

The Lenovo logo is displayed in white text on a black rectangular background.

GPU Configuration Guide for Lenovo ThinkAgile MX Solutions

Last Update: February 2023

Provides detailed steps to install and configure GPU drivers for use with Lenovo ThinkAgile MX solutions

Covers GPU configuration using the Windows Admin Center GPUs Extension

Discusses the process to assign host GPUs to VMs using DDA

Provides details for managing GPU Pools using Microsoft Windows Admin Center

Dave Feisthammel
Hussein Jammal
Laurentiu Petre



Table of Contents

1	Introduction.....	1
1.1	Supported GPUs.....	1
2	GPU driver installation	3
2.1	Download GPU driver installer	3
2.2	GPU driver installation on host.....	5
2.3	Install required INF file	8
3	Manage GPUs using WAC.....	10
3.1	Install GPUs extension.....	10
3.2	Create GPU pool.....	10
3.3	GPU assignment using DDA.....	13
3.4	GPU driver installation on virtual machine	15
4	Summary	17
5	Additional resources	18
6	Trademarks and special notices.....	19

1 Introduction

Graphics Processing Unit (GPU) virtualization technologies enable GPU acceleration in a virtualized environment, typically within virtual machines. If a workload is virtualized with Hyper-V, then graphics virtualization can be employed in order to provide GPU acceleration from the physical GPU to the virtualized apps or services.

Starting in Azure Stack HCI version 21H2, GPUs can be included in an Azure Stack HCI cluster to provide GPU acceleration to workloads running in clustered virtual machines. This document discusses the basic prerequisites of this capability and how to use GPUs with clustered virtual machines running the Azure Stack HCI operating system to provide GPU acceleration to workloads running in the cluster.

GPU acceleration is provided using Discrete Device Assignment (DDA), also known as GPU pass-through, which allows you to dedicate one or more physical GPUs to a virtual machine. Clustered virtual machines can take advantage of GPU acceleration and clustering capabilities such as high availability via failover. Although Live Migration of a virtual machine that has one or more GPUs assigned to it is not currently supported, virtual machines can be automatically restarted and placed where GPU resources are available in the event of a failure.

Note: Microsoft does not yet support GPU partitioning (GPU-P) in the Azure Stack HCI operating system. Currently, it is only possible to assign the entire GPU from the host to a single virtual machine using DDA.

This document provides instructions and examples to configure GPUs for use by the Azure Stack HCI operating system. We include information for installing GPU device drivers, configuring Windows Admin Center to manage GPUs, creating GPU Pools, and assigning virtual machines to GPUs in a pool.

1.1 Supported GPUs

Several GPUs are supported for Lenovo ThinkAgile MX solutions based on the particular solution. The GPUs supported by each ThinkAgile MX solution are shown in this section. Solutions can be easily identified by their Machine Type (MT) shown in the list.

ThinkAgile MX3520 Appliance (MT 7D5R) and ThinkAgile MX Certified Node (MT 7Z20) based on the ThinkSystem SR650 Rack Server:

- ThinkSystem NVIDIA Tesla T4 16GB PCIe Passive GPU (Feature Code B4YB)
- ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU (Feature Code BP05)
- ThinkSystem NVIDIA A10 24GB PCIe Gen4 Passive GPU (Feature Code BFTZ)
- ThinkSystem NVIDIA A30 24GB PCIe Gen4 Passive GPU (Feature Code BJHG)
- ThinkSystem NVIDIA A100 40GB PCIe Gen4 Passive GPU (Feature Code BEL5)

ThinkAgile MX1020 Appliance (MTs 7D5S and 7D5T) and MX1021 Certified Node (MTs 7D1B and 7D2U) based on the ThinkSystem SE350 Edge Server:

- ThinkSystem NVIDIA Tesla T4 16GB PCIe Passive GPU (Feature Code B4YB)
- ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU (Feature Code BP05)

Note: Since a GPU consumes the only available PCIe slot in an SE350 Edge Server, only 4 SSD or NVMe devices can be configured in any SE350 that includes a GPU

ThinkAgile MX3330 Appliance (MT 7D19) and ThinkAgile MX3331 Certified Node (MT 7D67) based on the ThinkSystem SR630 V2 Rack Server:

- ThinkSystem NVIDIA Tesla T4 16GB PCIe Passive GPU (Feature Code B4YB)
- ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU (Feature Code BQZT)

For ThinkAgile MX3530 Appliance (MT 7D6B) and ThinkAgile MX3531 Certified Node (MT 7D66) based on the ThinkSystem SR650 V2 Rack Server:

- ThinkSystem NVIDIA Tesla T4 16GB PCIe Passive GPU (Feature Code B4YB)
- ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU (Feature Code BP05)
- ThinkSystem NVIDIA A10 24GB PCIe Gen4 Passive GPU (Feature Code BFTZ)
- ThinkSystem NVIDIA A30 24GB PCIe Gen4 Passive GPU (Feature Code BJHG)
- ThinkSystem NVIDIA A40 48GB PCIe Gen4 Passive GPU (Feature Code BEL4)
- ThinkSystem NVIDIA A100 40GB PCIe Gen4 Passive GPU (Feature Code BEL5)

Note: The ThinkSystem NVIDIA Quadro RTX 6000 24GB PCIe Passive GPU is supported only for the Windows Server 2019 operating system. Since Windows Server 2022 and Azure Stack HCI operating systems are not supported, we do not show it in the list above.

2 GPU driver installation

Before GPUs can be used in a system, the appropriate device driver must be installed. This includes both host and guest systems. The driver is installed on the host first in order to assign the GPU(s) to a virtual machine, but must also be installed in any virtual machines that will use the GPU(s).

All currently supported GPUs for ThinkAgile MX solutions are NVIDIA GPUs. Therefore, the driver can be downloaded directly from NVIDIA at the following URL:

<https://www.nvidia.com/Download/index.aspx?lang=en-us>

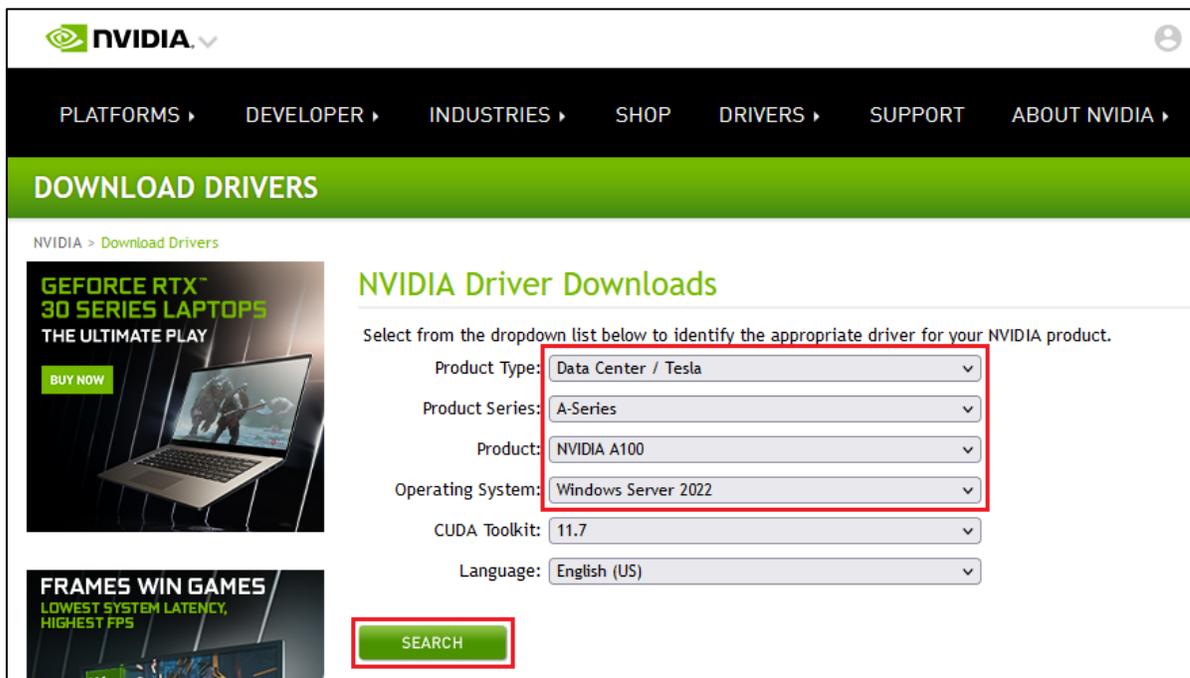
For NVIDIA GPUs, in addition to the standard GPU driver, an additional INF file needs to be installed on host systems. This INF file informs Hyper-V on how to correctly reset the GPU during VM reboots. This guarantees the GPU is in a clean state when the VM boots up. More information and a link to download the required INF files is available from NVIDIA at the following URL:

<https://docs.nvidia.com/datacenter/tesla/gpu-passthrough/index.html#introduction>

2.1 Download GPU driver installer

To download the appropriate driver, follow these steps:

1. Navigate to <https://www.nvidia.com/Download/index.aspx?lang=en-us> and then choose the appropriate options for the GPU installed. This includes the Product Type (“NVIDIA RTX / Quadro” or “Data Center / Tesla”), Product Series (“Quadro RTX Series” or “A-Series”), Product (choose appropriate GPU model), and Operating System (“Windows Server 2019” or “Windows Server 2022”).



The screenshot shows the NVIDIA website's driver download page. The navigation bar includes 'PLATFORMS', 'DEVELOPER', 'INDUSTRIES', 'SHOP', 'DRIVERS', 'SUPPORT', and 'ABOUT NVIDIA'. The main heading is 'DOWNLOAD DRIVERS'. Below this, there's a section for 'NVIDIA Driver Downloads' with a sub-heading 'Select from the dropdown list below to identify the appropriate driver for your NVIDIA product.' The form contains the following fields:

- Product Type: Data Center / Tesla
- Product Series: A-Series
- Product: NVIDIA A100
- Operating System: Windows Server 2022
- CUDA Toolkit: 11.7
- Language: English (US)

A green 'SEARCH' button is located at the bottom of the form. There are also promotional banners for 'GEFORCE RTX 30 SERIES LAPTOPS' and 'FRAMES WIN GAMES' on the left side of the page.

2. Once the correct values have been selected, click Search.
3. On the page that opens, click the SUPPORTED PRODUCTS tab to ensure that the correct GPU has been selected before clicking Download.

DOWNLOAD DRIVERS

NVIDIA > DRIVERS > DATA CENTER DRIVER FOR WINDOWS

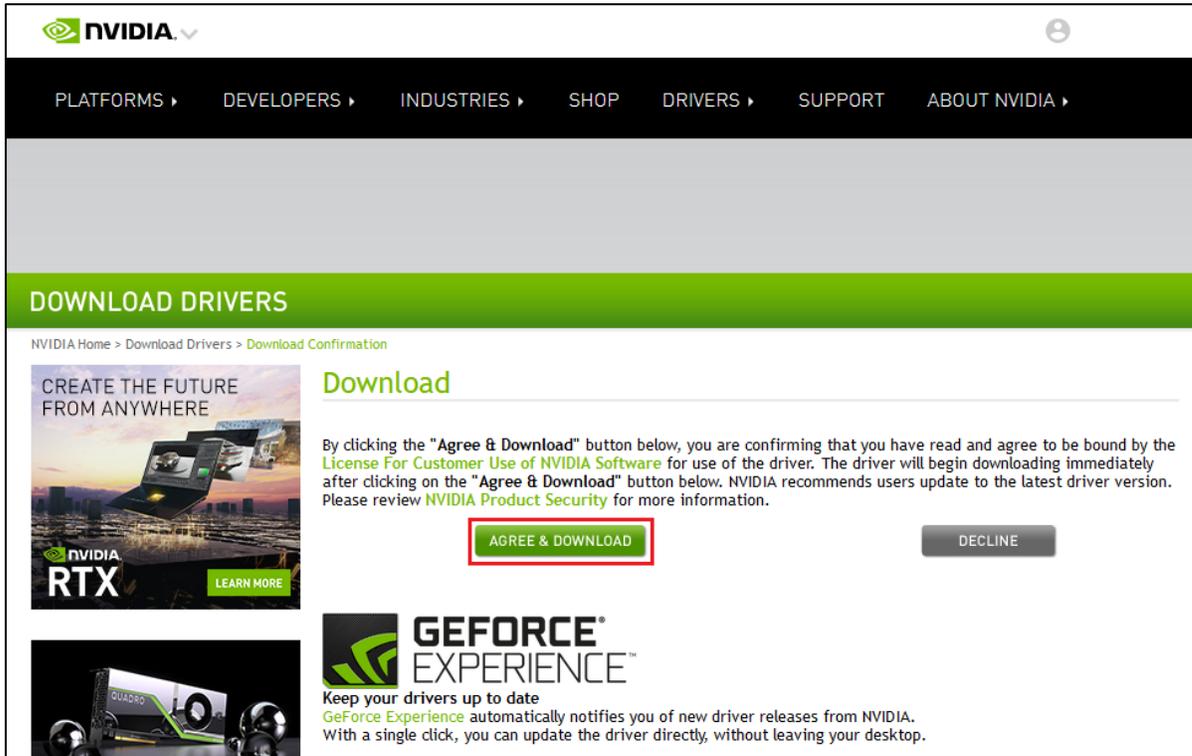
DATA CENTER DRIVER FOR WINDOWS

Version: 516.31 ^{WHQL}
 Release Date: 2022.6.6
 Operating System: Windows Server 2016, Windows Server 2019, Windows Server 2022
 CUDA Toolkit: 11.7
 Language: English (US)
 File Size: 637.19 MB

DOWNLOAD

RELEASE HIGHLIGHTS	SUPPORTED PRODUCTS	ADDITIONAL INFORMATION
	<p>A-Series: NVIDIA A100, NVIDIA A40, NVIDIA A30, NVIDIA A16, NVIDIA A10, NVIDIA A2</p> <p>RTX-Series: RTX 8000, RTX 6000, NVIDIA RTX A6000, NVIDIA RTX A5000, NVIDIA RTX A4000, NVIDIA T1000, NVIDIA T600, NVIDIA T400</p> <p>HGX-Series: HGX A100</p> <p>T-Series: Tesla T4</p> <p>V-Series: Tesla V100</p> <p>P-Series: Tesla P100, Tesla P40, Tesla P6, Tesla P4</p> <p>M-Class: M60, M40, M6, M4</p>	

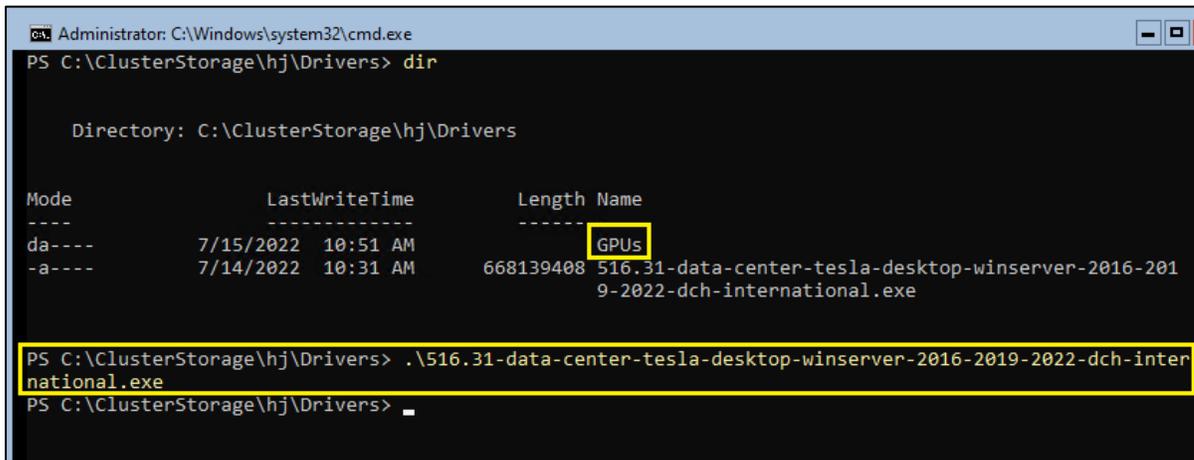
4. Click the AGREE & DOWNLOAD button to begin the driver download.



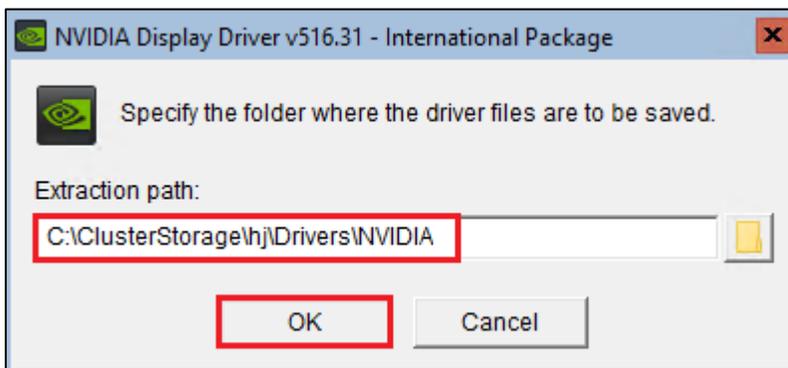
2.2 GPU driver installation on host

Once downloaded, the driver installation file can be copied to each node that contains one or more GPUs and installed. Alternatively, if the cluster is already configured, the installer can be copied to a CSV where it can be accessed by all nodes in the cluster. To install the driver from a CSV, follow these steps on each cluster node:

1. At the SConfig screen, type 15 and press Enter to leave SConfig and enter Windows PowerShell.
2. At the PowerShell prompt, navigate to the directory that contains the GPU driver installation file.
3. On the first node only, create a directory into which you will copy the extracted driver installation files. This will allow skipping the extraction process on all the other nodes in the cluster when installing the driver. In our example, the directory containing the downloaded driver installer is named "C:\ClusterStorage\hj\Drivers" and the directory that will be used for the extracted installation files is named "C:\ClusterStorage\hj\Drivers\GPUs".
4. Once the directory has been created, run the downloaded installer from PowerShell, as shown in the following example screenshot.



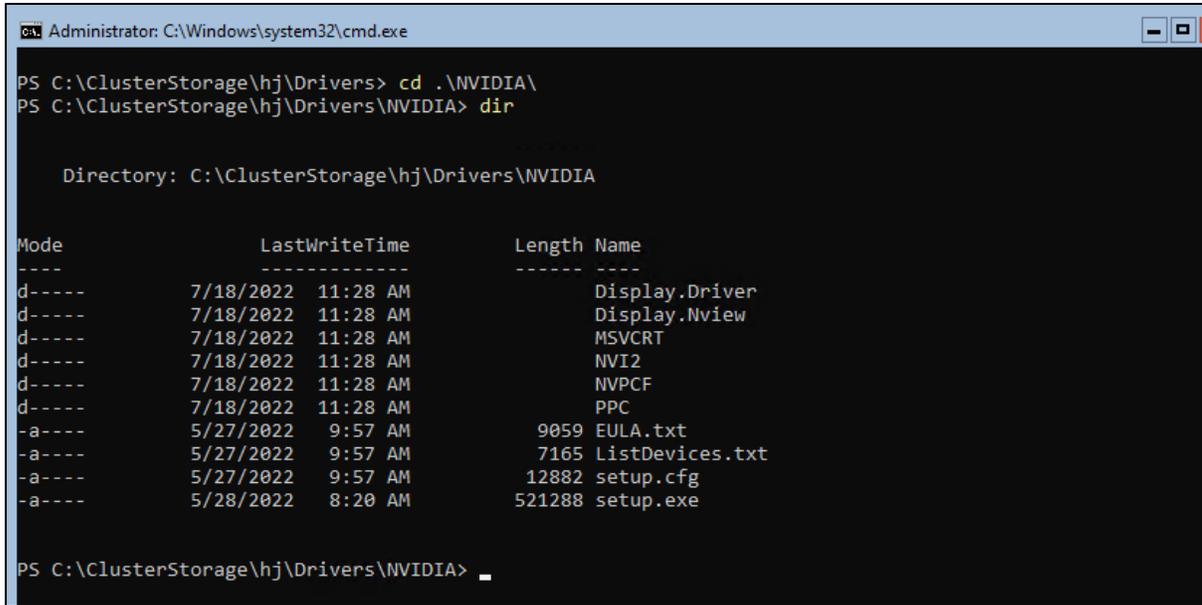
5. In the window that opens, enter the path for the installer to extract its files. This directory will be automatically deleted after driver installation.



6. The installer will take a minute or so to extract its files to the Extraction path specified and then present the main driver installation window, as shown in the following screenshot.



7. At this point, return to the PowerShell window (do not proceed with driver installation) and copy the content from the Extraction path specified to the directory created to hold the extracted content for other nodes in Step 3 above. In our example, we copied the extracted files from “C:\ClusterStorage\hj\Drivers\NVIDIA” to “C:\ClusterStorage\hj\Drivers\GPUs” resulting in the following content in both directories:



```
Administrator: C:\Windows\system32\cmd.exe
PS C:\ClusterStorage\hj\Drivers> cd .\NVIDIA\
PS C:\ClusterStorage\hj\Drivers\NVIDIA> dir

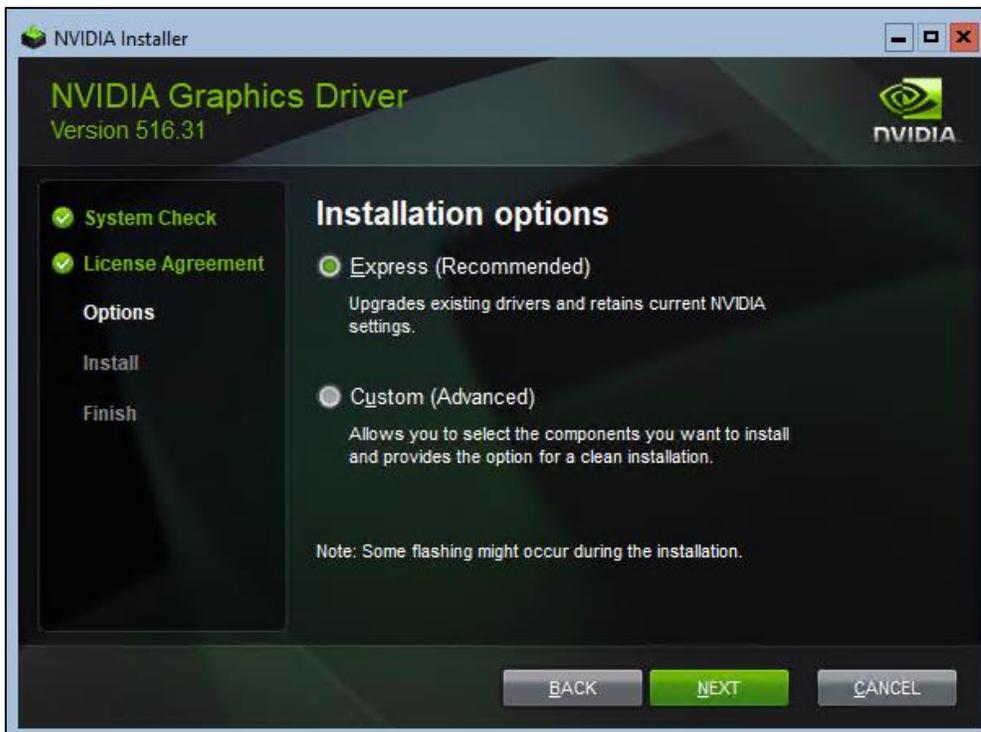
Directory: C:\ClusterStorage\hj\Drivers\NVIDIA

Mode                LastWriteTime         Length Name
----                -
d-----            7/18/2022  11:28 AM             Display.Driver
d-----            7/18/2022  11:28 AM             Display.Nview
d-----            7/18/2022  11:28 AM             MSVCRT
d-----            7/18/2022  11:28 AM             NVI2
d-----            7/18/2022  11:28 AM             NVPCF
d-----            7/18/2022  11:28 AM             PPC
-a-----            5/27/2022   9:57 AM           9059 EULA.txt
-a-----            5/27/2022   9:57 AM           7165 ListDevices.txt
-a-----            5/27/2022   9:57 AM          12882 setup.cfg
-a-----            5/28/2022   8:20 AM          521288 setup.exe

PS C:\ClusterStorage\hj\Drivers\NVIDIA>
```

The driver can be installed on additional nodes from this directory without waiting for content to be extracted. Simply run “setup.exe” from this location.

8. Back in the driver installation window, click through the wizard to install the driver. For general use, the Express installation option is suitable.



9. Install the driver on each of the other nodes. If the extracted installer content was copied as instructed in Step 7, simply run “setup.exe” from the copied directory (in our example, we copied the content to “C:\ClusterStorage\hj\Drivers\GPUs”).



After the GPU driver has been installed on all nodes, the required INF file must be installed on all nodes in order to assure that Hyper-V can correctly reset the GPU during VM reboots.

2.3 Install required INF file

To install the required INF file to allow proper assignment of a GPU to a VM using DDA, follow these steps:

1. Download the ZIP archive that contains the INF files from NVIDIA at the following URL:

<https://docs.nvidia.com/datacenter/tesla/gpu-passthrough/index.html#introduction>

NVIDIA GPU Passthrough Support

This document describes support for passthrough of NVIDIA GPUs.

Introduction

This bulletin describes the changes required to support GPU passthrough (Discrete Device Assignment or DDA) to pass an entire NVIDIA GPU PCIe device into a VM. The information contained in this bulletin is specifically targeted towards systems that use NVIDIA GPUs with [Azure Stack HCI version 21H2](#).

To enable NVIDIA GPUs in passthrough with Azure Stack HCI, an additional INF file needs to be installed on the system. An FAQ is provided below.

Frequently Asked Questions

INF File

An INF file is required to enable DDA functionality with NVIDIA GPUs. The INF informs Hyper-V on how to correctly reset the GPUs during VM reboots. This guarantees the GPU is in a clean state when the VM boots up, and avoids issues such as the driver not being able to initialize the GPU after VM reboot.

For any issues with using the INF, contact your OEM via their support channels.

Download INF

The required INF file (contained in a zip archive) for Azure Stack HCI can be downloaded from [here](#).

OS Support

The INF is intended to be used only for Azure Stack HCI, version 21H2.

GPUs Supported

Currently T4, A2, A10, A16, A30, A40 and A100 GPUs are supported. The list will be updated if more GPUs require the INF to be installed.

Installing the INF

Please follow this guide to install the INF: <https://docs.microsoft.com/en-us/windows-hardware/drivers/devtest/pnputil-examples>.

2. Extract the INF files from the archive.
3. After locating the appropriate INF files for the GPU installed, copy them to each of the nodes that has one or more GPUs installed.
4. On each node, use the following PowerShell command to install the INF file (our example uses the INF file for the NVIDIA A30 GPU):

```
pnputil /add-driver *A30*.inf /install
```

After the appropriate INF files have been installed on all nodes, Windows Admin Center (WAC) can be configured to manage the GPUs.

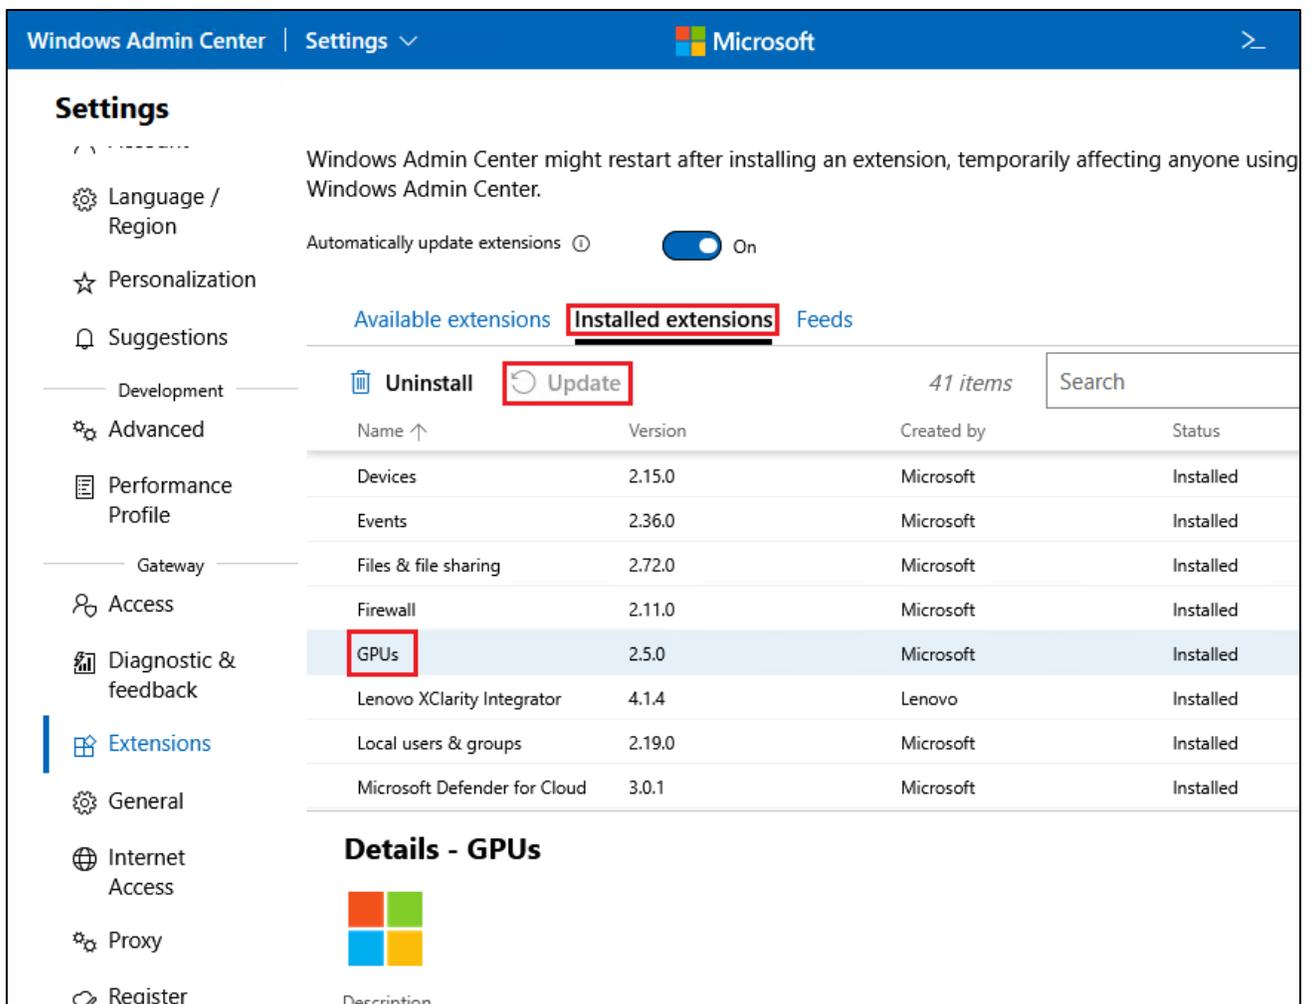
3 Manage GPUs using WAC

If you are using Windows Admin Center (WAC) to manage the Azure Stack HCI cluster, check to ensure that WAC can see the GPUs installed in the nodes.

3.1 Install GPUs extension

You will need to install the GPUs extension in WAC if not already done. To do this, follow these steps:

1. In WAC, navigate to Settings > Extensions.
2. If the GPUs extension is shown in the list of Available extensions, select it and then click Install.
3. If the GPUs extension is not shown in the list of Available extensions, click the Installed extensions heading. Check to make sure the latest version of the GPUs extension is installed. That is, with the GPUs extension selected, the Update button should be grayed out.

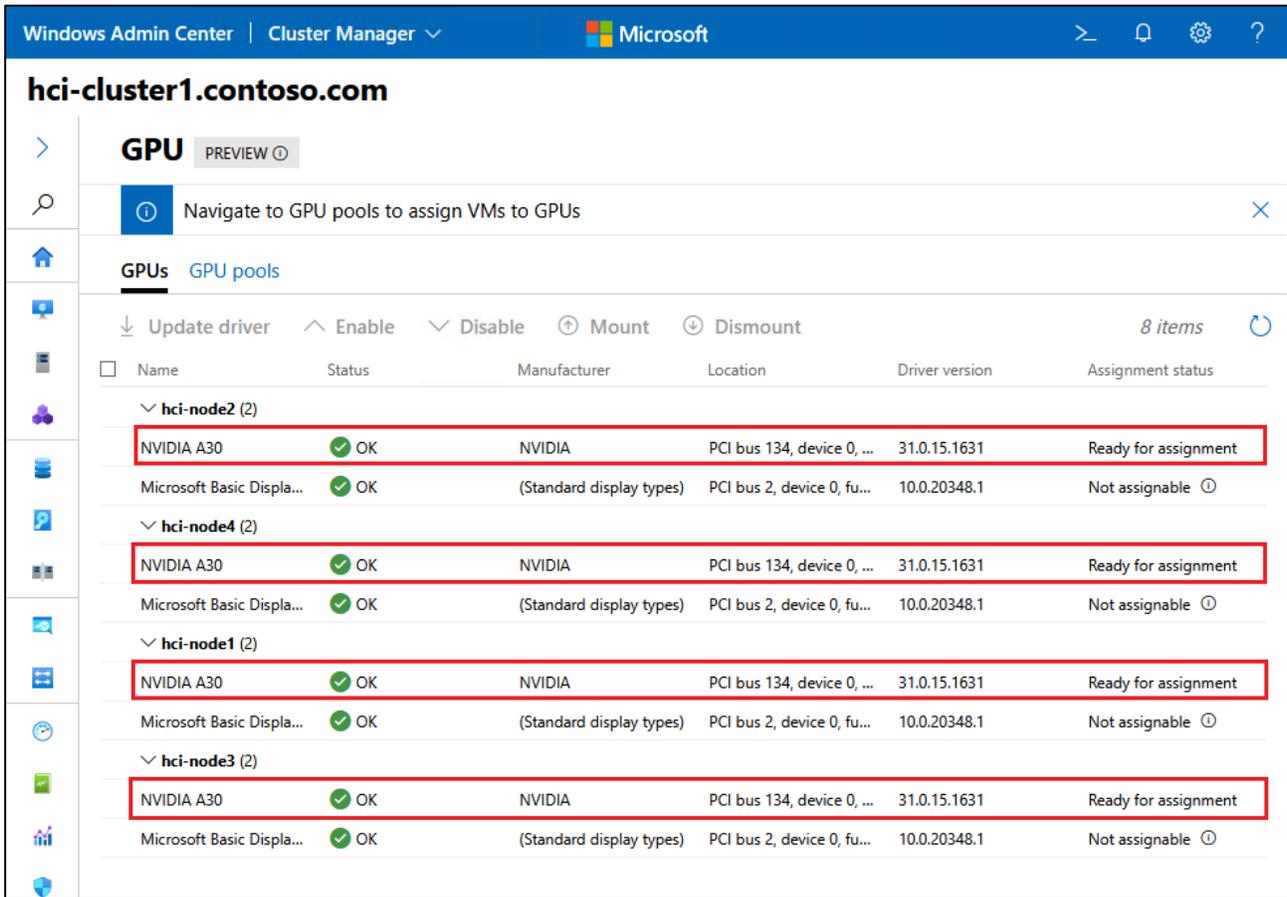


3.2 Create GPU pool

Once the current version of the GPUs extension has been installed in WAC, the GPUs installed in the HCI cluster nodes can be managed and configured using WAC. This includes some basic GPU management functions, such as the ability to mount and unmount GPUs from the host, and to create GPU pools. Once a

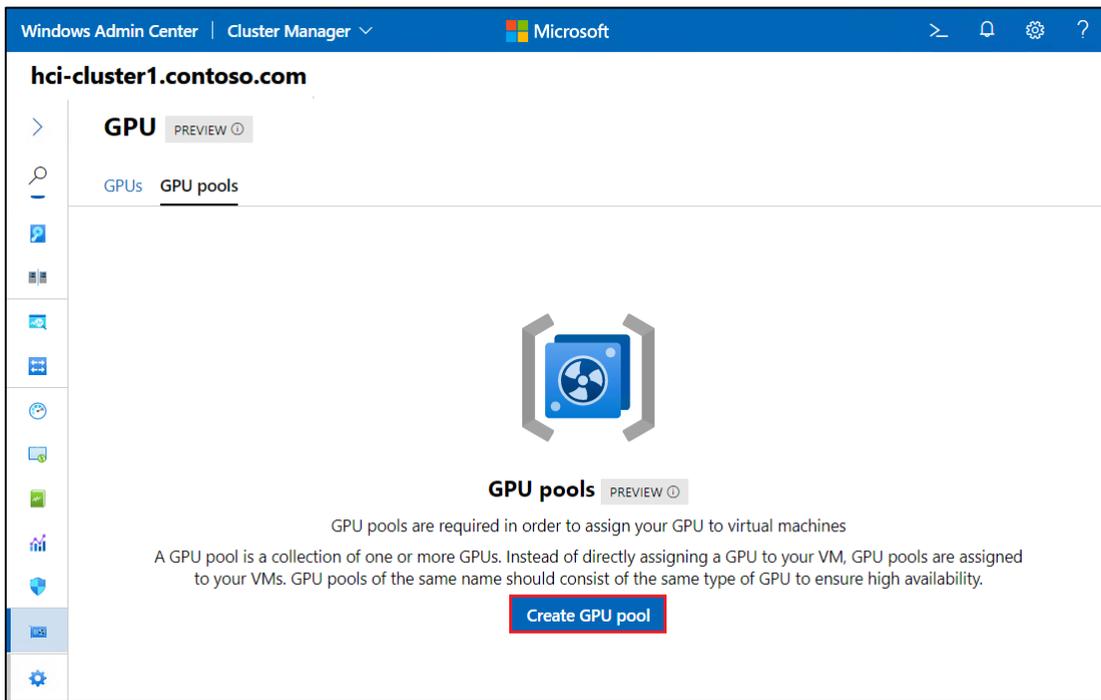
GPU pool has been created, assignment of GPUs to virtual machines using Discrete Device Assignment (DDA) can be done, which makes the entire GPU available to a specified virtual machine.

To verify that all GPUs have been identified in WAC, connect to the cluster in WAC and then use the left navigation pane to select the GPUs extension. Each node should be shown, including installed GPUs by name. The following screenshot shows our 4-node cluster with an NVIDIA A30 24GB PCIe Gen4 Passive GPU installed in each of the nodes.



GPU Pools are created to allow GPUs to be assigned to specific virtual machines. To create a GPU Pool, follow these steps:

1. In WAC, with the GPU extension showing, click the GPU pools heading, and then click on the Create GPU pool button.



2. Select the appropriate servers from the list presented, provide a name for the new GPU pool, choose the GPUs to place in the pool, and then click Save. It should NOT be necessary to check the Assign without host INF file (not recommended) checkbox since the appropriate INF file was installed in the *Install required INF file* section above.

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*

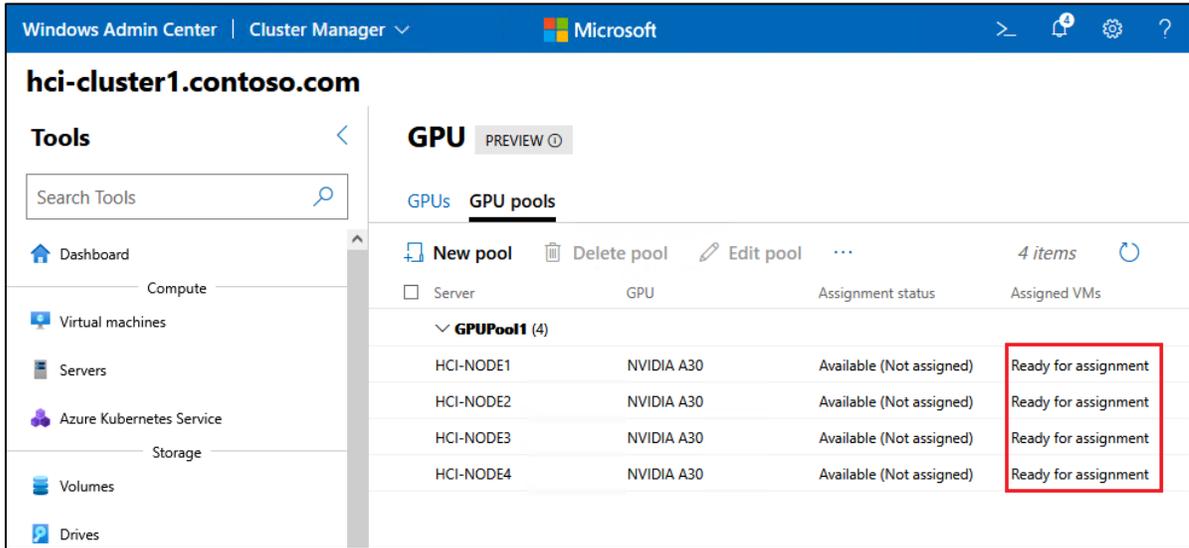
Select GPUs *

Name	Status	Manufacturer
hci-node2	OK	
NVIDIA A30	OK	NVIDIA
hci-node3	OK	
NVIDIA A30	OK	NVIDIA
hci-node4	OK	
NVIDIA A30	OK	NVIDIA

Advanced >

Assign without host INF file (not recommended)

- Once the GPU pool has been created, ensure that all GPUs in the pool show “Ready for assignment” in the Assigned virtual machines column.



The environment is now ready for GPUs to be assigned to individual virtual machines.

3.3 GPU assignment using DDA

Once a GPU pool has been created, virtual machines can be assigned to the pool. Assigning one or more entire GPUs to a single virtual machine via Discrete Device Assignment (DDA) is also known as GPU pass-through and is supported by all GPUs available for ThinkAgile MX solutions.

To assign a virtual machine to the pool, follow these steps:

- In WAC, with the GPU extension showing and click the + Assign VM to pool button to open the VM assignment screen. If the + Assign VM to pool button is not visible, it is likely hidden behind the ellipsis (“...”). Currently, the Type of assignment cannot be changed, since only DDA is currently supported by Microsoft. For Server, choose a node that is the host for the virtual machine needing the GPU assignment (in our example “hci-node1.contoso.com”). For GPU pool, choose the appropriate GPU pool (in our example “GPUPool01”). For Virtual machine, choose an appropriate virtual machine that is running on the selected host (in our example “GPU-VM4”). In the Advanced area, set the High memory mapped IO space (in MB) setting to “66560” and choose whether or not to check the Configure offline action to force shutdown checkbox, depending on your needs. Once all settings have been specified, click Assign.

For more about high memory mapped IO space and other considerations for using DDA, see the Microsoft reference article here:

<https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/plan/plan-for-deploying-devices-using-discrete-device-assignment>

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 by using the Advanced options below. [Read more about DDA assignment](#)

Server* hci-node1.contoso.com ▾

GPU pool* GPUPool1 ▾

The pool named GPUPool1 consists of the following GPUs: NVIDIA A30

Virtual machine* GPU-VM4 ▾

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

Advanced >

Low memory mapped IO space (in MB) 3072

High memory mapped IO space (in MB) 66560

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 relating to not enough resources, you'll likely need to modify these values. Consult [Plan for Deploying Devices using Discrete Device Assignment to 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. Thus, this VM will automatically shutdown prior to a failover to ensure GPU high availability. If this is unchecked, you will have to manually shut down your VM to successfully migrate.

Assign Cancel

2. Ensure that the virtual machine was assigned properly to the GPU pool.

Windows Admin Center | Cluster Manager | Microsoft

hci-cluster1.contoso.com

Tools

GPU PREVIEW

GPUs GPU pools

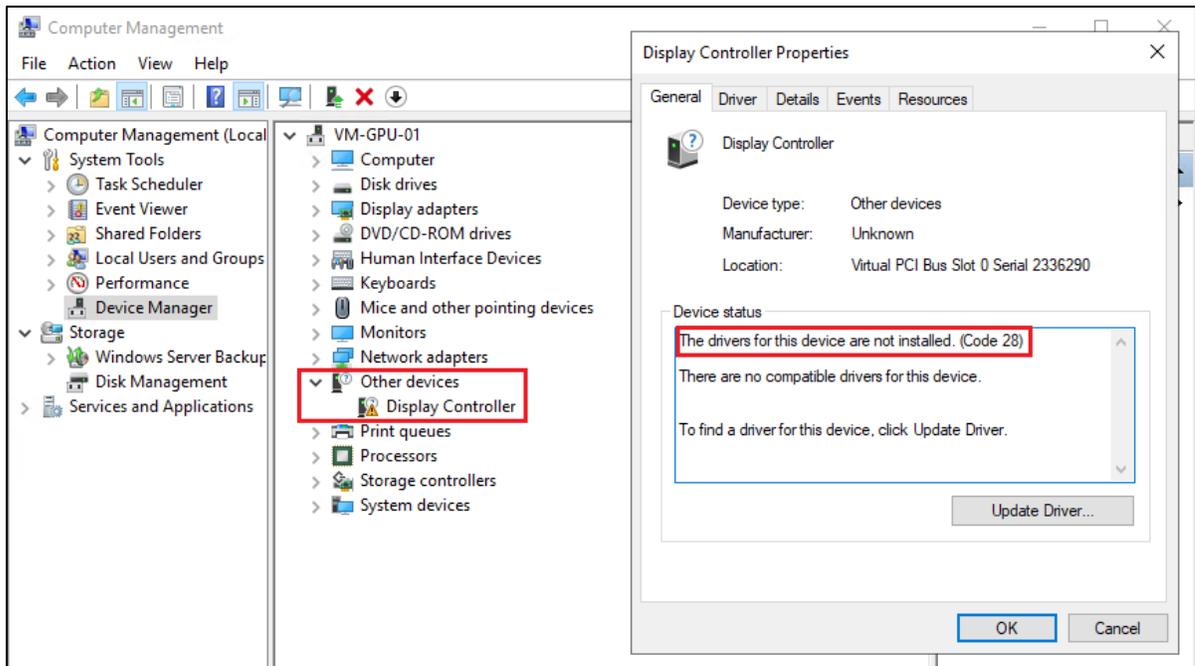
New pool Delete pool Edit pool ... 4 items

Server	GPU	Assignment status	Assigned VMs
GPUPool1 (4)			
HCI-NODE1	NVIDIA A30	Assigned	GPU-VM4
HCI-NODE2	NVIDIA A30	Available (Not assigned)	Ready for assignment
HCI-NODE3	NVIDIA A30	Available (Not assigned)	Ready for assignment
HCI-NODE4	NVIDIA A30	Available (Not assigned)	Ready for assignment

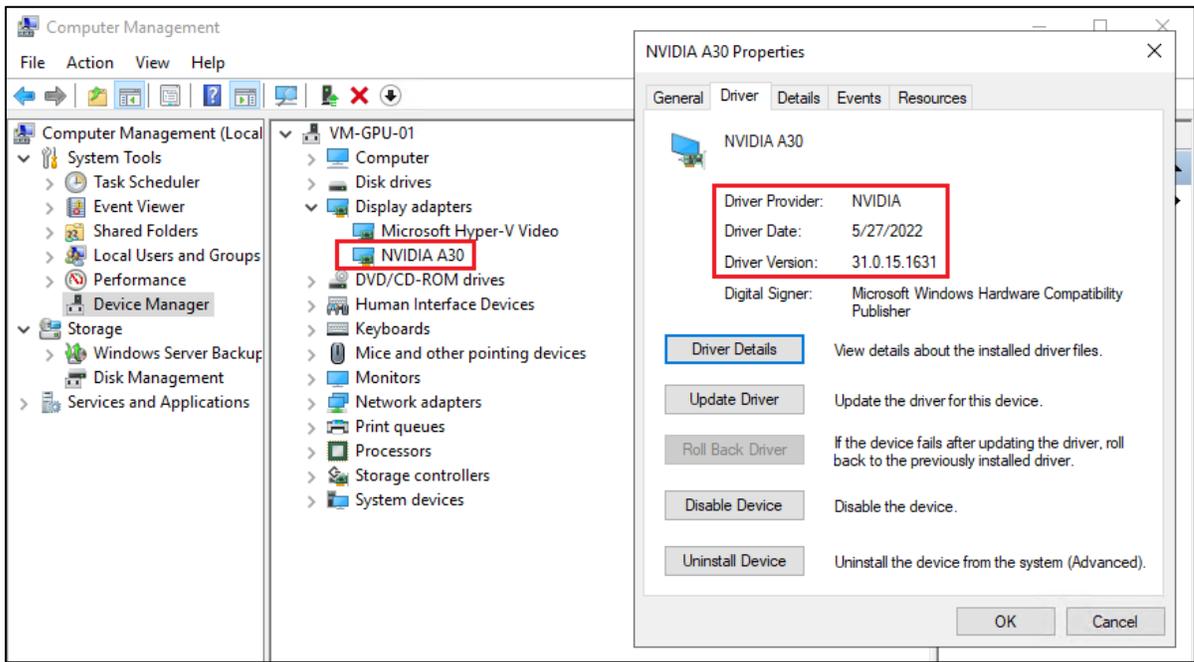
3.4 GPU driver installation on virtual machine

Even though the GPU driver has already been installed in all cluster nodes with a GPU installed, the driver must also be installed in any virtual machine that will use a GPU from the pool. To install the GPU driver in a virtual machine, follow these steps:

1. Access the virtual machine by whatever method is suitable. For our example, we simply use Failover Cluster Manager to connect to the virtual machine.
2. To verify that the GPU driver is not yet installed in the virtual machine, open Device Manager. With no GPU device driver installed, an item “Display Controller” should be listed under “Other devices”. Looking at the Properties for this device will show that no driver is installed.



3. Copy the GPU driver installer to the virtual machine and run it, exactly as was done for the host.
4. After the driver is installed, check Device Manager again to ensure the driver is recognized.



The virtual machine now has complete access and use of the GPU, which can be used for workloads as desired.

4 Summary

GPU virtualization technologies enable GPU acceleration in a virtualized environment, typically within virtual machines. If a workload is virtualized with Hyper-V, then graphics virtualization can be employed in order to provide GPU acceleration from the physical GPU to the virtualized apps or services. In order for a virtual machine to use a GPU installed in its Hyper-V host, several tasks must be accomplished. This document has provided the steps used to perform the following tasks:

- Install the GPU device driver in the host.
- Use the WAC GPUs Extension to create a GPU Pool.
- Assign virtual machines to GPUs in the Pool.
- Install the GPU device driver in any virtual machines that have been assigned.

5 Additional resources

The following resources might be useful in working with Lenovo ThinkAgile MX solutions.

Resources for Lenovo ThinkAgile MX Series solutions

<https://lenovopress.com/servers/thinkagile/mx-series>

Lenovo Press document: **Microsoft Storage Spaces Direct (S2D) Deployment Guide**

<https://lenovopress.com/lp0064>

Lenovo Press document: **ThinkAgile MX1021 on SE350 Azure Stack HCI (S2D) Deployment Guide**

<https://lenovopress.com/lp1298>

Lenovo Press document: **How to Deploy Azure Stack HCI clusters via Microsoft Windows Admin Center**

<https://lenovopress.com/lp1524>

Lenovo Press document: **Lenovo Certified Configurations for Microsoft Azure Stack HCI – V1 Servers**

<https://lenovopress.com/lp0866>

Lenovo Press document: **Lenovo Certified Configurations for Microsoft Azure Stack HCI – V2 Servers**

<https://lenovopress.com/lp1520>

Lenovo ThinkAgile MX Best Recipe landing page

<https://datacentersupport.lenovo.com/us/en/solutions/HT507406>

Lenovo ThinkAgile MX Series

<https://www.lenovo.com/au/en/data-center/software-defined-infrastructure/ThinkAgile-ThinkAgile-MX-Certified-Node/p/WMD00000377>

Microsoft article: *Plan for GPU acceleration in Windows Server*

<https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/plan/plan-for-gpu-acceleration-in-windows-server>

Microsoft article: *Use GPUs with clustered VMs*

<https://docs.microsoft.com/en-us/azure-stack/hci/manage/use-gpu-with-clustered-vm>

Microsoft article: *Deploy graphics devices using Discrete Device Assignment*

<https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/deploy/deploying-graphics-devices-using-dda>

6 Trademarks and special notices

© Copyright Lenovo 2022.

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo, the Lenovo logo, ThinkSystem, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.