out the labels to create the namespace, and then rebuilds the SCM stacks to expose the new logical device. The new logical persistent disk is added in Device Manager under **Persistent memory disks**. **Get-PhysicalDisk** displays the storage device as **MediaType SCM**.
- **Remove-PmemDisk**—Removes the given persistent memory disk. This cmdlet accepts the output of **Get-PmemDisk**. It deletes the namespace labels and then rebuilds the SCM stacks to remove the logical device.

## Configuring persistent memory as SCM logical devices

On each server node, verify unused persistent memory regions and configure them as new SCM logical devices:

1. Run **Get-PmemPhysicalDevice** to verify that 12 INVDIMM physical devices are available and healthy:

```
PS C:\> Get-PmemPhysicalDevice
DeviceId DeviceType HealthStatus OperationalStatus PhysicalLocation FirmwareRevision Persistent memory
-----
1 008906320000 INVDIMM device Healthy {Ok} A7 102005395 126 GB 0 GB
1001 008906320000 INVDIMM device Healthy {Ok} B7 102005395 126 GB 0 GB
101 008906320000 INVDIMM device Healthy {Ok} A10 102005395 126 GB 0 GB
1011 008906320000 INVDIMM device Healthy {Ok} B8 102005395 126 GB 0 GB
1021 008906320000 INVDIMM device Healthy {Ok} B9 102005395 126 GB 0 GB
11 008906320000 INVDIMM device Healthy {Ok} A8 102005395 126 GB 0 GB
1101 008906320000 INVDIMM device Healthy {Ok} B10 102005395 126 GB 0 GB
111 008906320000 INVDIMM device Healthy {Ok} A11 102005395 126 GB 0 GB
1111 008906320000 INVDIMM device Healthy {Ok} B11 102005395 126 GB 0 GB
1121 008906320000 INVDIMM device Healthy {Ok} B12 102005395 126 GB 0 GB
121 008906320000 INVDIMM device Healthy {Ok} A12 102005395 126 GB 0 GB
21 008906320000 INVDIMM device Healthy {Ok} A9 102005395 126 GB 0 GB
```

2. Run **Get-PmemUnusedRegion** to verify that two unused Pmem regions are available, one region for each physical CPU:

```
PS C:\> Get-PmemUnusedRegion

RegionId      TotalSizeInBytes      DeviceId
-----
-----
```

1	811748818944	{1, 111, 21, 101...}
3	811748818944	{1001, 1111, 1021, 1101...}

3. Run the `Get-PmemUnusedRegion | New-PmemDisk` script to create two Pmem disks, one for each Pmem region:

```
PS C:\> Get-PmemUnusedRegion | New-PmemDisk
Creating new persistent memory disk. This may take a few moments.
Creating new persistent memory disk. This may take a few moments.
```

4. Run the `Get-PhysicalDisk | ?{$_.MediaType -eq "SCM"}` script to verify that both Pmem disks are available as physical disk devices:

```
PS C:\> Get-PmemDisk
DiskNumber    Size      HealthStatus    AtomicityType    CanBeRemoved
PhysicalDeviceIds
-----
-----
11            756 GB    Healthy         None             True             {1, 111, 21, 101...}
12            756 GB    Healthy         None             True             {1001, 1111, 1021, 1101...}
```

When you run the `Enable-ClusterS2D` command to enable Storage Spaces Direct, the SCM logical devices are automatically detected and used as cache for NVMe and SSD capacity devices.