NVIDIA vGPU Software for Citrix Virtual Apps & Desktops on Citrix Hypervisor

Deployment Guide

Document History

nv-quadro-vgpu-deployment-guide-citrixonxenserver-v1-112020

Version	Date	Authors	Description of Change
01	Aug 17, 2020	AFS, JJC, EA	Initial Release
02	Dec 16, 2020	AFS, JJC, EA	Formatting update

Table of Contents

Chapter 1. Getting Started

NVIDIA virtual GPU (vGPU) allows multiple virtual machines (VMs) to have simultaneous, direct access to a single physical GPU, using the same NVIDIA graphics drivers that are deployed on non-virtualized operating systems. It also enables multiple GPUs to be aggregated and allocated to a single virtual machine to power the most demanding workloads. This gives VMs unparalleled graphics performance and application compatibility, together with cost-effectiveness and scalability brought about by sharing a GPU among multiple workloads.

This chapter covers how NVIDIA vGPU solutions fundamentally alters the landscape of desktop virtualization and enables users and applications of all levels of complexity and graphics requirements to utilize said solutions. It also describes the NVIDIA vGPU architecture, the GPUs recommended for virtualization, the three virtual GPU software editions, and key standards supported by NVIDIA virtual GPU technology.

1.1 Why NVIDIA vGPU?

The promise of desktop virtualization, realized for server workloads years ago, is flexibility and manageability. Initially, desktop virtualization was used where the flexibility and security were the primary drivers due to cost considerations. Over the last five years, the cost of desktop virtualization has been coming down quickly due to advances in storage, multi-core processors, and software advances from virtualization leaders like Citrix.

The big remaining challenge for desktop virtualization is providing a cost-effective yet rich user experience. There have been attempts to solve this problem with software graphics or shared GPU technologies, but those technologies do not support the rich applications needed to be successful and ensure end user adoption. This compared to dedicated GPU pass-through that provide 100% application compatibility, but only for the highest end user cases due to the high cost with limited density of virtual machines per host server.

Due to the lack of scalable, sharable, and cost-effective per user GPUs that provide 100% application compatibility, providing a cost-effective rich user experience has been challenging for broad use cases in desktop virtualization. Meanwhile, high-end 3D applications simply did not work in a virtualized environment or were so expensive to implement with pass-thru it was reserved for only the most limited of circumstances.

Today, this is no longer true thanks to the NVIDIA vGPU solution combined with Citrix Virtual Desktops and Apps running on Citrix Hypervisor. NVIDIA vGPU gives you the best of both worlds where multiple virtual desktops or applications share a single physical GPU, and multiple GPUs on a single physical PCI card, all providing the 100% application compatibility of pass-through graphics, but with lower cost since multiple virtual session hosts can share a single graphics card to provide a rich,

yet more cost-effective user experience. With Citrix Virtual Apps and Desktops, you are able to centralize, pool, and more easily manage traditionally complex and expensive, distributed workstations and desktops. Now all your user groups can take advantage of the promise of virtualization.

1.2 NVIDIA vGPU Architecture

A high-level architecture of NVIDIA vGPU is illustrated below. The NVIDIA virtual GPU enabled VDI environment is illustrated below in Figure 1.1. Here, we have GPUs in the server, and the NVIDIA vGPU manager software (VIB) is installed on the host server. This software enables multiple VMs to share a single GPU or if there are multiple GPU's in the server, they can be aggregated so that a single VM can access multiple GPUs. This GPU enabled environment, provides not only unprecedented performance, but it also enables support for more users on a server because work that was typically done by the CPU, can be offloaded to the GPU. Physical NVIDIA GPUs can support multiple *virtual* GPUs (vGPUs) and be assigned directly to guest VMs under the control of NVIDIA's Virtual GPU Manager running in a hypervisor.

Guest VMs use the NVIDIA vGPUs in the same manner as a physical GPU that has been passed through by the hypervisor. In the VM itself, vGPU drivers are installed which pertain to the different license levels that are available. The NVIDIA RTX Virtual Workstation (RTX vWS) license pertains to the NVIDIA RTX Enterprise driver, whereas the NVIDIA Virtual PC (vPC) and NVIDIA Virtual Applications (NVIDIA vApps) pertain to the NVIDIA graphics driver.

Figure 1.1 NVIDIA vGPU Platform Solution Architecture

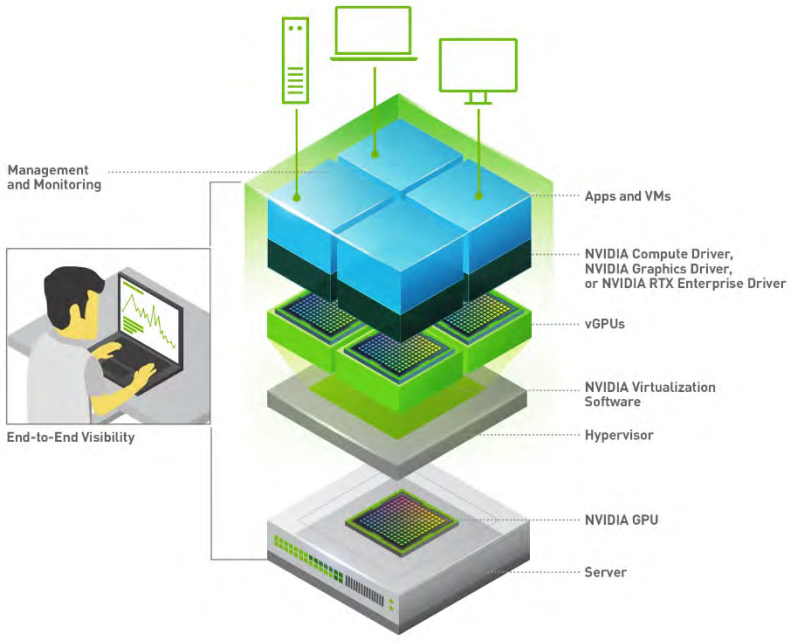

NVIDIA vGPUs are comparable to conventional GPUs in that they have a fixed amount of GPU-Memory and one or more virtual display outputs or *heads*. Multiple heads support multiple displays. Managed by the NVIDIA vGPU Manager installed in the hypervisor, the vGPU Memory is allocated out of the physical GPU frame buffer at the time the vGPU is created. The vGPU retains exclusive use of that GPU Memory until it is destroyed.

All vGPUs resident on a physical GPU share access to the GPU's engines, including the graphics (3D) and video decode and encode engines. Figure1.2 shows the vGPU internal architecture. VM's guest OS leverages direct access to the GPU for performance and critical fast paths. Non-critical performance management operations use a para-virtualized interface to the NVIDIA Virtual GPU Manager.

Figure 1.2 NVIDIA vGPU Internal Architecture

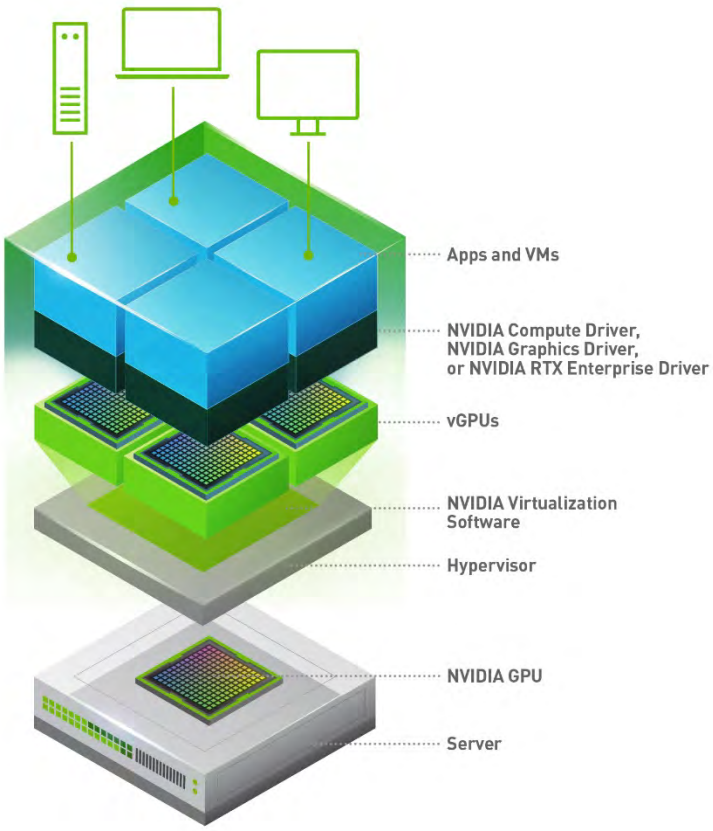

1.3 NVIDIA vGPU Software Licensed Products

NVIDIA virtual GPU software divides NVIDIA GPU resources so the GPU can be shared across multiple virtual machines running any application.

- The portfolio of NVIDIA virtual GPU software products for desktop virtualization is as follows:
- NVIDIA RTX Virtual Workstation (vWS)
- NVIDIA Virtual PC (NVIDIA vPC)
- NVIDIA Virtual Applications (NVIDIA vApps)

To run these software products, you need an NVIDIA GPU and software license that addresses your specific use case. For NVIDIA Virtual Applications (NVIDIA vApps) you can use Citrix Virtual Apps and for NVIDIA Virtual PC (NVIDIA vPC) and NVIDIA RTX Virtual Workstation (vWS) you can use Citrix Virtual Desktop.

For further details on vGPU licensing, please refer to the vGPU Packaging and Licensing-Guide.

1.4 Supported Graphics Protocols

This version of NVIDIA vGPU software includes support for:

- Full DirectX 12, Direct2D, and DirectX Video Acceleration (DXVA)
- OpenGL 4.6
- NVIDIA vGPU SDK (remote graphics acceleration)
- Vulkan 1.1
- OpenCL and CUDA applications WITHOUT Unified Memory are supported on supported GPUs.
- NVIDIA CUDA Toolkit and OpenCL Support on NVIDIA vGPU Software

Note: Unified Memory and CUDA tools are **NOT** supported on NVIDIA vGPU

1.5 Before You Begin

This section describes the general prerequisites and some general preparatory steps that must be addressed before proceeding with the deployment.

Note: This deployment guide assumes you are building an environment as a proof of concept and is not meant to be a production deployment, as a result, choices made are meant to speed up and ease the process. See the corresponding guides for each technology, and make choices appropriate for your needs, before building your production environment.

1.5.1 Server BIOS Settings

Configure the BIOS as appropriate for your physical hosts, as described below:

- Hyperthreading – Enabled
- Power Setting or System Profile– High Performance
- CPU Performance (if applicable) – Enterprise or High Throughput
- Memory Mapped I/O above 4-GB - Enabled (if applicable)
- VT-d or AMD IOMMU – Enabled

1.5.2 Citrix GPU Utilization Patch

KB4586830 & KB458639 address an issue with incorrect Canonical Display Driver (CDD) buffer flushing, which degrades performance in Remote Desktop Protocol (RDP) Windows 2000 Display Driver Model (XDDM) scenarios. This issue affects applications that use graphics processing units (GPU) to operate, such as Microsoft Teams, Microsoft Office, and web browsers.

- Server 2016 – KB4586830
- Server 2019 – KB4586839

Please follow the steps below to enable KB4586830 on Server 2016. It is not enabled by default post installation.

- To Enable the fix - reg add HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Policies\Microsoft\FeatureManagement\Overrides /v 1826589834 /t REG_DWORD /d 1 /f
- To Disable the fix - reg add HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Policies\Microsoft\FeatureManagement\Overrides /v 1826589834 /t REG_DWORD /d 0 /f

1.5.3 Citrix Virtual Machine Requirements

Before getting started with a deployment, please review the following Guest OS requirements below.

- Delivery Controller
- Virtual Delivery Agent (VDA) for single-session OS
- Virtual Delivery Agent (VDA) for multi-session OS

1.5.3.1 Additional Sizing Resources for Citrix

- Database sizing tool for XenDesktop 7
- Citrix VDI Best Practices for XenApp and XenDesktop 7.15 LTSR
- Citrix Virtual Apps and Desktops Single-Server Scalability
- Complete system requirements

1.5.4 Virtual GPU Evaluation licenses

In order to run a PoC/trial of NVIDIA Virtual GPU, a vGPU license is required. An evaluation license is available at the vGPU Evaluation site.

Chapter 2. Installing Citrix Hypervisor

This chapter covers the following Citrix Hypervisor installation topics:

- Choosing an install method
- Preparing the USB boot media
- Installing Citrix Hypervisor from the USB media
- Initial host configuration

Note: This deployment guide assumes you are building an environment as a proof of concept and is not meant to be a production deployment, as a result, choices made are meant to speed up and ease the process. See the corresponding guides for each technology, and make choices appropriate for your needs, before building your production environment.

- For the purpose of this guide, Citrix Hypervisor 8.2.0 is used as the hypervisor version.

2.1 Choosing the Installation method

With the ability to install from USB drive, Citrix Hypervisor offers flexibility verses local hard drive install. Please see documentation regarding best practices for logs when booting from USB or similar. This lab used the server's remote console and virtual media to boot from ISO file.

2.2 Installing Citrix Hypervisor

Tip: Throughout the installation, quickly advance to the next screen by pressing F12. Use Tab or the arrow keys to move between options and Space or Enter to select. Press F1 for general help.

After bootup through USB Boot Media, the **Welcome to Citrix Hypervisor** screen will display.

1. Press **[Enter]** to begin preparing installation. The system will scan your hardware to make sure an installation is compatible.

The following **Select Keymap** window displays.

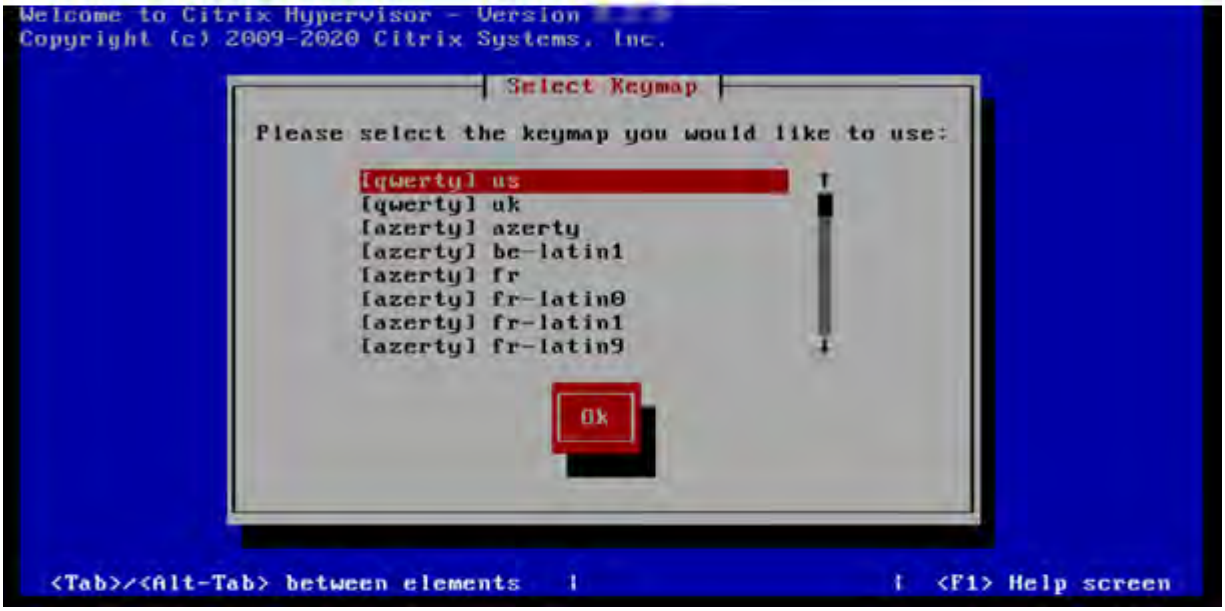

2. Select your desired keyboard layout using the arrow keys and press **[Enter].**

The **Welcome to Citrix Hypervisor Setup** window displays.

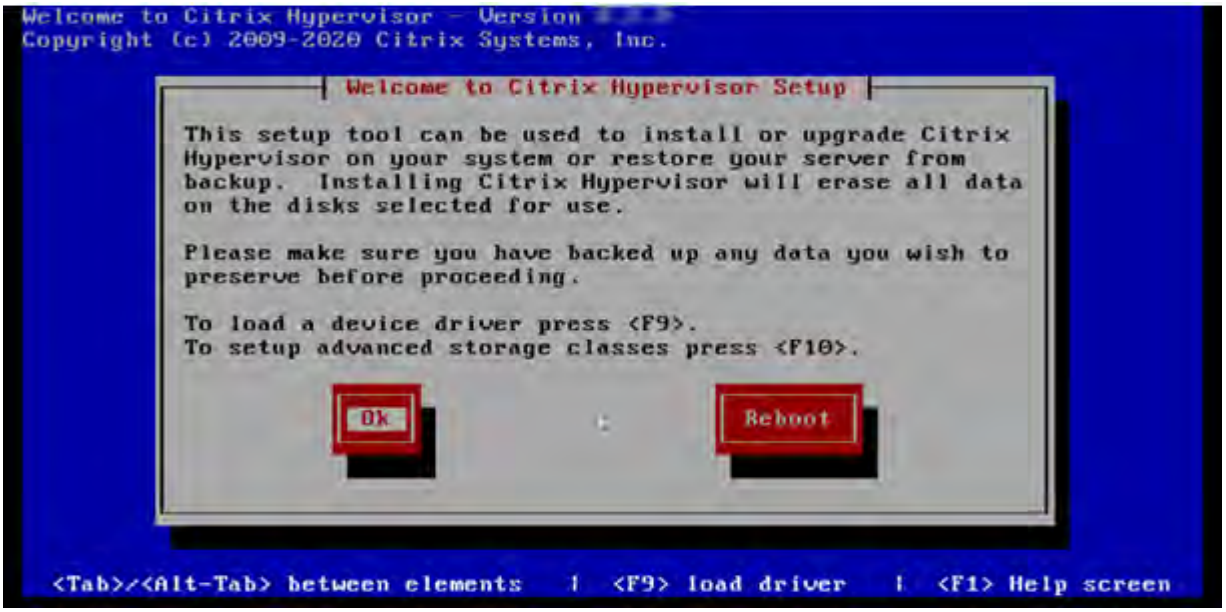

3. Back up any data that you want to preserve. Installing Citrix Hypervisor overwrites data on any hard drives that you select to use for the installation. Select **OK** by pressing **[Enter].**

The **End User License Agreement** (EULA) displays.

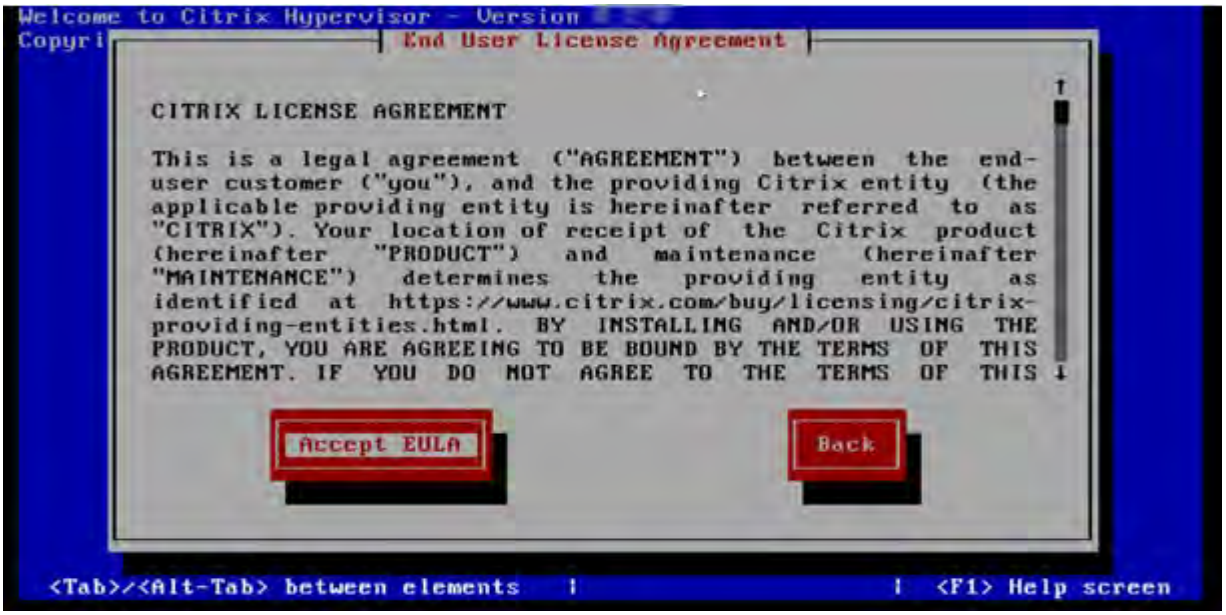

4. Read the EULA and then accept by pressing **[Enter]** to proceed.

> A system hardware warning may appear if hardware virtualization has not been enabled in the System BIOS. While installation can still proceed, make sure to enable hardware virtualization in the System BIOS after installation reboot to avoid any complications with XenCenter.

The **Select Primary Disk** window displays.

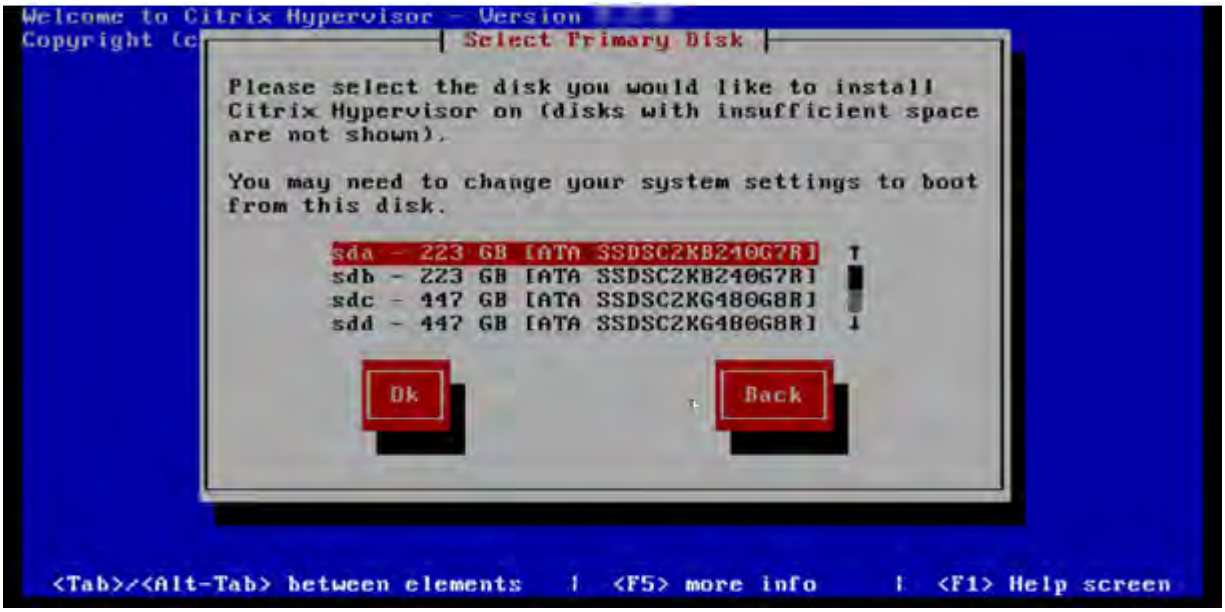

5. If you have multiple hard disks, choose a Primary Disk for the installation. This will overwrite any data upon the Primary Disk. Select **Ok** and press **[Enter].**

The **Virtual Machine Storage** window displays.

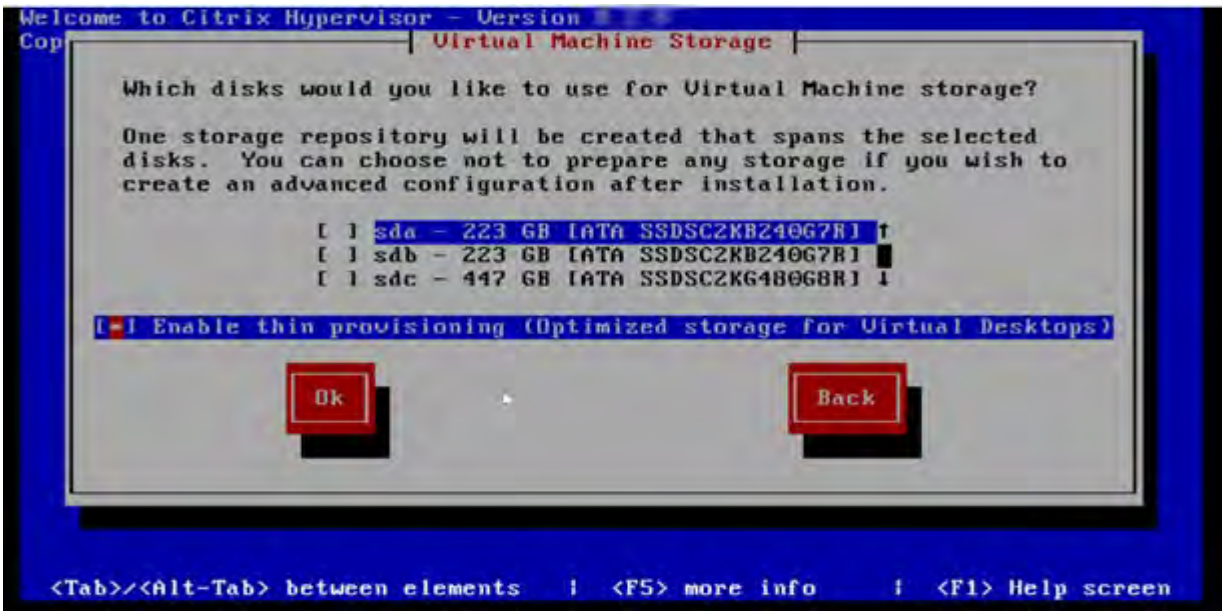

6. Choose which disks you want to use for Virtual Machine Storage. Enable thin provisioning for optimized storage if you would like to do so. Select and **Ok** and press **[Enter].**

Note: See Citrix documentation for support on thin provisioning: https://support.citrix.com/article/CTX233865

The **Select Installation Source** window displays.

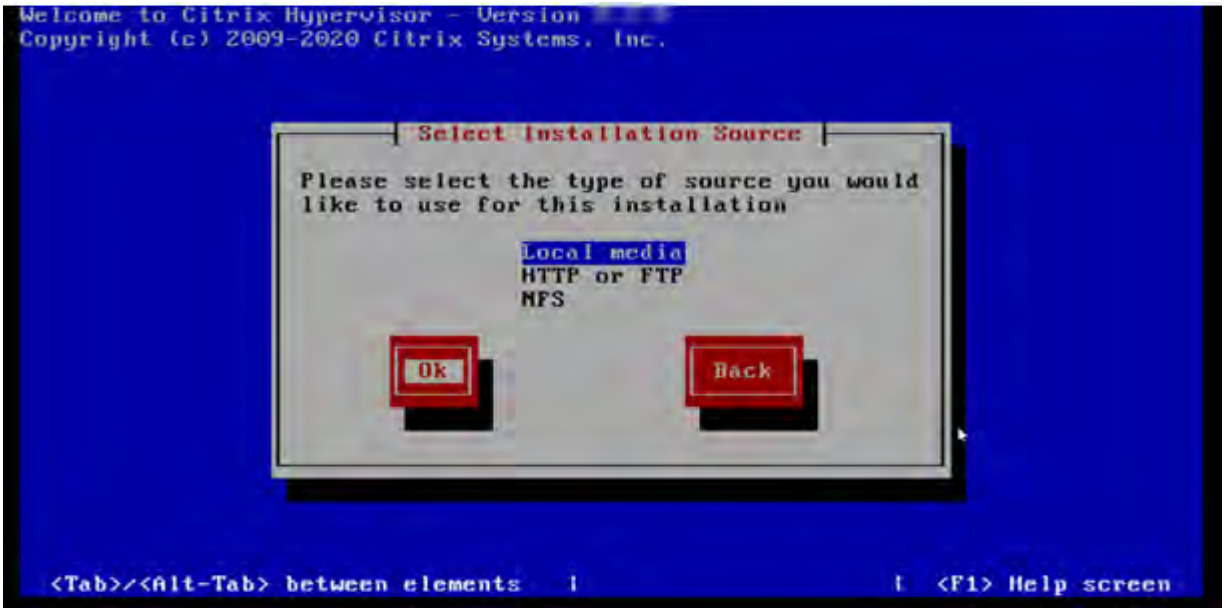

7. Select **Local Media** as your Installation Source. Then navigate to **OK** and press **[Enter].**

The **Verify Installation Source** window displays.

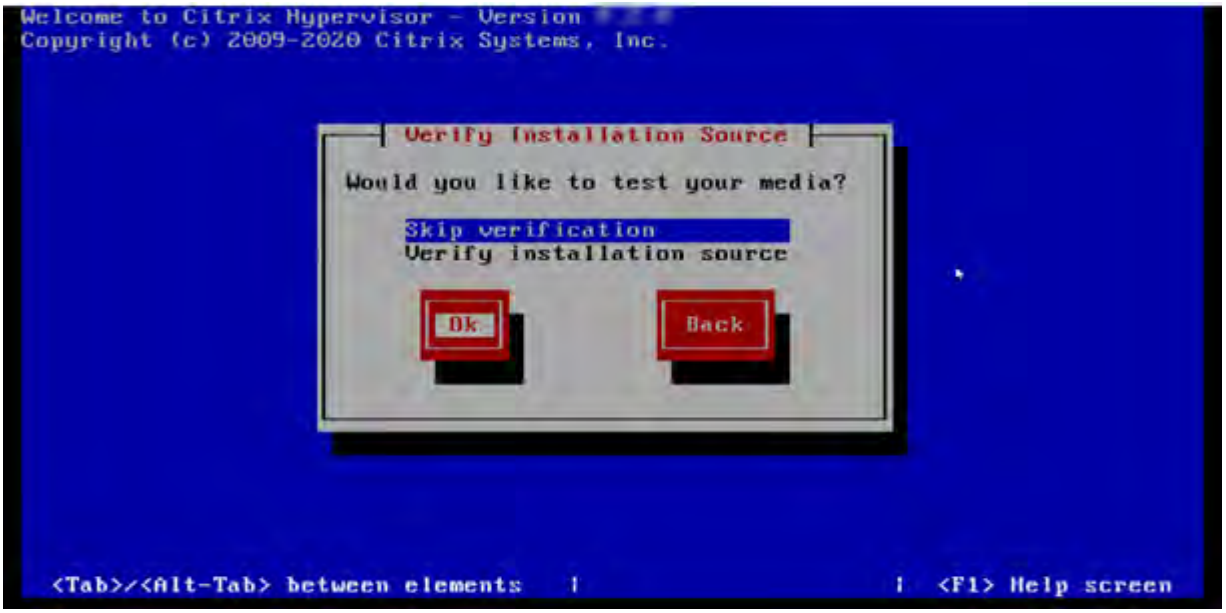

8. Select **Skip verification**. Then navigate to **OK** and press **[Enter].**

The **Set Password** window displays.

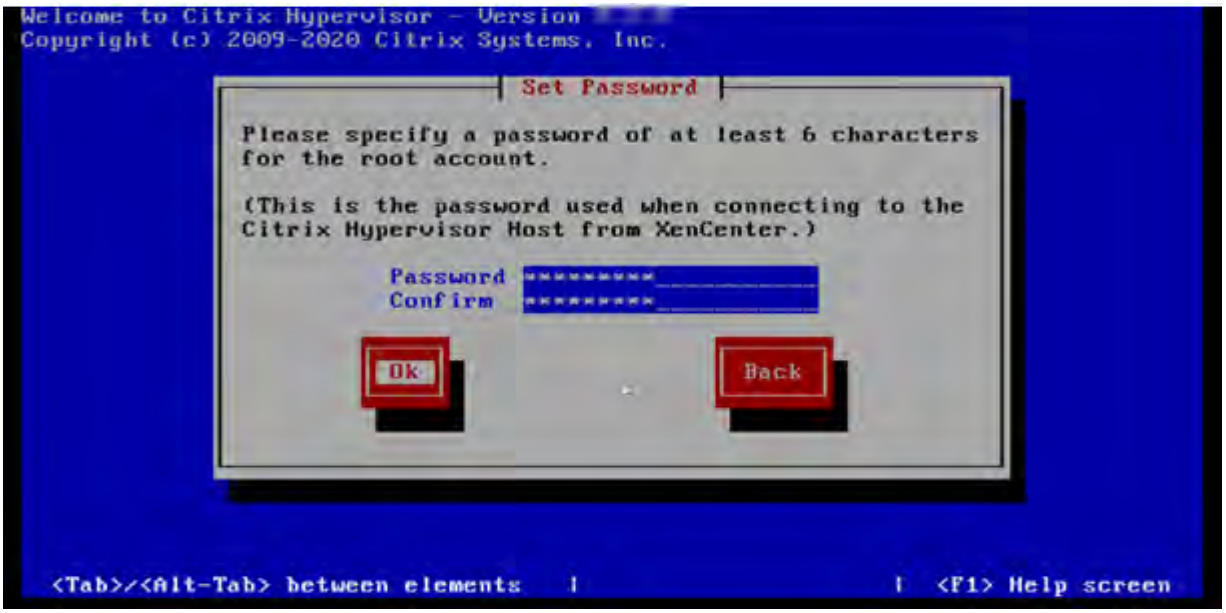

9. Create and confirm a root password for the Citrix Hypervisor Host in the password field. This is the password used when connecting to the Citrix Hypervisor Host from XenCenter. Then navigate to **OK** and press **[Enter].**

> ! CAUTION: To prevent unauthorized access, your selected root password should contain at least eight (8) characters and consist of a mix of lowercase and capital letters, digits, and special characters.

The **Networking** window displays.

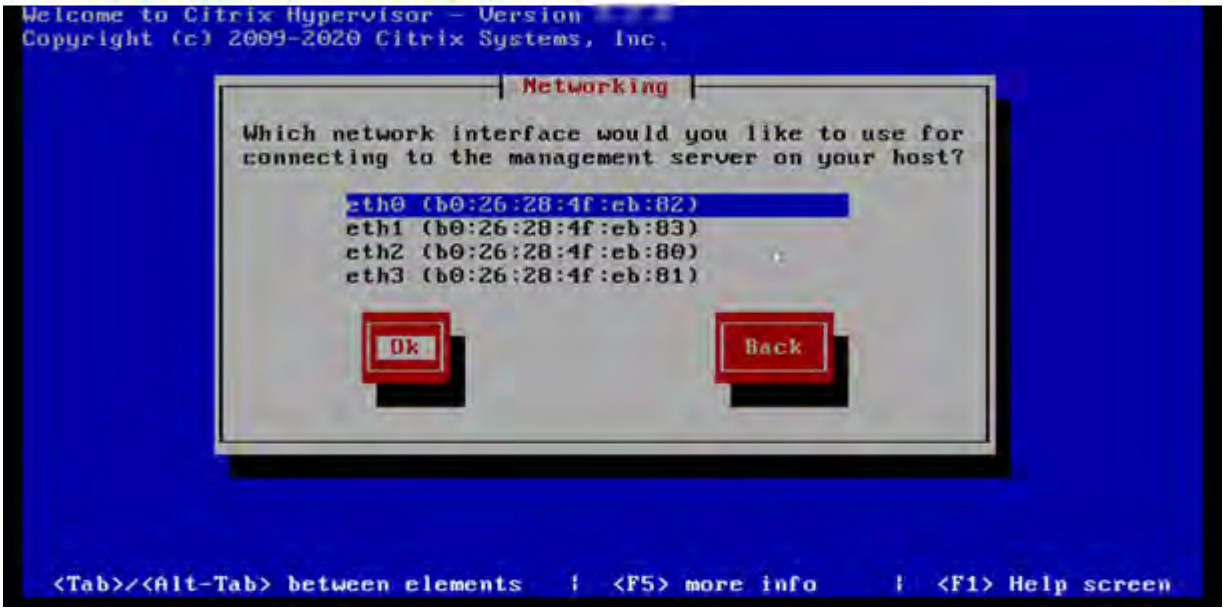

10. Specify the Network Interface or Ethernet connection you want to use for the management server on your host. Select **Ok** and press **[Enter].**

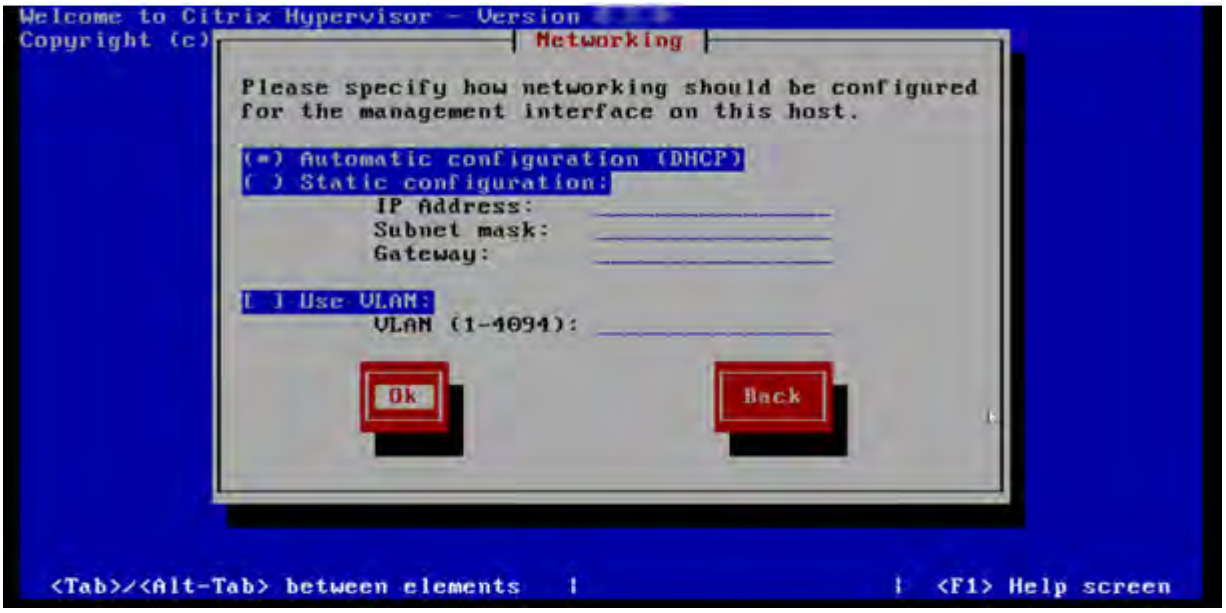

11. Specify automatic configuration (DHCP) or Static Configuration. If selecting Static Configuration, enter the required information into the field. Navigate to **OK** and press **[Enter].**

The **Hostname and DNS Configuration** windows displays.

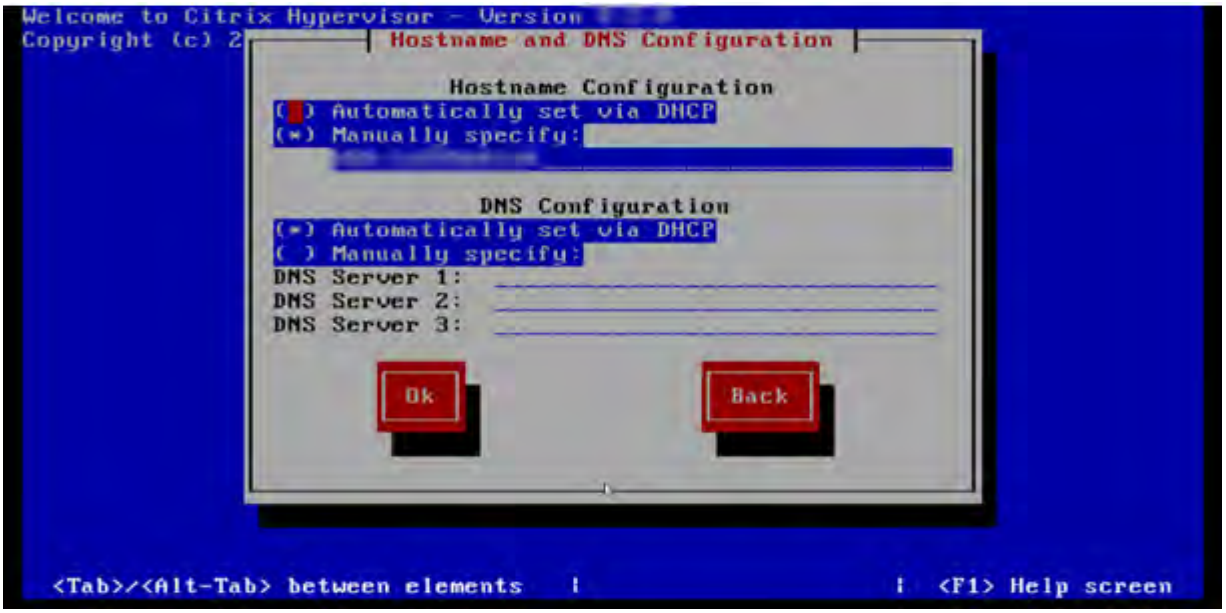

12. Specify the host name and the DNS configuration manually or automatically through DHCP. If you manually configure the DNS, enter the IP addresses of your primary DNS server as well as any secondary servers if needed. Navigate to **OK** and press **[Enter].**

The **Select Time Zone** window displays.

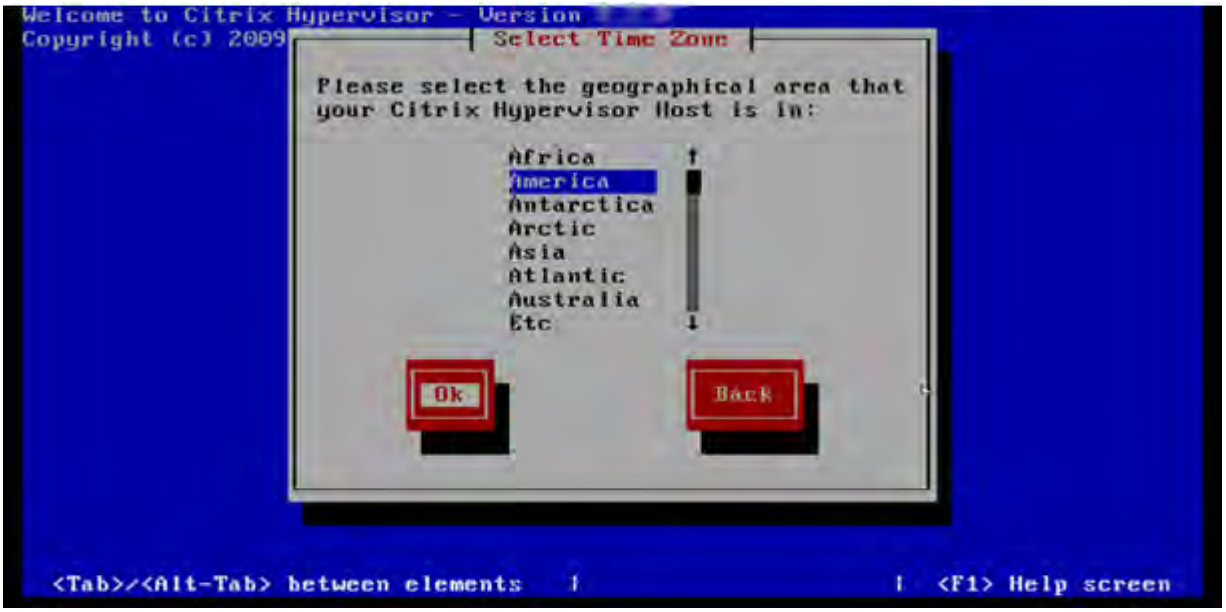

13. Configure Time Zone for Citrix Hypervisor. Select the appropriate listing. Navigate to **OK** and press **[Enter].**

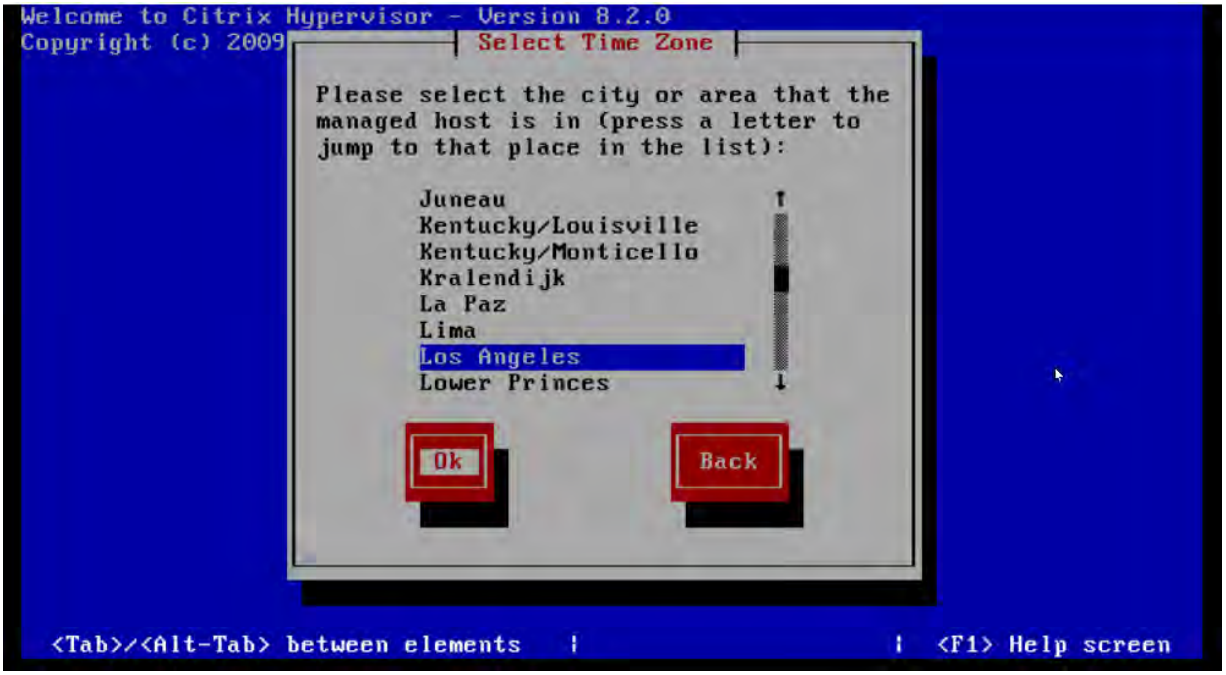

14. Select the appropriate city for local time zone. Navigate to **OK** and press **[Enter].**

The **System Time** window displays.

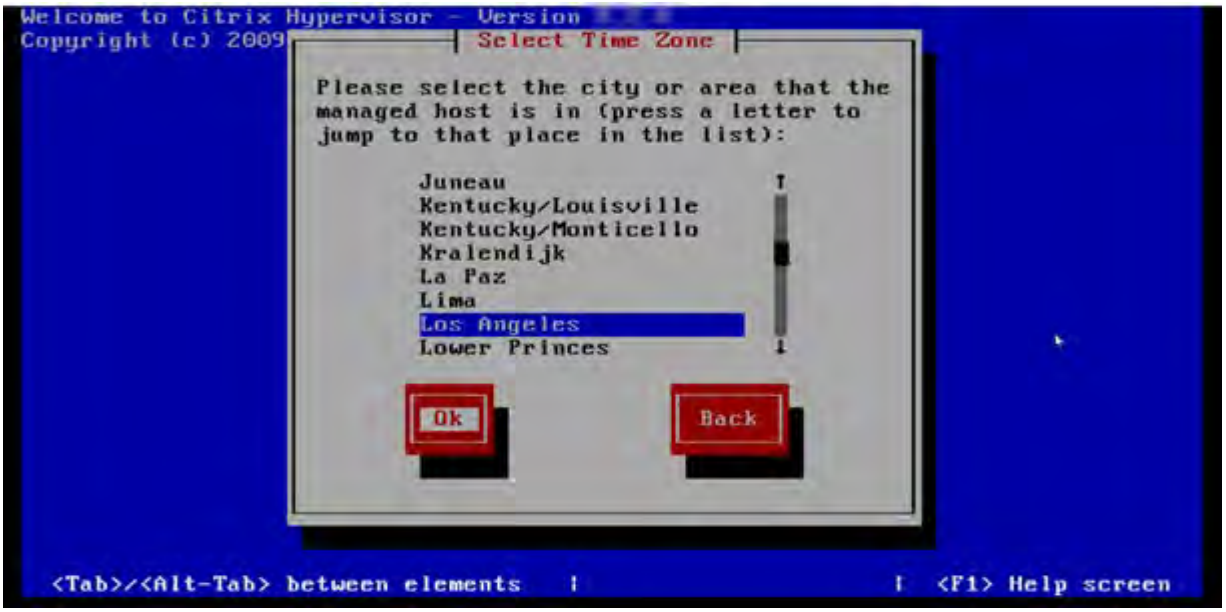

15. Configure the system time with NTP servers. Select Using NTP. Navigate to **OK** and press **[Enter].**

The **NTP Configuration** window displays.

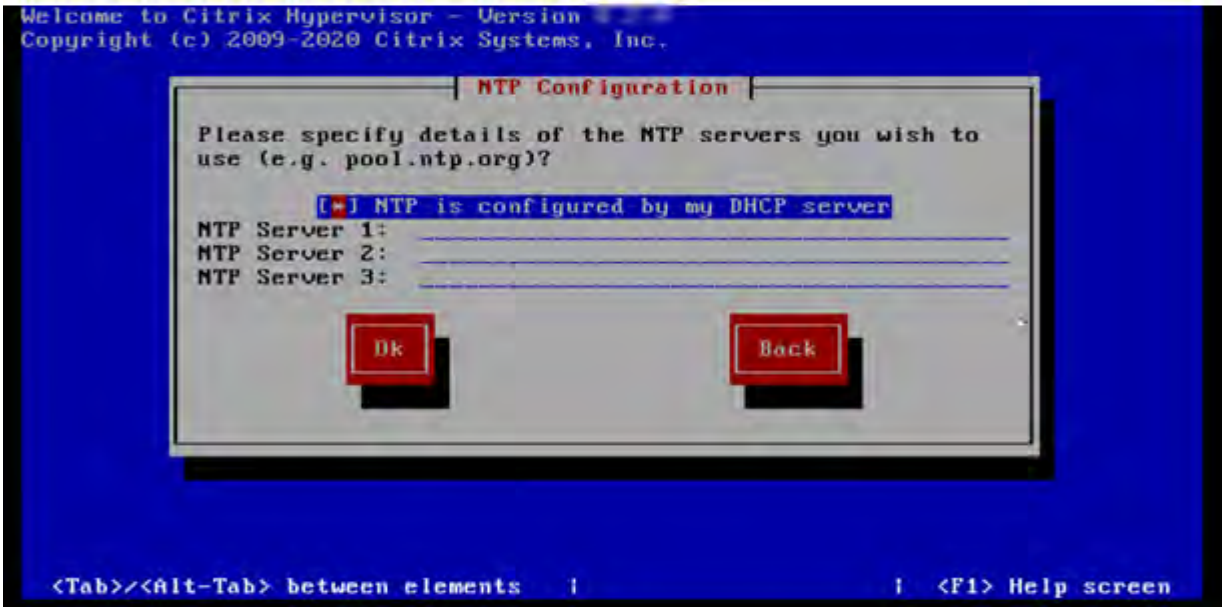

16. Select NTP is configured by the DHCP server or conversely enter at least one NTP server name in the fields below. Navigate to **OK** and press **[Enter].**

The **Confirm Installation** window displays.

17. Confirm the Installation by selecting **Install XenServer.** Then press **[Enter]** to proceed with installation. Wait for the installation to finish.

The **Installation Complete** window displays.

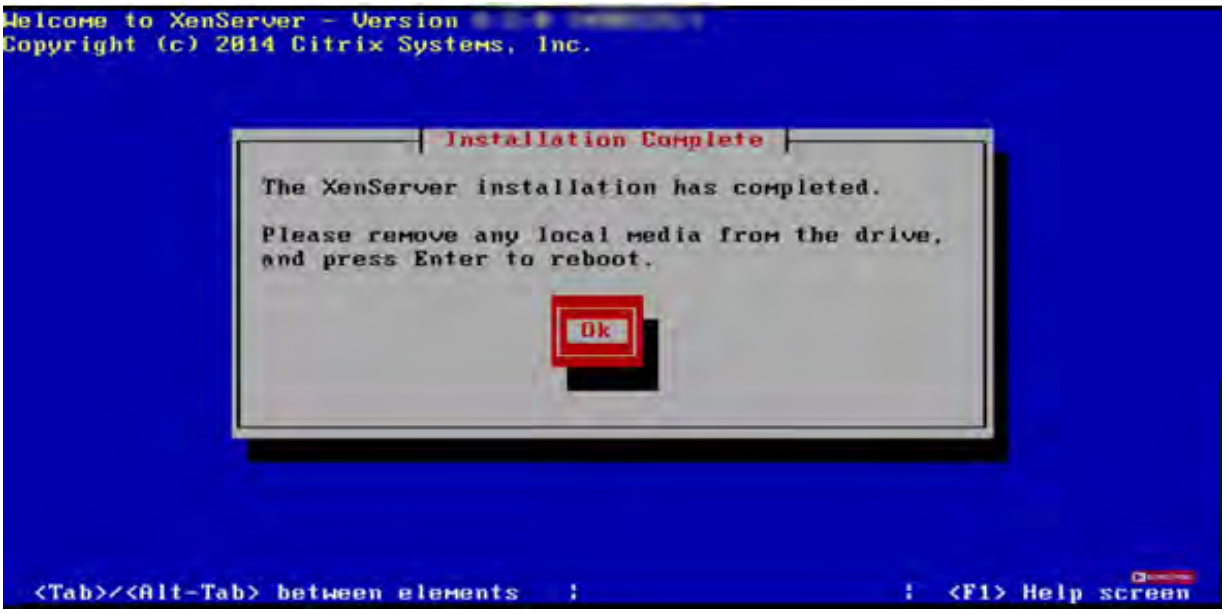

18. Eject the installation media. Press **[Enter]** to reboot the server.

The **Citrix Hypervisor configuration console** will display upon reboot.

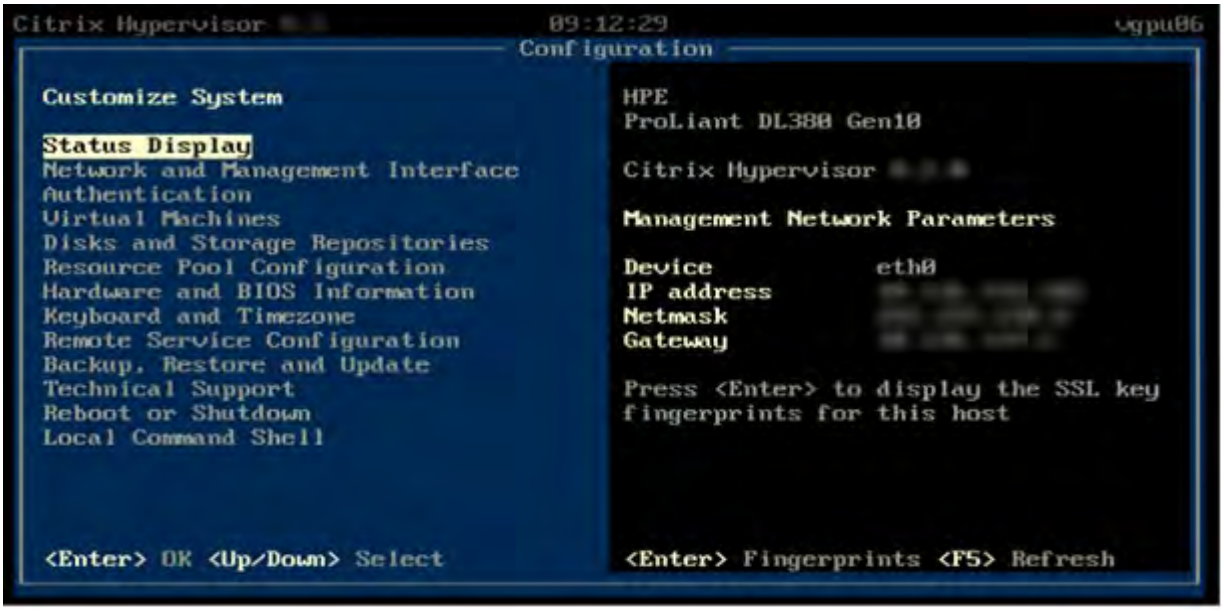

Record the IP address from the Management Network Parameters for use in the next section, Installing Citrix XenCenter.

Note: If you properly set everything during installation, nothing should be required from this console. All additional configurations can be performed through XenCenter which is covered in a later section.

Chapter 3. Installing Citrix XenCenter

This chapter covers installing Citrix XenCenter, including:

- Citrix XenCenter Setup and Installation
- Adding a Server with XenCenter
- Installing Licenses for Citrix Hypervisor with XenCenter

3.1 Installing Citrix XenCenter

1. Download the XenCenter installer from the Citrix download site at www.citrix.com or enter the IP address of the server into your browser.
2. Install Citrix XenCenter to a computer that you want to manage your servers. Select the Installer and open the Citrix XenCenter Setup Wizard.
3. Follow the Setup wizard, select **Next** to continue.

4. Choose an install location and select either the **All Users** or **Just Me** radio button as needed. Select **Next** to continue.

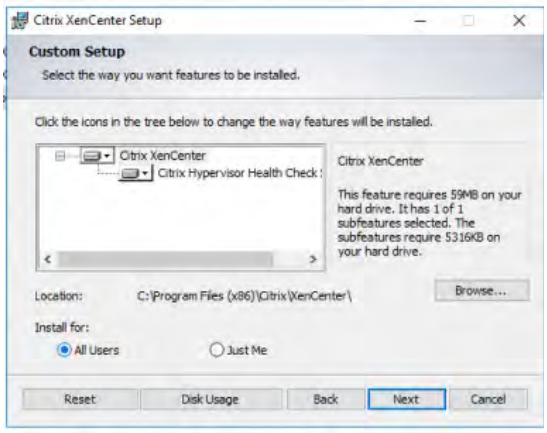

5. Select **Install** to proceed.

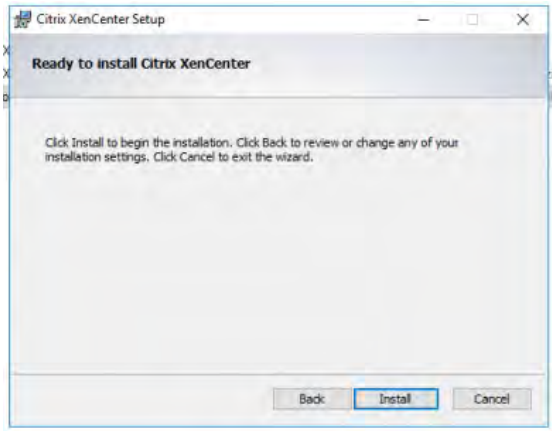

6. Select **Finish** to conclude installation of XenCenter Appliance.

3.2 Post Installation

This section describes post install and configuration for XenCenter.

3.2.1 Adding a Server with XenCenter

1. Launch XenCenter from the Start Menu. A prompt will appear to periodically check for updates and new versions of Citrix Hypervisor and XenCenter if available. Select **Yes** or **No**, as needed.

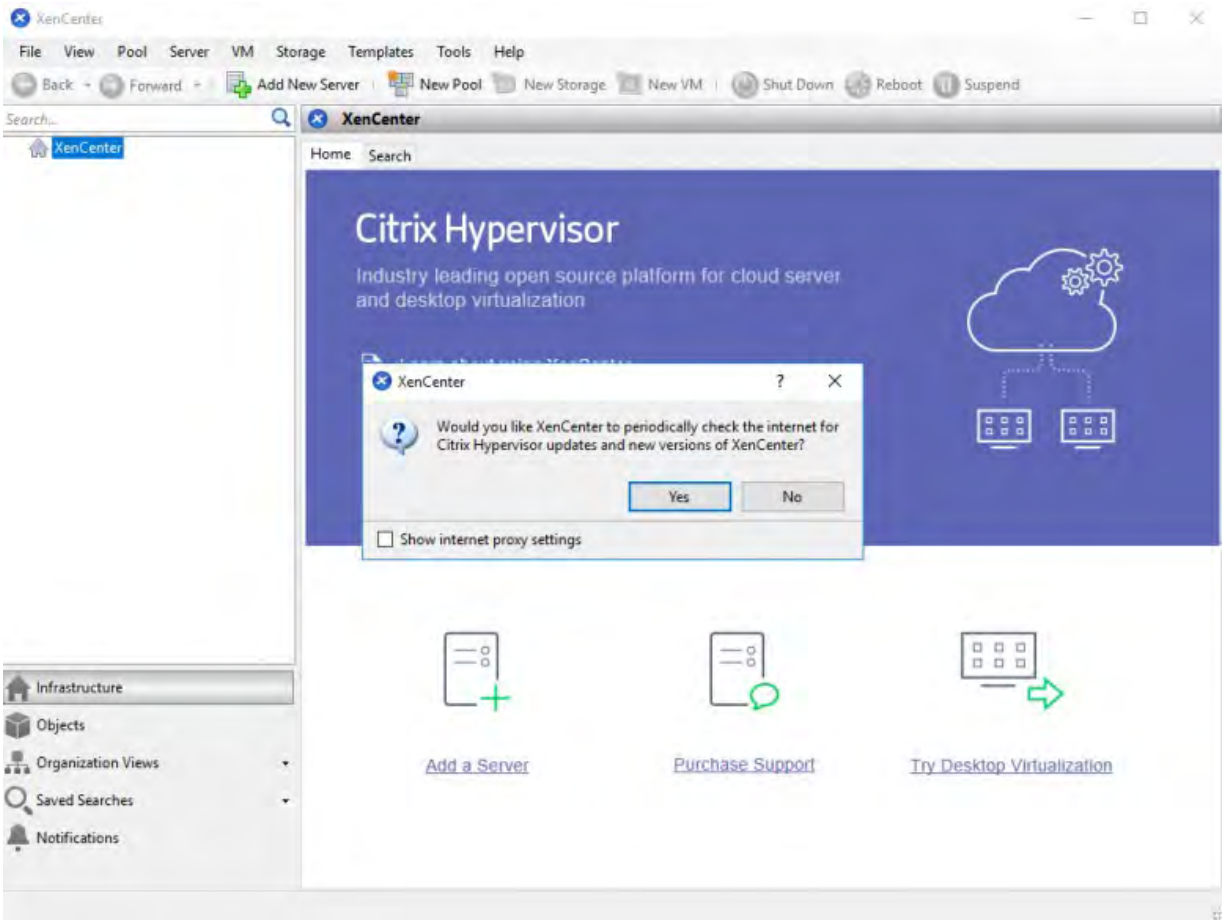

2. Click the **Add New Server** icon to open the Add New Server dialog box. In the Server field, enter the IP address of the Citrix Hypervisor server that was previously recorded in section 3.1. Enter the root username and password that was set during Citrix Hypervisor installation. Select **Add** to continue.

Note: An SSL security certificate may pop up. Click **accept** to continue.

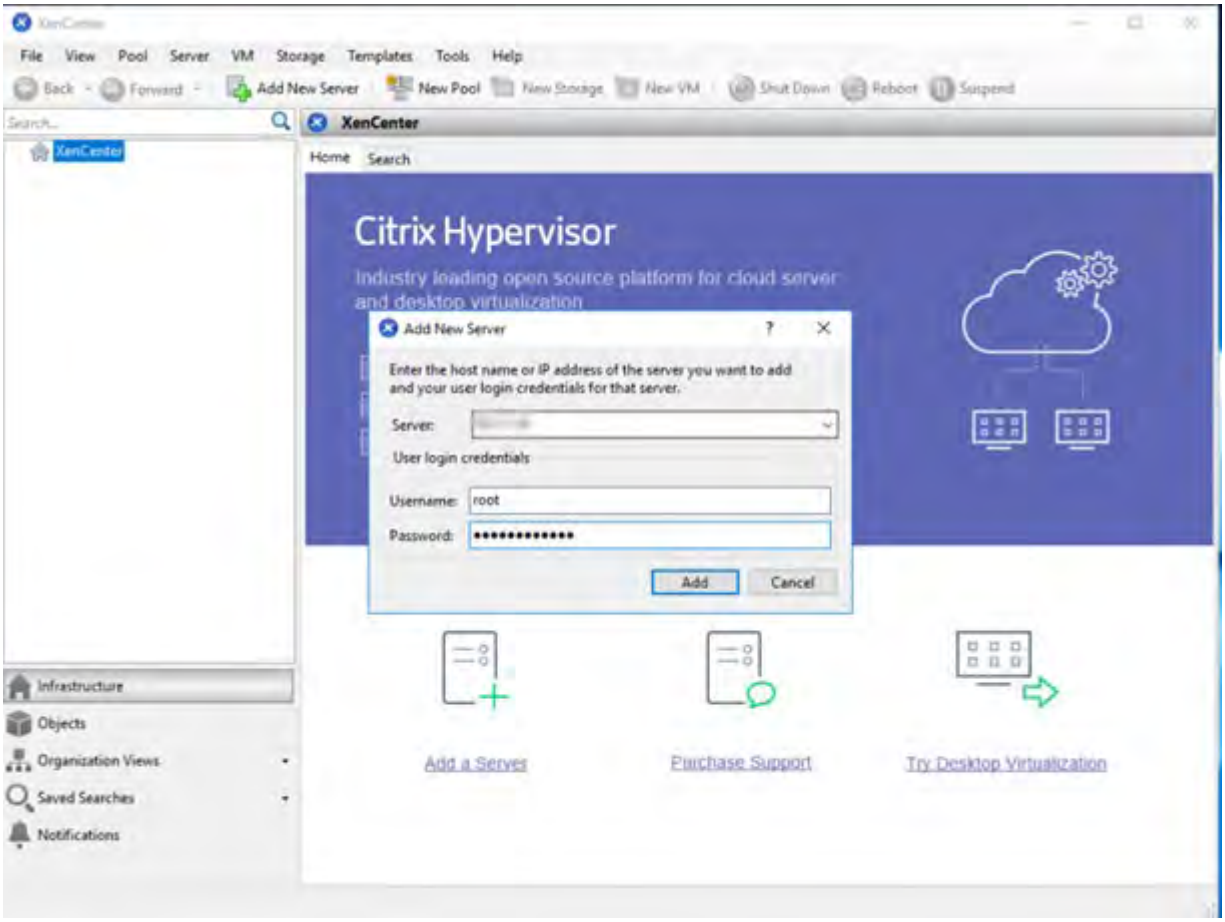

3. A Save and Restore Connection State dialog box will appear. Click **Save and restore connection state on startup** when XenCenter is launched if appropriate. Select **OK** to continue.
4. The **Health Check Overview** window will display. Enroll for Health Checks if appropriate.

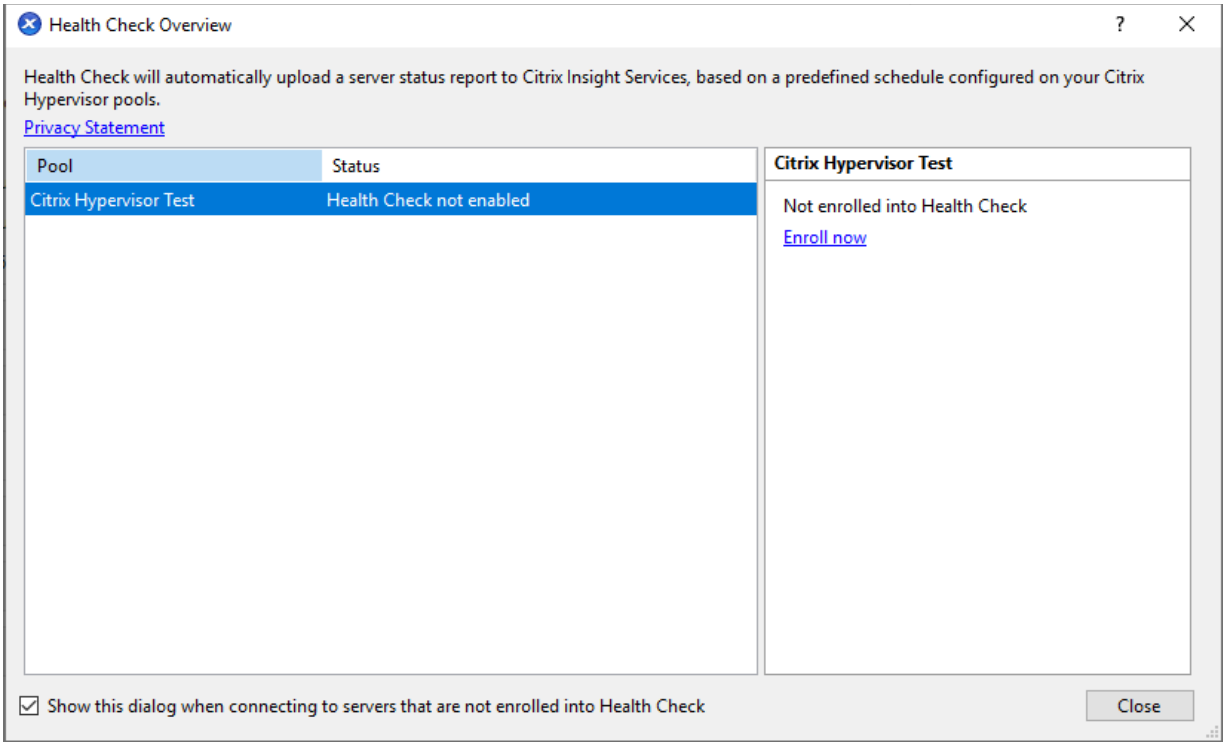

3.2.2 Adding Licenses to Your XenCenter

1. Within the XenCenter **Tools** pulldown, select **License Manager**.

> Note: You can use Citrix Hypervisor without a license (Free Edition). However, this edition will provide a restricted set of features.

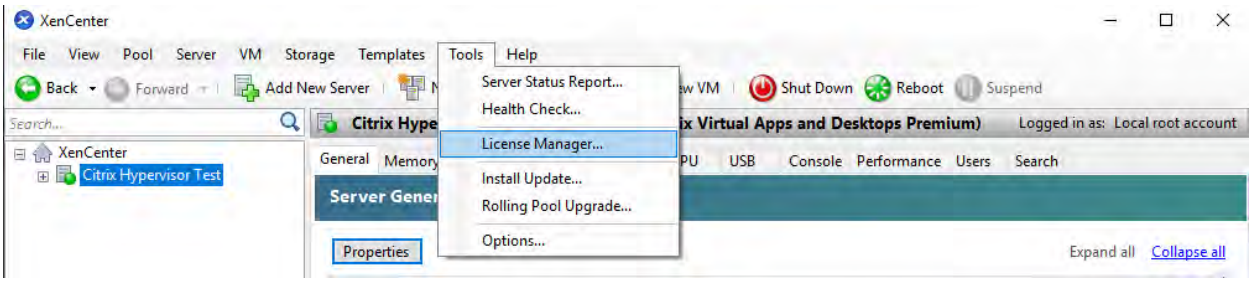

2. The License Manager window will display. Select **Assign License.**

> Note: For further information on Citrix Licensing see the Licensing Product Documentation.

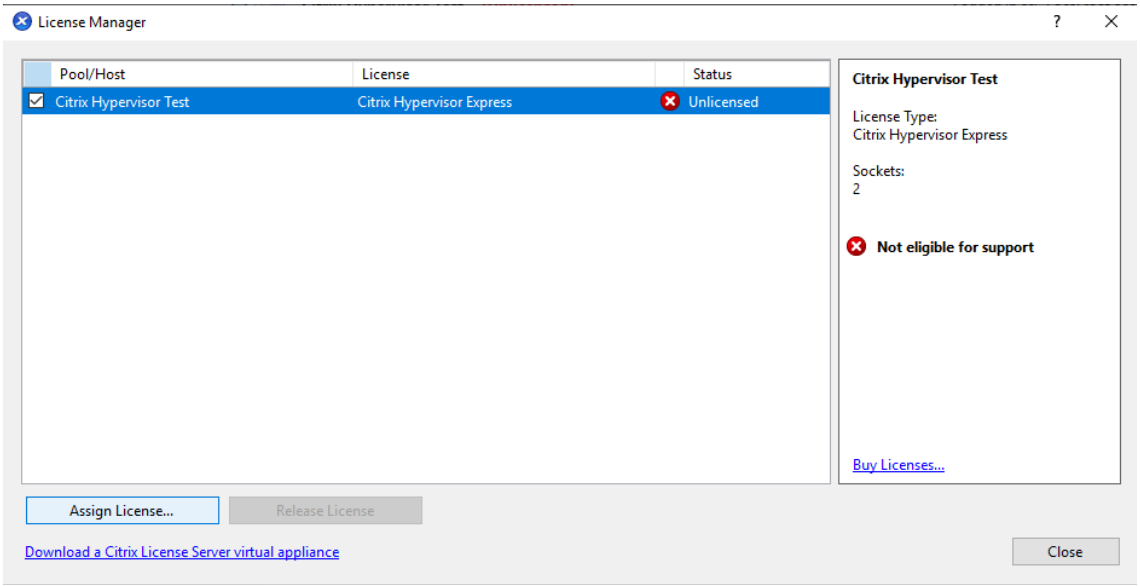

3. Enter the IP Address and appropriate License version in accordance with the License Server. Select **Ok** to continue.

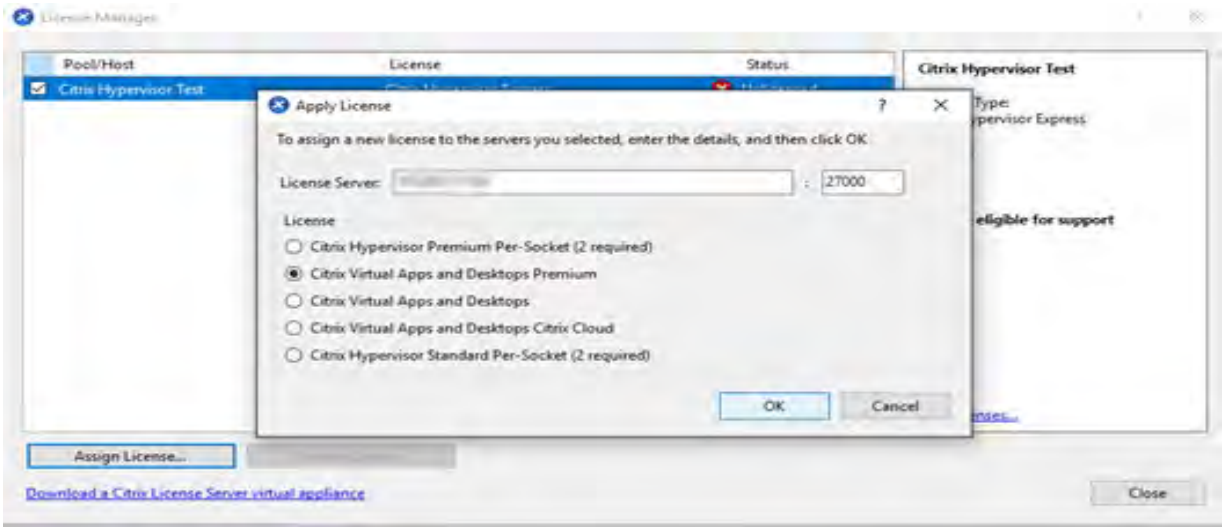

4. If configured correctly, the License Manager will update the status as Licensed within the Status column. Click close.

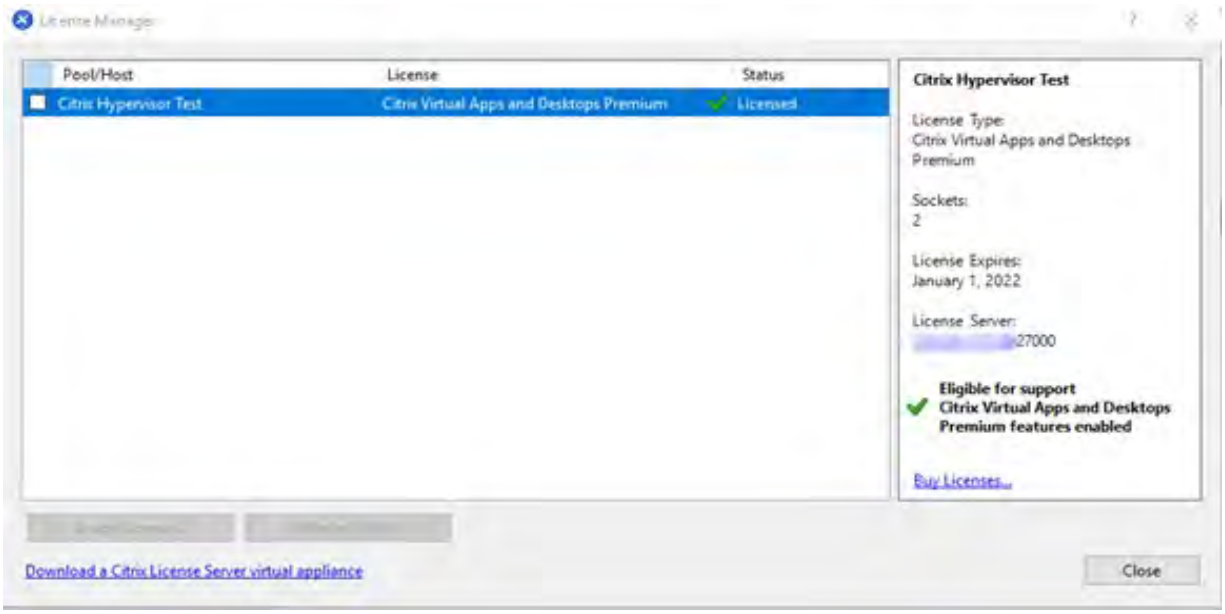
License Manager
Pool/Host
License
Status
Citrix Hypervisor Test
Citrix Virtual Apps and Desktops Premium
Licensed
Citrix Hypervisor Test
License Type:
Citrix Virtual Apps and Desktops Premium
Sockets:
2
License Expires:
January 1, 2022
License Server:
27000
Eligible for support
Citrix Virtual Apps and Desktops Premium features enabled
Buy Licenses...
Download a Citrix License Server virtual appliance
Close

Chapter 4. Building Citrix Virtual Apps & Desktops

This chapter covers installation and configuration of the core components for a Citrix Virtual Apps & Desktops environment, including:

- Citrix Delivery Controller, Citrix Studio, Citrix License Server

Note: For the latest install and configure information refer to Citrix's Install Core Components Section of the Citrix Virtual Apps and Desktops Product Documentation. Additionally, Citrix offers solutions for deploying Delivery Controllers to the Cloud. Citrix Cloud is outside the scope of this document, so please consult Citrix and refer to the Citrix Cloud Product Documentation if you choose to have any components of your Citrix infrastructure in the Cloud.

Note: A VM with a Windows Server OS was pre-created. This VM will host the Delivery Controller, Citrix Studio and License Server. The following sections describe the installation and configuration of these core components.

4.1 Installing the Citrix Delivery Controller

The Citrix Delivery Controller Server must meet the requirements listed in section 1.5.3 Citrix Virtual Machine Requirements.

Use the following procedure to install Citrix Delivery Controller:

1. Attach the iso file to the server OS and open it via File Explorer.
2. Launch the ***Auto Select*** Application and accept the Windows User Account Control Popup.
3. Click **Start** for ***Virtual Apps and Desktops*** section.

4. Select **Delivery Controller** to launch the Citrix Virtual Apps and Desktops installer.

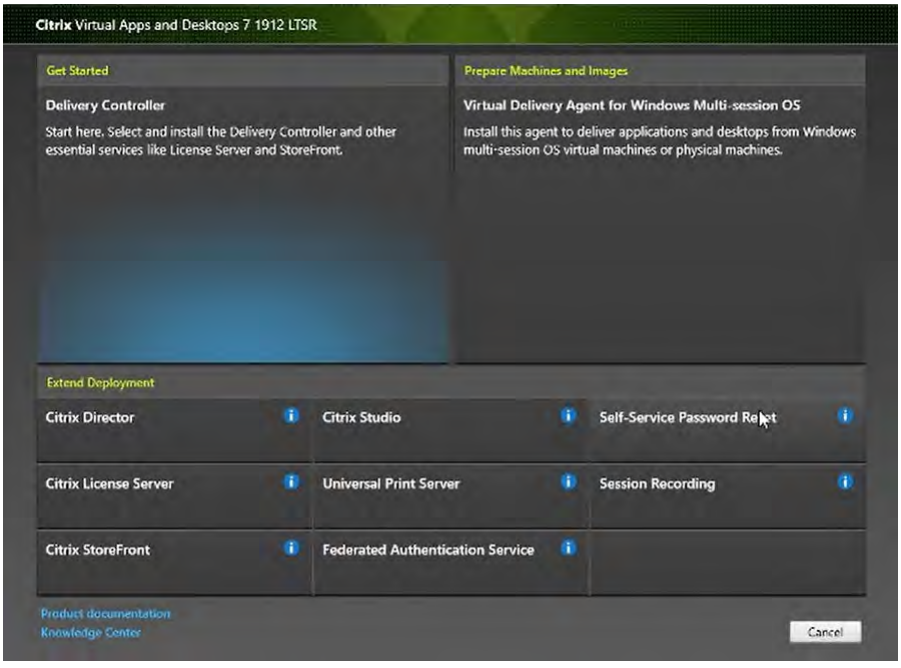

5. Scroll and read and through the Software License Agreement. Click the radio button, **I have read, understand, and accept the terms of the license agreement** to accept the agreement. Select **Next** to continue.

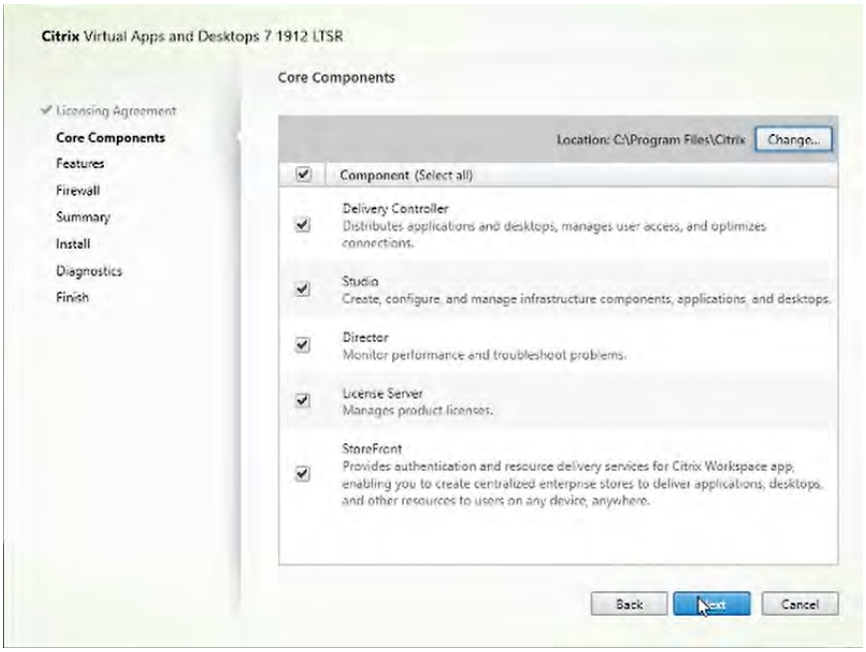

6. The Core Components window allows you to choose an install location as well as the core components which can be installed. For the purposes of POC/trial, ensure all components are selected and click **Next**.

Note: In a production environment, only Delivery Controller and Studio should be checked. Director, License Server, & StoreFront, should all reside on their own isolated servers. See the Citrix Virtual Apps & Desktops Install & Configure Product Documentation for production deployment instructions.

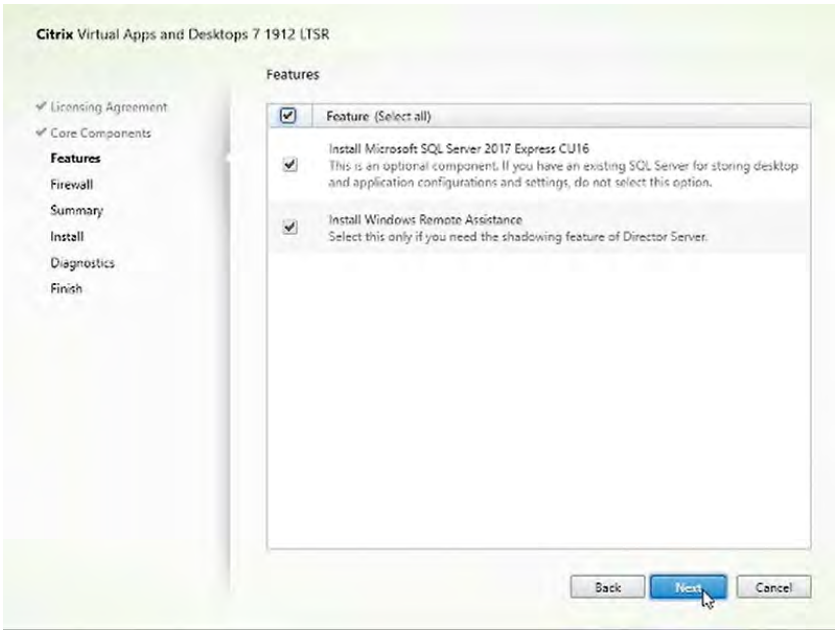

7. The Features window allows you to choose which Features to install. Ensure all features are selected and click **Next**.

8. The Firewall window allows you to configure Windows Firewall. Select the **Automatically** radio button and click **Next**.

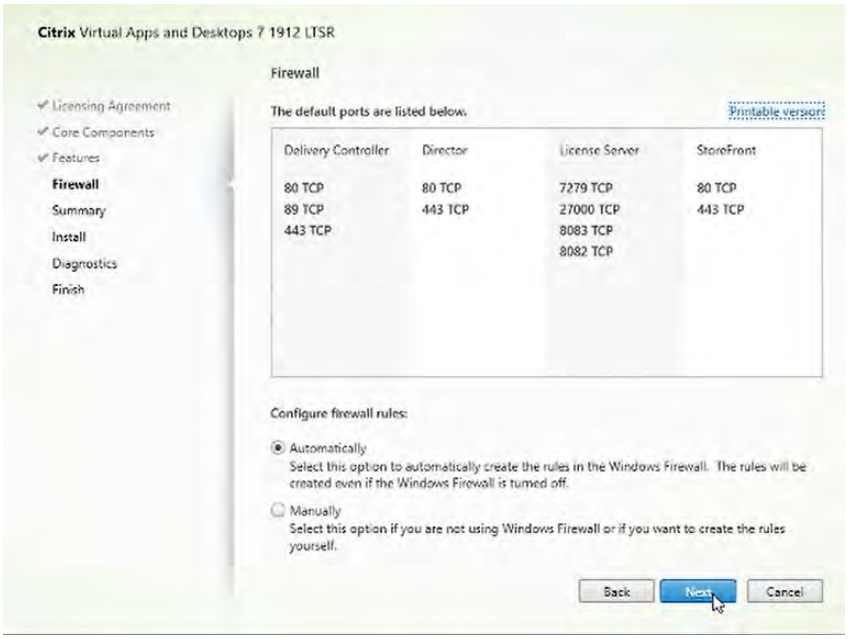

9. On the Summary window select **Install**.

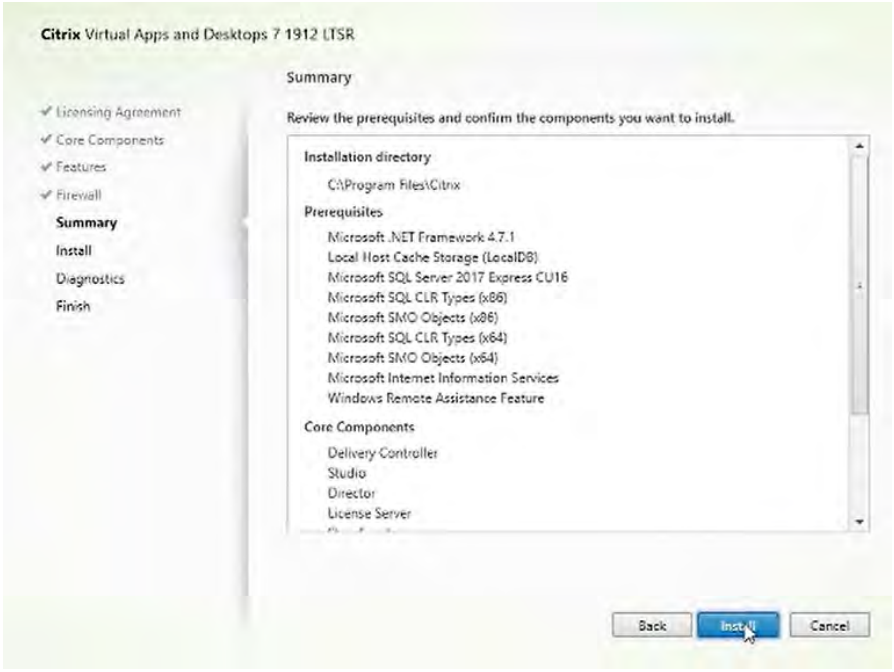

Accept any ***Reboot Prompts*** and reconnect to the server.

10. In the Diagnostics window, select the appropriate option to ***Collect diagnostic information*** **according to your organization policies**.

11. Select **Next** to continue.
12. Select **Finish** to complete the install.

4.2 Configuring the Citrix Delivery Controller

The following steps use Citrix Studio to configure the Citrix Delivery Controller.

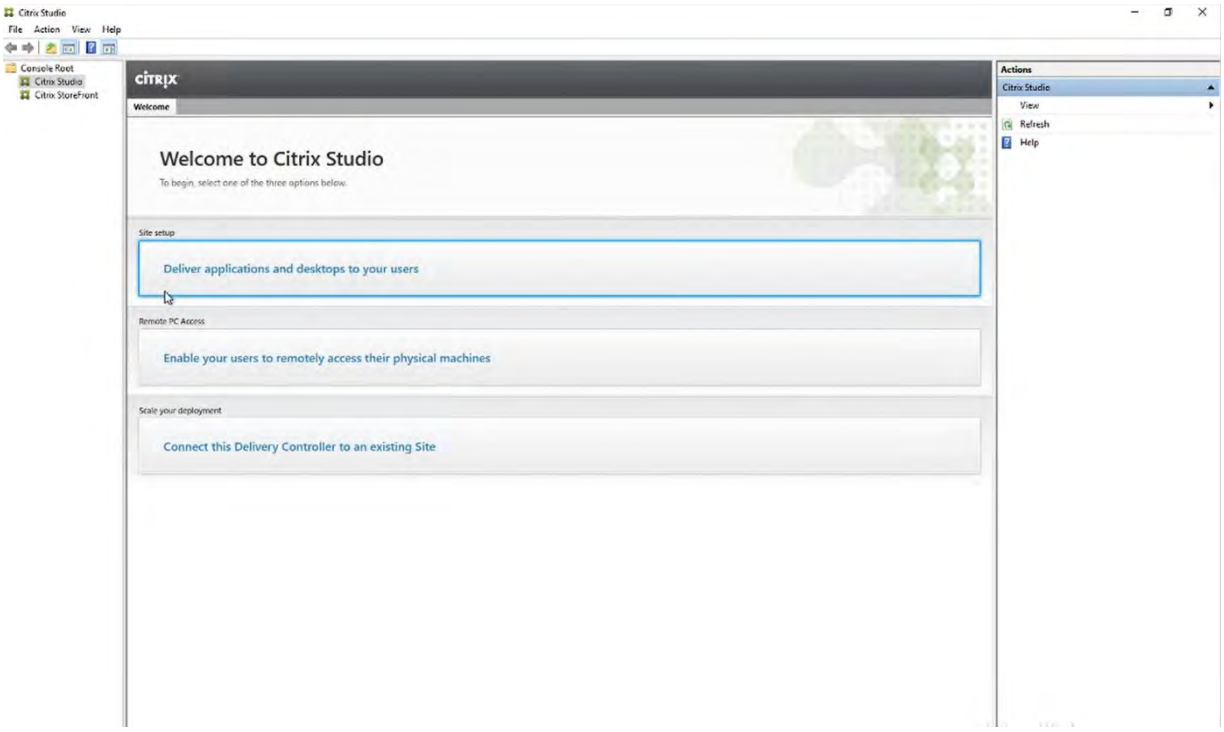

1. If Citrix Studio was not launched automatically after the install, launch **Citrix Studio** from the Windows Start Menu.
2. Select **Deliver applications and desktops to your users**.
3. On the Introduction window, ensure the "**A fully configured, production-ready Site (recommended for new users)**" radio button is selected.

4. In the **Site name:** type in a Site name.

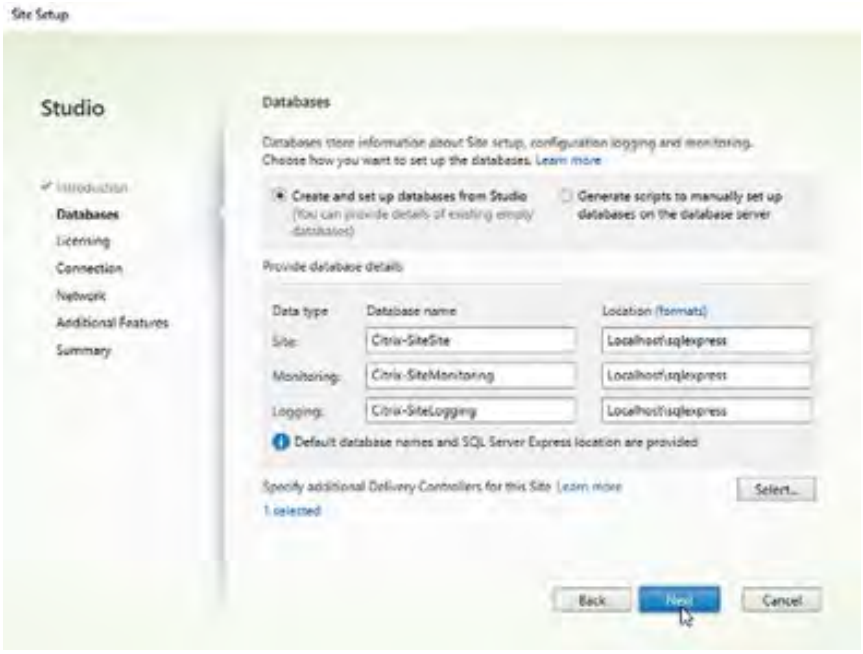

5. On the Licensing window in the "**License server address:**" ensure **localhost:27000** is specified.

> Note: In a production environment, you should provide the FQDN of your separate License Server, & correct port number.

6. Next, ensure the "**Use the free 30-day trial**" radio button is selected and click **Next**.

> Note: In a production environment, you should use the existing license from the FQDN License Sever.

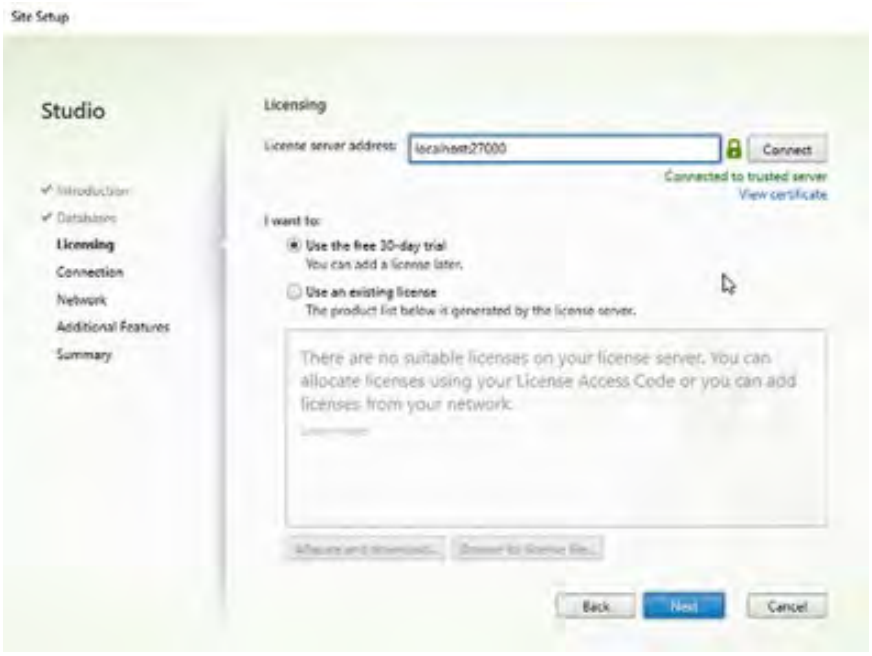

7. On the Connection Window perform the following tasks.
   a) Under the **Connection type:** drop down menu, select **Citrix Hypervisor.**
   b) In the **Connection address:** enter the URL for your Citrix Hypervisor.
   c) In the **Username:** enter your Citrix Hypervisor administrator's username.
   d) In the **Password:** enter the Citrix Hypervisor administrator's password.
   e) In the **Connection name:** enter a name for this connection.

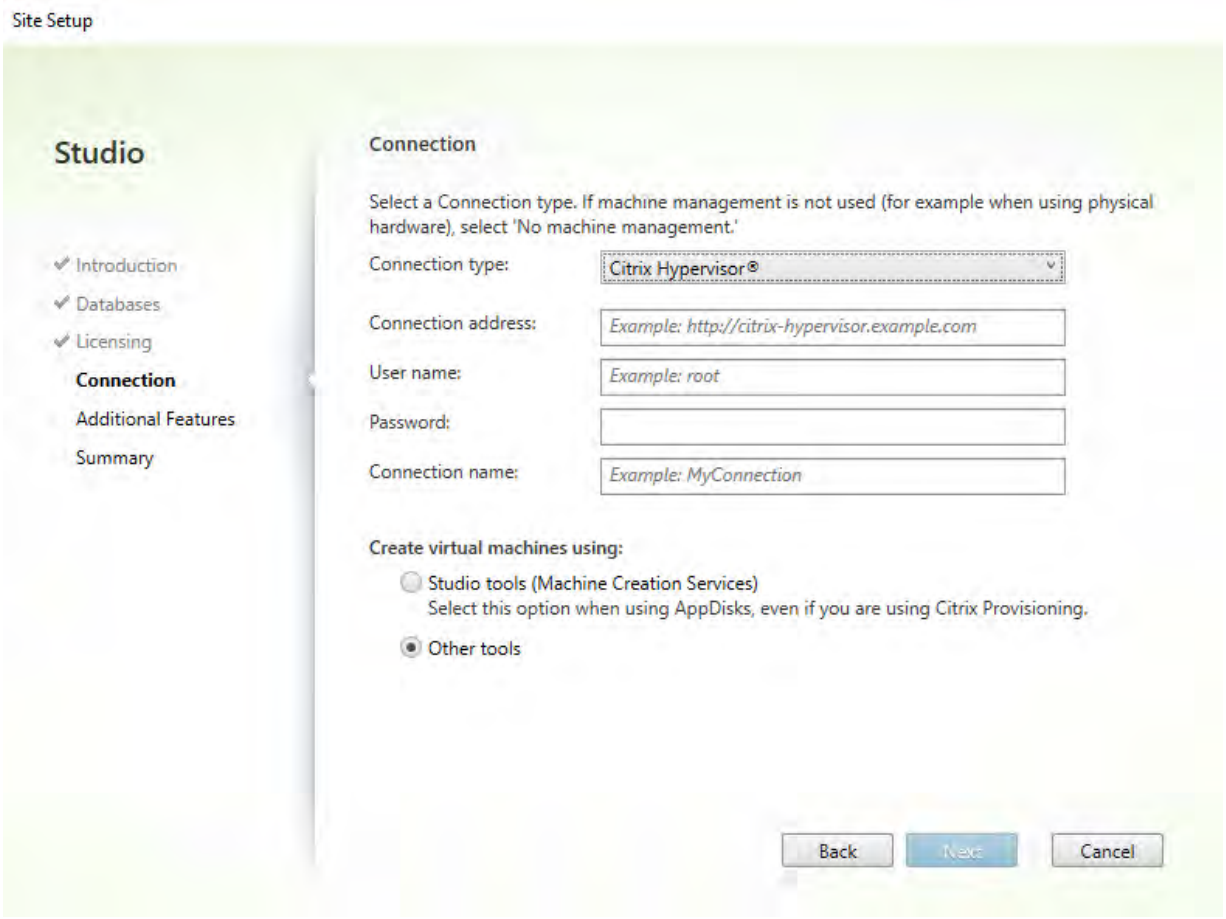

8. Ensure the **Other tools** radio button is selected and click **Next**.

9. On the Additional Features window, click **Next**.

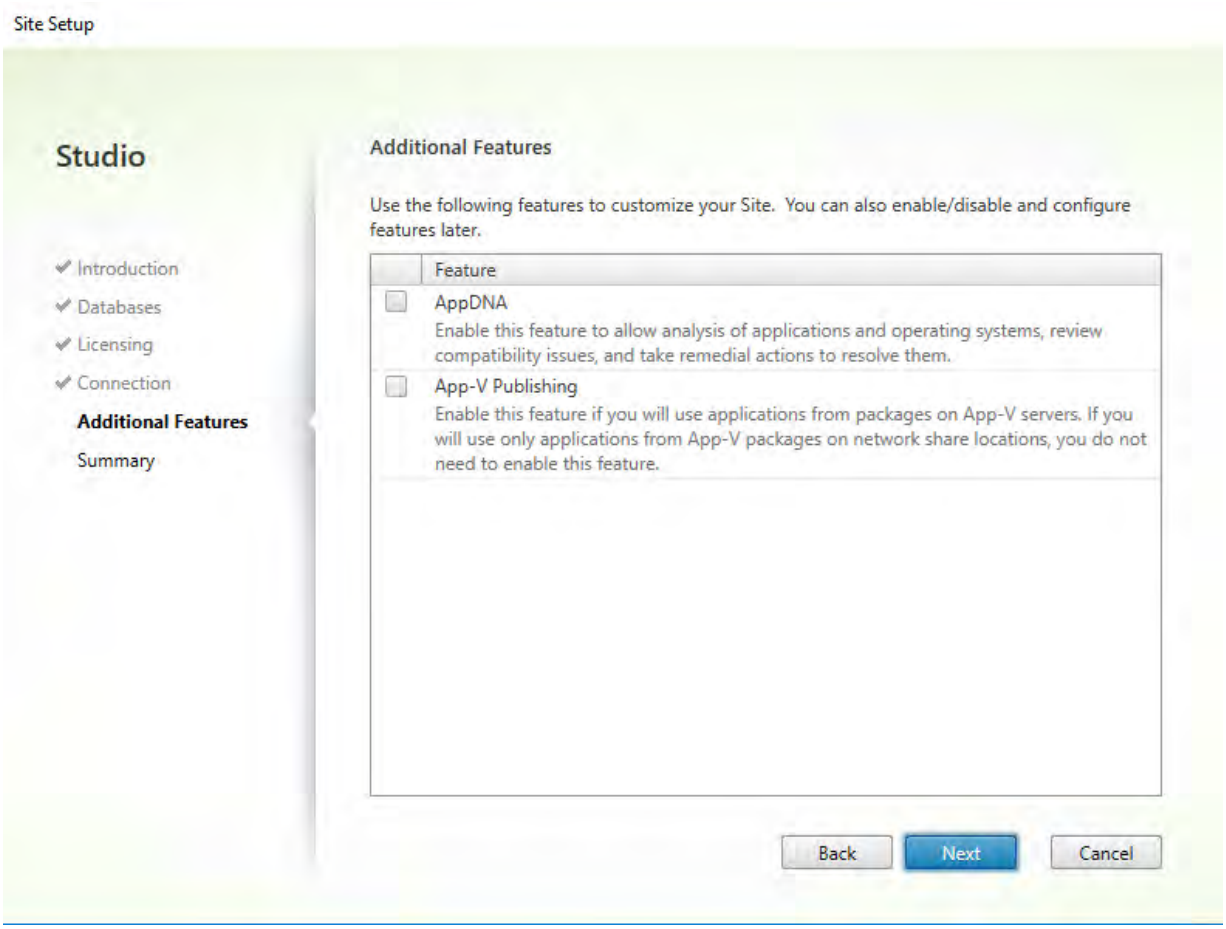

10. On the Summary window, click **Finish**.

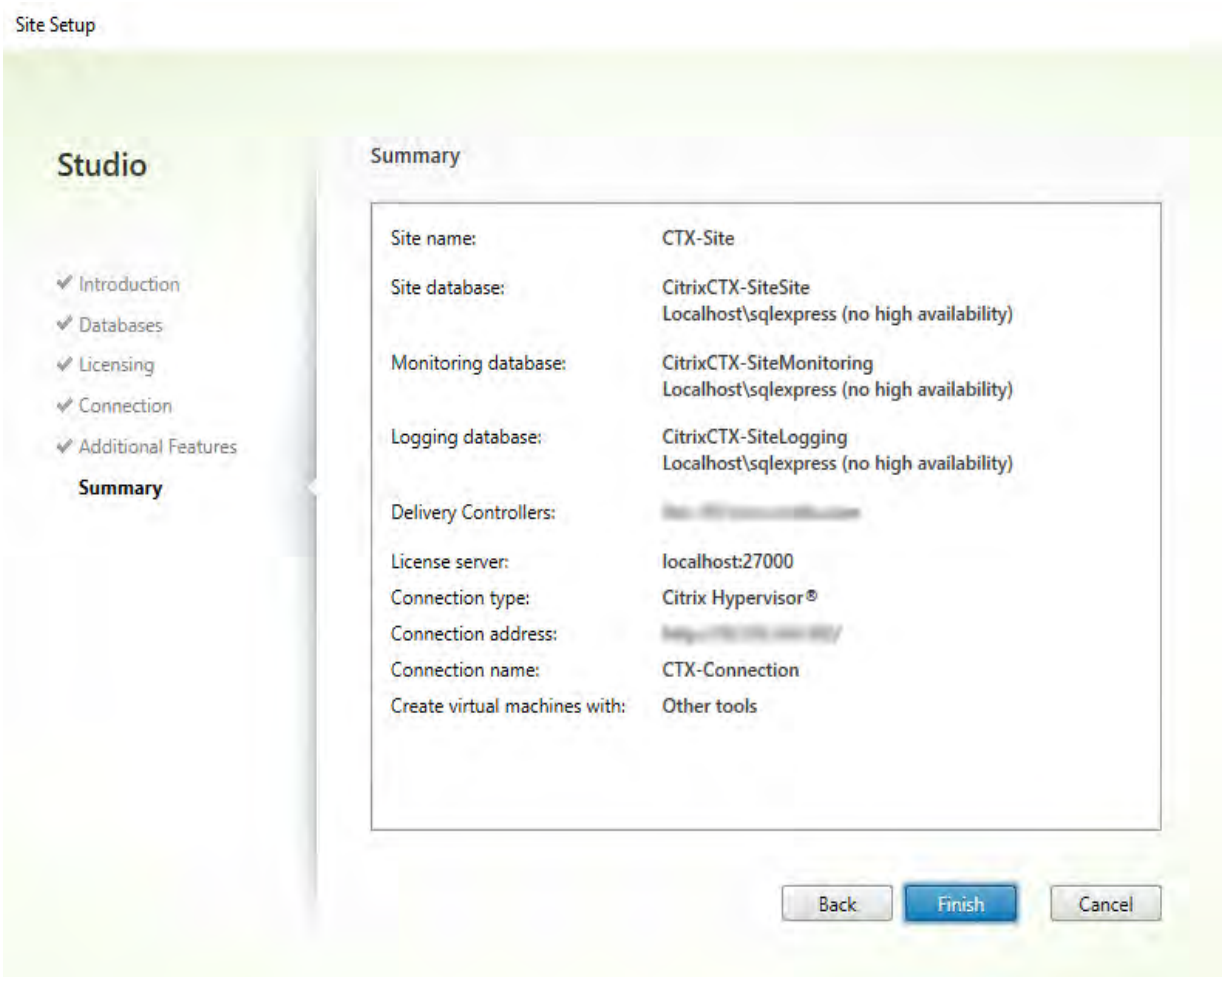
Site Setup
Studio
Introduction
Databases
Licensing
Connection
Additional Features
Summary
Summary
Site name:
CTX-Site
Site database:
CitrixCTX-SiteSite
Localhost\sqlexpress (no high availability)
Monitoring database:
CitrixCTX-SiteMonitoring
Localhost\sqlexpress (no high availability)
Logging database:
CitrixCTX-SiteLogging
Localhost\sqlexpress (no high availability)
Delivery Controllers:
License server:
localhost:27000
Connection type:
Citrix Hypervisor®
Connection address:
Connection name:
CTX-Connection
Create virtual machines with:
Other tools
Back
Finish
Cancel

Chapter 5. NVIDIA vGPU Manager Installation

This chapter covers the installation of the Nvidia vGPU Manager via two methods:

- Copying the .RPM file to Dom0 Server through a command line
- Installing the .ISO file through Citrix XenCenter

5.1.1 Copying the RPM file to a Citrix Hypervisor Host

Note: The update, .iso names, and versions are examples and may be different in your environment. As an example, the file will look similar to [Nvidia-vGPU-CitrixHypervisor-8.2-367.64x86_64.rpm].

1. For this guide, PuTTY will be used to copy files from a remote Windows system to the Citrix Hypervisor.
2. Use the following command to copy the RPM file to dom0

```
pscp somefile root@<serverip>:/tmp
```

3. Use the following rpm command to install the package:

```
rpm –iv NVIDIA-vGPU-xenserver-8.2-450.80.x86_64.rpm
```

4. Restart the Citrix Hypervisor server:

```
shutdown -r now
```

5. After you restart the Citrix Hypervisor server, verify that the vGPU package has been installed and loaded correctly by checking the NVIDIA kernel driver:

```
[root@xenserver ~]#lsmod |grep nvidia
nvidia 8152994 0
i2c_core 20294 2 nvidia, i2c_i801
```

6. Verify that the NVIDIA kernel driver can successfully communicate with the vGPU physical GPUs in your host. Run the nvidia-smi command to produce a listing of the GPUs in your platform similar to:

```
[root@xenserver:~] nvidia-smi
Wed Jan 13 19:48:05 2021
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 450.55       Driver Version: 450.55       CUDA Version: N/A      |
|-------------------------------+----------------------+----------------------+
```

```
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|                               |                      |               MIG M. |
|===============================+======================+======================|
|   0  Tesla T4            On   | 00000000:81:00.0 Off |                  Off |
| N/A   33C    P8    15W /  70W |     79MiB / 16383MiB |      0%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+
|   1  Tesla T4            On   | 00000000:C5:00.0 Off |                  Off |
| N/A   31C    P8    15W /  70W |     79MiB / 16383MiB |      0%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                                  |
|  GPU   GI   CI        PID   Type   Process name                  GPU Memory |
|        ID   ID                                                   Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
[root@xenserver:~]
```

5.1.2 Installing the NVIDIA vGPU manager from XenCenter

XenCenter can be used to install or update Supplemental Packs on the Citrix Hypervisor hosts. The NVIDIA Virtual GPU Manger supplemental pack is provided as a .iso file.

1. Navigate to **Tools** and select **Install Update.**

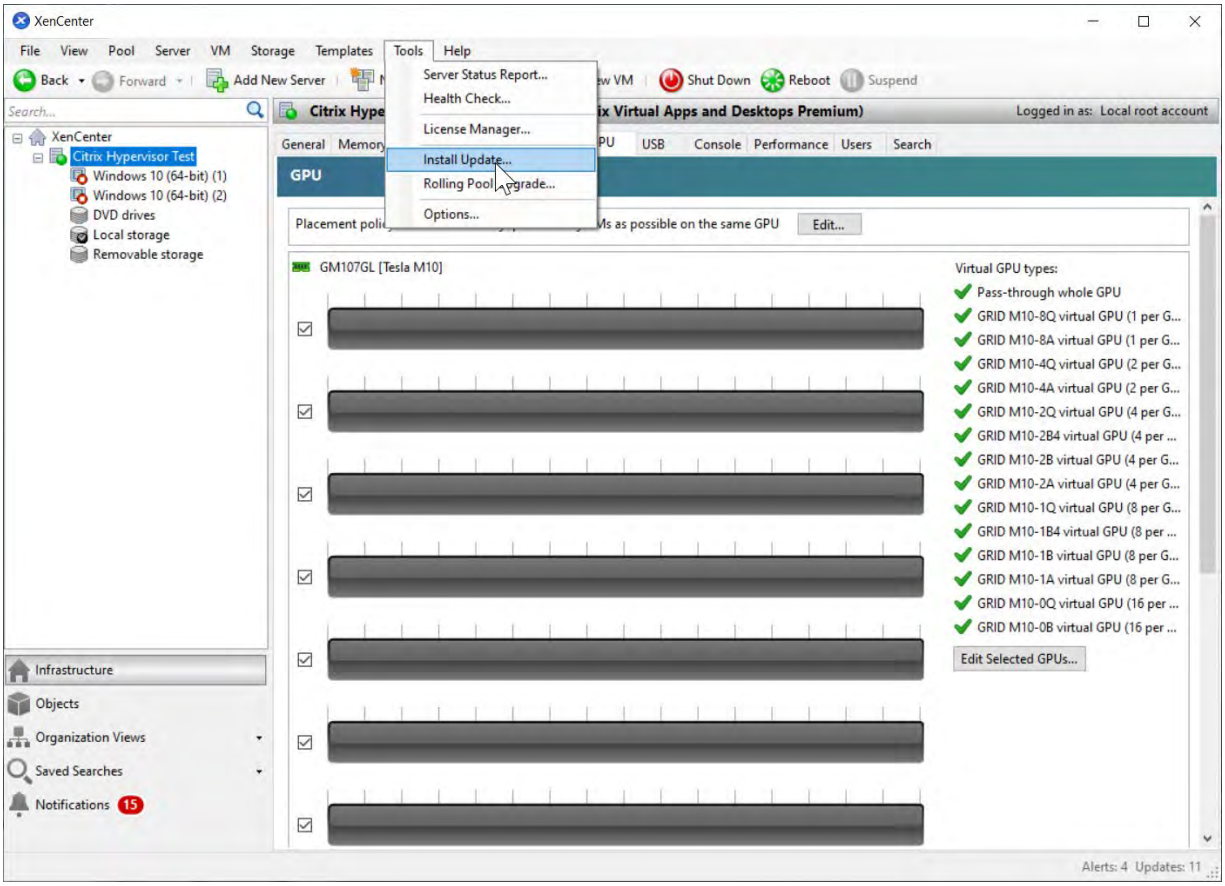

2. At the **Before You Start** window, read through the instructions, and select **Next** to continue.

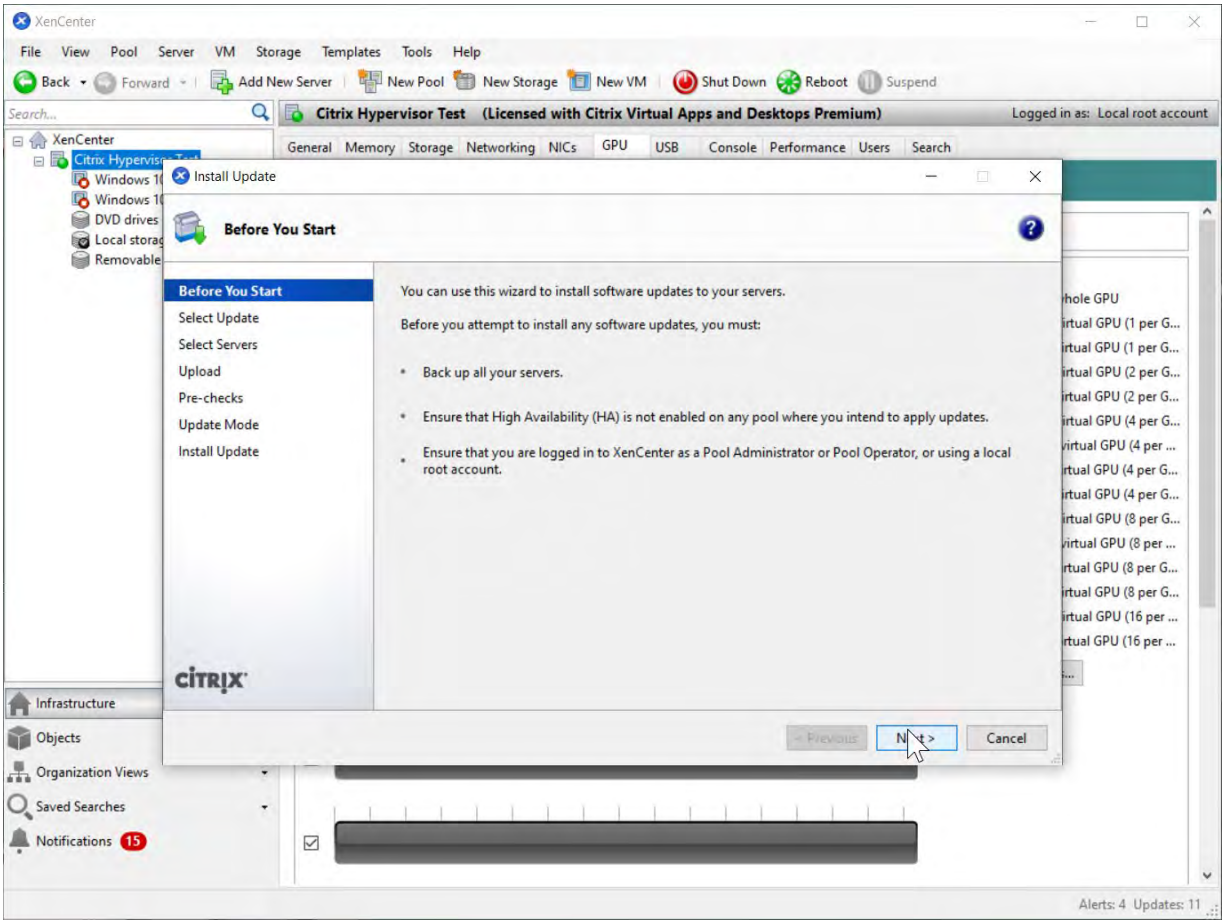

3. Select the bottom option for **Select update or supplemental pack from disk**. Proceed to the file path where the NVIDIA Citrix Hypervisor Supplemental Pack .iso is located. Select **Next** to continue.

Note: The update, .iso names, and versions are examples and may be different in your environment. As an example, the file will look similar to [NVIDIA-vGPU-xenserver-8.2-450.80.x86_64.iso]

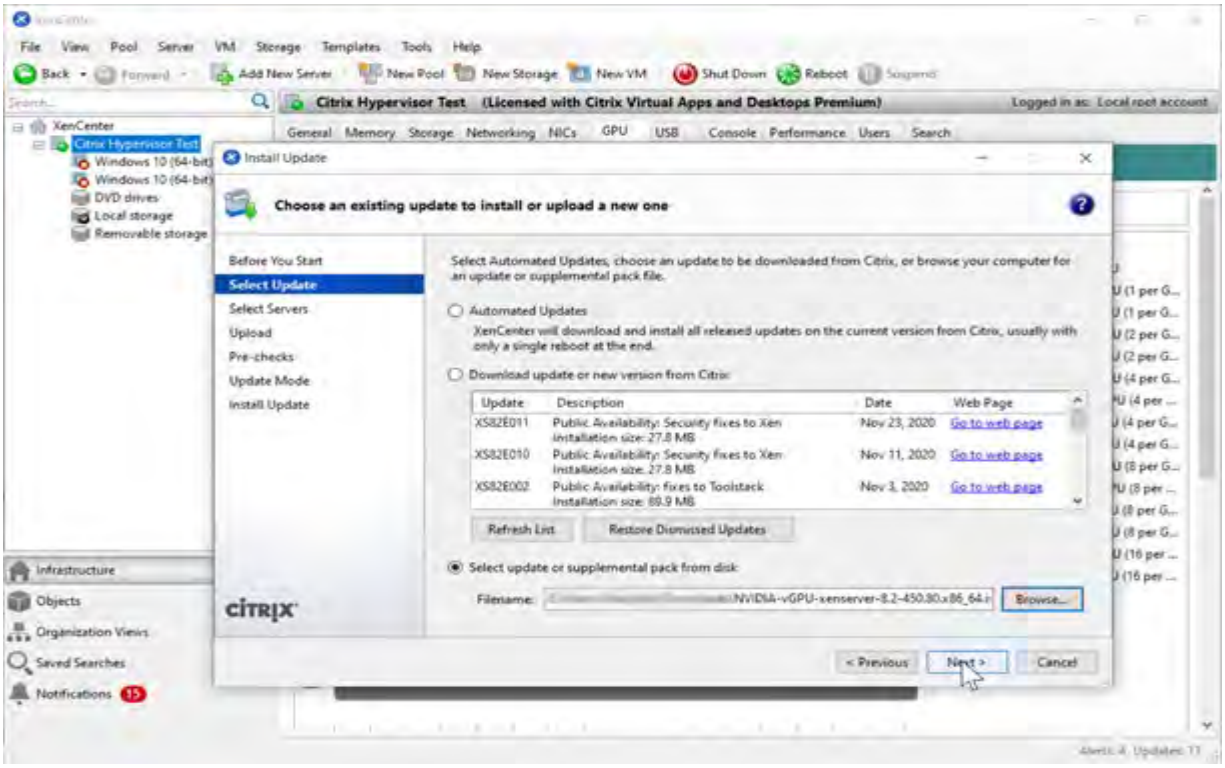

4. Select the servers on which to install the .iso file. Navigate to **Next** to continue.

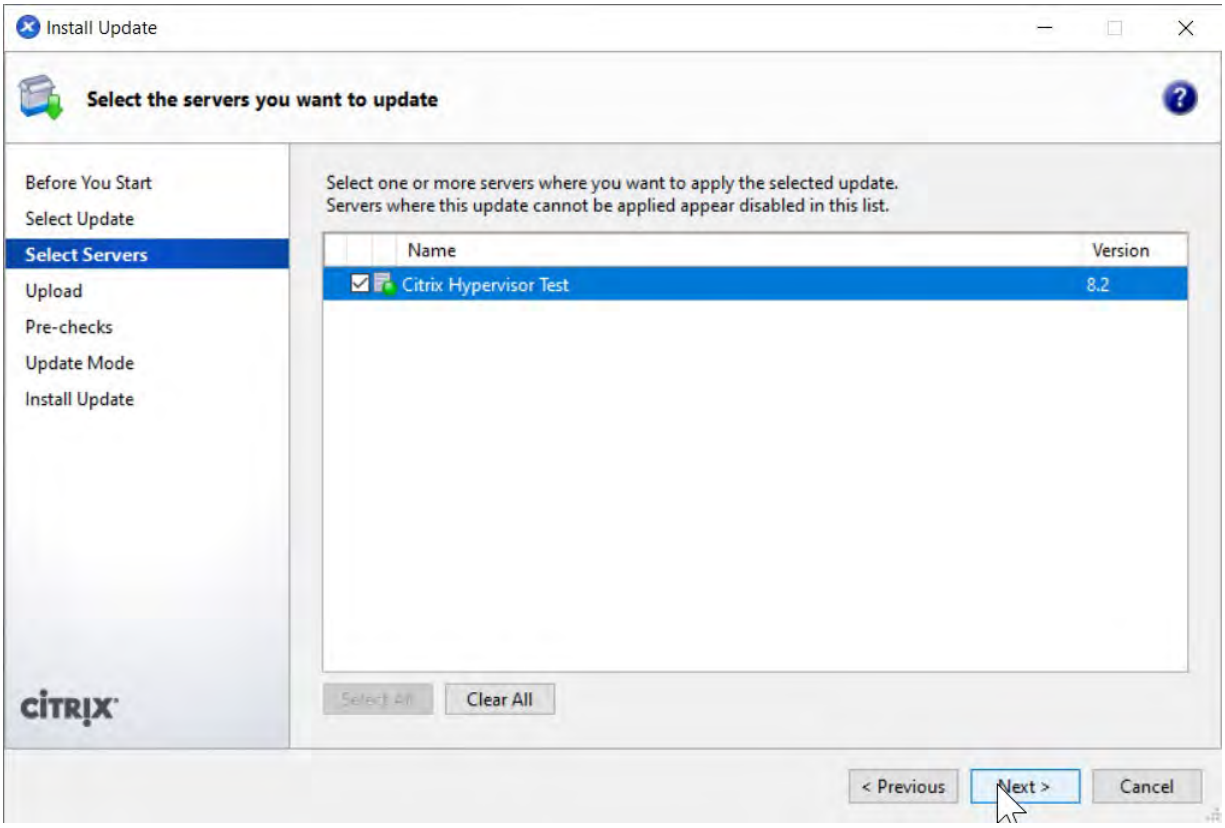

5. The file will be uploaded from the local machine disk to the Citrix Hypervisor's selected storage location. After successful completion, select **Next** to continue.

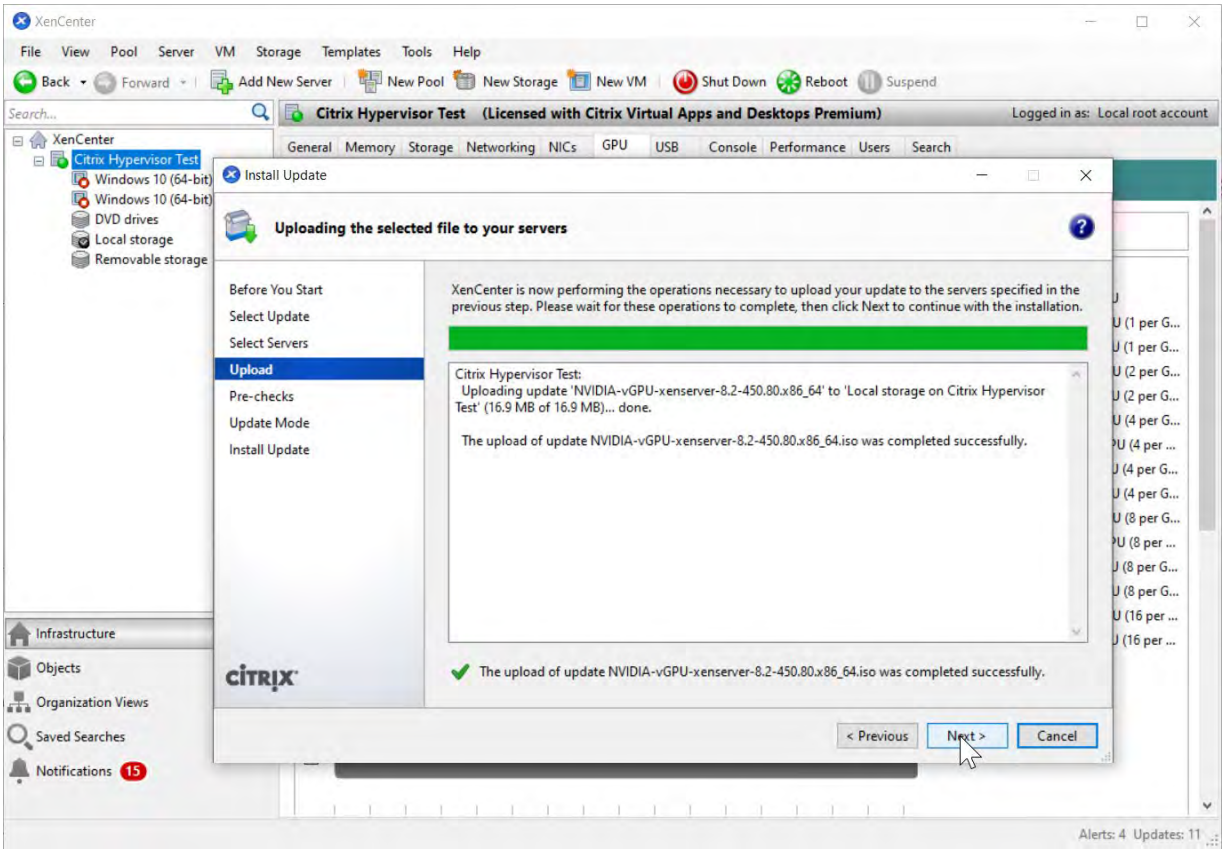

6. Installation prechecks will begin. Once all prechecks are resolved, select **Next** to continue.

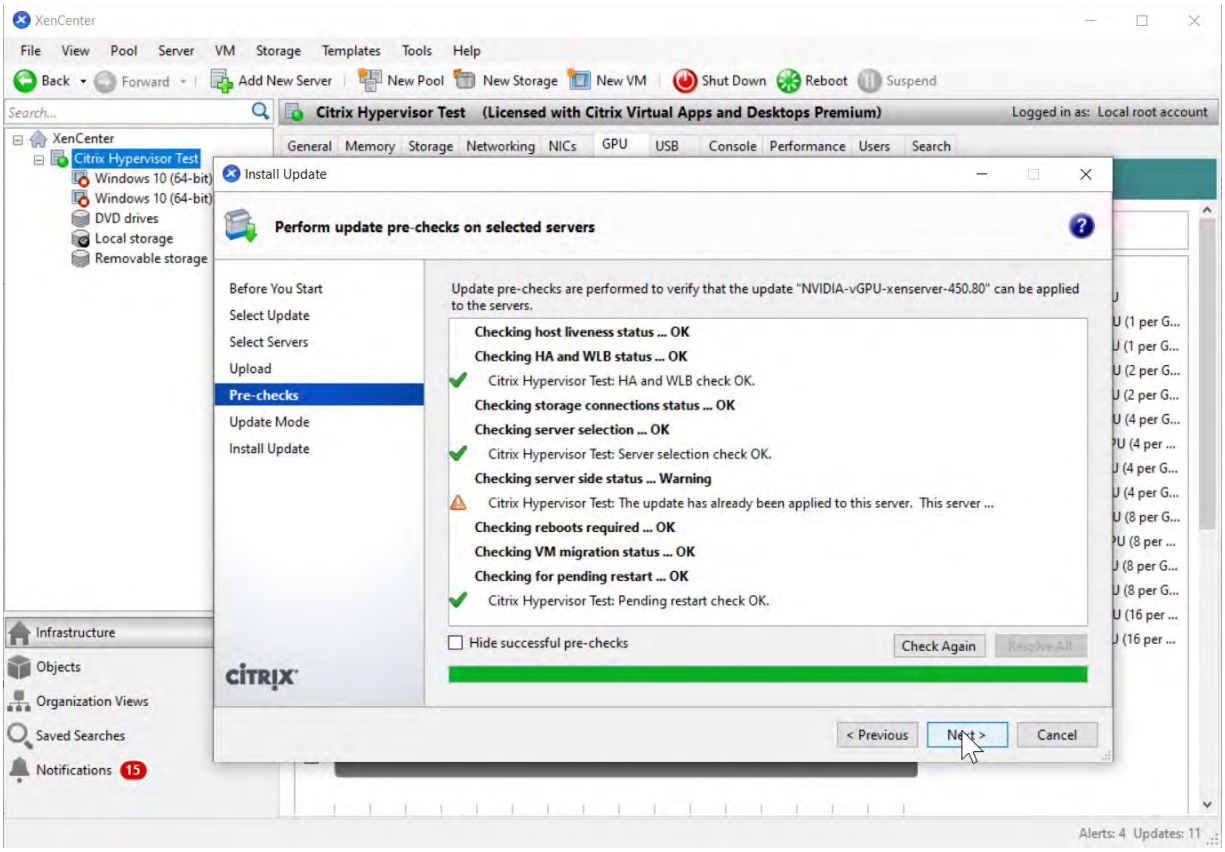

7. In the **Update Mode** window, select Allow XenCenter to carry out the post-update tasks as soon as the update has been applied. Then select **Install Update.**

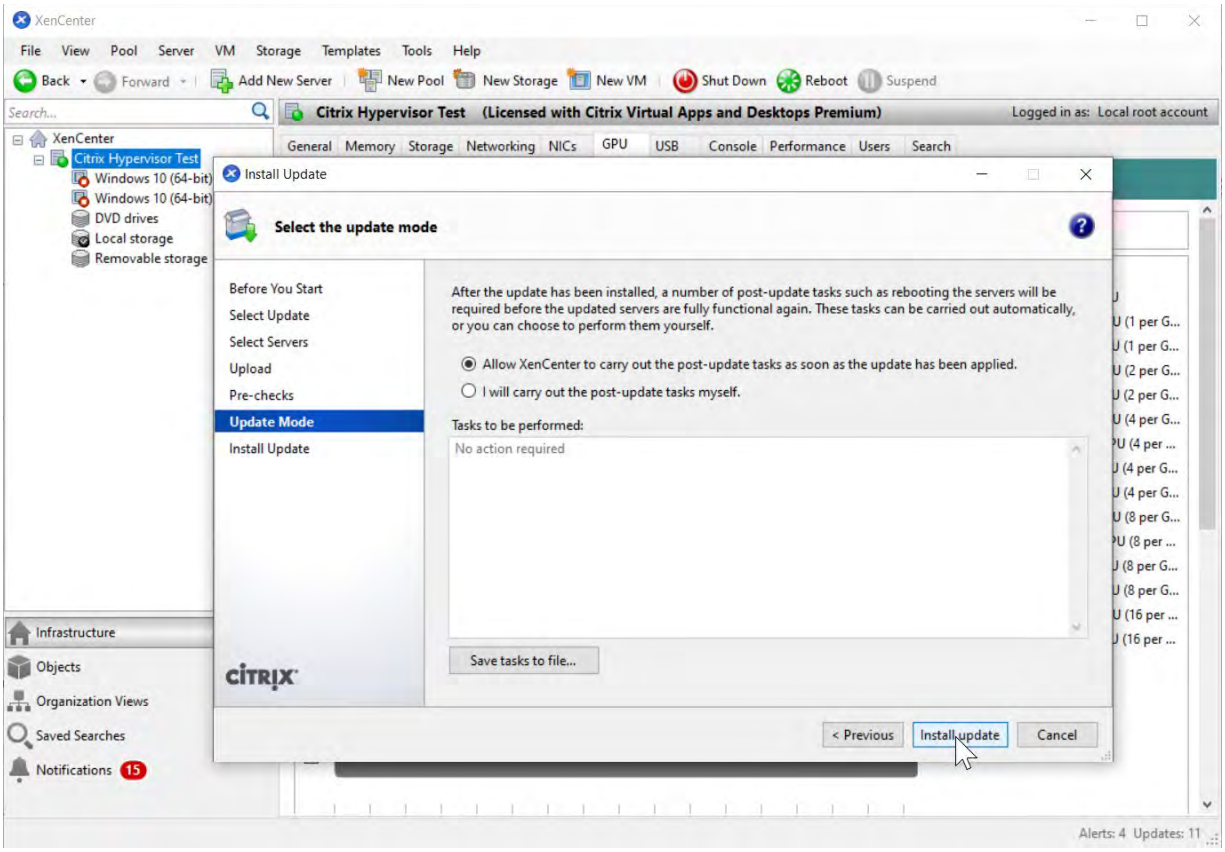

8. Once the installation is complete, select **Finish** to continue. Reboot the server to conclude installment of the NVIDIA vGPU Manager via .iso file.

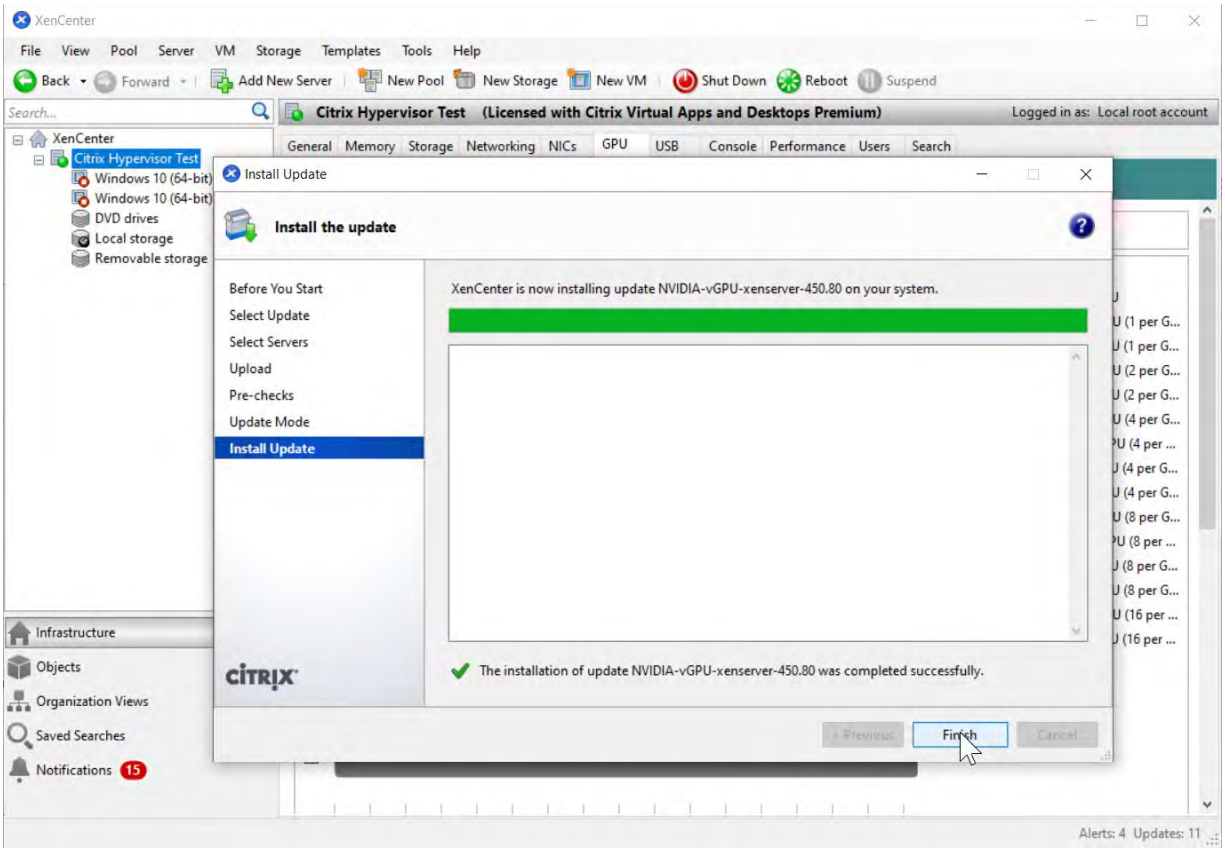
XenCenter
File View Pool Server VM Storage Templates Tools Help
Back Forward Add New Server New Pool New Storage New VM Shut Down Reboot Suspend
Search...
Citrix Hypervisor Test (Licensed with Citrix Virtual Apps and Desktops Premium)
Logged in as: Local root account
XenCenter
Citrix Hypervisor Test
Windows 10 (64-bit)
Windows 10 (64-bit)
DVD drives
Local storage
Removable storage
General Memory Storage Networking NICs GPU USB Console Performance Users Search
Infrastructure
Objects
Organization Views
Saved Searches
Notifications 15
Alerts: 4 Updates: 11
Install Update
Install the update
Before You Start
Select Update
Select Servers
Upload
Pre-checks
Update Mode
Install Update
XenCenter is now installing update NVIDIA-vGPU-xenserver-450.80 on your system.
The installation of update NVIDIA-vGPU-xenserver-450.80 was completed successfully.
CITRIX
Previous
Finish
Cancel

Chapter 6. Deploying the NVIDIA vGPU Software License Server

This chapter covers deployment of the NVIDIA vGPU software license server, including:

- Platform Requirements
- Installing the Java Runtime Environment on Windows
- Installing the License Server Software on Windows

6.1 Platform Requirements

Before proceeding, ensure that you have a platform suitable for hosting the license server.

6.1.1 Hardware and Software Requirements

- The hosting platform may be a physical machine, an on-premises virtual machine (VM), or a VM on a supported cloud service. NVIDIA recommends using a host that is dedicated solely to running the license server.
- The recommended minimum configuration is 2 CPU cores and 4 GB of RAM. A high-end configuration of 4 or more CPU cores with 16 GB of RAM is suitable for handling up to 150,000 licensed clients.
- At least 1 GB of hard drive space is required.
- The hosting platform must run a supported operating system.
- On Window platforms, .NET Framework 4.5 or later is required.

6.1.2 Platform Configuration Requirements

- The platform must have a fixed (unchanging) IP address. The IP address may be assigned dynamically by DHCP or statically configured but must be constant.
- The platform must have at least one unchanging Ethernet MAC address, to be used as a unique identifier when registering the server and generating licenses in the NVIDIA Licensing Portal.
- The platform's date and time must be set accurately. NTP is recommended.

6.1.3 Network Ports and Management Interface

The license server requires TCP port 7070 to be open in the platform's firewall, to serve licenses to clients. By default, the installer will automatically open this port. The license server's management interface is web-based and uses TCP port 8080. The management interface itself does not implement access control; instead, the installer does not open port 8080 by default, so that the management interface is only available to web browsers running locally on the license server host. Access to the management interface is therefore controlled by limiting remote access (via VNC, RDP, etc.) to the license server platform.

Note: If you choose to open port 8080 during license server installation, or at any time afterwards, the license server's management interface is unprotected.

6.2 Installing the NVIDIA vGPU Software License Server on Windows

The license server requires a Java runtime environment, which must be installed separately before you install the license server.

6.2.1 Installing the Java Runtime Environment on Windows

If a suitable Java runtime environment (JRE) version is not already installed on your system install a supported JRE before running the NVIDIA license server installer.

1. Download a supported 64-bit Oracle Java SE JRE or OpenJDK JRE.
   a) Download Oracle Java SE JRE from the Java Downloads for All Operating Systems page.
      i. Download Oracle Java SE JRE from the java.com: Java + You page.
   b) Download OpenJDK JRE from the Community builds using source code from OpenJDK project on GitHub.
2. Install the JRE that you downloaded.
   a) Oracle Java SE JRE installation:

b) OpenJDK JRE installation:

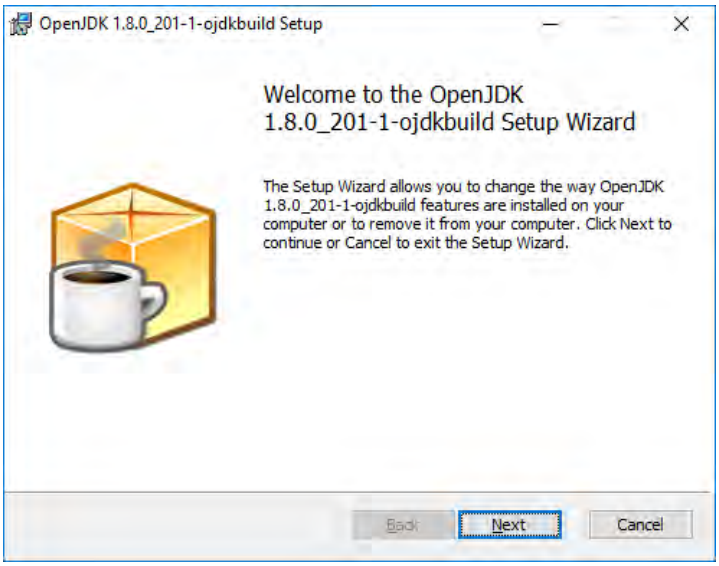

3. Set the JAVA_HOME system variable to the full path to the jre folder of your JRE installation.
   a) **For 64-bit Oracle Java SE JRE**: C:\Program Files\Java\jre1.8.0_191
   b) **For 64-bit OpenJDK JRE**: C:\Program Files\ojdkbuild\java-1.8.0-openjdk-1.8.0.201-1\jre

Ensure that the path does not include any trailing characters, such as a slash or a space.

If you are upgrading to a new version of the JRE, update the value of the JAVA_HOME system variable to the full path to the jre folder of your new JRE version.

4. Ensure that the Path system variable contains the path to the java.exe executable file.
   a) **For 64-bit Oracle Java SE JRE**: C:\Program Files\Java\jre1.8.0_191\bin

b) **For 64-bit OpenJDK JRE**: C:\Program Files\ojdkbuild\java-1.8.0-openjdk-1.8.0.201-1\bin

6.2.2 Installing the License Server Software on Windows

1. Unzip the license server installer and run setup.exe.
2. Accept the EULA for the license server software and the Apache Tomcat software used to support the license server's management interface.

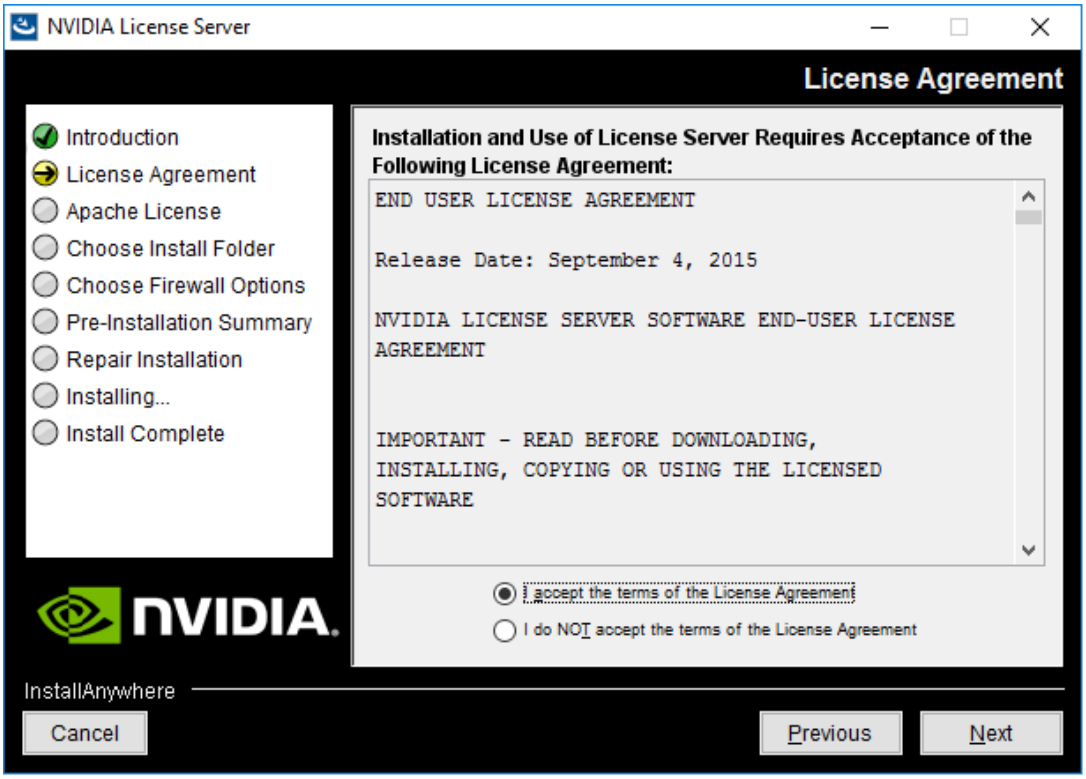

3. Choose the destination folder where you want the license server software to be installed.

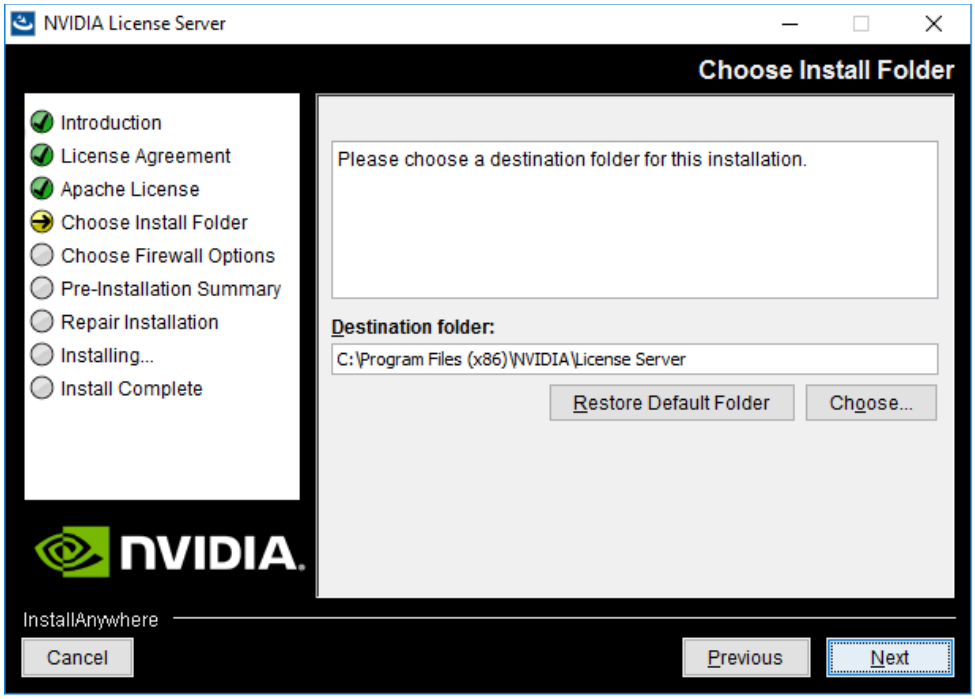

4. In the Choose Firewall Options dialog box, select the ports to be opened in the firewall.

   To enable remote clients to access licenses from the server and prevent remote access to the management interface, use the default setting, which sets ports as follows:

   a) Port 7070 is opened to enable remote clients to access licenses from the server.

   b) Port 8080 is closed to ensure that the management interface is available only through a web browser running locally on the license server host.

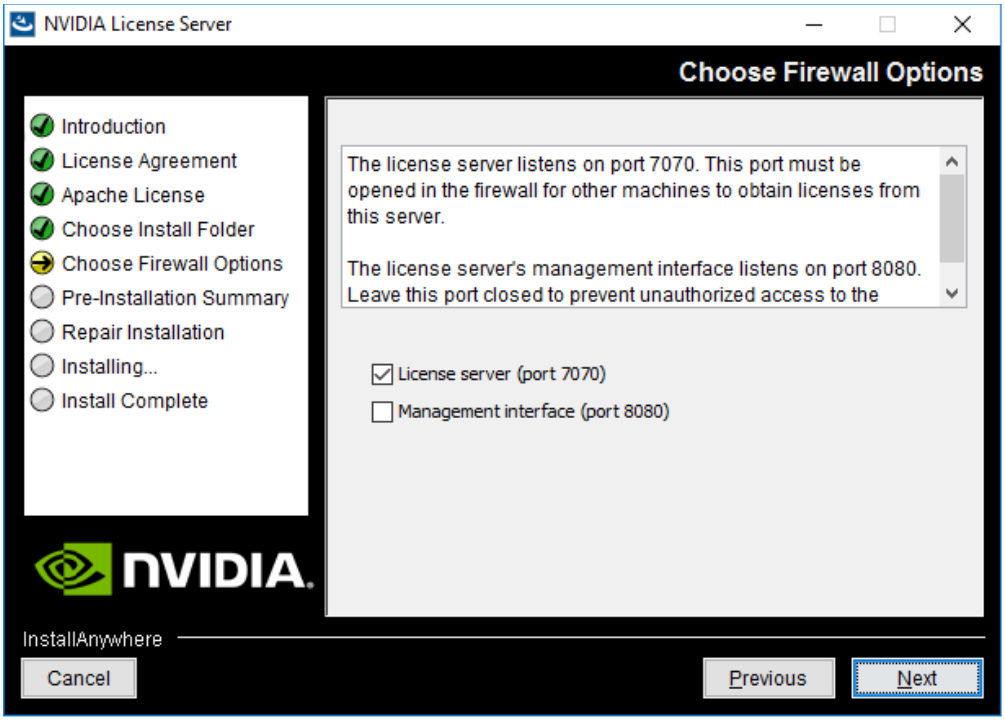

5. After installation has completed successfully, click Done to exit the installer.

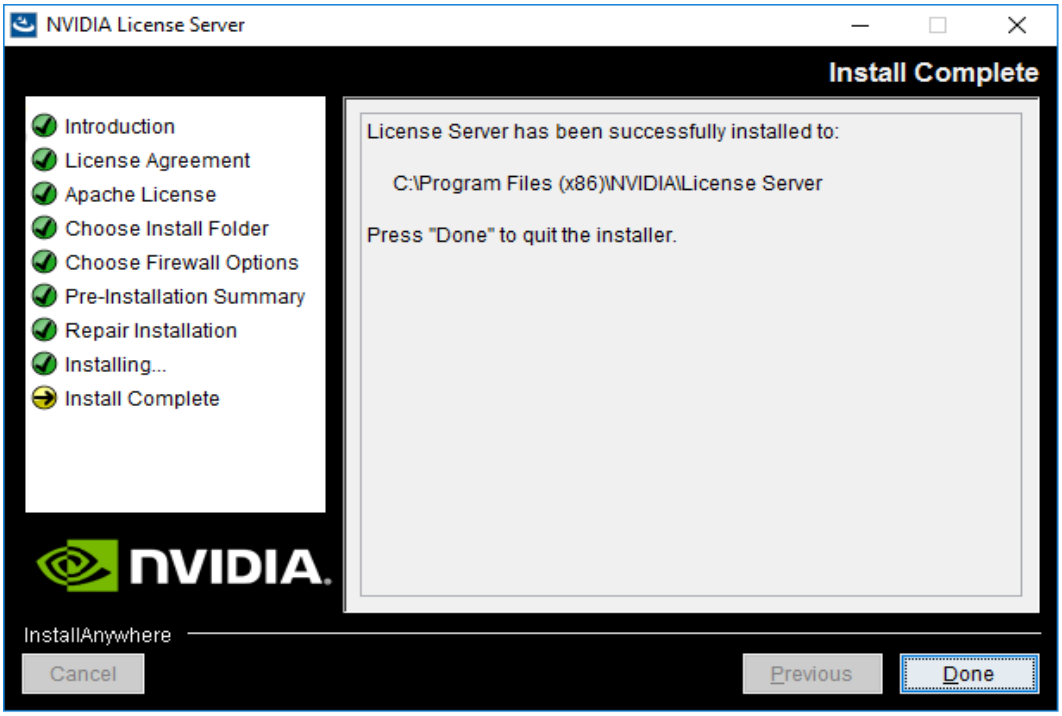

6.2.3 Obtaining the License Server's MAC Address

The license server's Ethernet MAC address uniquely identifies your server to the NVIDIA Licensing Portal. You will need this address to register your license server with the NVIDIA Licensing Portal to generate license files.

1. Open a web browser on the license server host and connect to the URL http://localhost:8080/licserver.
2. In the license server management interface, select **Configuration**.

3. On the License Server Configuration page that opens, in the **Server host ID** drop-down list, select the platform's ETHERNET address.

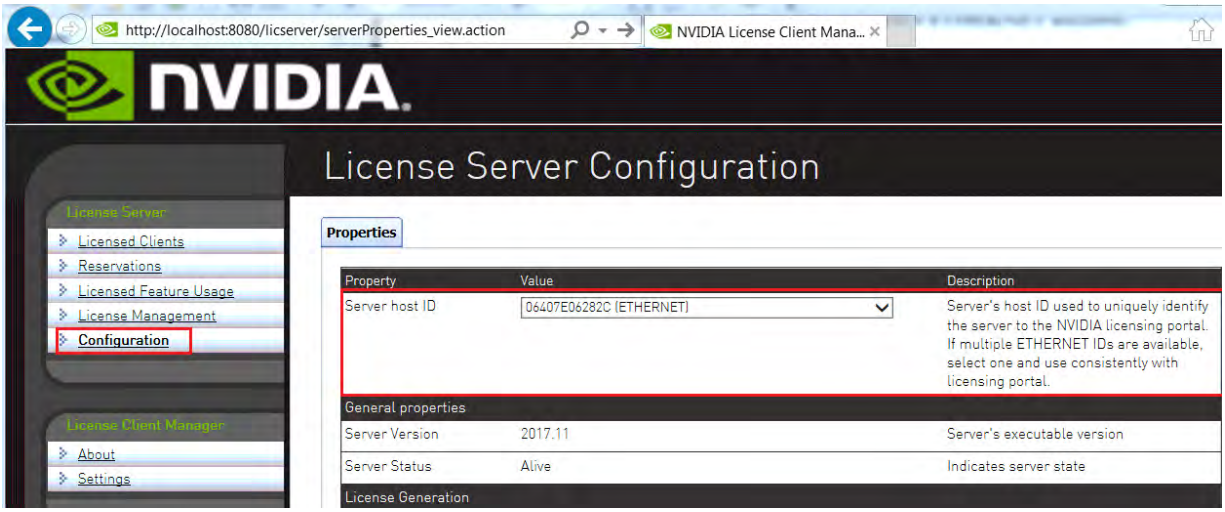

6.2.4 Managing your License Server and Getting your License Files

To be able to download NVIDIA vGPU software licenses, you must create at least one license server on the NVIDIA Licensing Portal and allocate licenses to the server. After creating a license server and allocating licenses to it, you can download your license file.

6.2.4.1 Creating a License Server on the NVIDIA Licensing Portal

1. In the NVIDIA Licensing Portal, navigate to the organization or virtual group for which you want to create the license server.

   a) If you are not already logged in, log in to the NVIDIA Enterprise Application Hub and click **NVIDIA LICENSING PORTAL** to go to the NVIDIA Licensing Portal.

b) **Optional:** If your assigned roles give you access to multiple virtual groups, select the virtual group for which you are creating the license server from the list of virtual groups at the top right of the page.

If no license servers have been created for your organization or virtual group, the NVIDIA Licensing Portal dashboard displays a message asking if you want to create a license server.

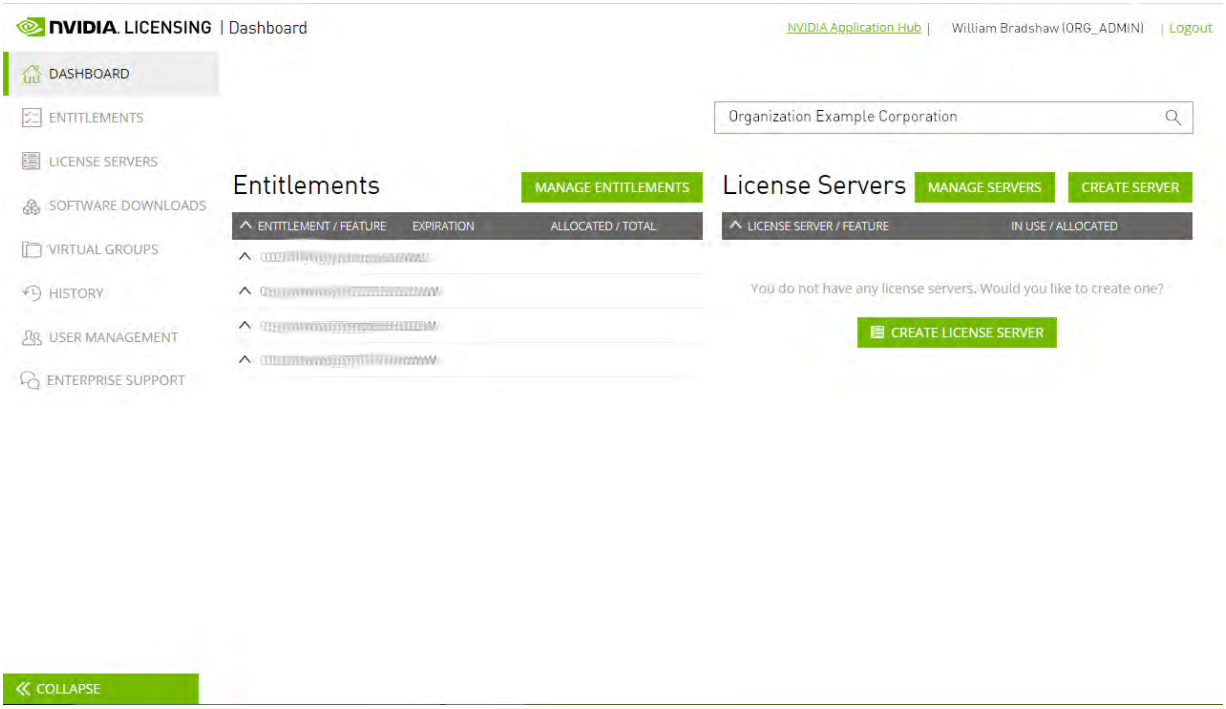

2. On the NVIDIA Licensing Portal dashboard, click **CREATE LICENSE SERVER**.

   The Create License Server pop-up window opens.

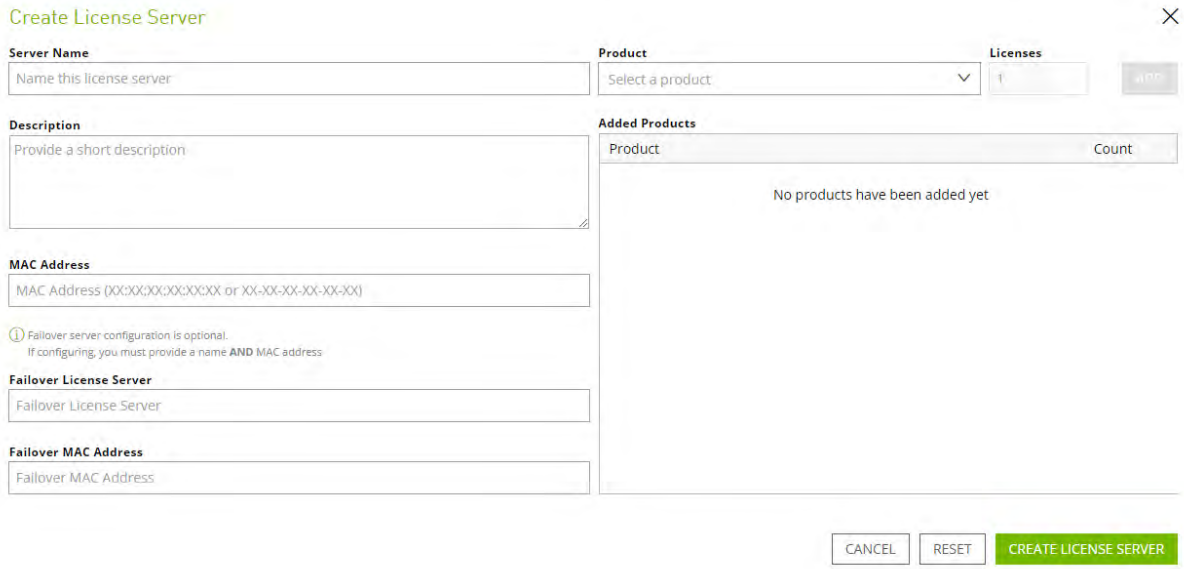

3. Provide the details of your license server.
   a) In the **Server Name** field, enter the host name of the license server.
   b) In the **Description** field, enter a text description of the license server. This description is required and will be displayed on the details page for the license server that you are creating.
   c) In the **MAC Address** field, enter the MAC address of your license server.
4. Add the licenses for the products that you want to allocate to this license server. For each product, add the licenses as follows:
   a) From the **Product** drop-down list, select the product for which you want to add licenses.
   b) In the **Licenses** field, enter the number of licenses for the product that you want to add.
   c) Click **ADD**.
5. Leave the **Failover License Server** and **Failover MAC Address** fields unset.
6. Click **CREATE LICENSE SERVER**.

6.2.4.2 Downloading a License File

Each license server that you create has license file associated with it. The license file contains all the licenses that you allocated to the license server. After downloading the license file, you can install it on the license server host associated with the license server on the NVIDIA Licensing Portal.

1. In the NVIDIA Licensing Portal, navigate to the organization or virtual group for which you want to download the license file.

   a) If you are not already logged in, log in to the NVIDIA Enterprise Application Hub and click **NVIDIA LICENSING PORTAL** to go to the NVIDIA Licensing Portal.

b) **Optional:** If your assigned roles give you access to multiple virtual groups, select the virtual group for which you are downloading the license file from the list of virtual groups at the top right of the page.

2. In the list of license servers on the NVIDIA Licensing Portal dashboard, select the license server whose associated license file you want to download.

3. In the License Server Details page that opens, review the licenses allocated to the license server.

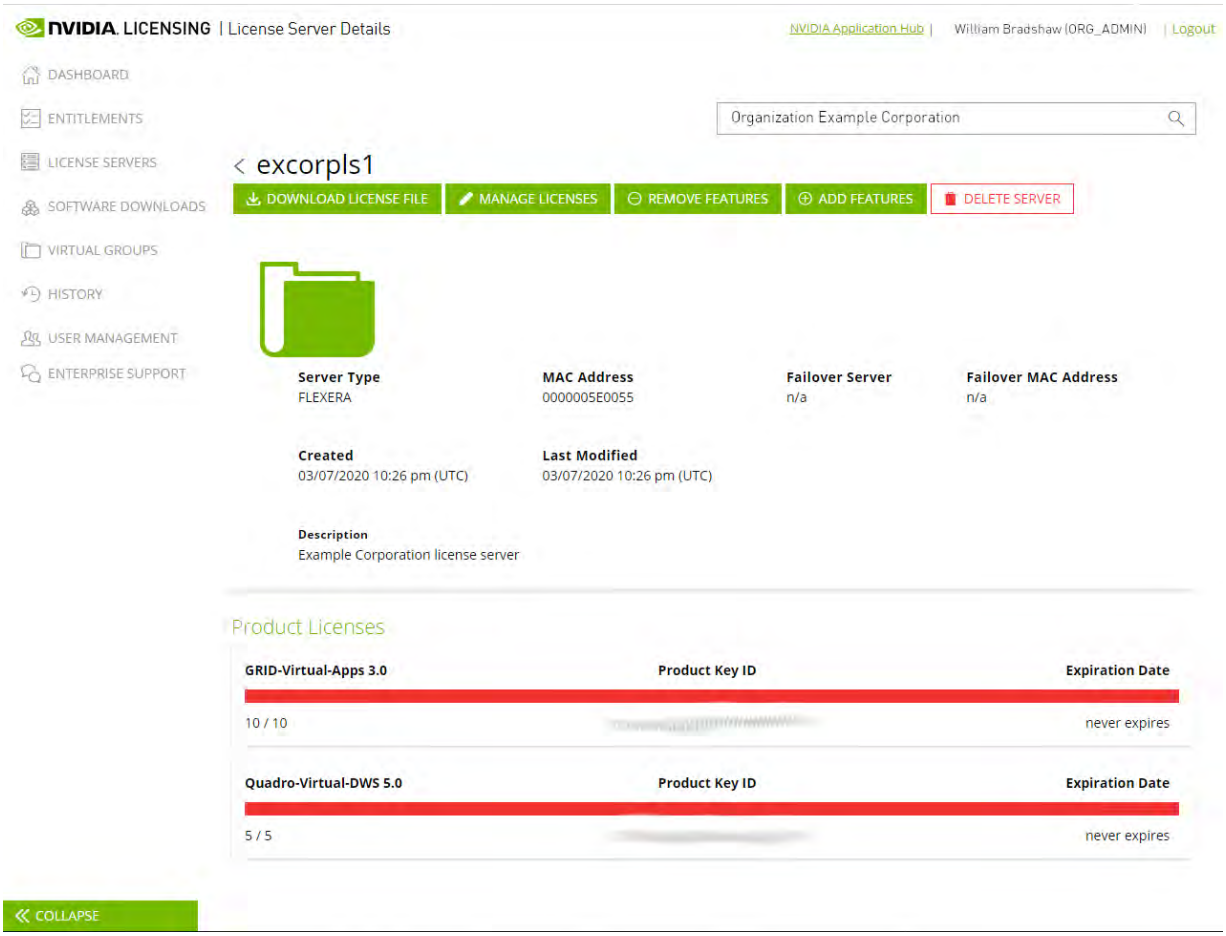

4. Click **DOWNLOAD LICENSE FILE** and save the .bin license file to your license server for installation.

6.2.5 Installing a License

NVIDIA vGPU software licenses are distributed as .bin files for download from the NVIDIA Licensing Portal.

Before installing a license, ensure that you have downloaded the license file from the NVIDIA Licensing Portal.

1. In the license server management interface, select **License Management**.

2. On the License Management page that opens, click **Choose File**.

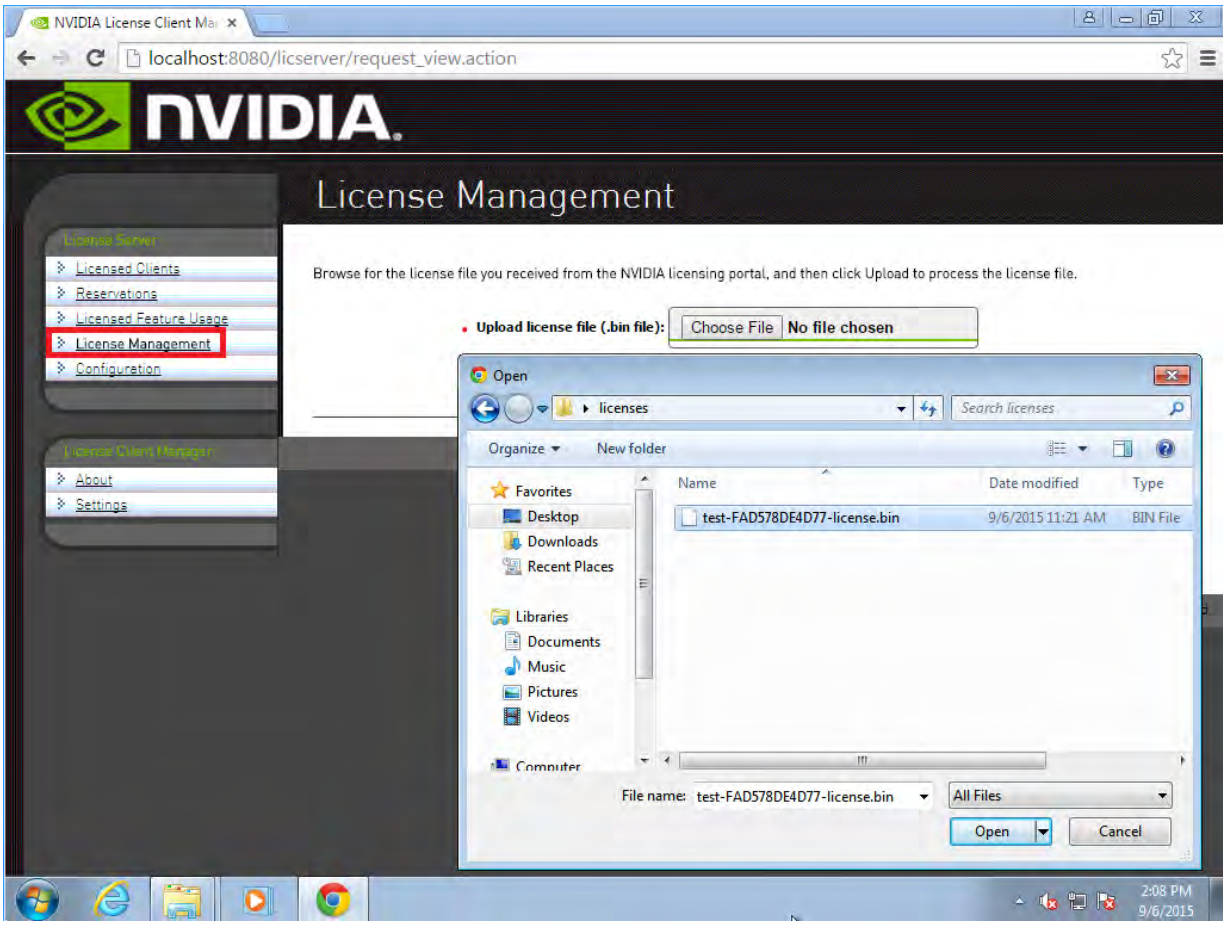

3. In the file browser that opens, select the .bin file and click **Open**.
4. Back on the License Management page, click **Upload** to install the license file on the license server. The license server should confirm successful installation of the license file.

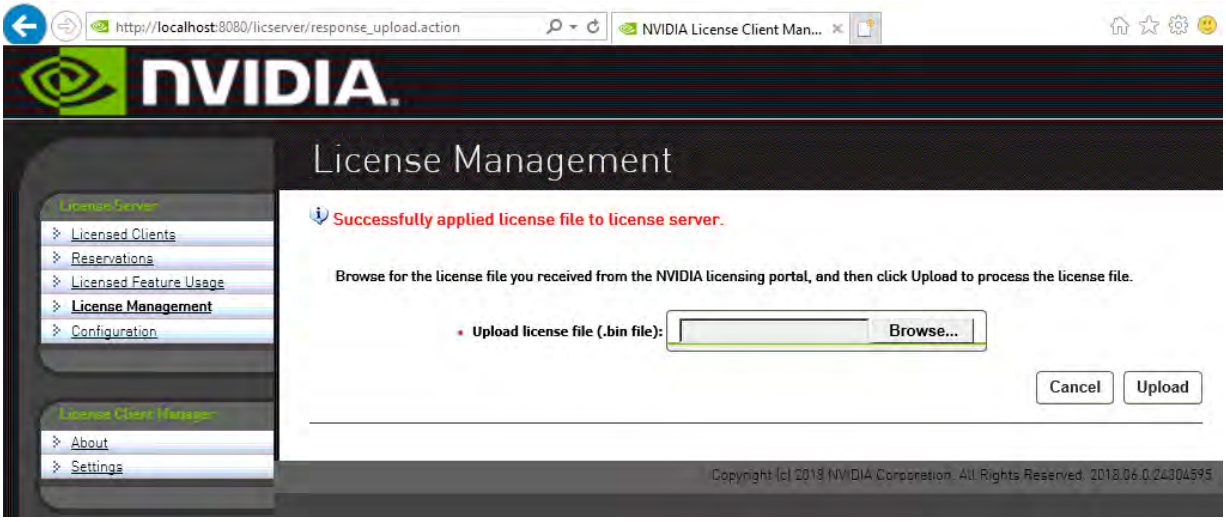

Note: For additional configuration options including Linux server deployment, securing your license server, and license provisioning, refer to the Virtual GPU Software License Server User Guide.

Chapter 7. Selecting the Correct vGPU Profiles

Choosing the right profile to maximize your stakeholders experience within the virtual instance is critical. Below, you will find guidance through the vGPU Manager and beyond to ensure your deployment is successful.

7.1 The Role of the vGPU Manager

NVIDIA vGPU profiles assign custom amounts of dedicated graphics memory to each virtual machine. NVIDIA vGPU Manager assigns the correct amount of memory to meet the specific needs within the workflow for the virtual machine user. Every virtual machine has dedicated graphics memory and must be assigned accordingly thus ensuring that it has the resources needed to handle the expected graphics load.

NVIDIA vGPU Manager allows multiple users to share each physical GPU by assigning the graphics resources of the available GPUs to virtual machines using a balanced approach. Depending on the number of GPUs within each NVIDIA card there can be multiple user types assigned.

7.2 The Full List of vGPU Profiles

vGPU profiles represent very flexible deployment options for virtual GPUs, varying the size of the allocated frame buffer memory depending on a number of factors, including the number and resolution of display heads. The division of frame buffer is what defines the number of users possible per GPU with that specific profile, while the number of heads defines the number of displays supported. Max resolution is consistent across all the profiles. The full list may also be found here: *https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html*

Series	Optimal Workload
Q-series	Virtual workstations for creative and technical professionals who require the performance and features of NVIDIA RTX Enterprise Drivers
C-series	Compute-intensive server workloads, such as artificial intelligence (AI), deep learning, or high-performance computing (HPC)
B-series	Virtual desktops for business professionals and knowledge workers

A-series	App streaming or session-based solutions for virtual applications users

Note: NVIDIA vGPU is a licensed product on all supported GPU boards. A software license is required to enable all vGPU features within the guest VM. The type of license required depends on the vGPU type.

- Q-series vGPU types require a NVIDIA RTX Virtual Workstation (vWS) license.
- C-series vGPU types require a NVIDIA Virtual Compute Server license but can also be used with a RTX vWS license.
- B-series vGPU types require a NVIDIA Virtual PC license but can also be used with a RTX vWS license.
- A-series vGPU types require a NVIDIA Virtual Applications license.

CAUTION: A NVIDIA Virtual Application license is required for Citrix Virtual Application deployments. A NVIDIA Virtual PC license is required for Citrix Virtual Desktop deployments.

Chapter 8. Creating Your First vGPU Virtual Desktop

This chapter describes how to:

- Create and configure a virtual machine in Citrix XenCenter
- Install Windows and Citrix VM Tools on the VM
- Prepare the VM for use as the golden image for Citrix MCS or PVS
- Install Citrix Virtual Delivery Agent on the VM
- Adjust additional VM settings and enable VM console access
- Enable the NVIDIA vGPU and finalizing the installation
- GPU Resource Allocation and Placement

8.1 Creating a Virtual Machine

The procedures detailed within this guide will use the simplest setup of Citrix Hypervisor; a single Citrix Hypervisor with local storage done through Windows 10 (64-bit). The default values may vary depending on your environment.

1. To create a Windows VM, navigate to the toolbar and select the **New VM** button to open the New VM wizard. This wizard allows you to configure the VM via parameters such as CPU, storage, and networking resources.

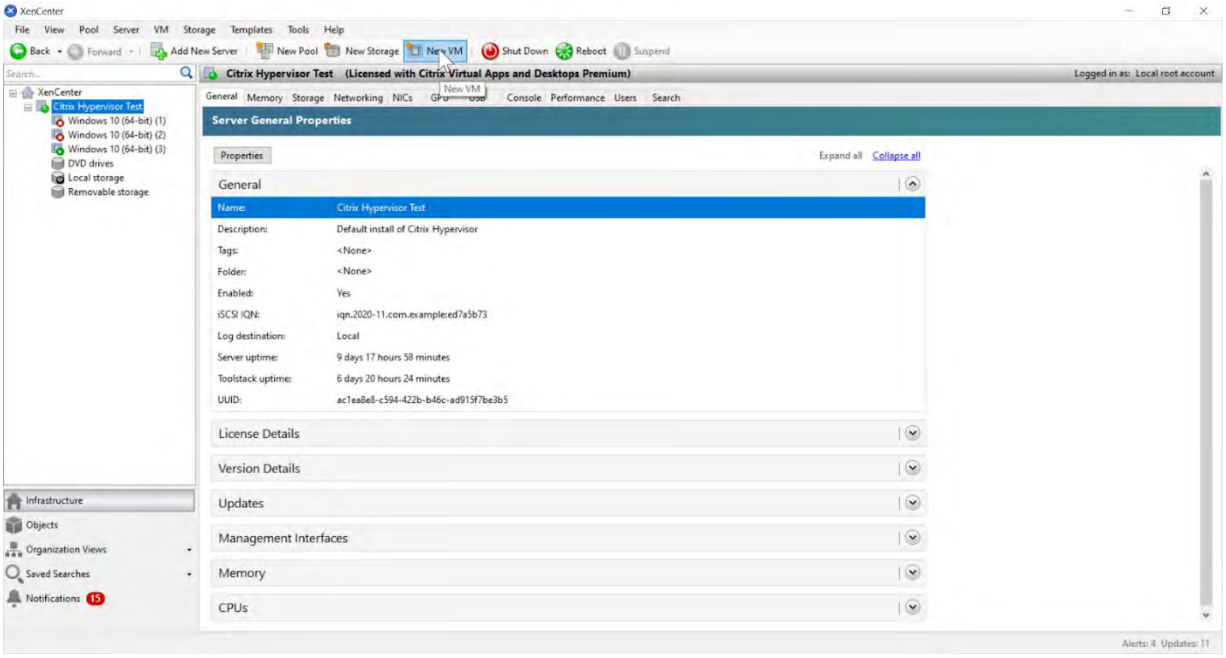

2. At the **Select a VM Template** window, navigate to the desired OS. For the purposes of this guide, select Windows 10 (64-bit). Select **Next** to continue.

Note: If the OS you are installing on your new VM is compatible only with the original hardware, check the Copy host BIOS strings to VM box. For example, use this option for an OS installation CD that was packaged with a specific computer. For more information go to Citrix.com.

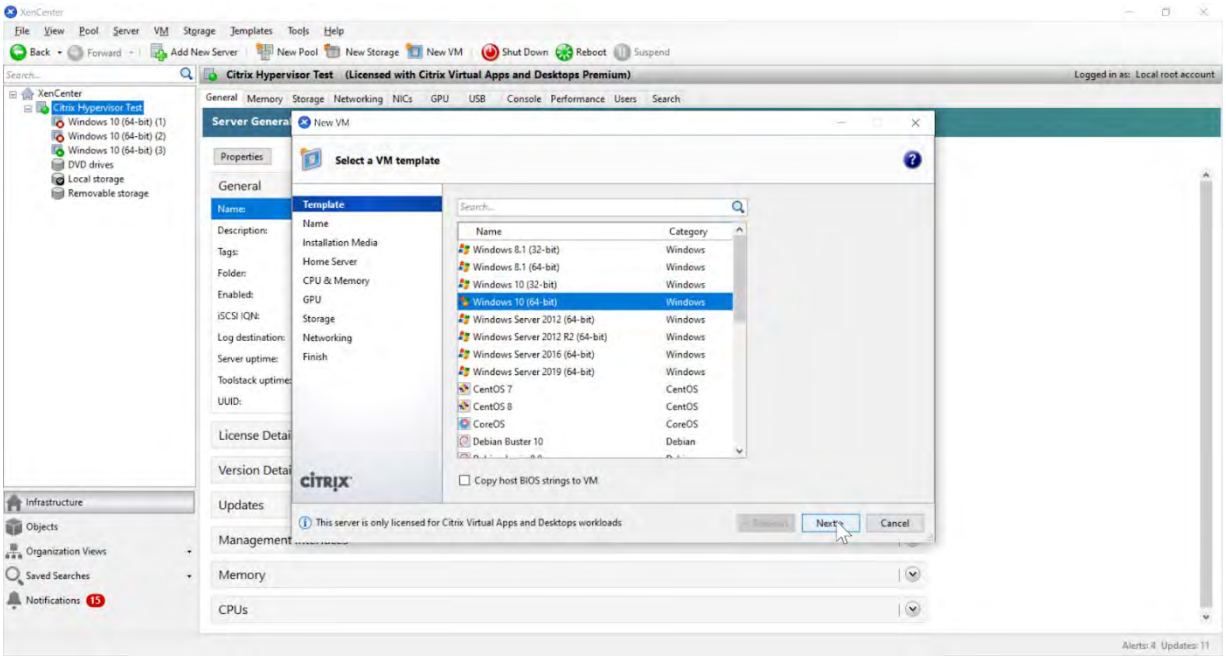

3. Provide a name and optimal description for the new VM. Select **Next** to continue.

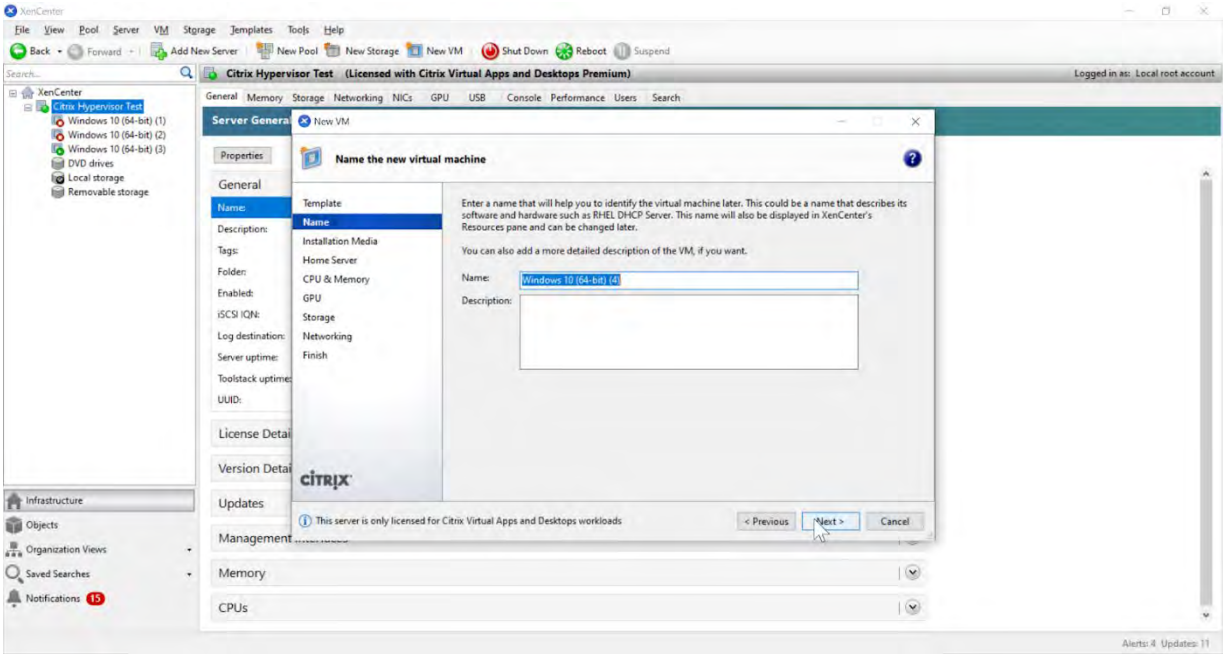

4. Choose the source of the OS media to install on the new VM. For the purpose of this guide, we will be selecting from a pre-existing ISO library and EUFI Boot mode. Find the ISO library in the drop-down list and select **Next** to continue.

Note: If you do not have a pre-existing ISO library, you can also attach an ISO to the host server's DVD drive and select it from the dropdown menu.

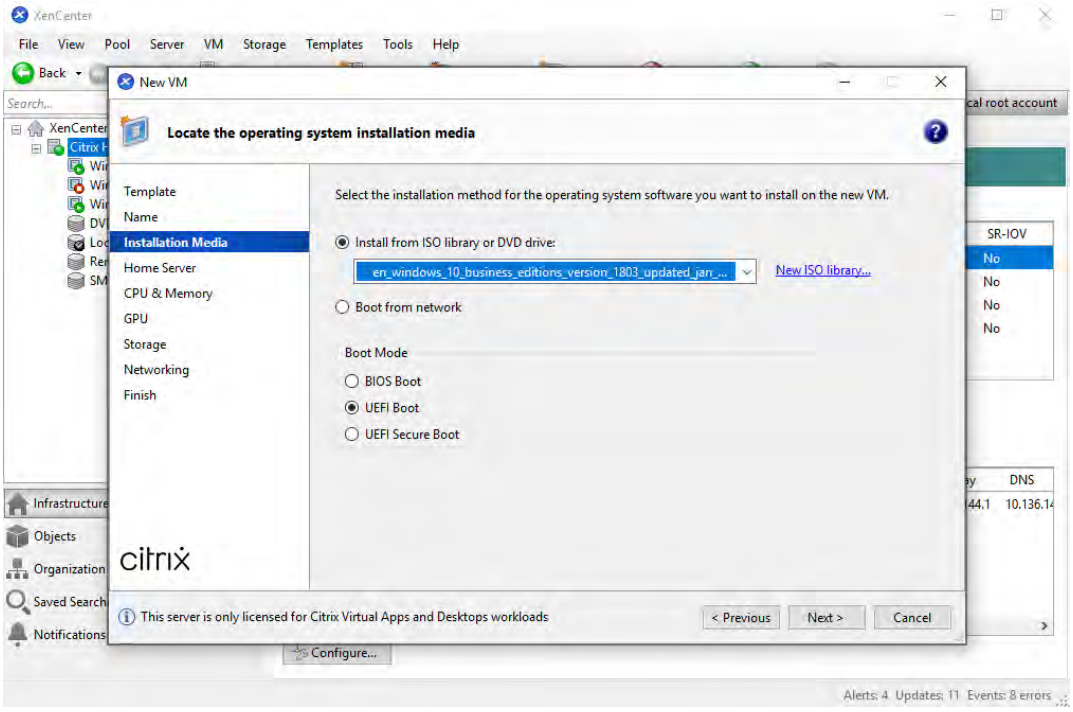

5. At the **Select a home server window**, select the appropriate Citrix Hypervisor server on which to place the VM. Choose **Next** to proceed.

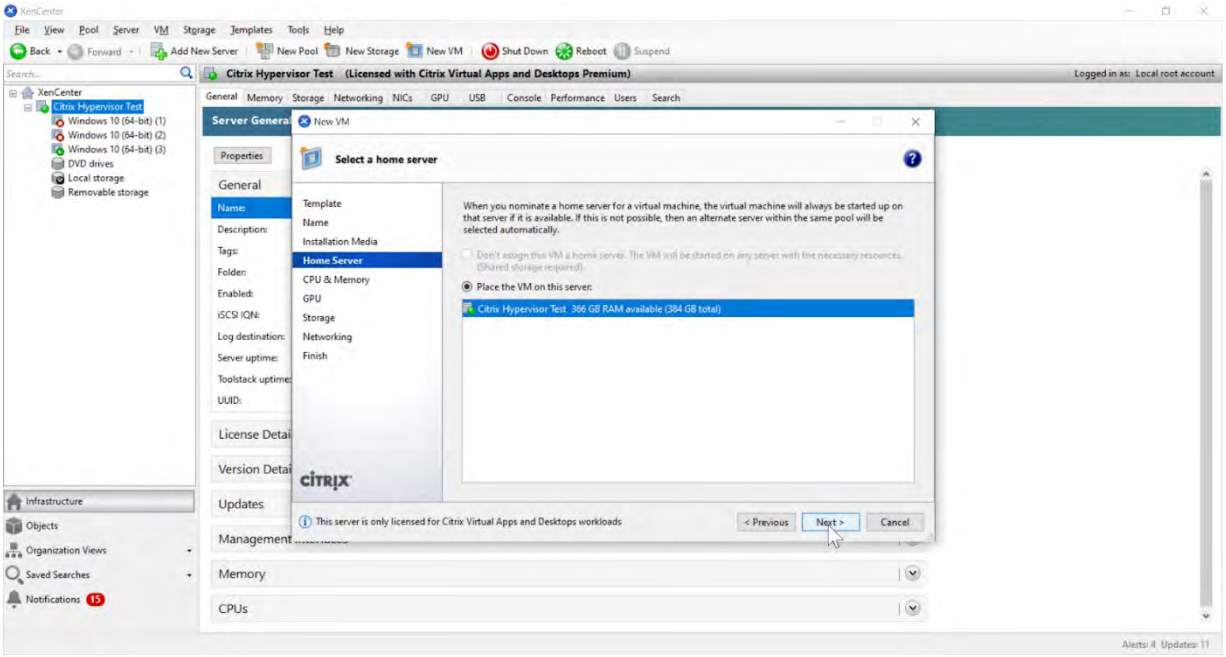

6. Change the VM specifications for CPU and Memory to preference. Select **Next** to continue.

Note: Each OS has different configuration requirements which are reflected in the templates.

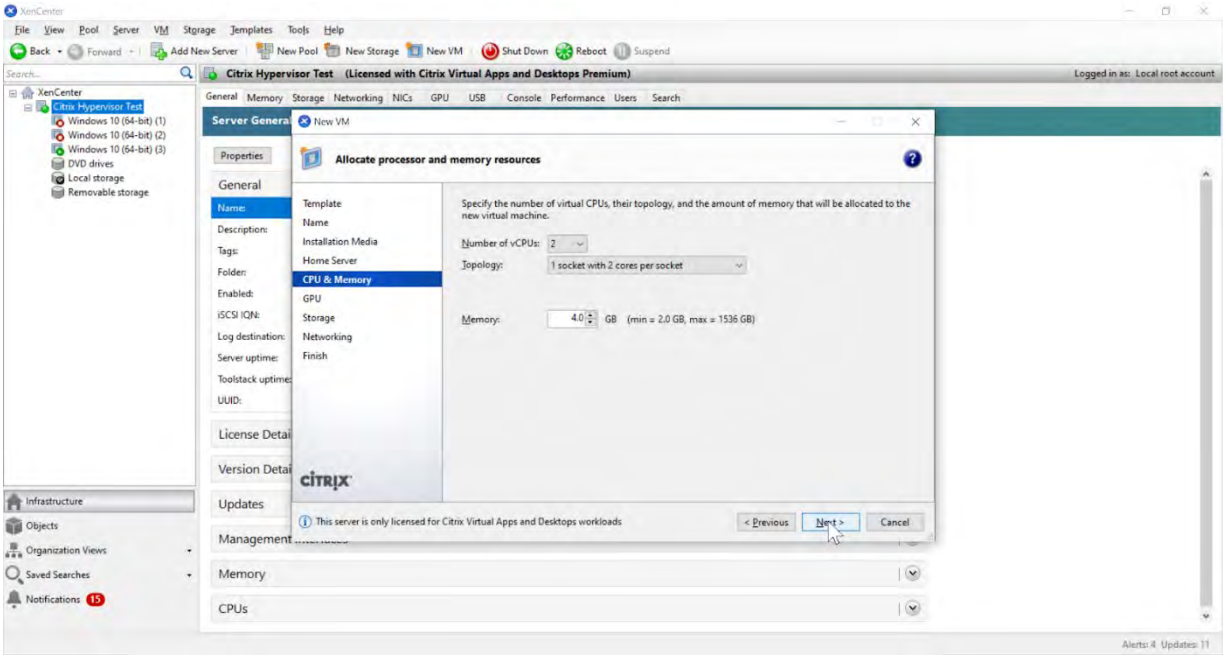

7. Attach a GPU to the new VM. The New VM wizard will allow designation of a dedicated GPU or virtual GPUs to the VM on a GPU enabled server via a configuration page. Select Add.

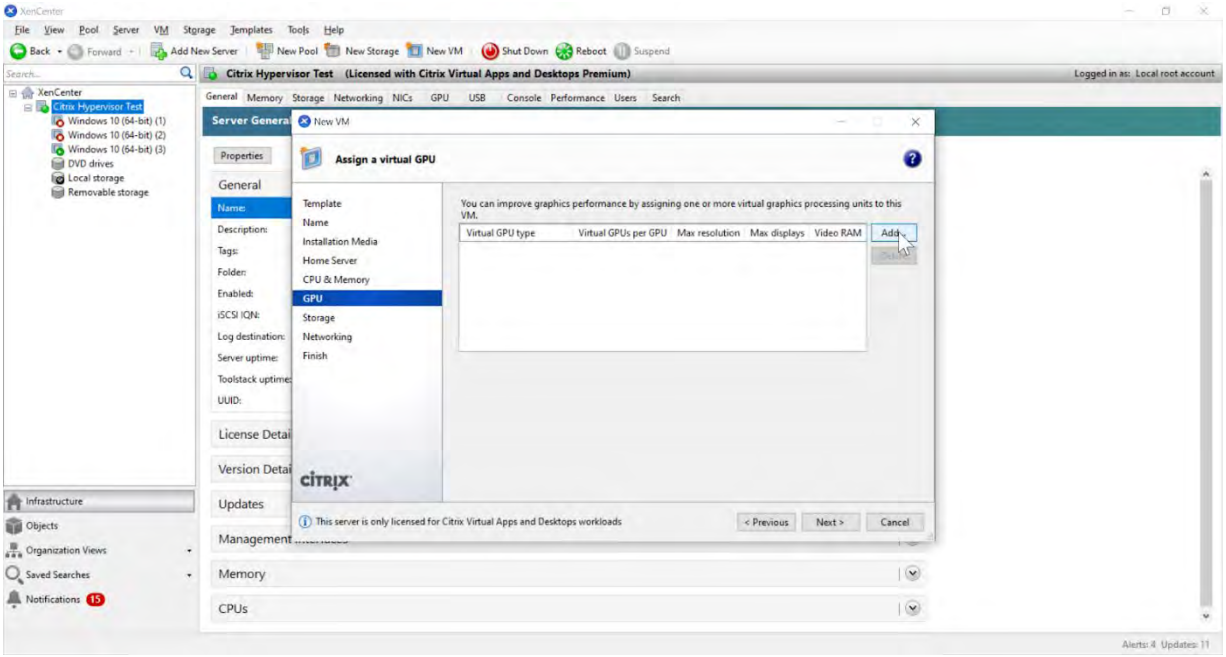

8. From the drop-down list select the desired GPU profile.

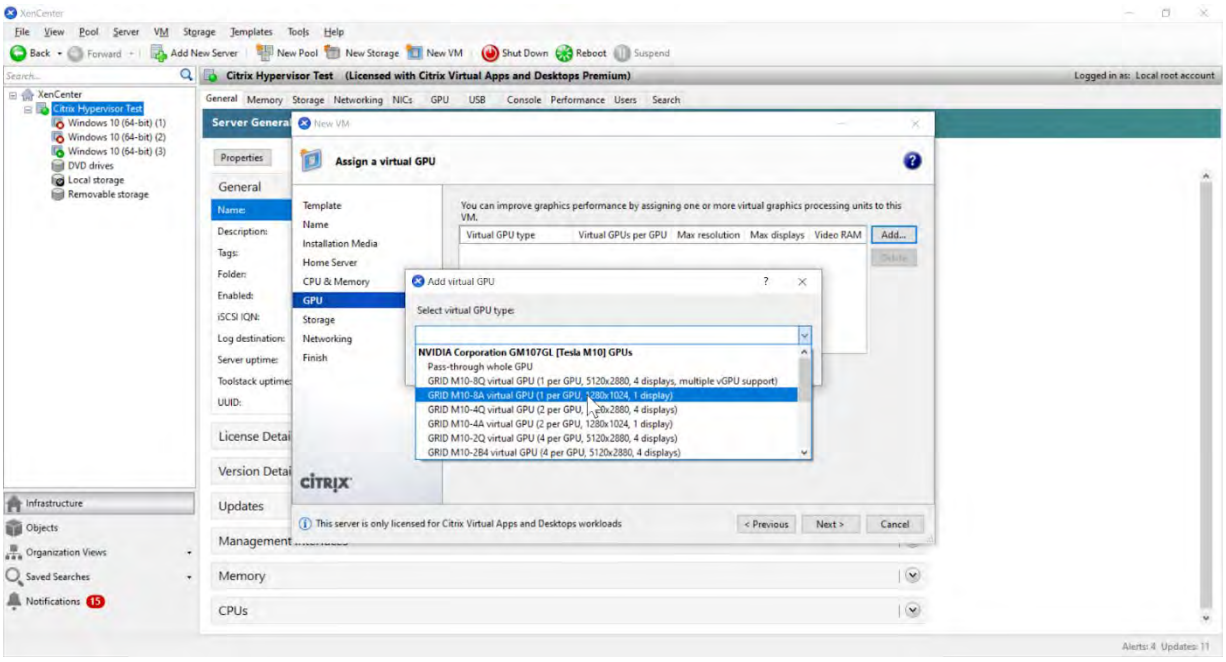

9. Navigate to Next to continue.

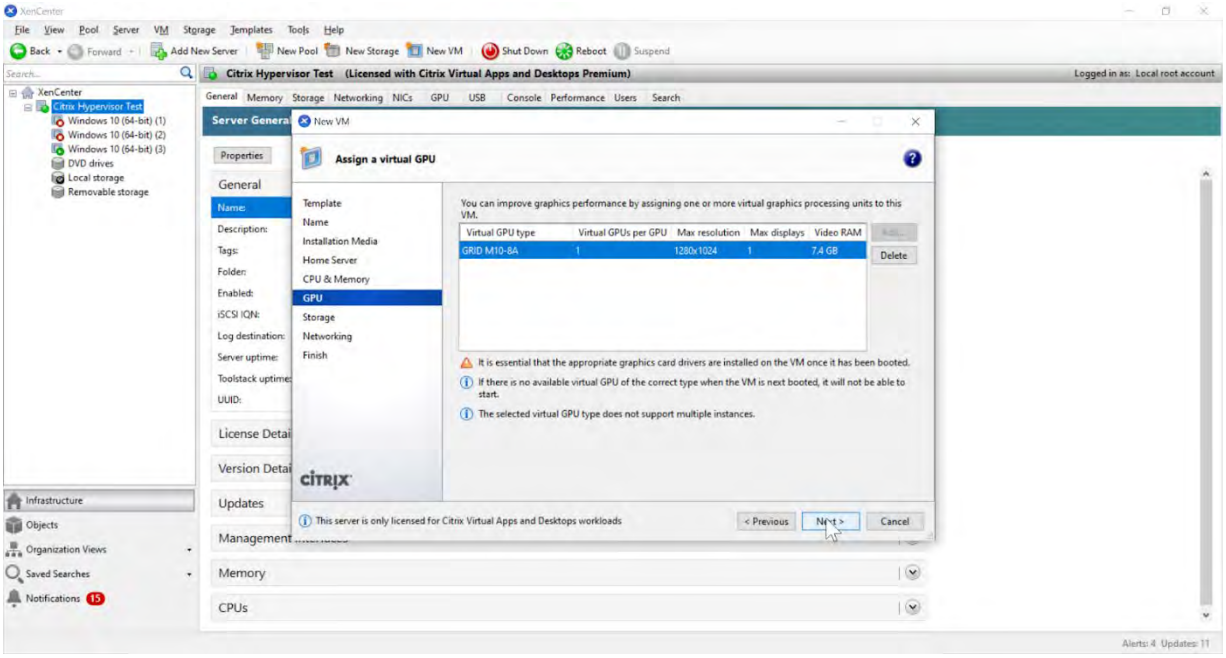

10. Configure storage for the new VM to desired specifications. You may also change the name, description, or size of your virtual disk by selecting Properties as well as add a new virtual disk by selecting Add. Select **Next** to continue after configuring storage appropriately for the VM.

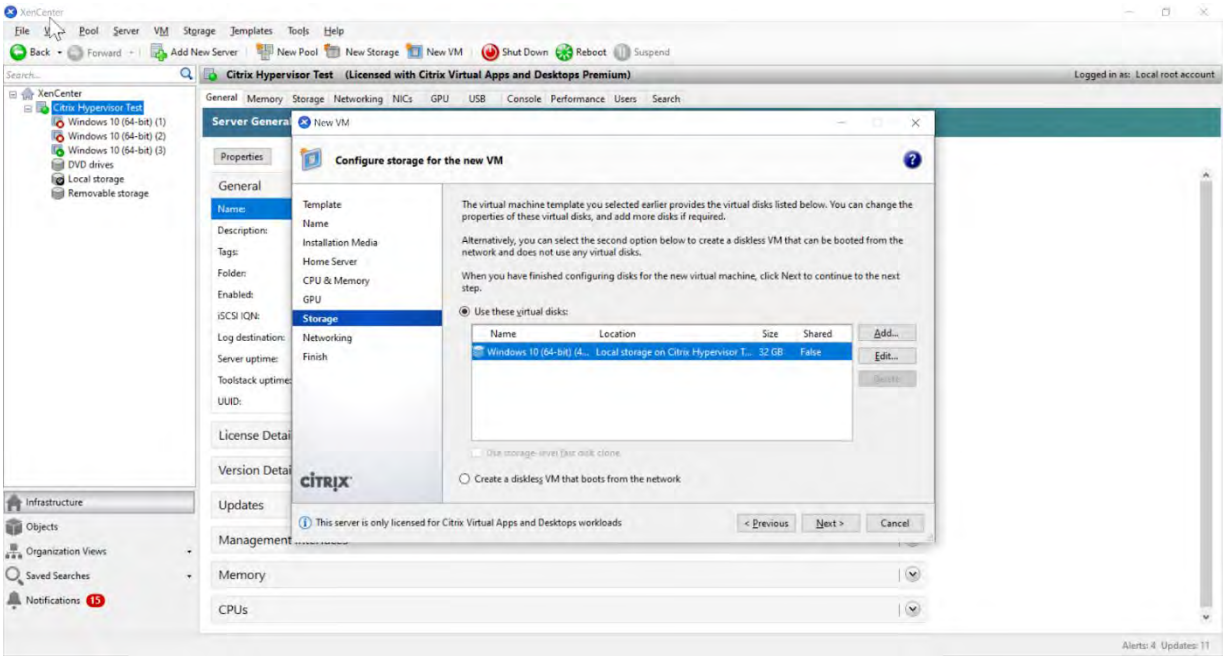

11. At the **Configure networking on the new VM** window, select the network that corresponds with the network that the VM requires. Select **Next** to continue.

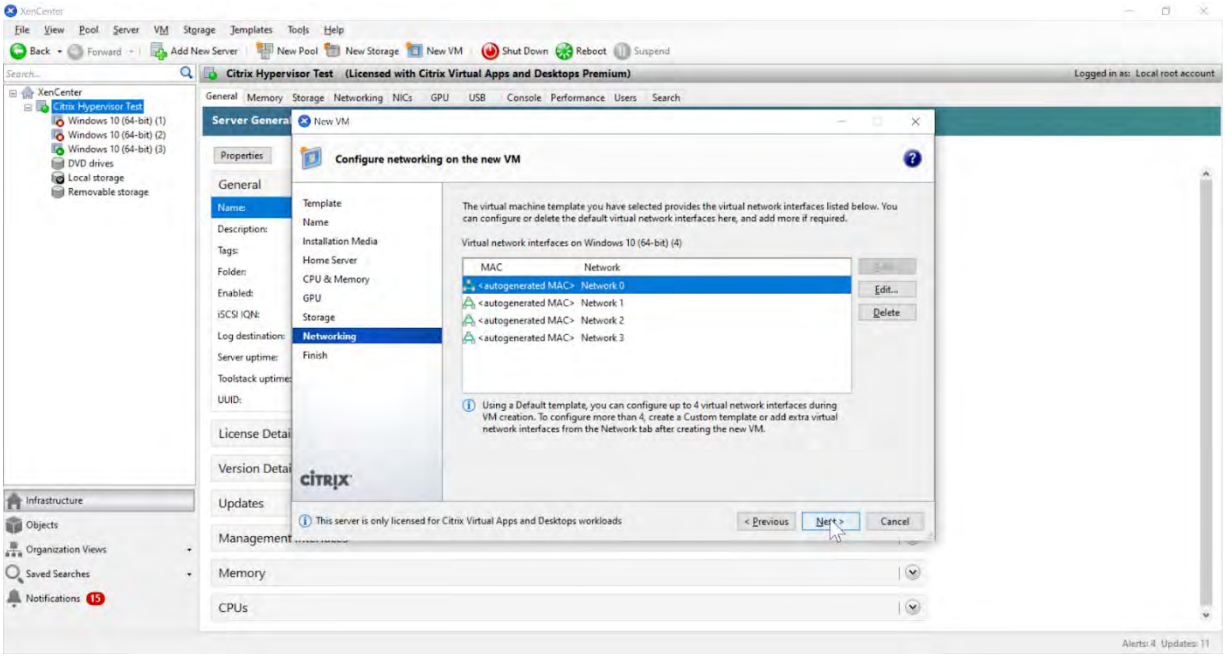

12. Review all settings and then select **Create Now** to create the VM. An icon for your new VM appears under the server in the left Resources pane.

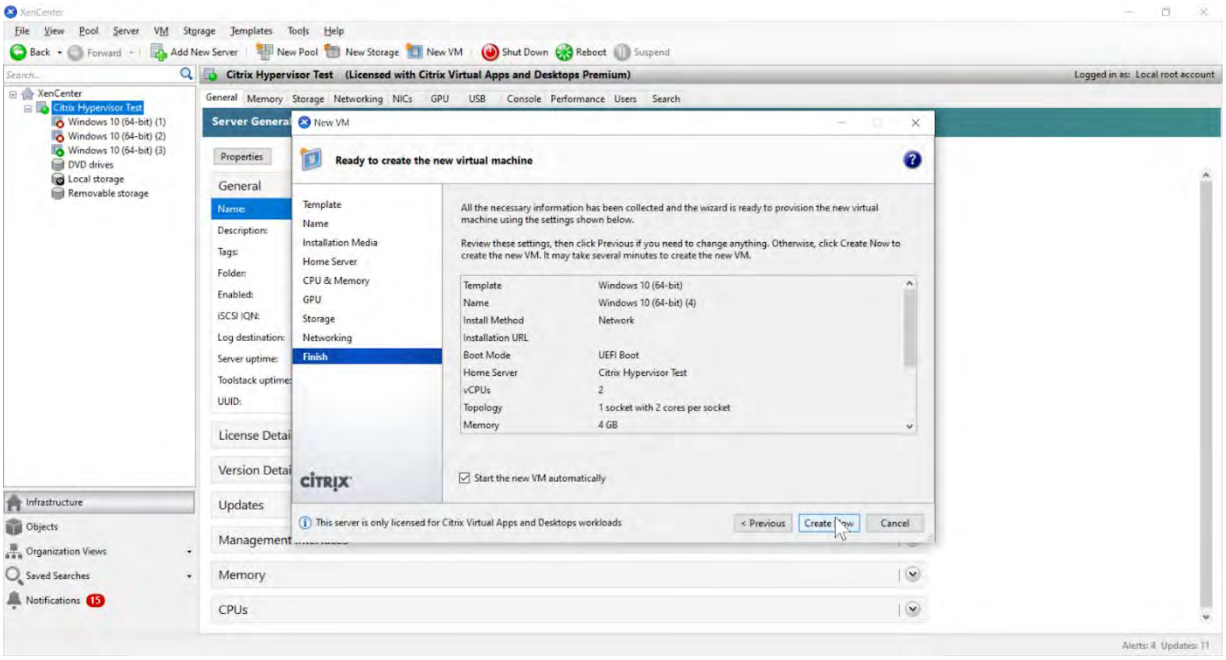

13. In the Resources pane, select the VM and then navigate to the **Console tab** to view the VM console.

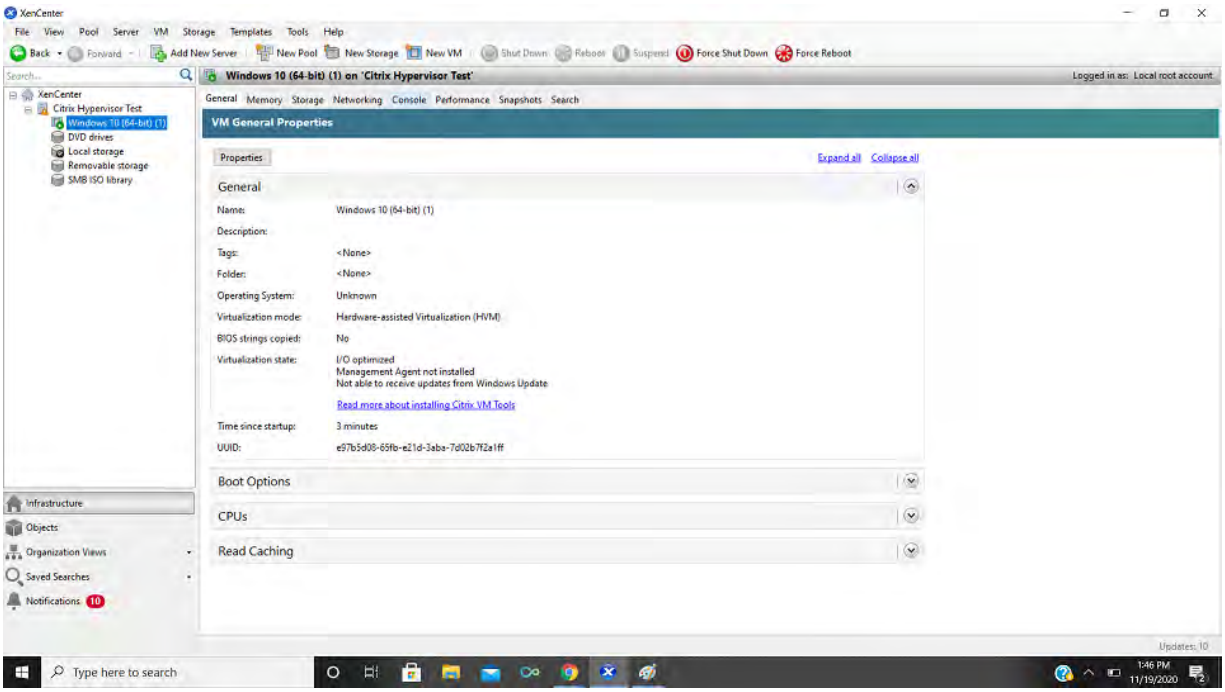

14. Install the Windows OS. Follow the installation screens and make your selections.

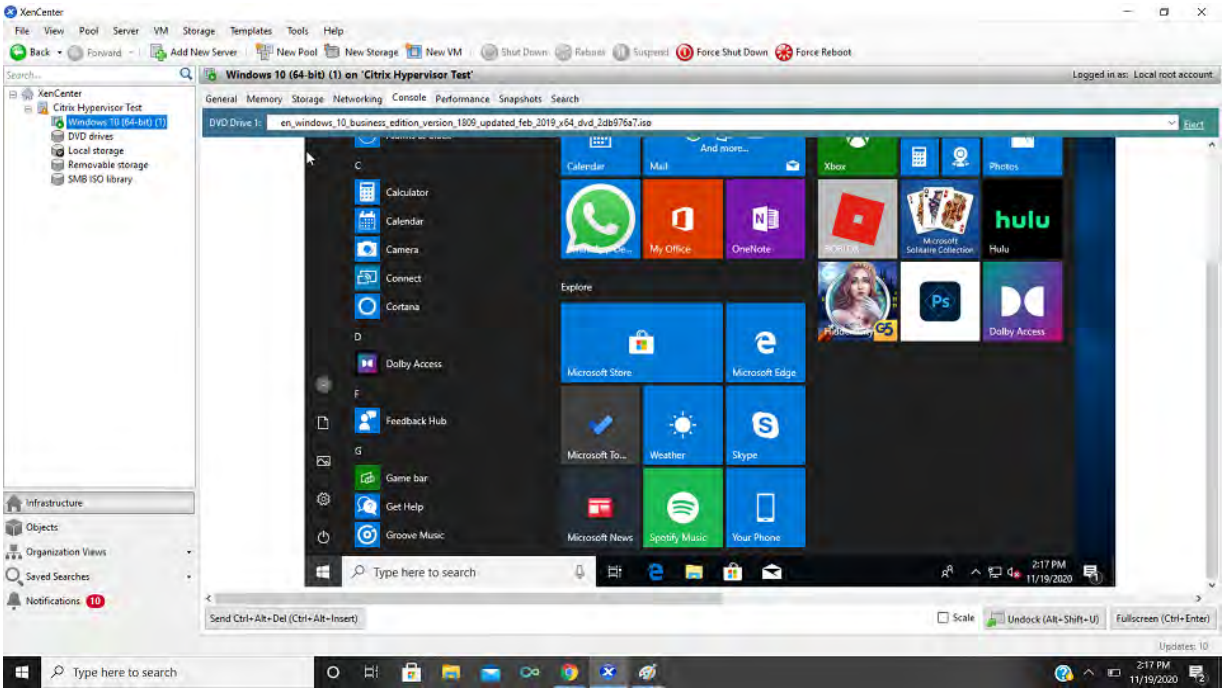

After the OS installation completes, reboot the VM.

8.2 Installing Citrix VM Tools

Citrix VM Tools for Windows must be installed on each Windows VM for the VM to have a fully supported configuration. A Windows VM will function without Citrix VM Tools, but performance is hampered and is not supported. Citrix VM Tools for Windows also enable certain functions and features, including cleanly shutting down, rebooting, suspending and live migrating VMs.

To install Citrix VM Tools for Windows:

1. Download the Citrix VM Tools for windows file onto the Windows VM. This file can be obtained from the Citrix Hypervisor Downloads page.
2. Run the managementagent.msi file to begin Citrix VM Tools installation. The VM Tools Installation wizard will display.

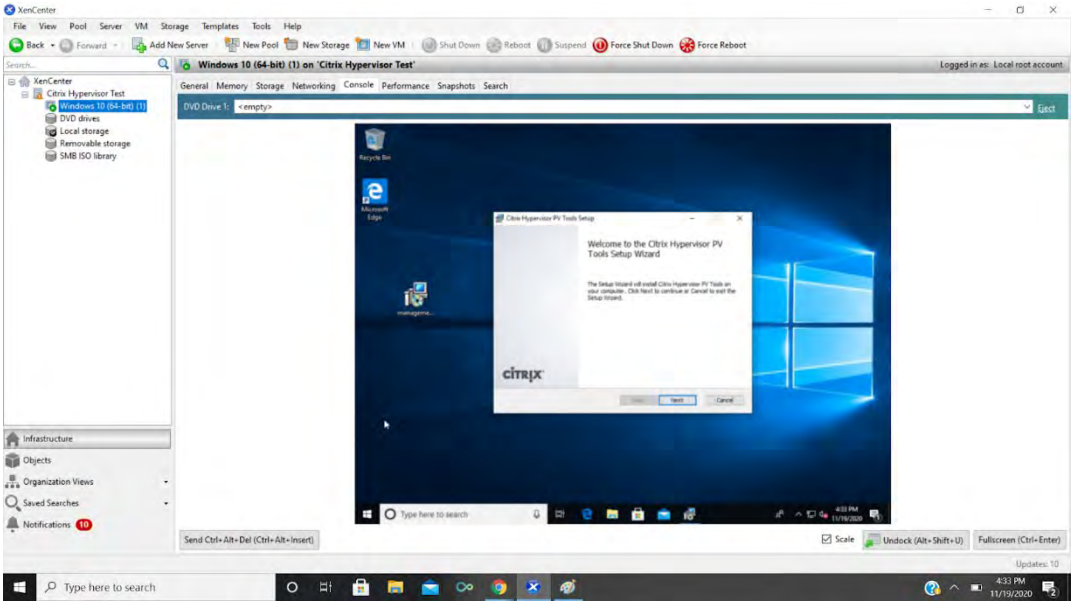

3. Select **Next** on the Welcome to the Citrix Hypervisor PV Tools Setup Wizard.
4. Read the End-User License Agreement and Check the **I accept the terms in the License Agreement** check box and client **Next.**

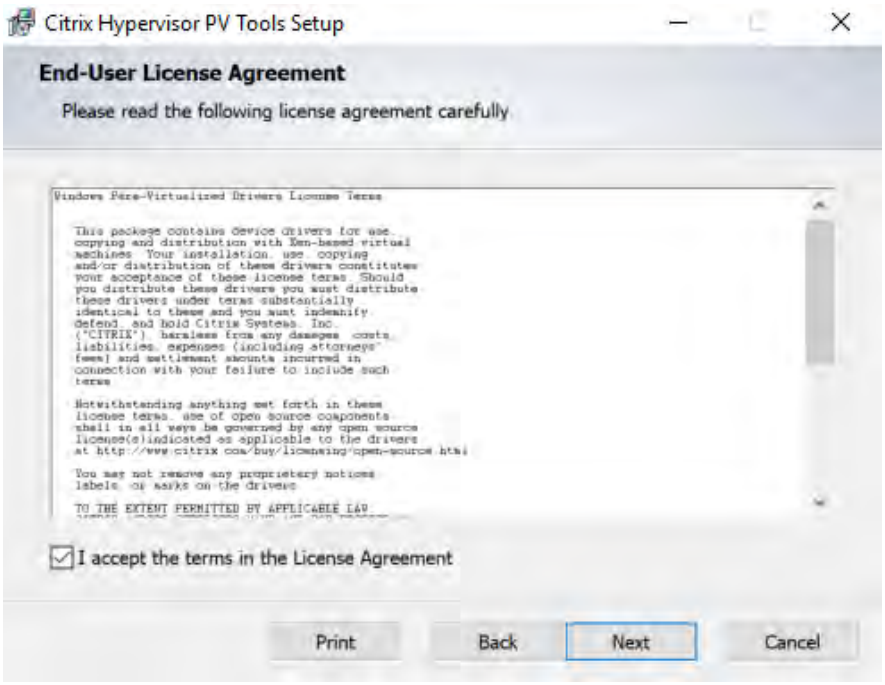

5. On the Destination Folder page, define an install location and click **Next**.

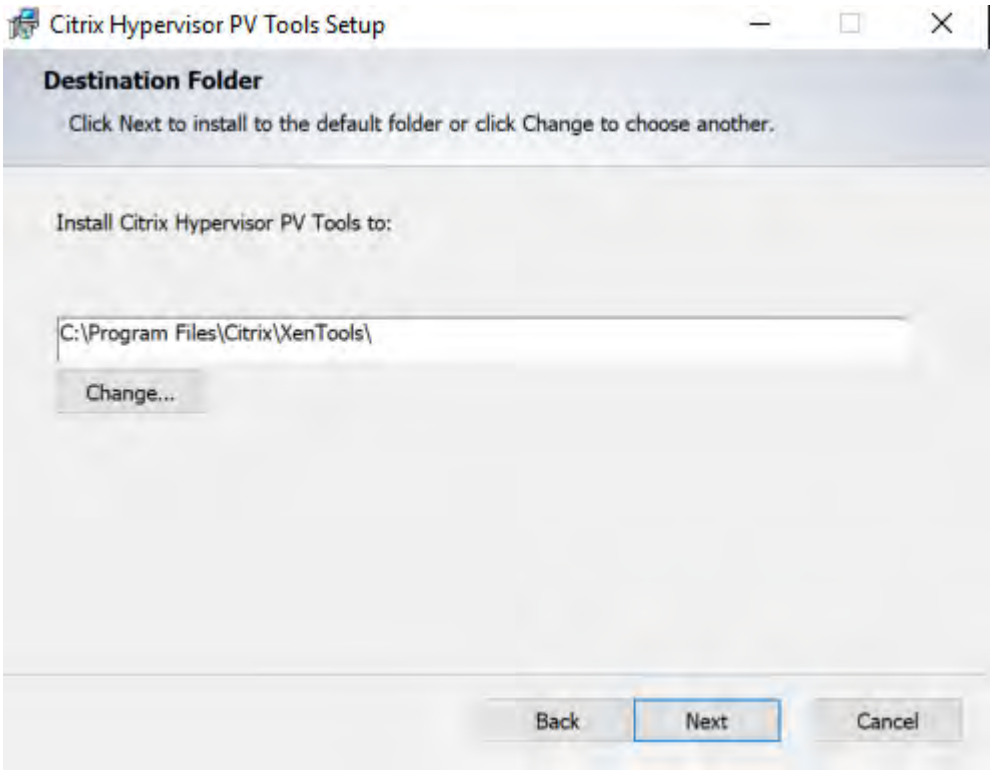

6. On the Installation and Update Setting Window, leave the Install I/O Drivers Now checkbox selected for optimal performance and functionality.
7. Choose the appropriate drop-down options for managing updates that are in line with your organizational needs and click, **Next**.

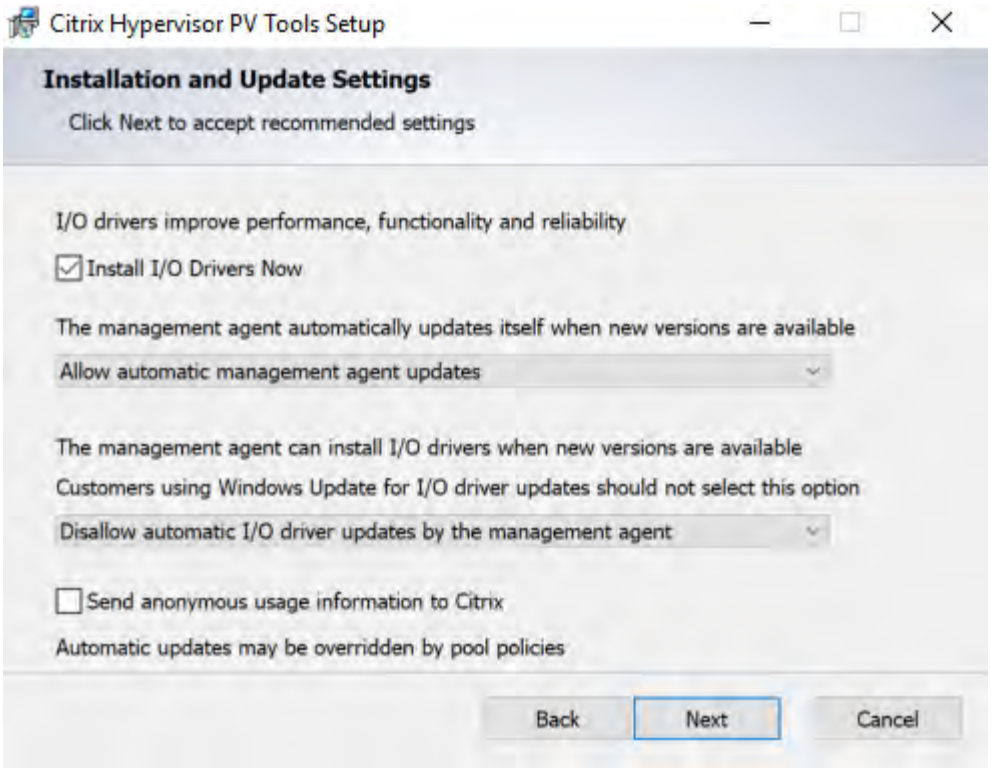

8. Select **Install**.

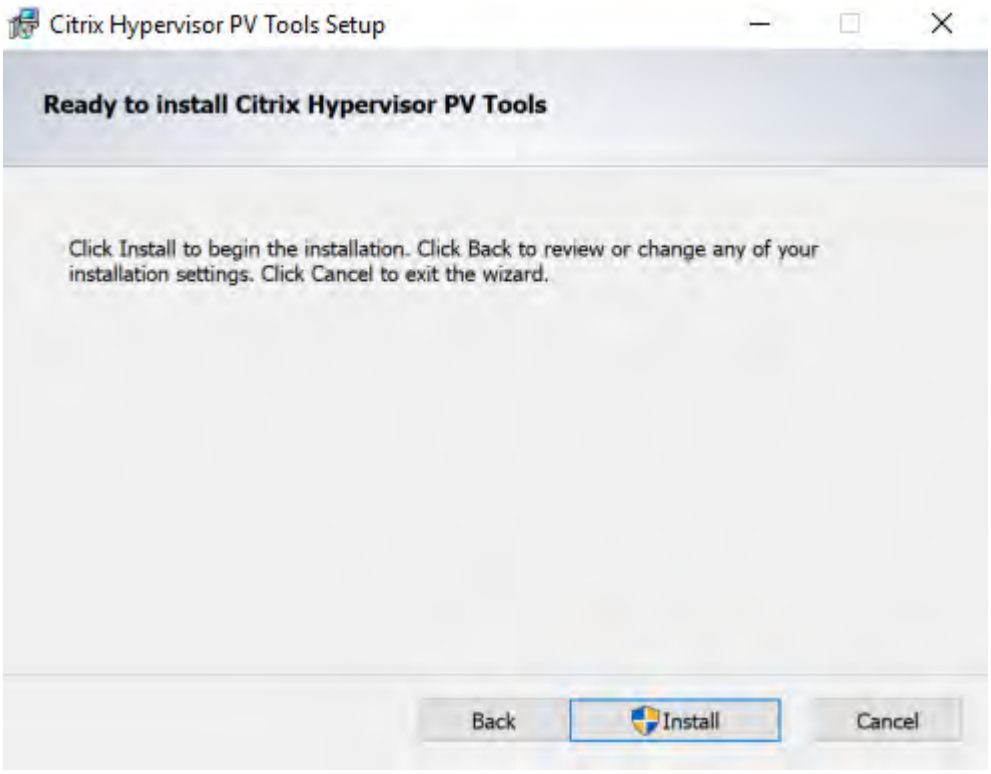

9. After Installation has finished, click **Finish**.
10. Restart the VM when prompted to complete the installation of Citrix VM Tools for Windows.

8.3 Adding the VM to the Domain

By joining the VM to the Windows Active Directory domain you are then able to manage it as you would any physical desktop in the domain.

Customize Windows on the virtual machine as follows:

- Join the domain
- Add appropriate Domain groups to Local Administrators
- Adding a VM to the domain:

1. On the VM, go to **Control Panel**, **System and Security**, **System**

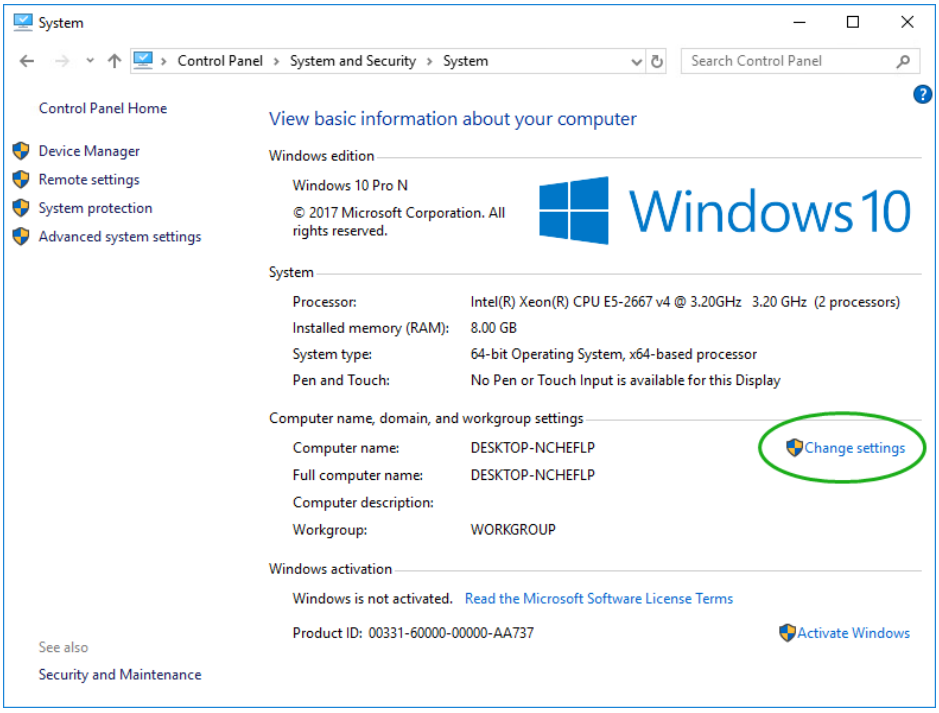

2. This brings up the System Properties window, on the Computer Name tab click Change:

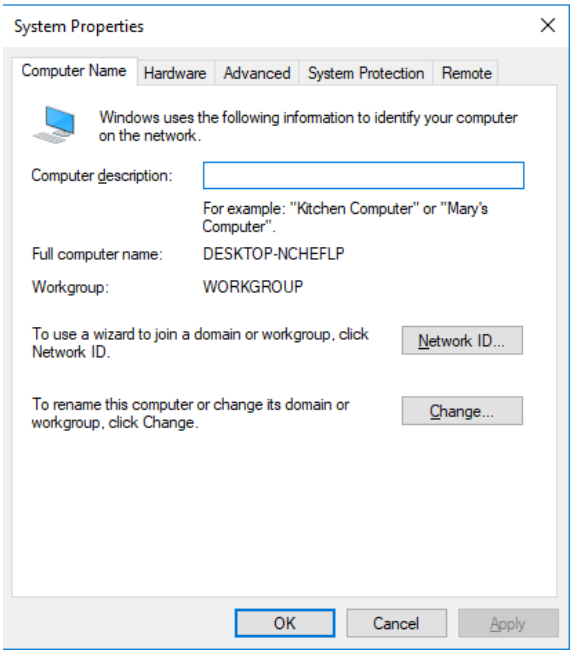

3. On the Computer Name/Domain Changes window, enter in an appropriate Computer name, then Domain name, and click OK. Our chosen naming is shown below, use what is appropriate for your POC/trial.

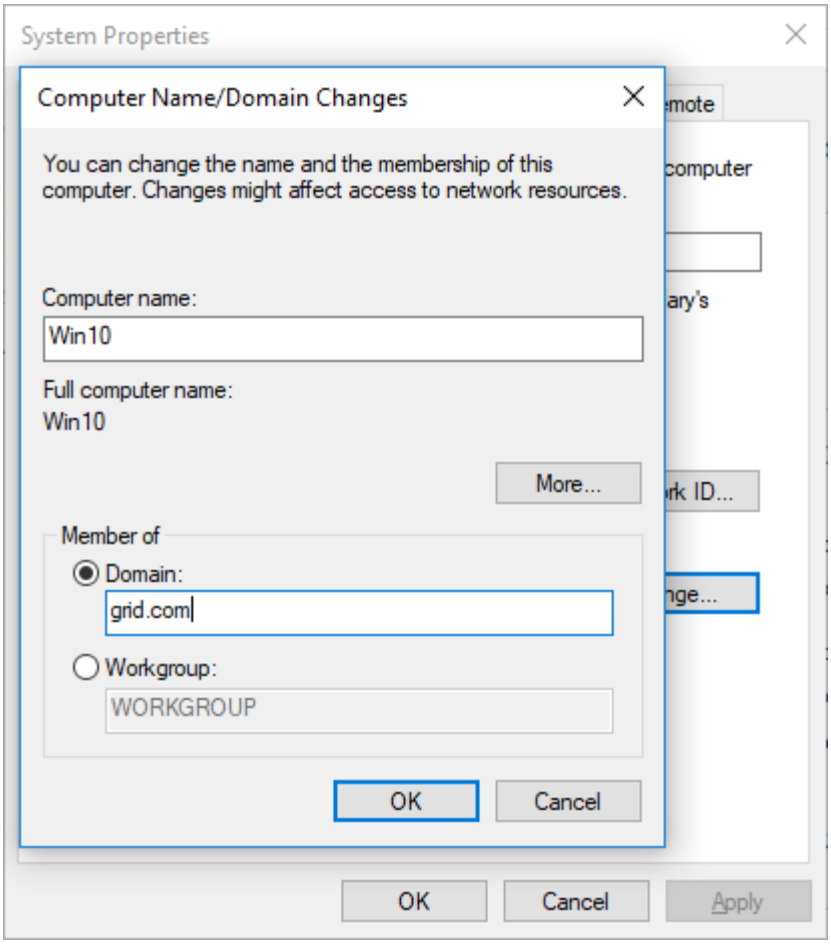

4. A security window pops up, fill in your specific domain administrator credentials and click OK:

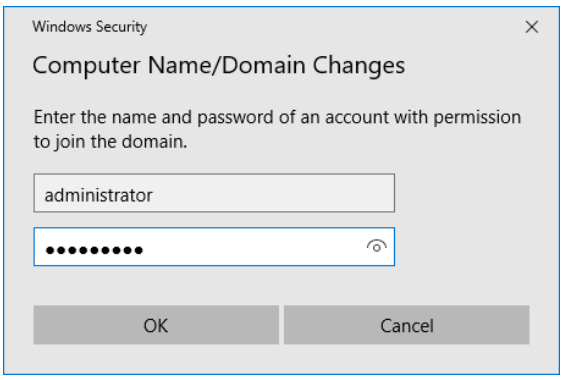

5. On successful authentication you will see the following welcome pop-up showing your VM is now on the domain (the domain name should reflect your domain information):

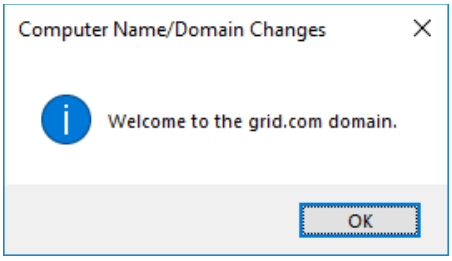

6. Click OK and the VM needs to reboot to complete the process, click OK again and the VM reboots immediately.

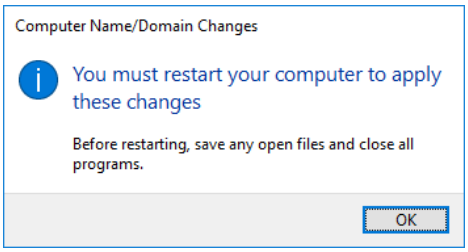

8.4 Installing the Citrix Virtual Delivery Agent

You need to install the correct version of the Citrix Virtual Delivery Agent (VDA) for your virtual machine. For the purpose of this guide, Citrix Virtual Apps and Desktop LTSR 7_1912 was used, therefore this guide uses VDA agent within the exact same LTSR service branch version.

Use the following procedure to install the Virtual Delivery Agent:

1. Attach the iso file to the server OS and open it via File Explorer.
2. Launch the ***Auto Select*** Application and accept the Windows User Account Control Popup.
3. Click **Start** for ***Virtual Apps and Desktops*** section.

4. Select **Virtual Delivery Agent for Windows Single-session OS** to launch the Citrix Virtual Delivery Agent installer.

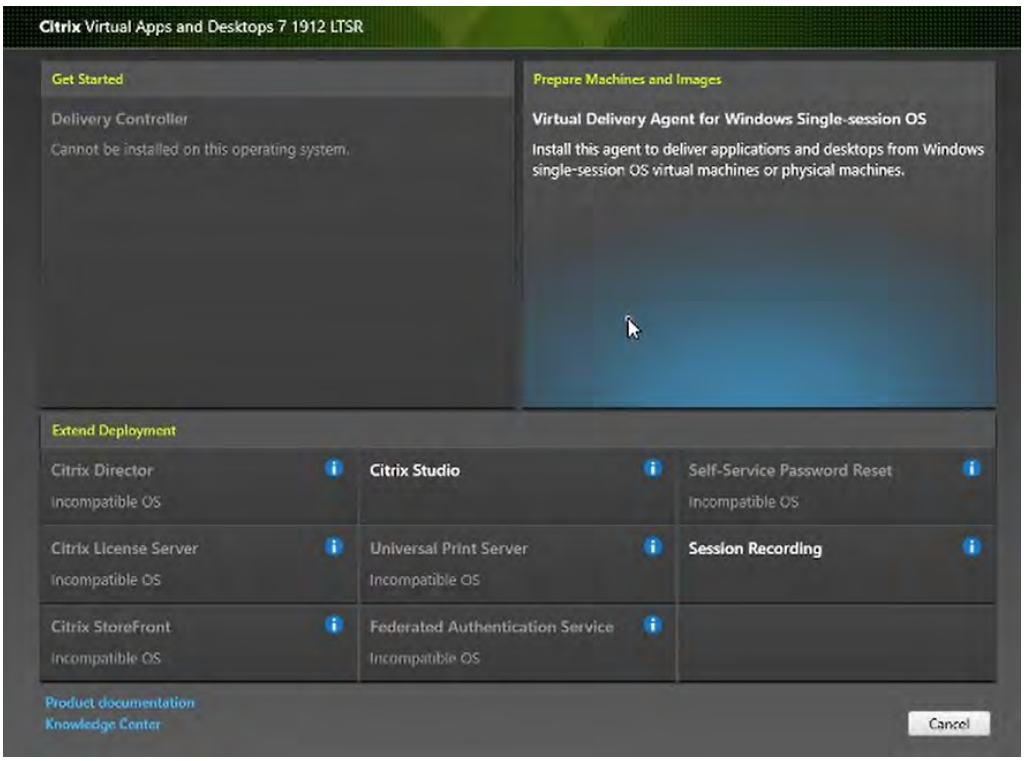

Note: If you are using an Operating System that supports multiple user sessions, you will select Virtual Delivery Agent for Windows Multi-session OS.

5. On the environment window, ensure the **Create a master MCS image** radio button is selected.

> Note: In a production environment you may use a different VM provisioning method.

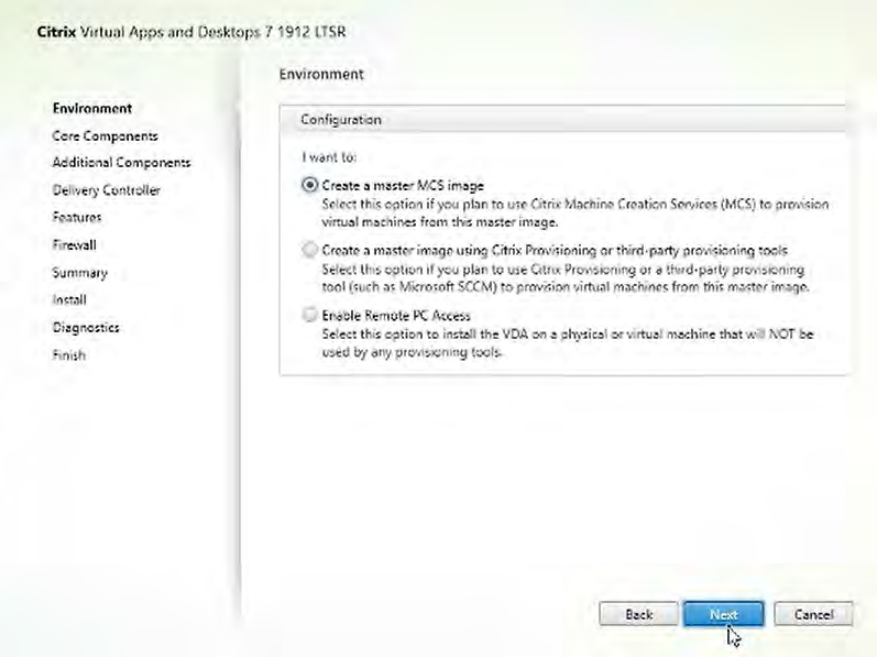

6. On the Core Components window, you can change the install location or leave the default location and click **Next.**

7. The Additional Components window allows you to choose which Components to install. Leave the defaults selected and click **Next.**

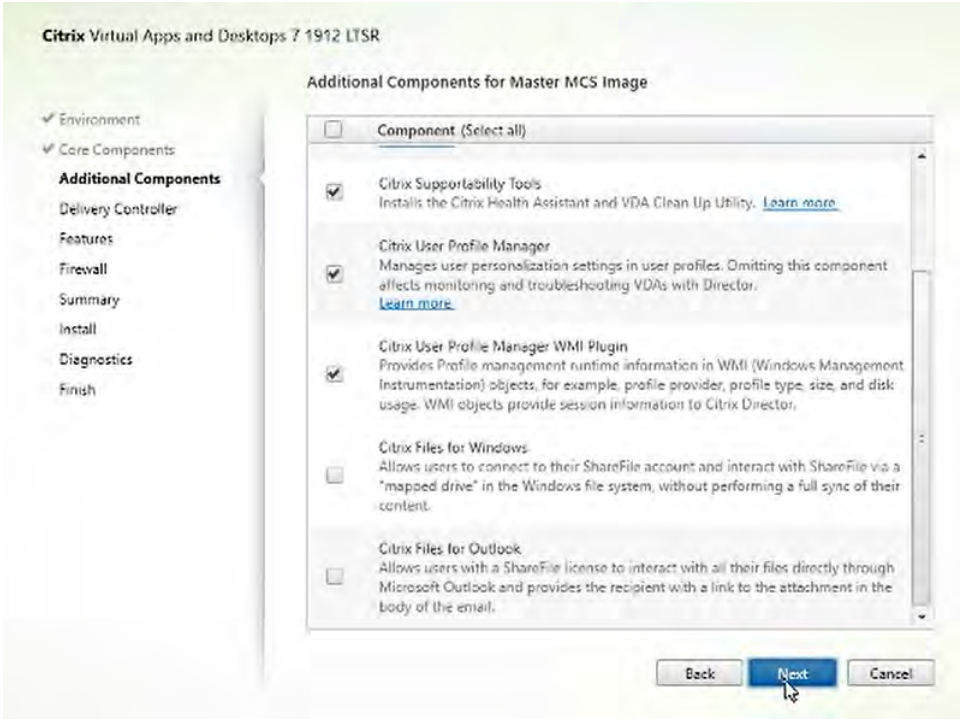

8. On the Delivery Controller window, type in the Fully Qualified Domain Name of your Delivery Controller in the **Controller address:** text field.

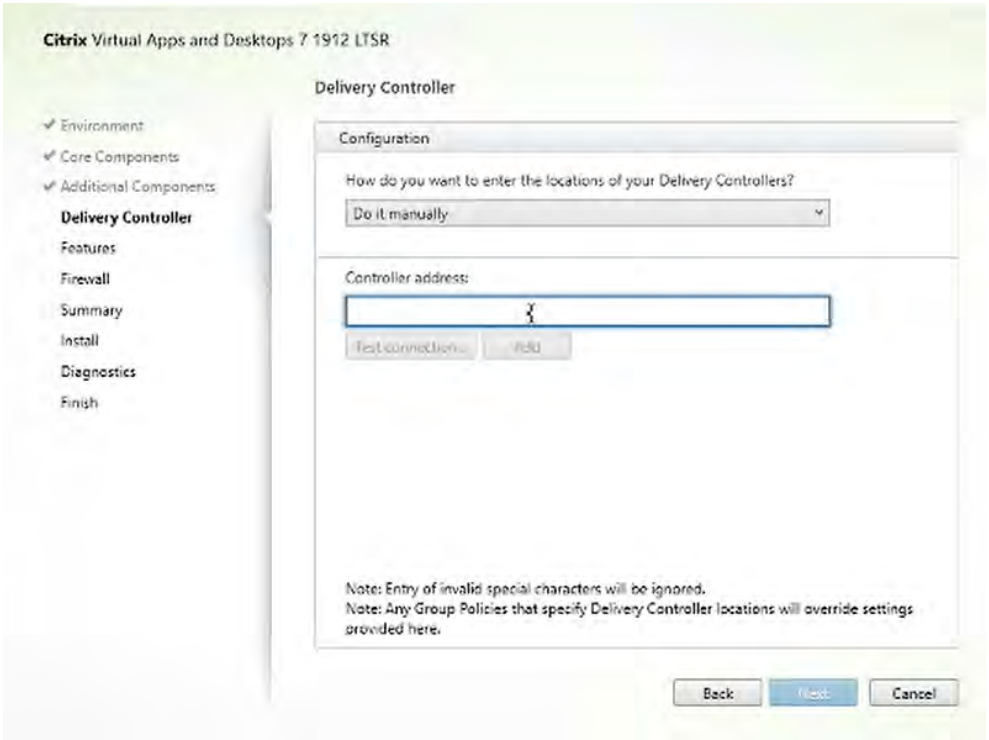

9. Click **Test connection..** then click **Add**.
10. Click **Next**.
11. The Features window allows you to choose which Features to install. Leave the defaults selected and click **Next**.

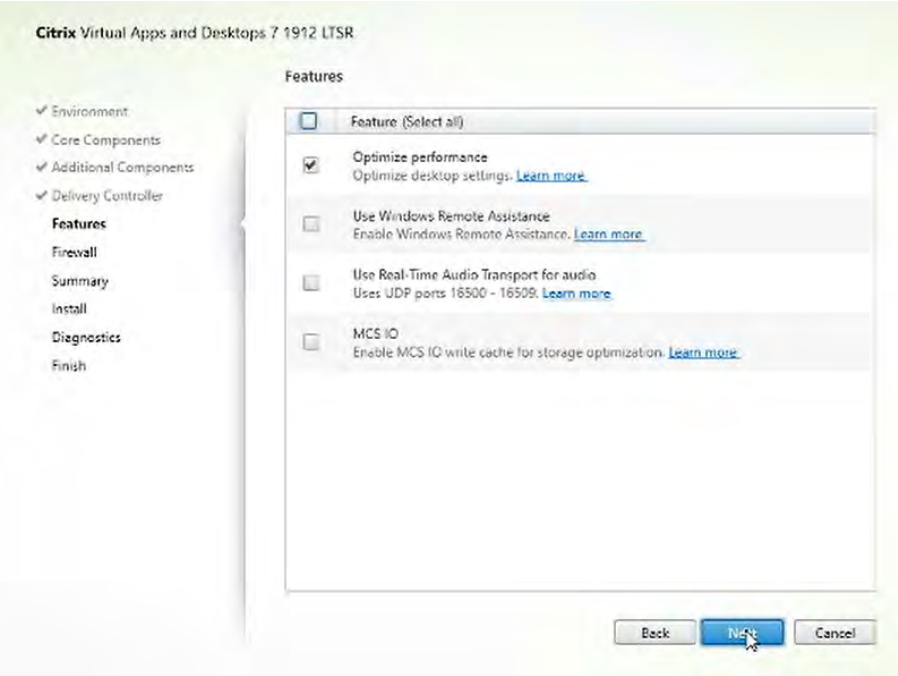

12. The Firewall window allows you to configure Windows Firewall. Select the **Automatically** radio button and click **Next**.

13. Review your configuration on the Summary window and click **Install.**

14. Accept any ***Reboot Prompts*** and reconnect to the server.

15. In the Diagnostics window, uncheck ***Collect diagnostic information***.

16. Click **Next** to continue.
17. On the Finish window, click **Finish** to complete the install and allow the server to restart.

8.5 Additional Virtual Machine Settings

Perform the following additional tasks on the virtual machine as required:

Note: To evaluate browser based HTML5 applications, consider using newer browsers that utilize hardware acceleration within virtual desktop environments.

- Turn Off Windows Firewall for all network types.

CAUTION: THESE INSTRUCTIONS ASSUME THAT THE VM IS BEING USED AS A PROOF-OF-CONCEPT ONLY AND THAT DISABLING THE FIREWALL WILL THEREFORE POSE ONLY A MINIMAL SECURITY BREACH. ALWAYS FOLLOW YOUR ESTABLISHED SECURITY PROCEDURES AND BEST PRACTICES WHEN SETTING UP SECURITY FOR A PRODUCTION MACHINE OR ANY ENVIRONMENT THAT CAN BE ACCESSED FROM OUTSIDE YOUR NETWORK.

- Shut down the virtual machine once this is completed.
- **Close the remote console; this will not be functional when vGPU is configured.**

Note: Take a snapshot of the virtual machine to preserve your work. Label this snapshot pre-vGPU and revert to it if you encounter any problems going forward, such as driver issues.

8.6 Installing NVIDIA Driver in Windows Virtual Desktop

Start the appropriate driver for the GPU inside the VM. The following example shows the specific case for installation of the NVIDIA drivers with a M10 GPU.

1. Copy the 32-bit or 64-bit NVIDIA Windows driver package to the VM, open the zip file, and run setup.exe.

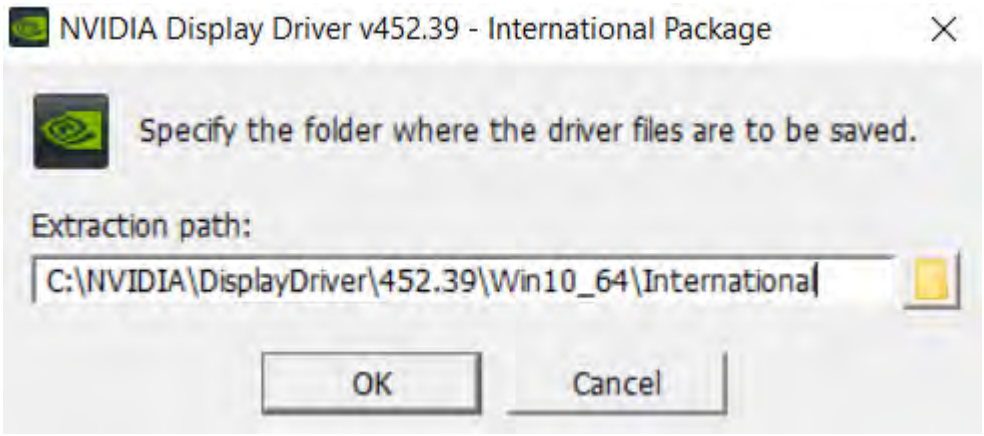

2. After a brief System Check, read through the EULA and select Agree and Continue.

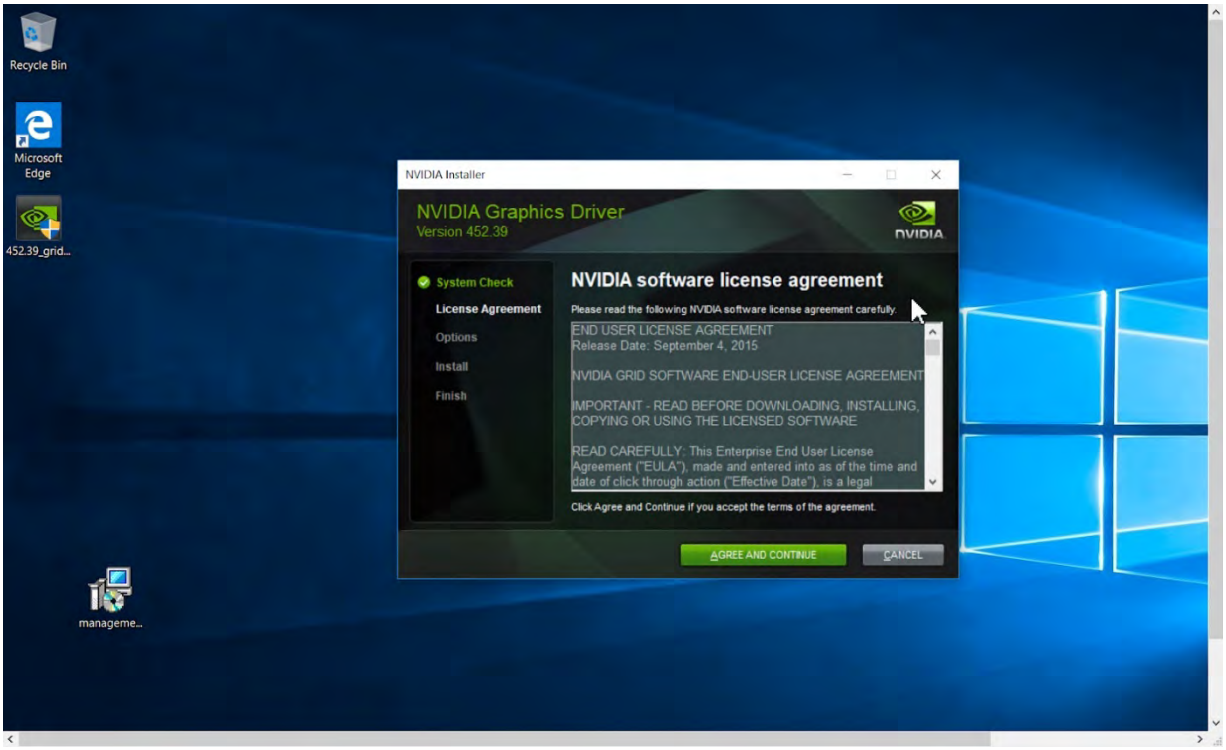

3. Select the desired installation options and navigate to Next to continue.

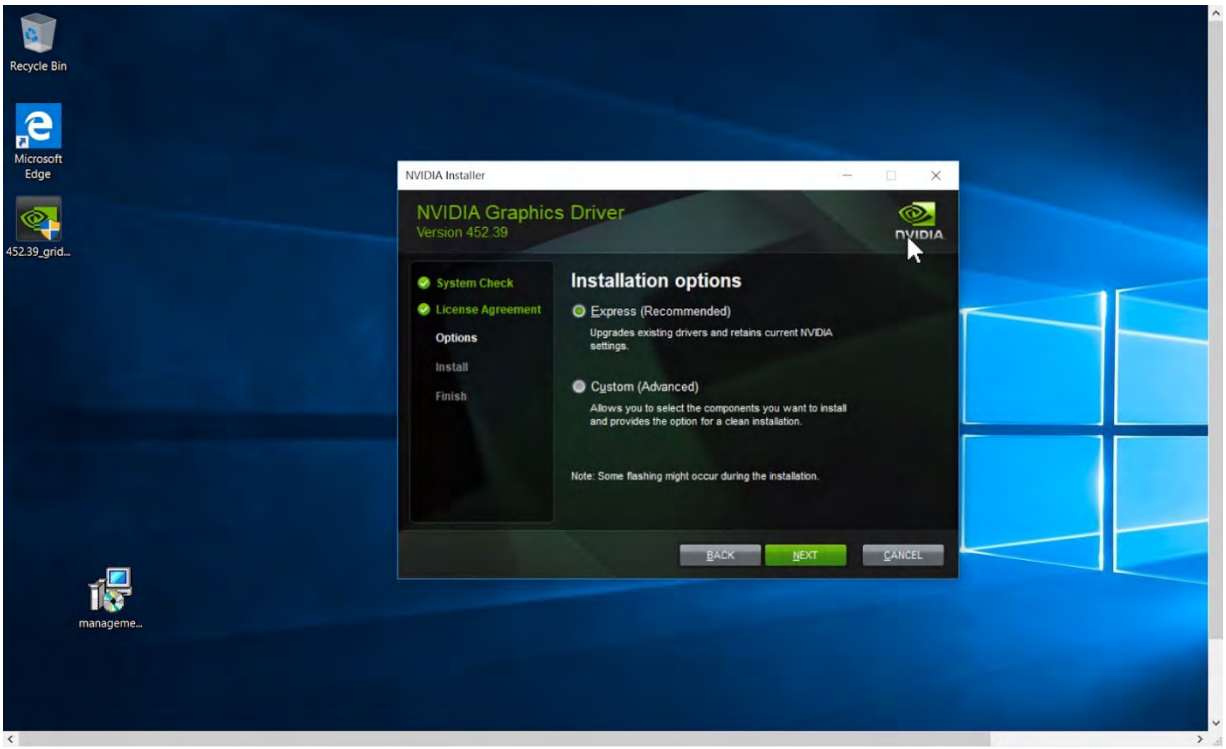

4. Wait for the installation to finish. Select Next to continue.

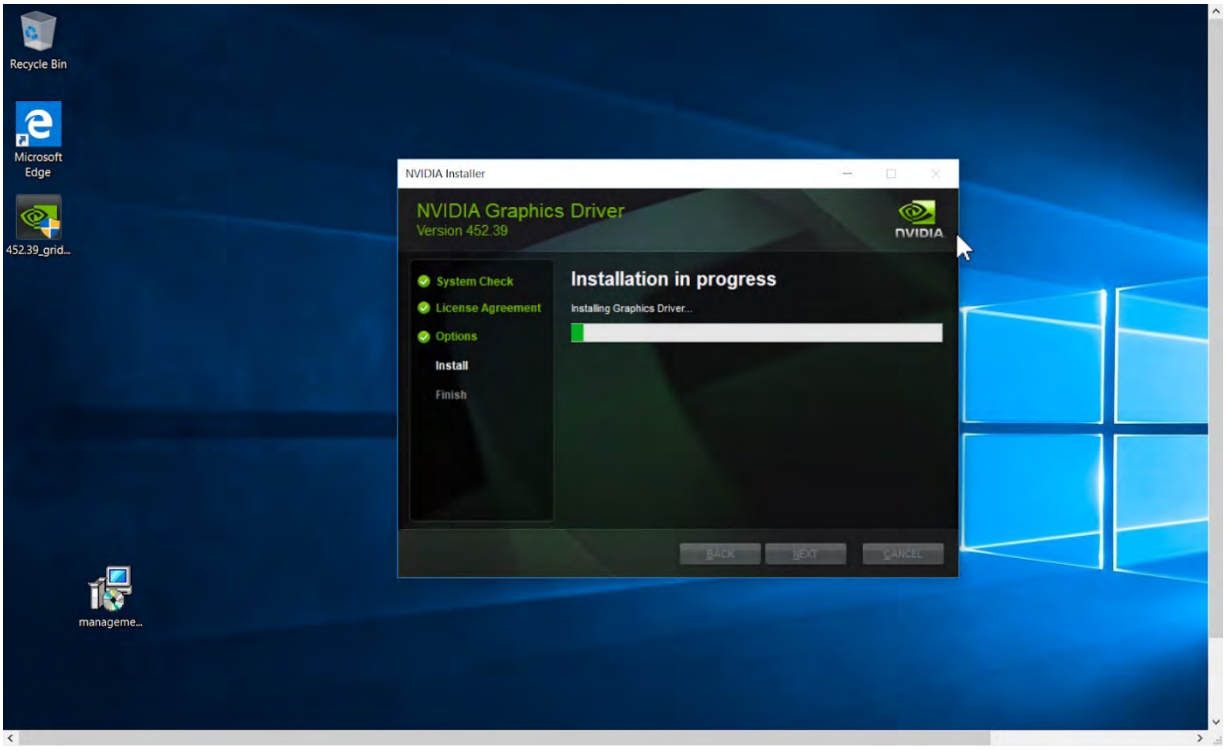

5. After the driver installation has completed, you may be prompted to reboot the VM. Select Restart Now to restart the VM. When the VM starts, it will boot to a Windows desktop.

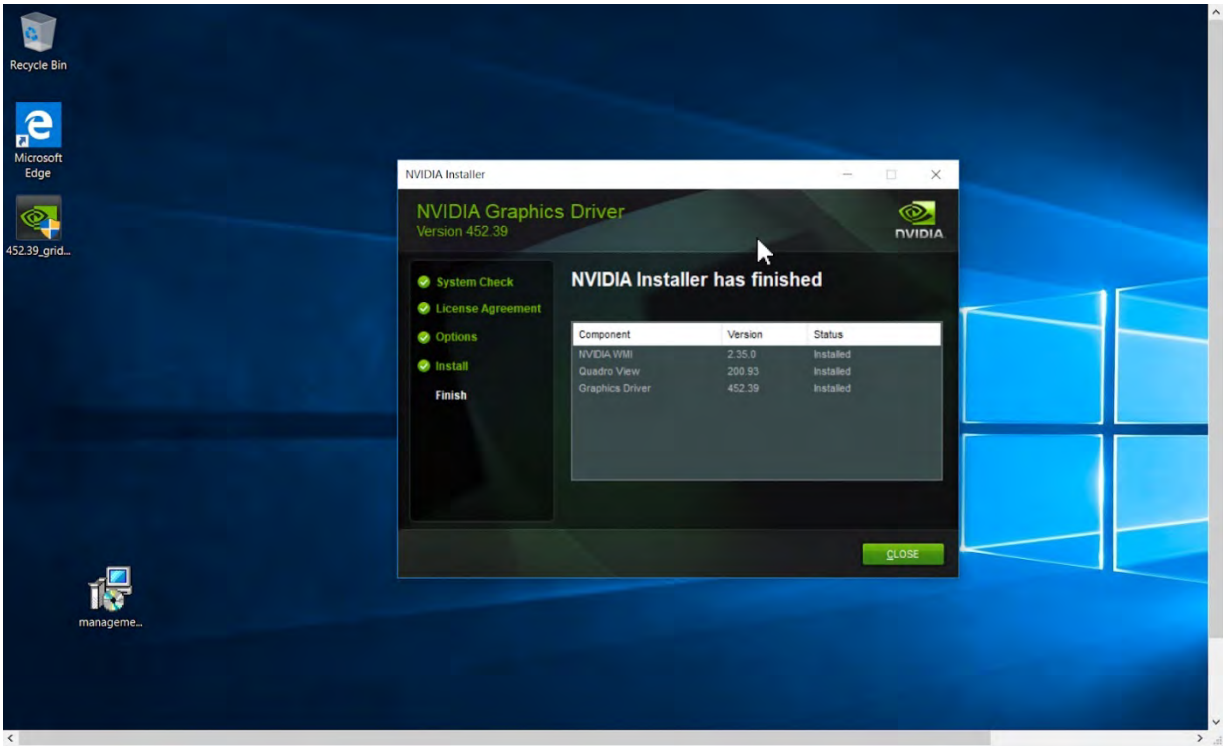

6. To verify that the NVIDIA driver is running, right-click on the desktop and select NVIDIA Control Panel.

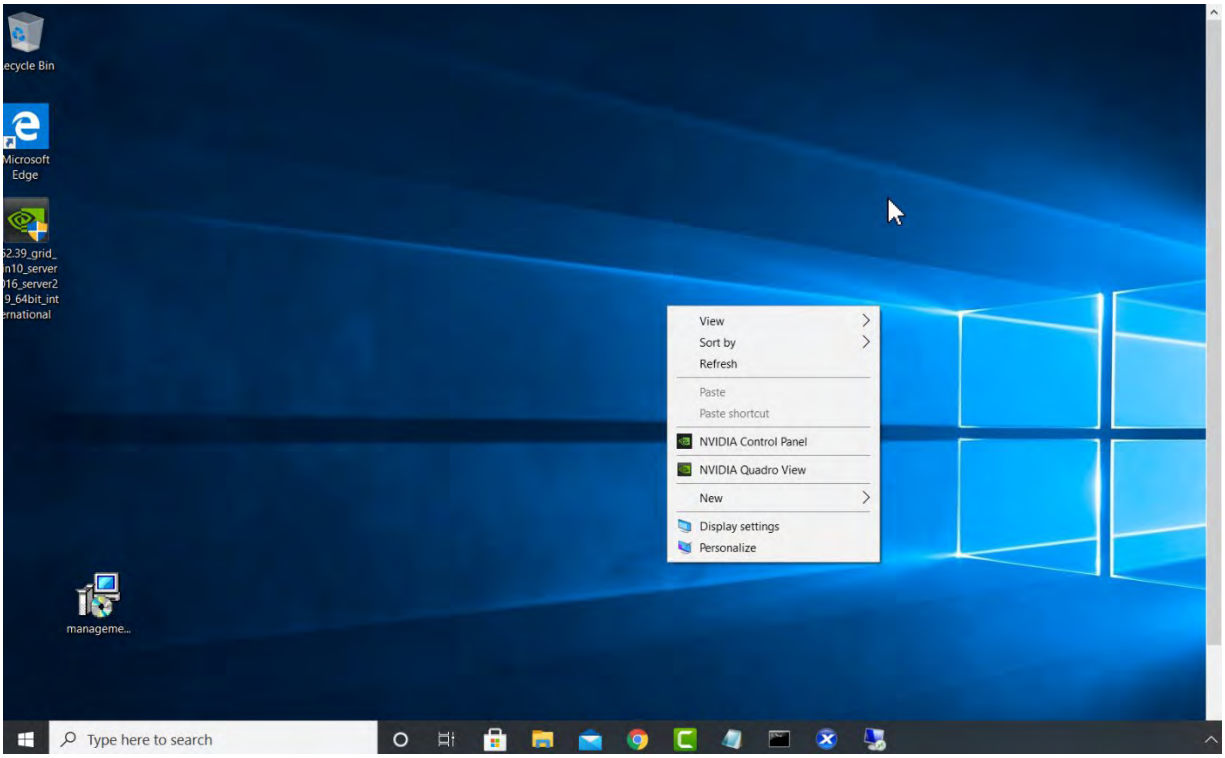

7. In the NVIDIA Control Panel, select System Information. This interface shows the GPU Type in use by the VM, its features, and the NVIDIA driver version in use.

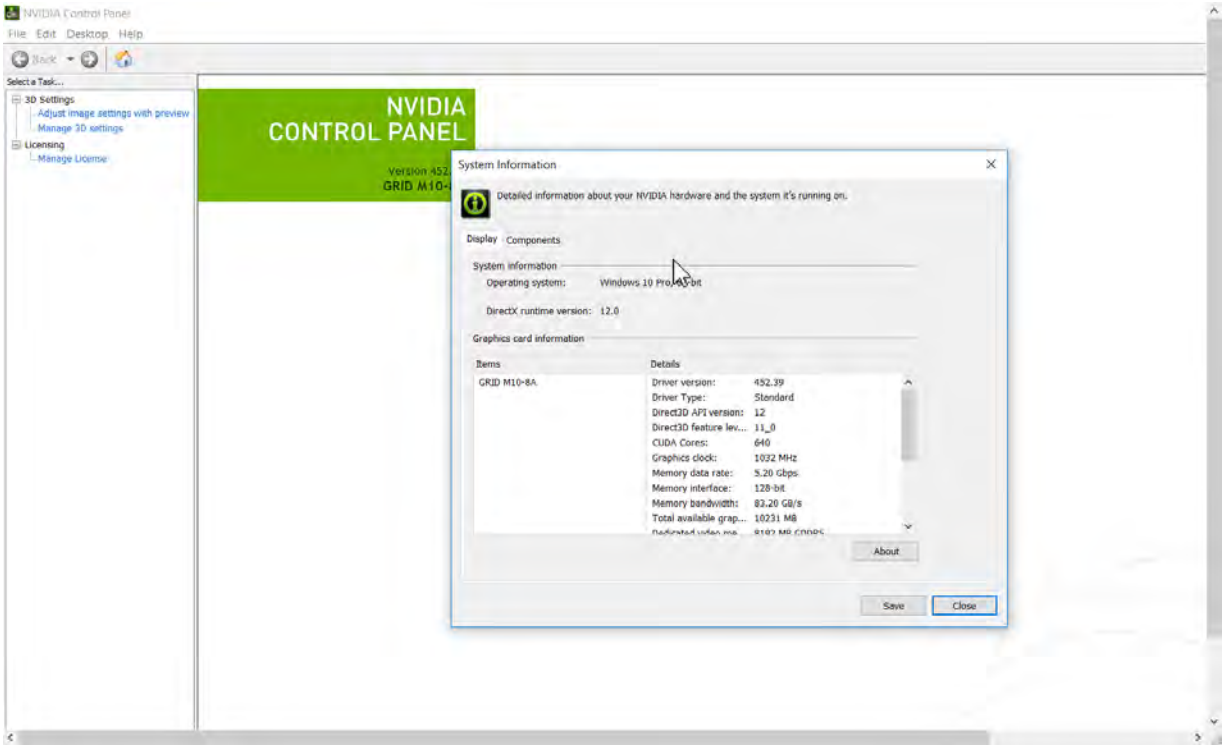

8.7 Licensing NVIDIA vGPU (Update 11.0)

NVIDIA vGPU is a licensed product. When booted on a supported GPU, a vGPU runs at reduced capability until a license is acquired. The performance of an unlicensed vGPU is restricted as follows.

Elapsed Time	Performance Degradation
20 minutes	• Frame rate is capped at 15 frames per second. • The performance of applications and processes that use CUDA is degraded.
24 hours	• Frame rate is capped at 3 frames per second. • CUDA stops working and CUDA API function calls fail. • GPU resource allocations for a vGPU are limited, which will prevent some applications from running correctly.

These restrictions are removed when a license is acquired. After you license NVIDIA vGPU, the VM that is set up to use NVIDIA vGPU is capable of running all DirectX (up to and including DirectX12 and DX12-Raytracing on Turing architecture cards),OpenGL & Vulkan graphics applications.

If licensing is configured, the virtual machine (VM) obtains a license from the license server when a vGPU is booted on these GPUs. The VM retains the license until it is shut down. It then releases the license back to the license server. Licensing settings persist across reboots and need only be modified if the license server address changes, or the VM is switched to running GPU pass through.

Note: For complete information about configuring and using NVIDIA vGPU software licensed features, including vGPU, refer to Virtual GPU Client Licensing User Guide.

8.7.1.1 Licensing NVIDIA vGPU on Windows

1. Open NVIDIA Control Panel:

   a) Right-click on the Windows desktop and select **NVIDIA Control Panel** from the menu.

   b) Open Windows Control Panel and double-click the **NVIDIA Control Panel** icon.

2. In NVIDIA Control Panel, select the **Manage License** task in the **Licensing** section of the navigation pane.

Note: If the **Licensing** section and Manage License task are not displayed in NVIDIA Control Panel, the system has been configured to hide licensing controls in NVIDIA Control Panel. For information about registry settings, see Virtual GPU Client Licensing User Guide.

The Manage License task pane shows that NVIDIA vGPU is currently unlicensed.

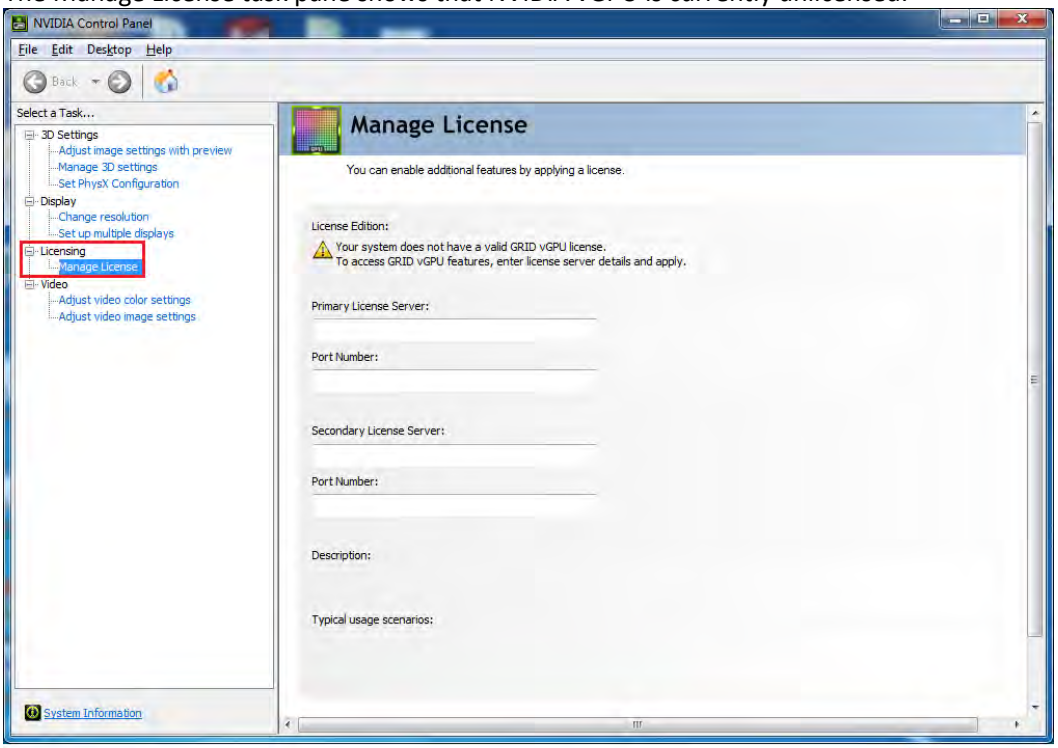

3. In the **Primary License Server** field, enter the address of your primary NVIDIA vGPU software License Server. The address can be a fully qualified domain name such as gridlicense1.example.com, or an IP address such as 10.31.20.45. If you have only one license server configured, enter its address in this field.
4. Leave the **Port Number** field under the **Primary License Server** field unset. The port defaults to 7070, which is the default port number used by NVIDIA vGPU software License Server.
5. In the **Secondary License Server** field, enter the address of your secondary NVIDIA vGPU software License Server. If you have only one license server configured, leave this field unset. The address can be a fully qualified domain name such as gridlicense2.example.com, or an IP address such as 10.31.20.46.
6. Leave the **Port Number** field under the **Secondary License Server** field unset. The port defaults to 7070, which is the default port number used by NVIDIA vGPU software License Server.
7. Click **Apply** to assign the settings. The system requests the appropriate license for the current vGPU from the configured license server.
8. The vGPU within the VM should now exhibit full frame rate, resolution, and display output capabilities. The VM is now capable of running the full range of DirectX and OpenGL graphics applications.
9. If the system fails to obtain a license, see Virtual GPU Client Licensing User Guide for guidance on troubleshooting.

Chapter 9. Creating a Citrix Machine Catalog

This chapter describes the following:

- Creating a Citrix Machine Catalog using Citrix Machine Creation Services (MCS) to deploy a Virtual Desktop
- Creating a Citrix Machine Catalog using Citrix Machine Creation Services (MCS) to deploy a Virtual Application

To create a pool of virtual machines for users to remotely access, a Citrix Machine Catalog must be configured and deployed via Citrix Studio. The Citrix Virtual Delivery Agent configured in the previous chapters to build out a pool of VMs. This chapter outlines how to create the pool of VMs by building a Citrix Machine Catalog with Citrix Machine Creation Services (MCS).

Note: Other machine deployment technologies are available like Citrix Provisioning (PVS) and is outside the scope of this document.

Please refer to Citrix product documentation for additional information regarding machine catalog and deliver group creation.

Part of creating a Citrix Machine Catalog involves choosing the correct type of operating system for your deployment. Citrix groups operating systems into three categories:

- Single-Session OS
  - Single-Session OS deployments allow for publishing only the Desktop or only the Application, but not within a single Citrix Machine Catalog.
- Multi-Session OS
  - Multi-Session OS deployments allow for the publishing of multiple desktops or multiple applications, as well as both desktops and applications within a single Citrix Machine Catalog.
- Remote PC Access.
  - Remote PC Access is outside the scope of a vGPU deployment.

Additionally, consult Citrix and your ISV partners to determine the best deployment method for your environment.

9.1 Creating a Citrix Machine Catalog for Virtual Desktops and Apps

1. Log on to the Citrix Delivery Controller and Launch **Citrix Studio** from the Windows Start Menu

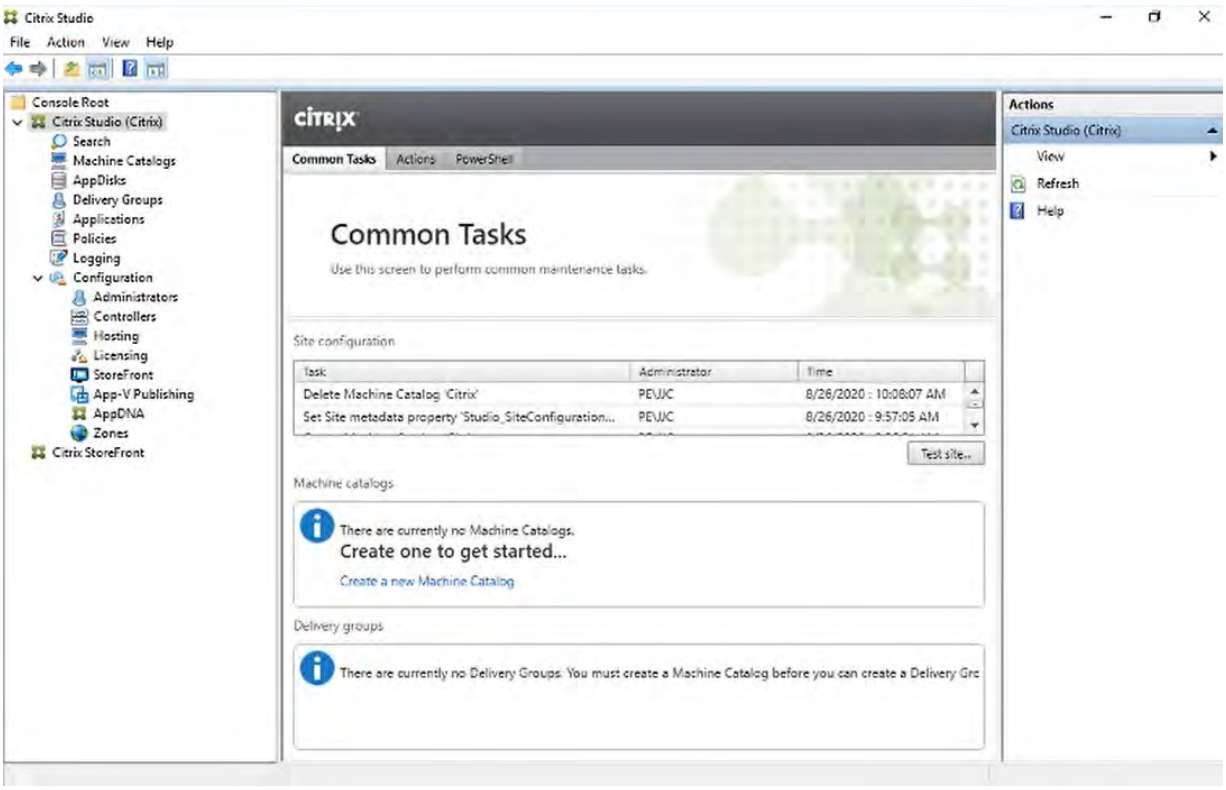

2. On the left menu panel, click **Machine Catalogs**.
3. Under the Actions Menu on the Right, click **Create Machine Catalog**.

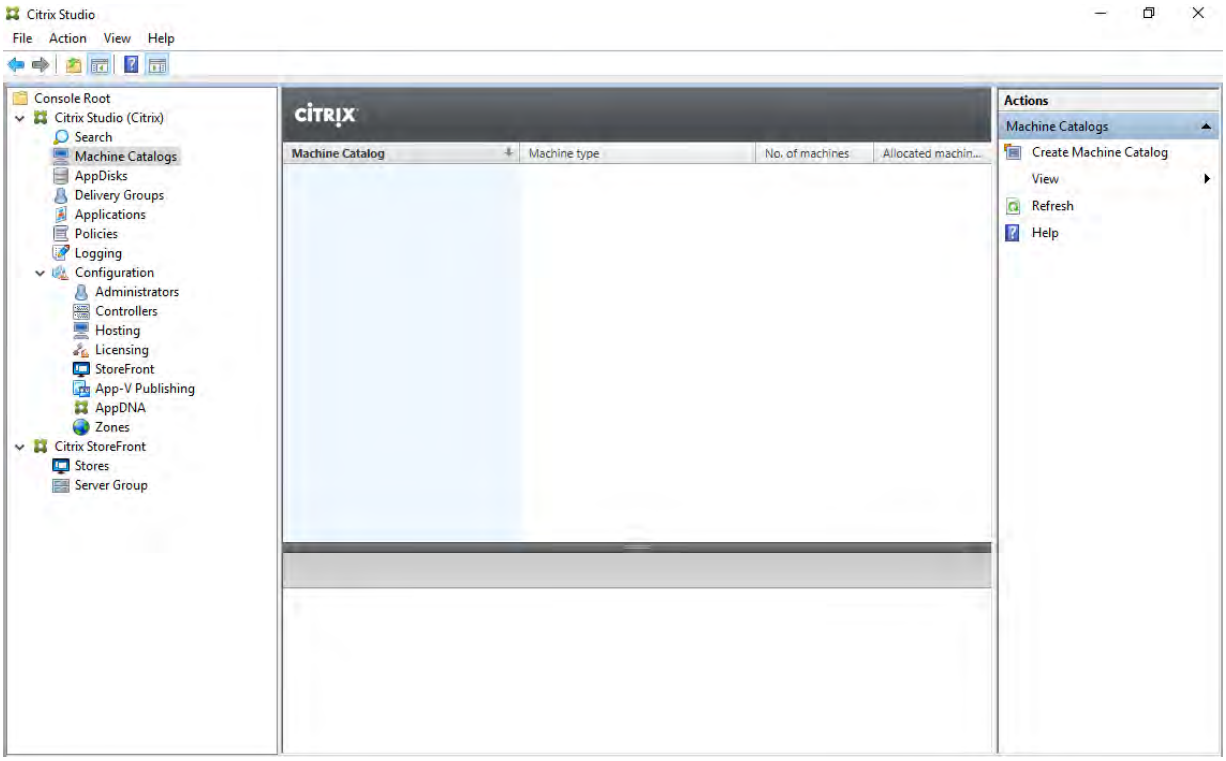

4. On the Introduction page click **Next**.
5. On the Operating System select either the **Single-session OS** radio button or the **Multi-session OS** radio button, in accordance with your VDA's operating system and NVIDIA vGPU licensing level.

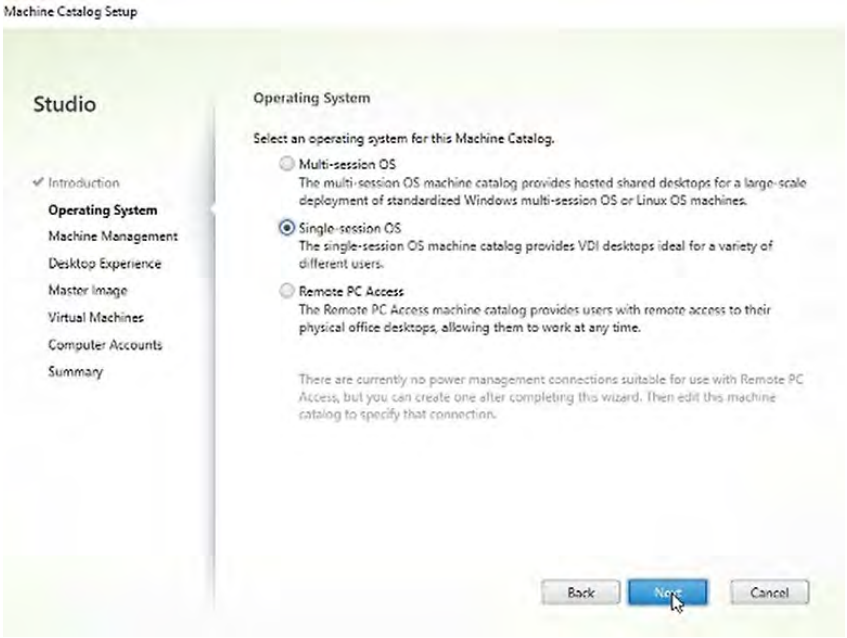

Citrix groups operating systems into three categories:

- Single-Session OS deployments allow for publishing only the Desktop or only the Application, but not within a single Citrix Machine Catalog.
- Multi-Session OS deployments allow for the publishing of multiple desktops or multiple applications, as well as both desktops and applications within a single Citrix Machine Catalog.
- Remote PC Access is outside the scope of a vGPU deployment. Additionally, consult Citrix and your ISV partners to determine the best deployment method for your environment.

Please refer to Citrix product documentation for additional information regarding machine catalog creation.

6. On the Machine Management Page, ensure the **Machines that are power managed (for example, virtual machines or blade PCs)** radio button is selected, as well as the **Another service or technology** radio button and click **Next**.

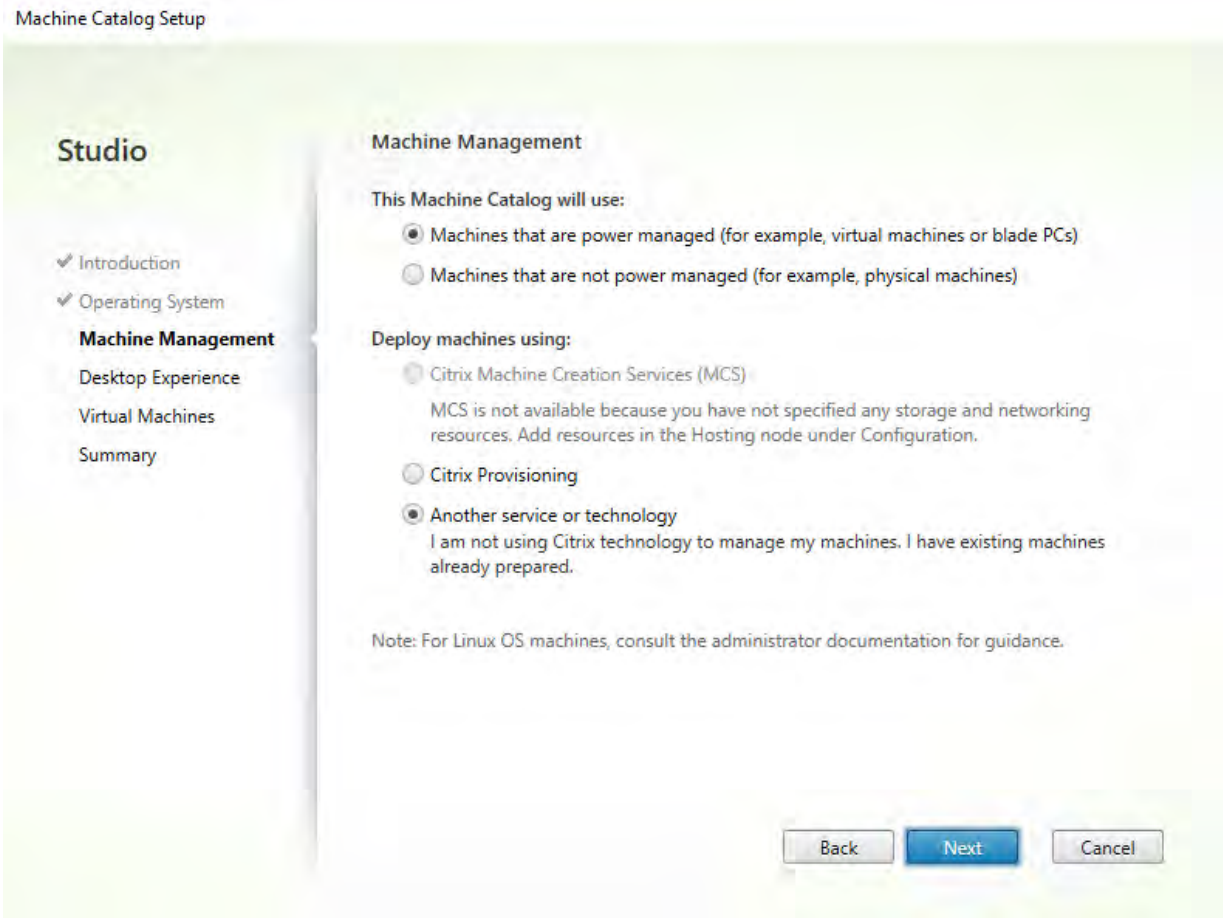

7. On the Desktop Experience window, choose which type of desktop experience you would like. For purposes of POC/trail, chose **I want users to connect to the same (static) desktop each time they log on.**

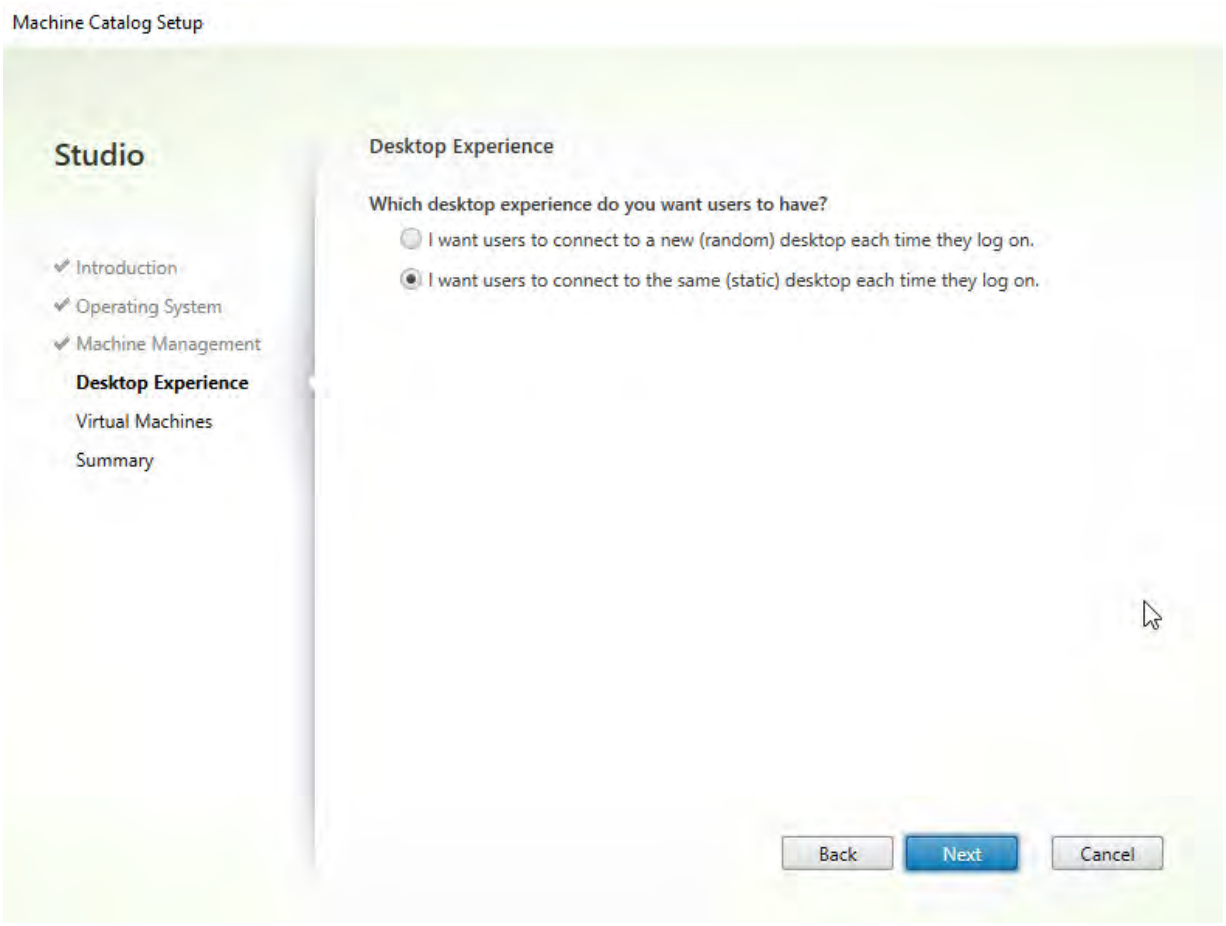

8. On the Virtual Machines and Users window, select your previously created master image VM with the Citrix VDA installed by clicking, **Add VMs…**.

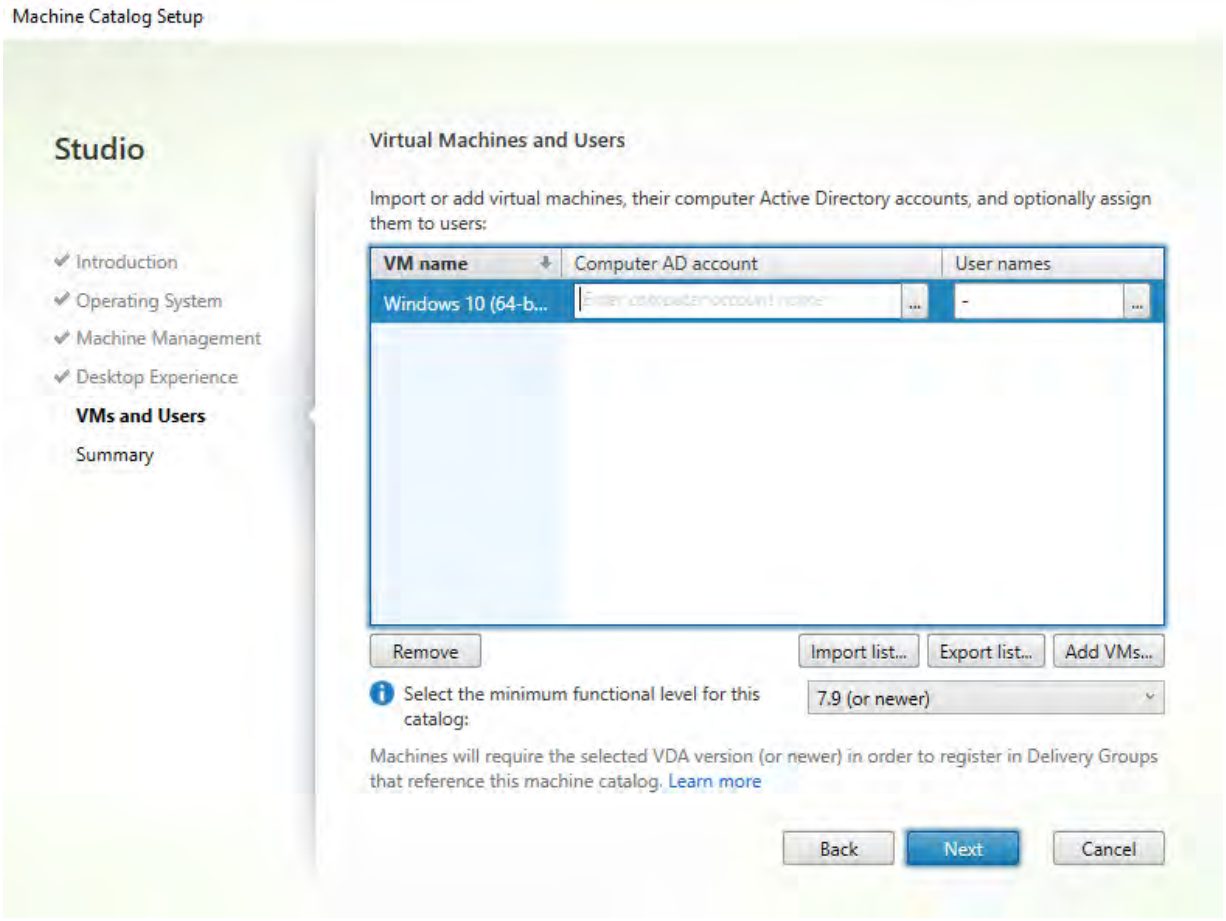

9. Provide the Computer AD account name and usernames or accounts that will require access.
10. In the **Select the minimum functional level for this catalog:** drop-down menu ensures the latest functional level available is selected and click, **Next**.

> Selecting a minimum function level lower than is available will result in a loss of new feature sets that may benefit or be required for you deployment needs.

11. On the Summary window, type a name for your Machine Catalog in the **Machine Catalog name:** text field.

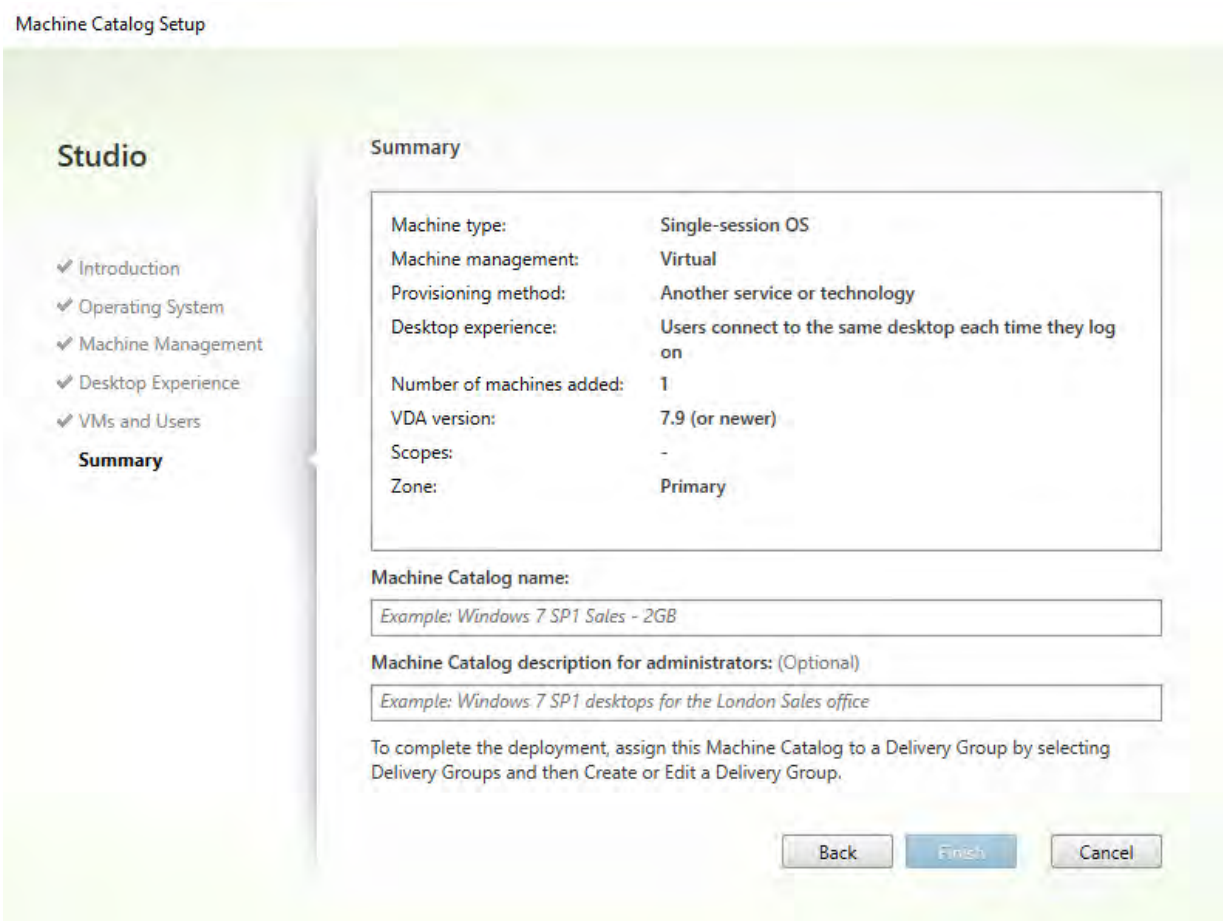

12. Click **Finish**.

Chapter 10. Creating a Citrix Delivery Group

This chapter describes the following:

- Creating a Citrix Delivery Group to Deploy a Virtual Desktop
- Creating a Citrix Delivery Group to Deploy a Virtual Application

For users to remotely access virtual applications and desktops, a Citrix Delivery Group must be configured and deployed via Citrix Studio. The Citrix Virtual Delivery Agent and the Citrix Machine Catalog created and configured in the previous chapters to build out the Delivery group.

10.1 Creating a Citrix Delivery Group for Virtual Desktops

Now that you have a Citrix Machine Catalog created, the next step is to create a Delivery Group so that users can access the resource, in this case a Virtual Desktop.

1. Log on to the Citrix Delivery Controller and Launch **Citrix Studio** from the Windows Start Menu
2. On the left menu pane, click **Delivery Groups**.
3. On the right Actions menu, click **Create Delivery Groups.**

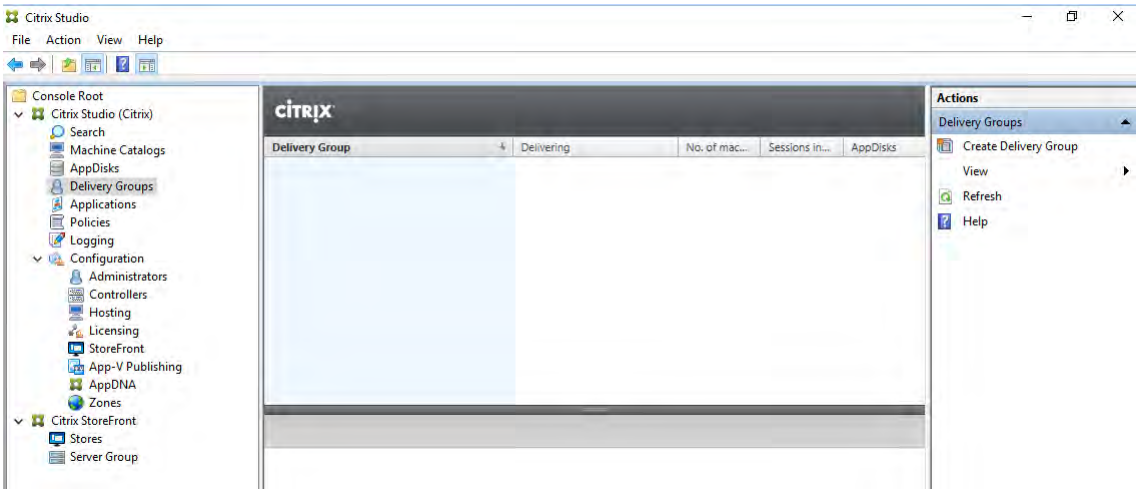

4. On the Introduction window, click **Next**.
5. On the Machines window, select the Machine Catalog that you created in the previous section and click **Next.**

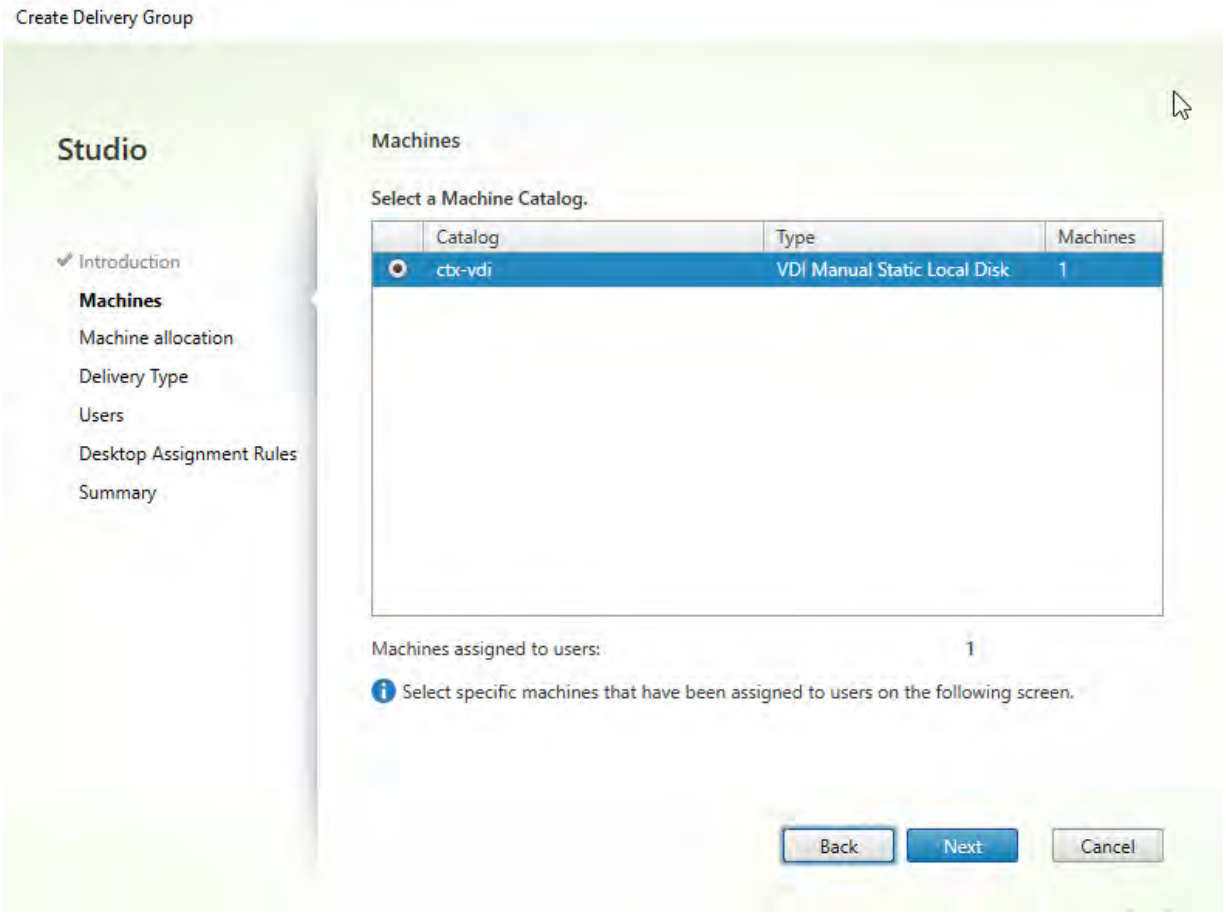

6. On the Machine allocation window, enter your machine name and the usernames of the accounts that will require access to the desktop and click **Next.**

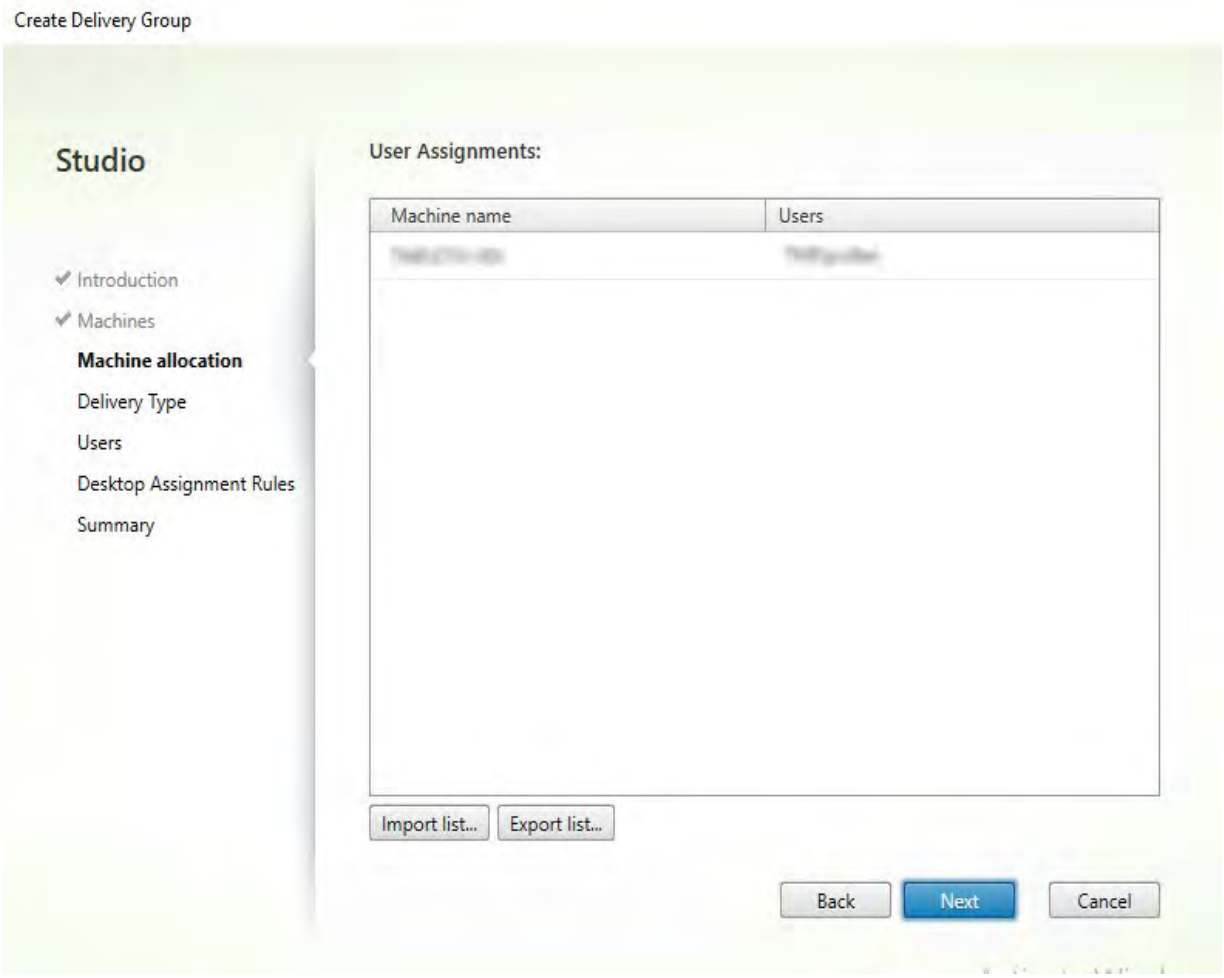

7. On the Delivery Type window, leave the **Desktops** radio button selected and click **Next.**

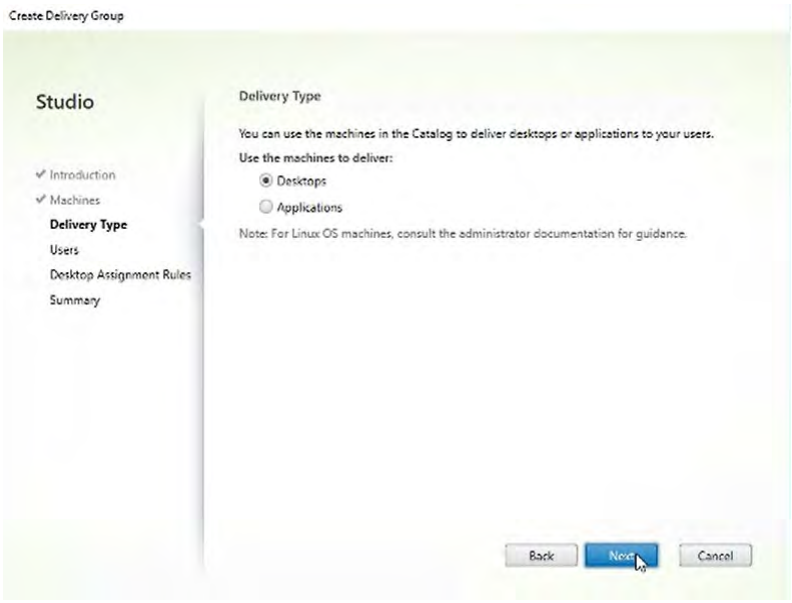

Because we previously chose a Machine Catalog type of "Single-Session OS" we must choose between deploying an application or a desktop. Had we created a Machine Catalog with a type of "Multi-Session OS," we would have had the option to deploy a desktop and an application.

Please refer to Citrix product documentation for additional information regarding machine catalog and deliver group creation.

Note: For NVIDIA vApps you can use Citrix Virtual Apps with a multi-session OS, and for NVIDIA vPC and RTX vWS you can use Citrix Virtual Desktop and is limited to single-session OS.

8. On the Users window, specify which users can access this delivery group. For POC/trial purposes leave the **Allow any authenticated users to use this Delivery Group**. radio button selected.

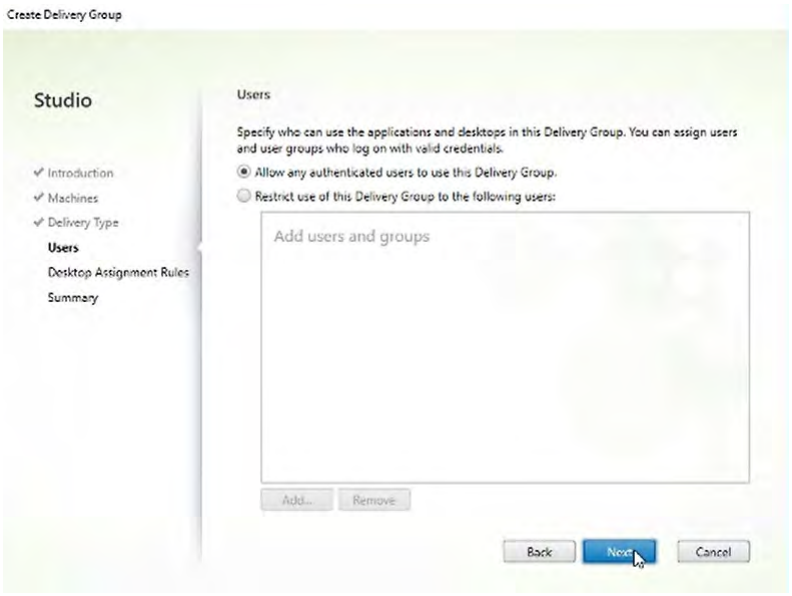

9. Click **Next**
10. On the Desktop Assignment Rules window, click **Add…**

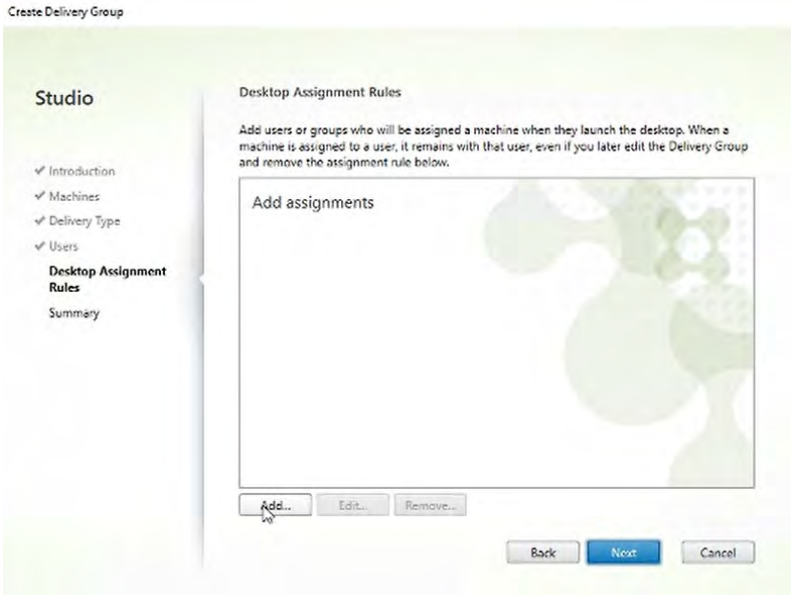

11. The Add Desktop Assignment Rule pops up. In the **Display name:** text field type a name for the Desktop that users will see.
12. For POC/trial purposes, leave the **Allow everyone with access to this Delivery Group to have desktop assigned** radio button selected.

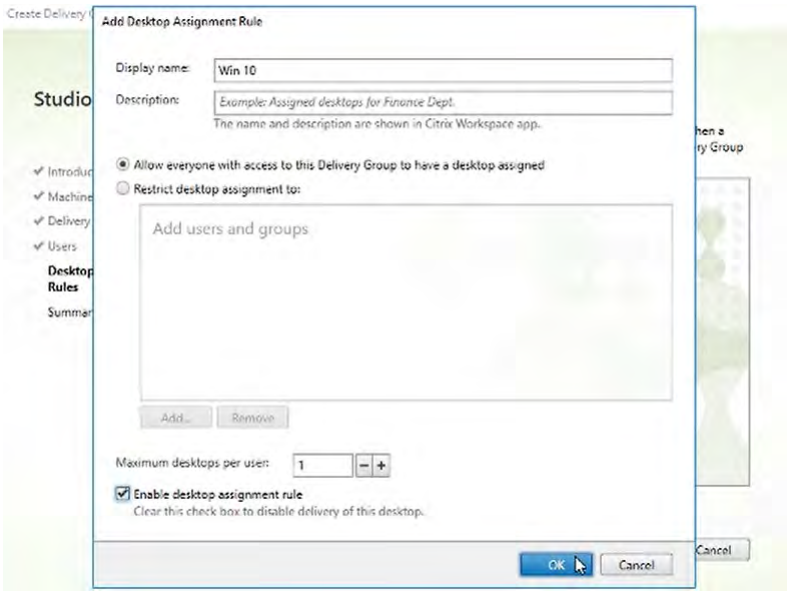

13. Select the **Enable desktop assignment rule** check box and click **OK.**
14. Click, **Next**.

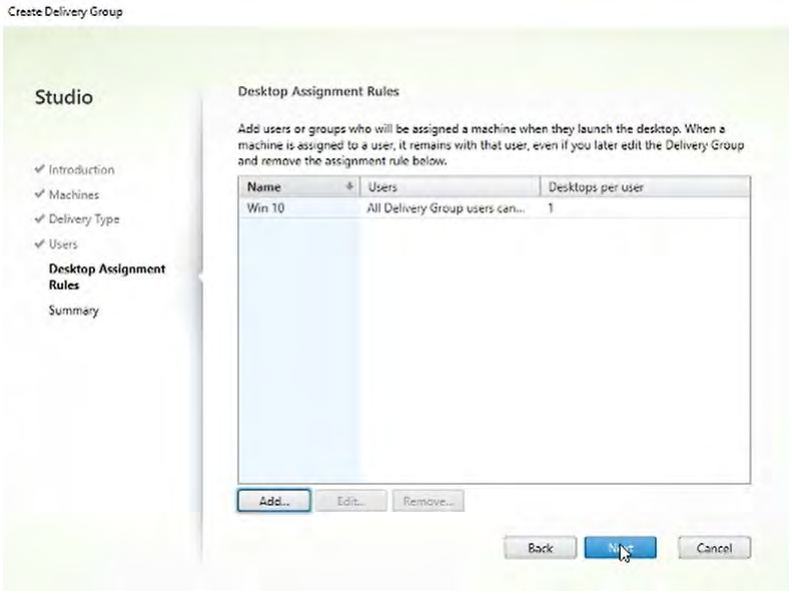

15. On the Summary window, type a name for the Delivery Group in the **Delivery Group name:** text field and click **Finish** to complete the creation of the Delivery Group.

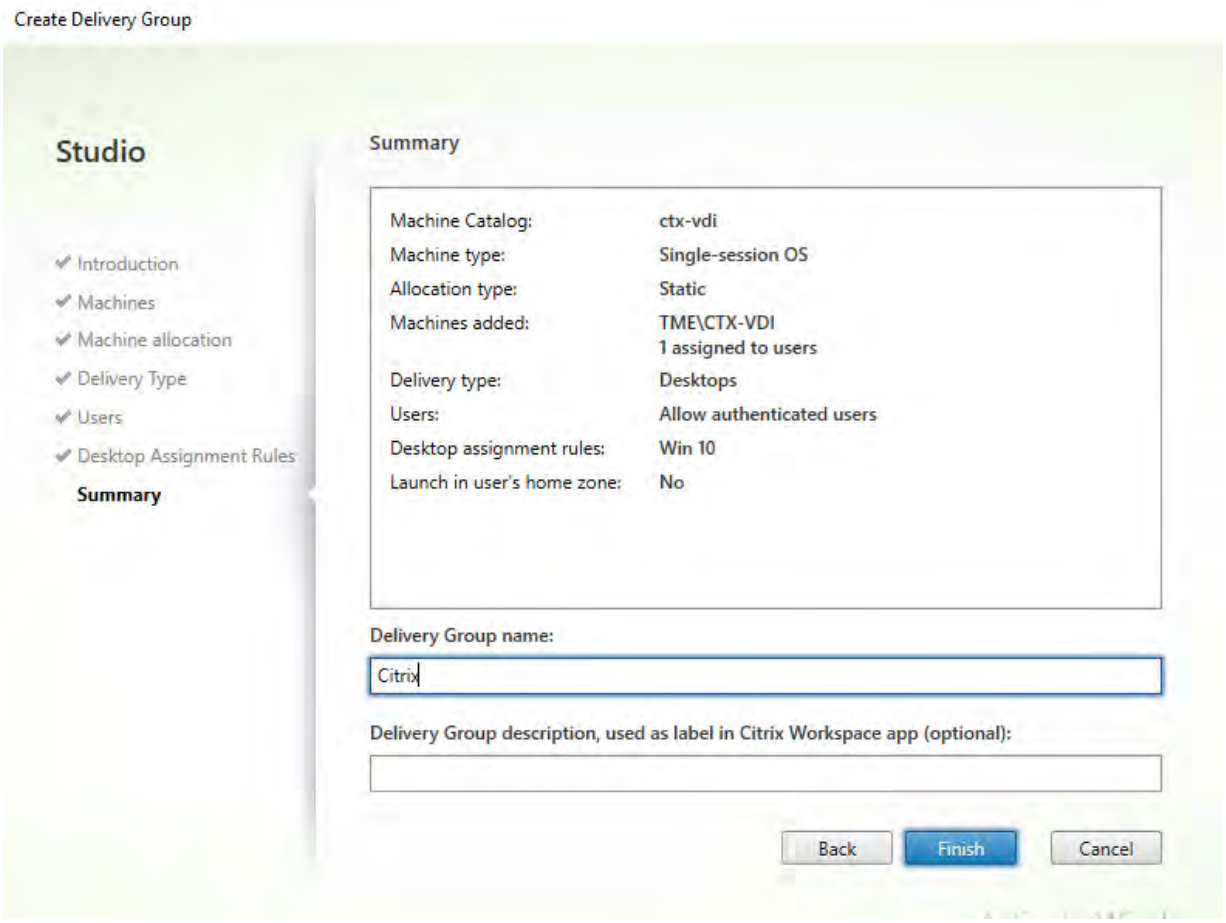

10.2 Creating a Citrix Delivery Group for Virtual Applications

Now that you have a Citrix Machine Catalog created, the next step is to create a Delivery Group so that users can access the resource, in this case a Virtual Application.

1. Log on to the Citrix Delivery Controller and Launch **Citrix Studio** from the Windows Start Menu
2. On the left menu pane, click **Delivery Groups**.
3. On the right Actions menu, click **Create Delivery Groups.**

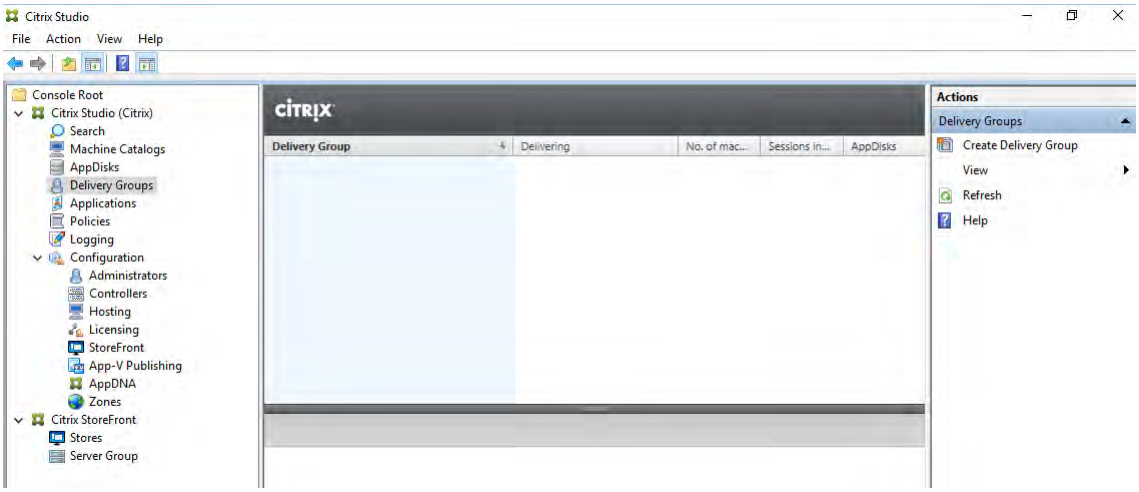

4. On the Introduction window, click **Next**.
5. On the Machines window, select the Machine Catalog that you created in the previous section and choose how many machines will be included in this delivery group. For POC/trial purposes, we choose 1 machine.

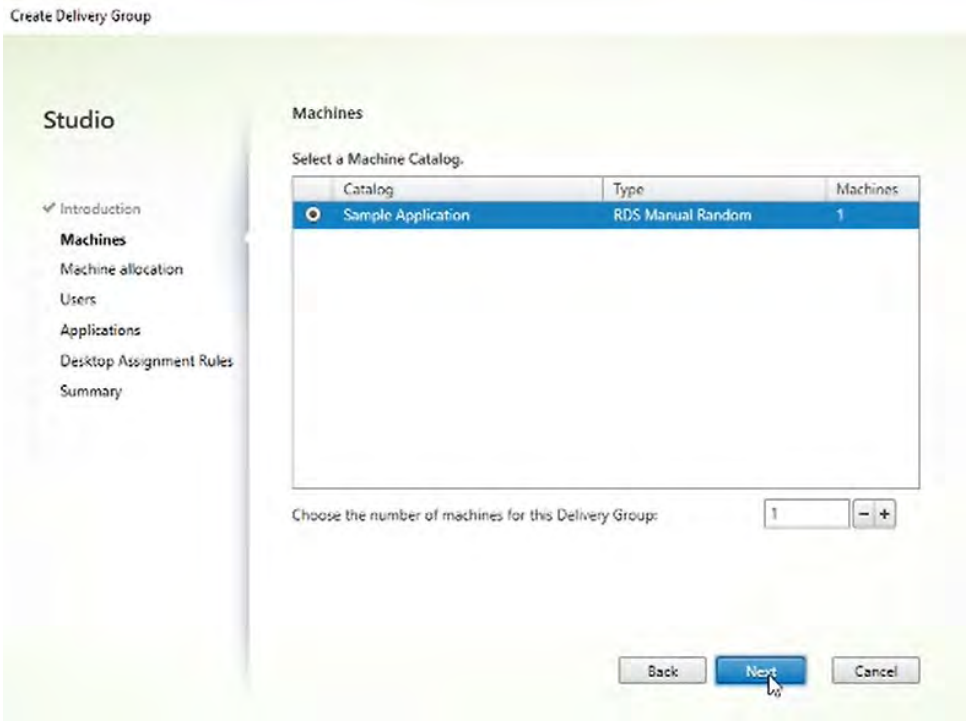

6. Click **Next.**
7. On the Users window, specify which users can access this delivery group. For POC/trial purposes leave the **Allow any authenticated users to use this Delivery Group**. radio button selected.

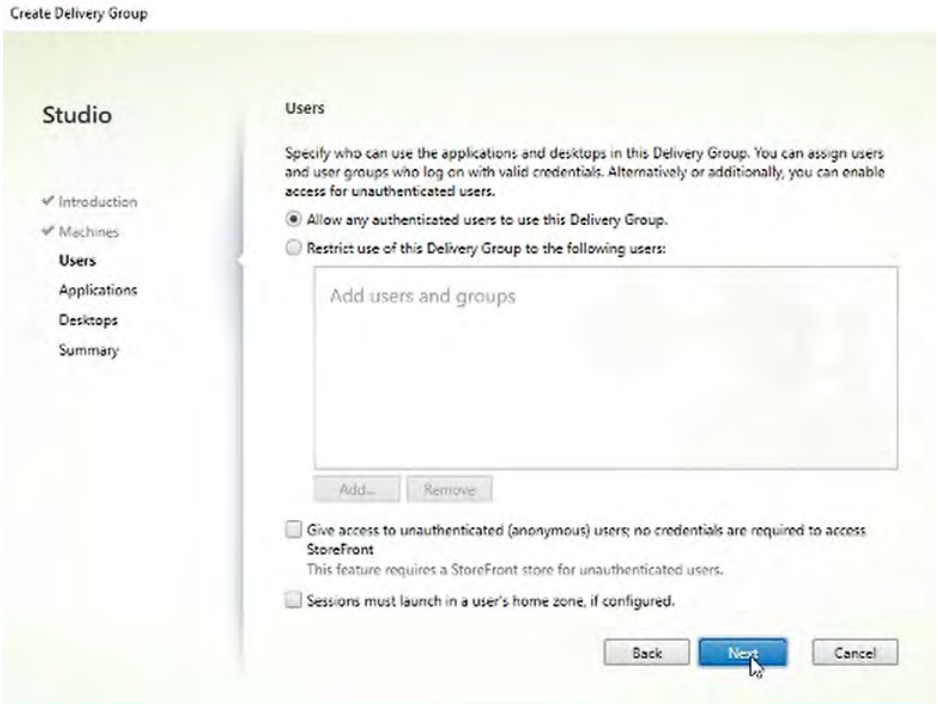

8. On the Application window, select the **Add…** dropdown menu.

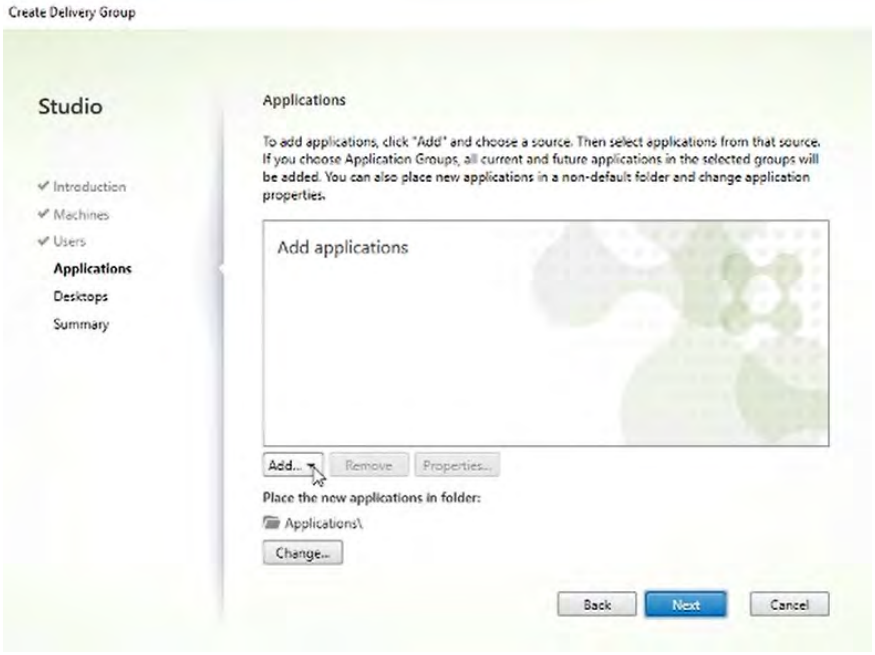

9. Select **Manually…**

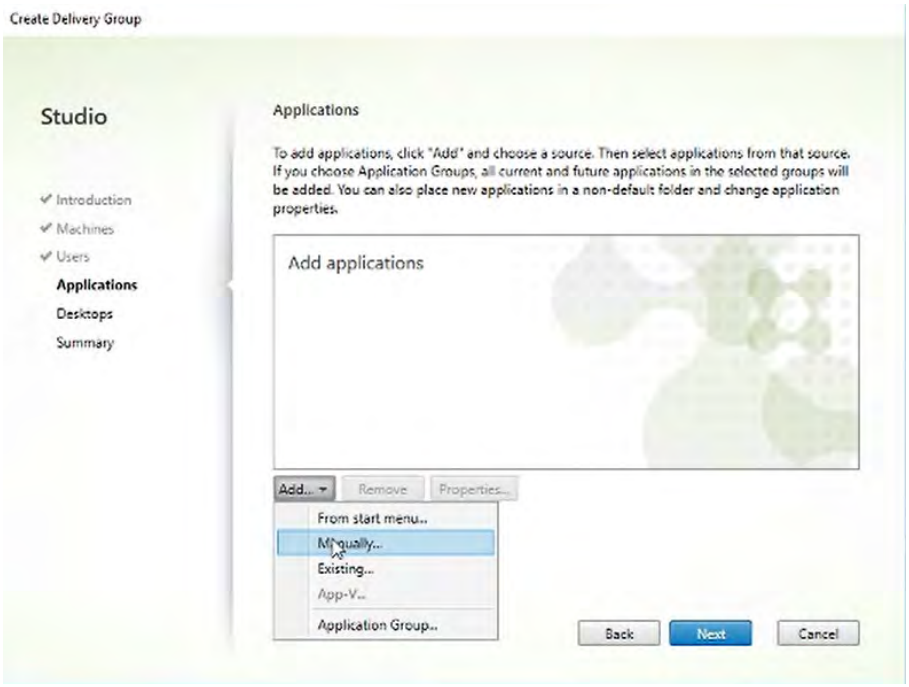

10. On the Add Applications Manually popup window, type a path to the executable, an application name for users and an application name for administrators, then Click **OK.**

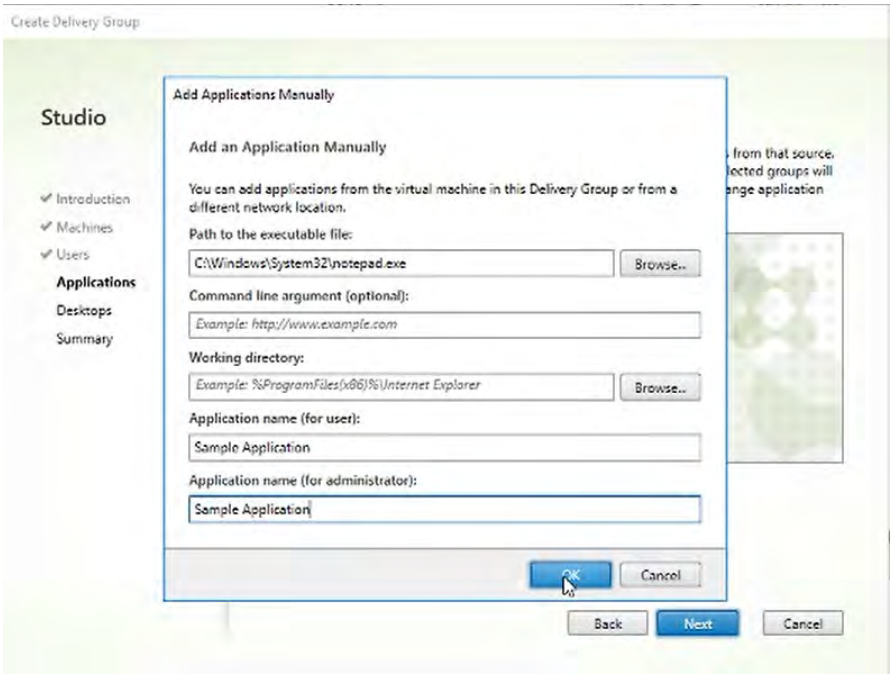

11. Click **Next**

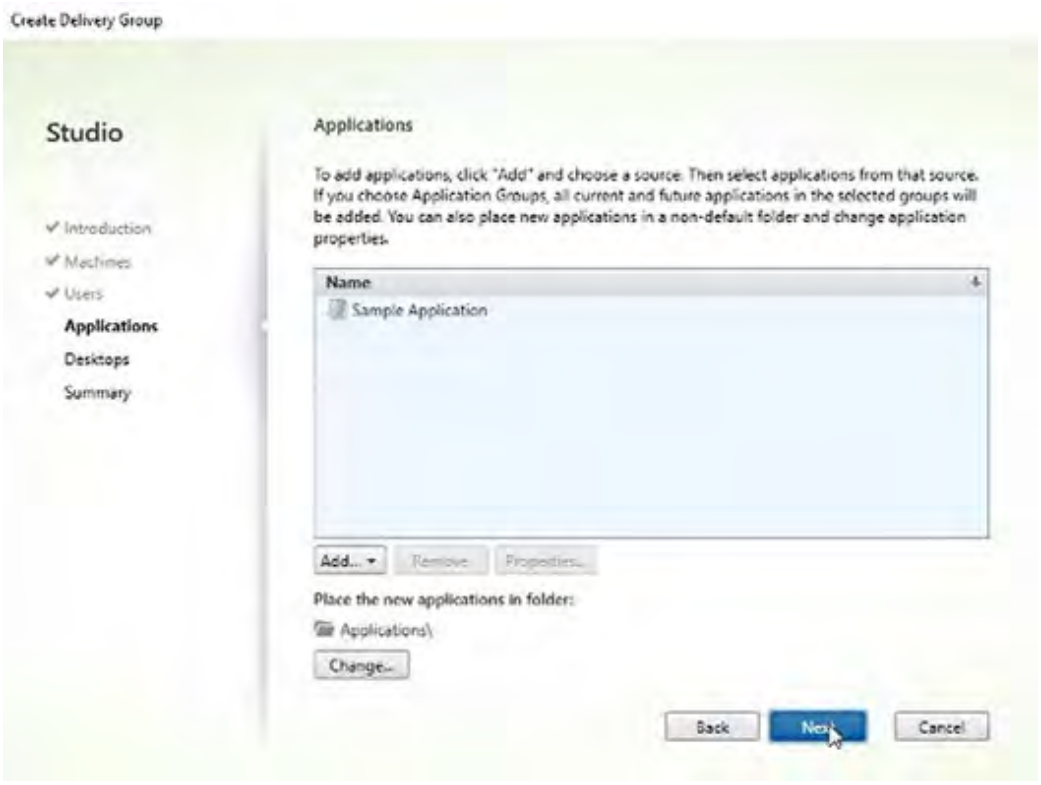

12. On the Desktops window, click **Next.**

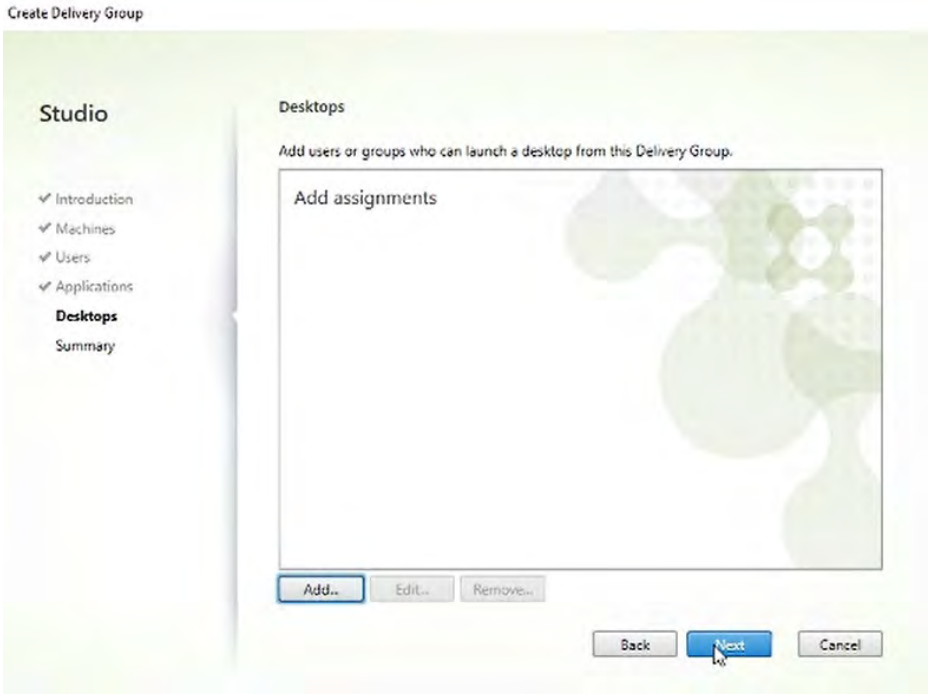

Because we previously chose a Machine Catalog type of “Multi-Session OS” we have the option to deploy both Applications and Desktops. Had we created a Machine Catalog with type of “Single-Session OS,” we would have only been able to deploy an application or a desktop. This section focuses on Application deployment, so we choose to not deploy a desktop.

Please refer to Citrix product documentation for additional information regarding machine catalog and deliver group creation.

13. On the Summary window, type a name for the Delivery Group in the **Delivery Group name:** text field and click **Finish** to complete the creation of the Delivery Group.

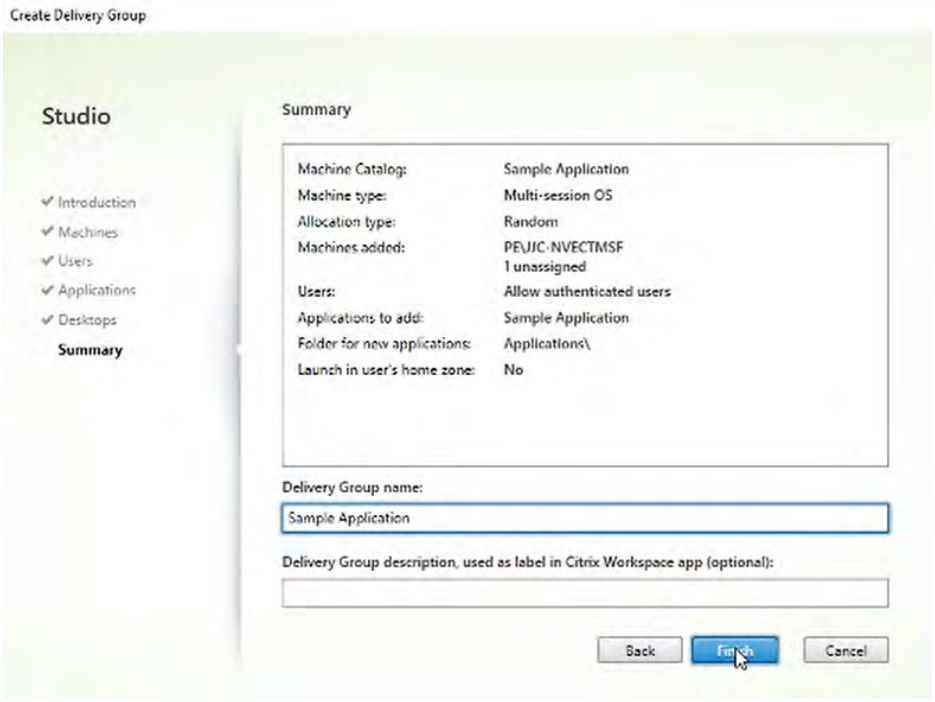

Chapter 11. Creating Citrix Policies for NVIDIA vGPU

This chapter describes the following:

- Creating a Citrix Policy for NVIDIA vGPU
- Creating a Microsoft Policy

A vGPU is not automatically accessible to a Citrix session. Both Citrix and Microsoft policies must be configured for a Citrix session to utilize the vGPU.

11.1 Creating a Citrix Policy for NVIDIA vGPU

The Citrix HDX 3D Pro protocol can utilize NVIDIA NVENC (Hardware-Accelerated Video Encoding). In order to utilize NVENC, the Optimize for 3D graphics workload Citrix Policy must be enabled.

Your full Citrix policy set depends on multiple parameters like application requirements, bandwidth requirements, & image quality requirements, etc. Only enabling the Optimize for 3D graphics workloads policy is only sufficient for a POC/trail purpose. Consult Citrix and your application partners to ensure you full Citrix Policy set is optimized for you deployment needs.

Additionally, refer to the Graphics Section of the Citrix Virtual Apps and Desktop Product Documentation for additional details.

1. Log on to the **Citrix Delivery Controller**
2. Launch **Citrix Studio** from the Windows Start Menu
3. Select **Policies** on the left menu pane, then select **Create Policy** on the right-side Action menu.

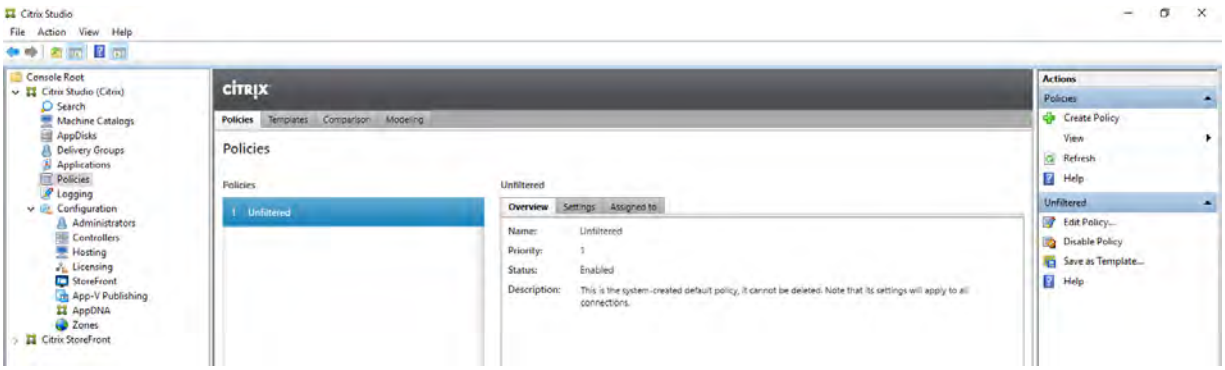

4. Search for **Optimize for 3D graphics workload** and click **Select**.

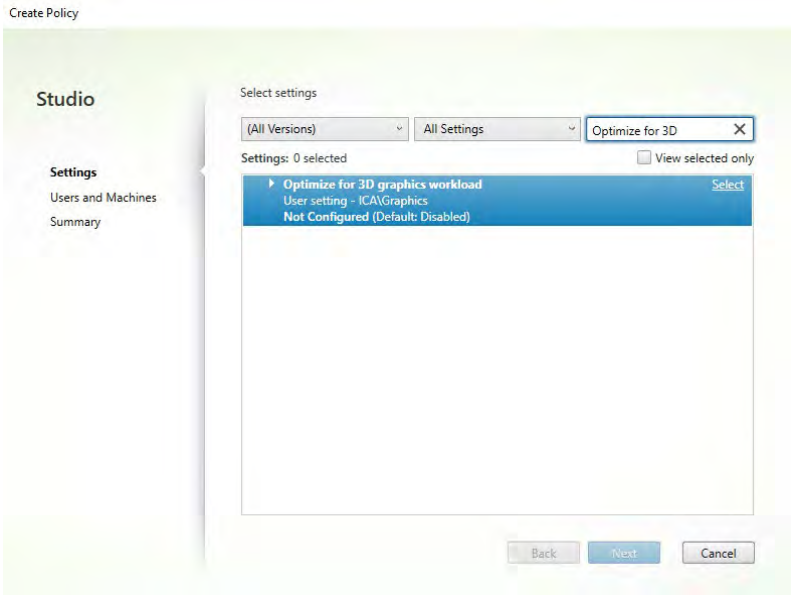

5. On the Edit Settings popup window, select the **Enabled** radio button and click **OK**.

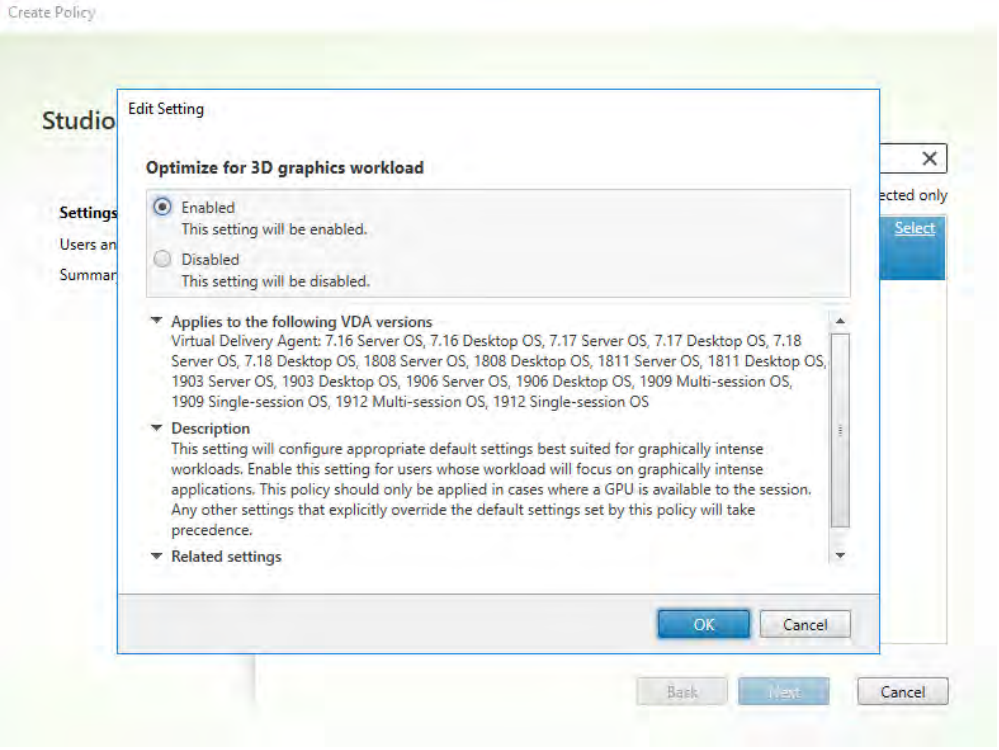

6. Click **Next**
7. Select the appropriate machines to apply the policy to. For POC/trial purposes, select the **All objects in the site** radio button and click **Next.**

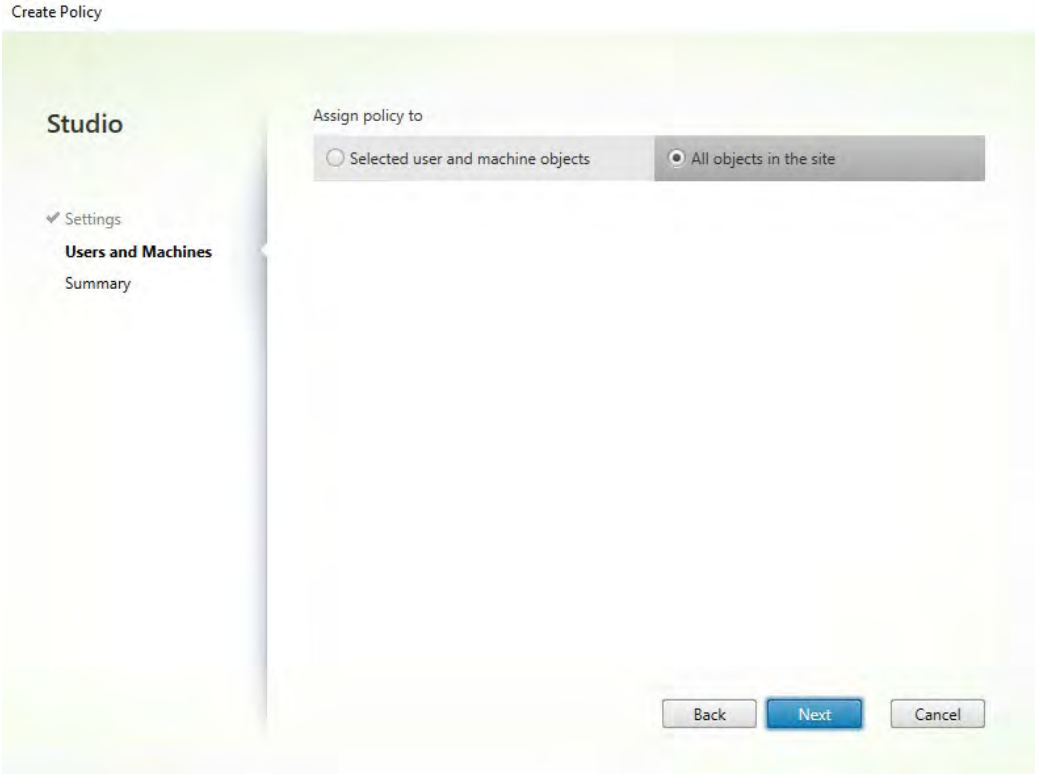

8. Check the **Enable policy** checkbox and provide a name for the policy.

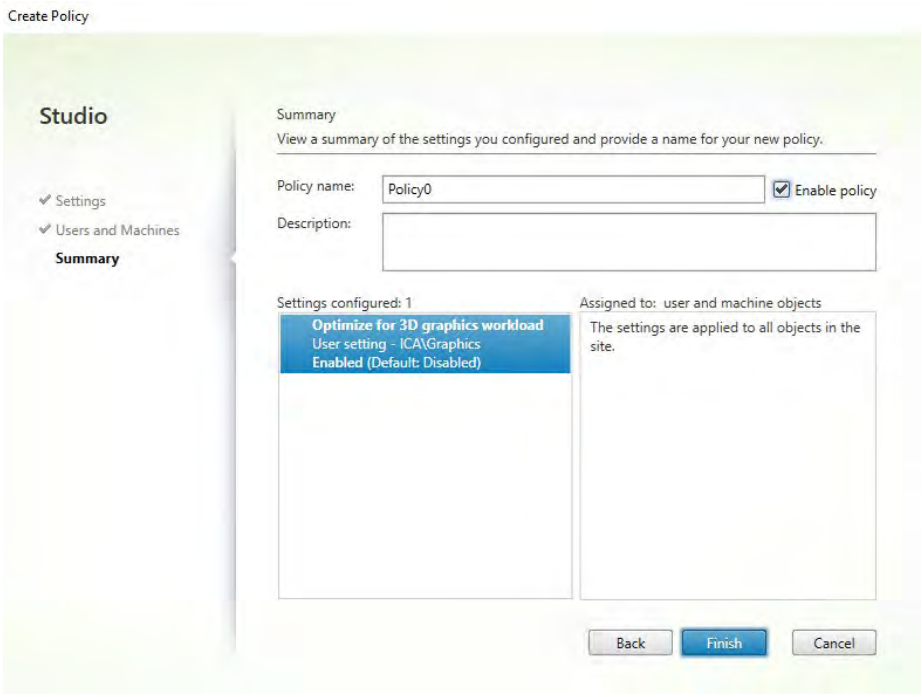

9. Click **Finish**

Note: There are several policies within Citrix Studio that affect performance and the efficient use of Virtual GPUs. Policies such as **Use Video Codec for Compression, Use hardware encoding for video codec** and **Target Frame Rate**.

Please refer to the following Citrix document which refers to these settings in more detail:

https://docs.citrix.com/en-us/tech-zone/design/design-decisions/hdx-graphics.html

11.2 Creating Microsoft Group Policy for NVIDIA vGPU

On Windows Server 2016 and Windows Server 2012 R2, Remote Desktop Services (RDS) sessions on a RD Session Host server use the Microsoft Basic Render Driver as the default adapter. To use the virtual GPU in RDS sessions and Citrix HDX 3D Pro sessions enable the **Use the hardware default graphics adapter for all Remote Desktop Services sessions** setting in the group policy.

Local Computer Policy > Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Remote Session Environment.

Chapter 12. Citrix Workspace App

Before connecting to a virtual application or desktop over a Citrix HDX connection, the Citrix Workspace App will need to be installed and configured onto a desktop or device which the virtual desktop will be accessed from. For this guide, connect to the Citrix StoreFront and use the web site which can detect and download the receiver.

12.1 Locating Citrix StoreFront Web Site

1. Locate the Citrix Storefront URL by opening **Citrix Studio** from the Windows menu on your Citrix Delivery Controller Server.
2. Expand **Citrix StoreFront** on the left menu pane and Select **Stores**.
3. Highlight the Store for your deployment in the top menu pane and select the **Receiver for Web Sites** tab under the Details section.
4. Right Click the URL and select **Copy URL.**

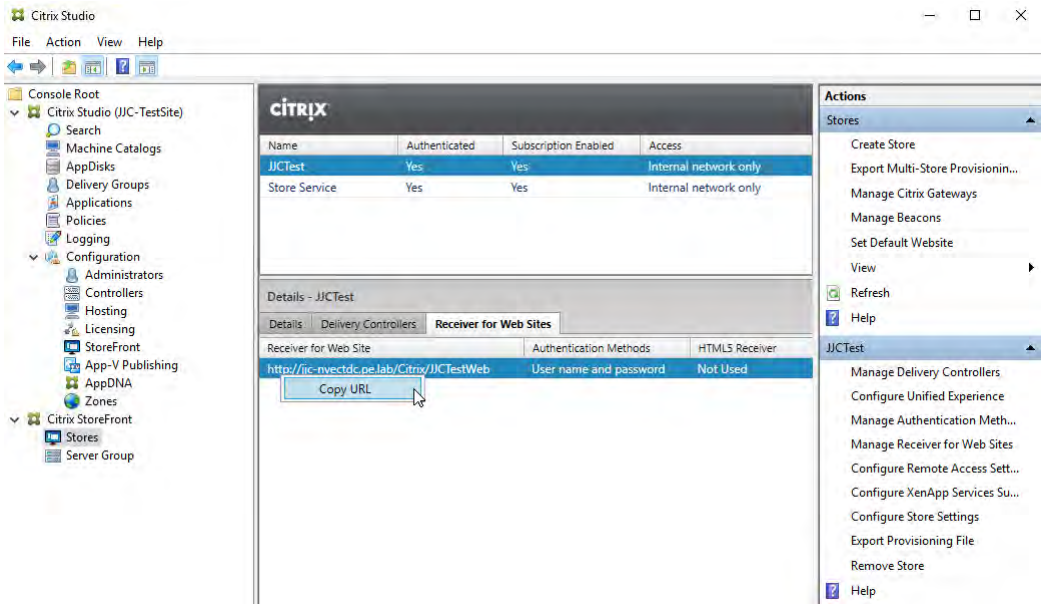

12.2 Installing Citrix Workspace App

1. Log into the physical device where you will be launching the virtual desktop from and open an internet browser. Navigate to the Citrix Storefront URL you previously copied in step 11.1.4. Click on **Detect Receiver**.

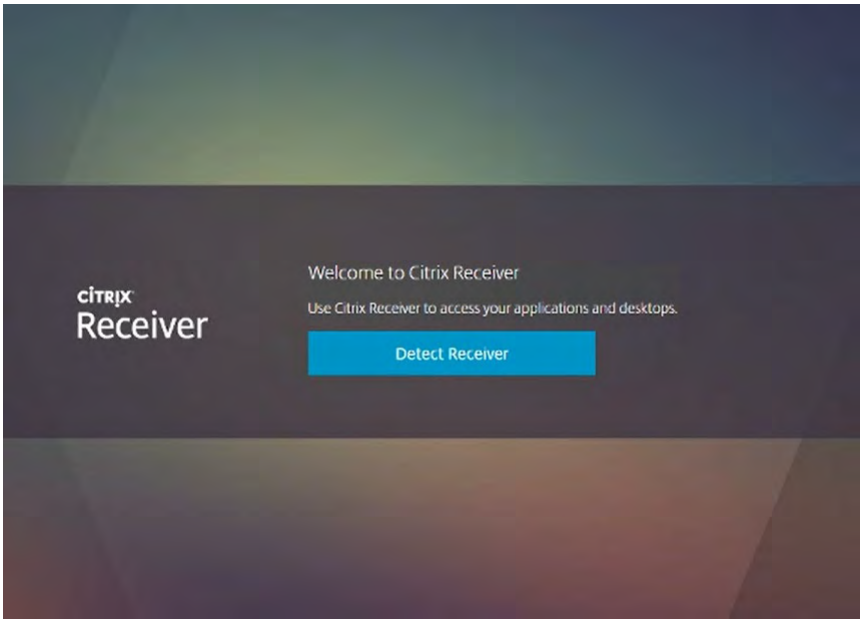

2. Review the Citrix license agreement and check the **I agree with the Citrix license agreement** checkbox and select **Download.**

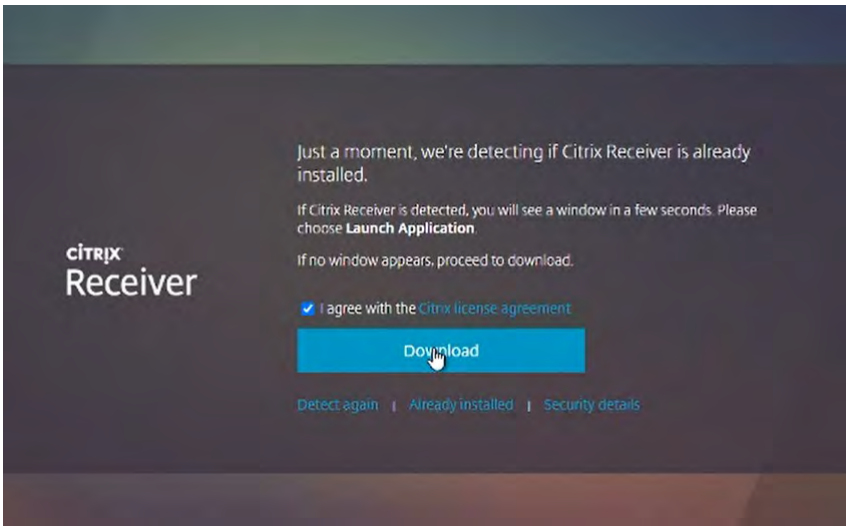

3. Locate the downloaded installer program and double click to begin installation.

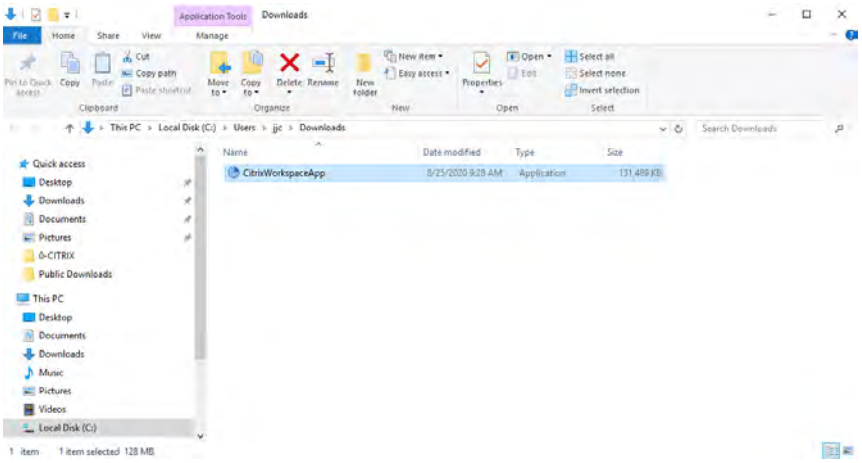

4. Click **Start** to begin the installation.

5. Select the **I accept the license agreement** check box and click **Next**.

6. Click **Install**

7. Click **Finish** to complete the install.

12.3 Launch a Citrix Virtual Desktop

1. Navigate back to the Citrix StoreFront web browser window and click **Continue**.

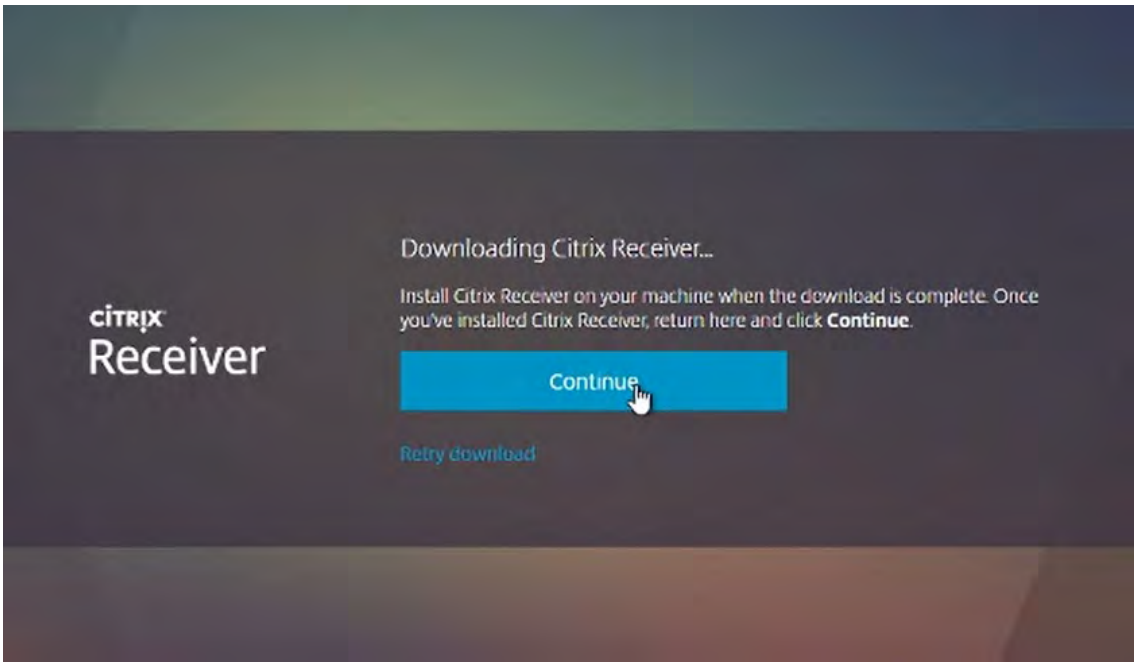

2. Enter a username and password.

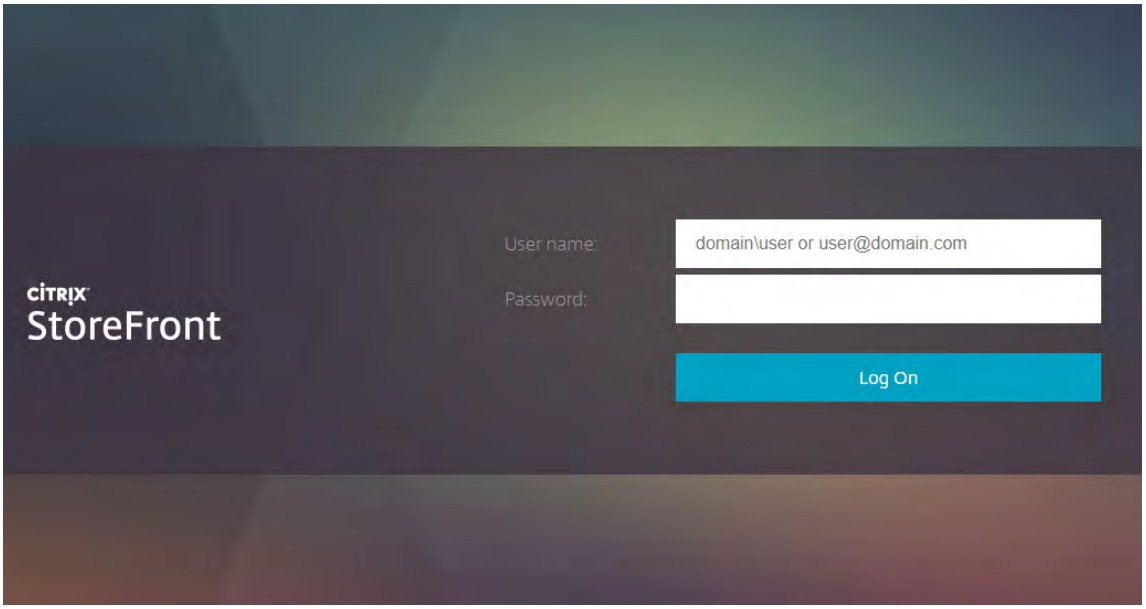

3. Click **Desktops** in the top menu and select your desktop.

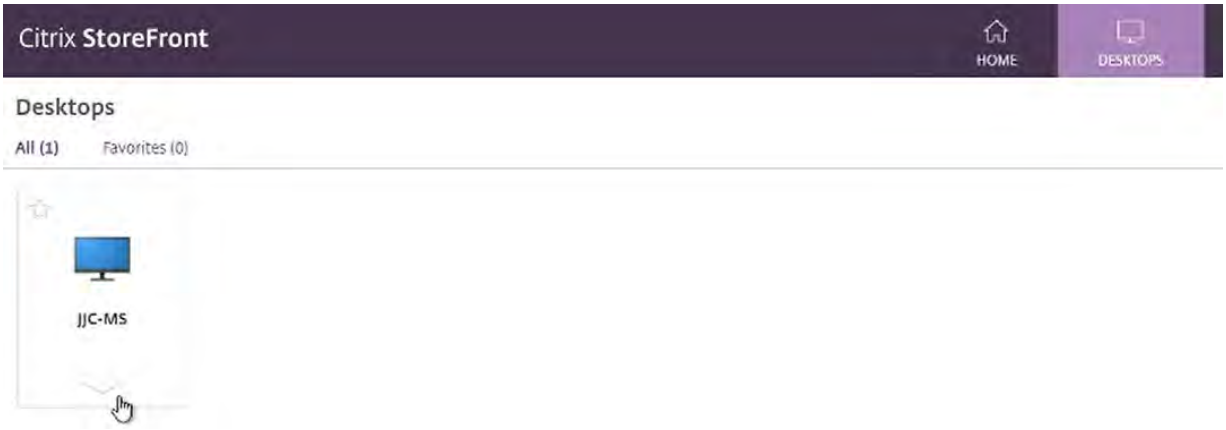
Citrix StoreFront
HOME
DESKTOPS
Desktops
All (1)
Favorites (0)
JJC-MS

Chapter 13. Troubleshooting

This section includes links and examples of where to explore and post in order to find solutions or assistance.

Note: Before troubleshooting or filing a bug report, review the release notes for information about known issues with the current release, and potential workarounds.

13.1 Forums

- NVIDIA forums are a very inclusive source of solutions to many problems that may be faced when deploying a virtualized environment. Please search on the NVIDIA forums located here first:
- https://gridforums.nvidia.com/
- You may also wish to look through the NVIDIA Enterprise Services Knowledgebase to further find support articles and links here:
- https://nvidia-esp.custhelp.com/app/answers/list/autologout/1
- Keep in mind that not all issues within your deployment may be answered in the NVIDIA vGPU forums. You may also have to reference forums from the hardware supplier, the hypervisor and application themselves. Some examples of key forums to look through are here:
  - Citrix Forums: https://discussions.citrix.com/
  - HPE ProLiant Server Forums: https://community.hpe.com/t5/ProLiant/ct-p/proliant
  - Dell Server Forums: https://www.dell.com/community/Servers/ct-p/ESServers
  - Lenovo Server Forums: https://forums.lenovo.com/t5/Datacenter-Systems/ct-p/sv_eg
  - Autodesk Knowledge Network: https://knowledge.autodesk.com/
  - Adobe Forums: https://forums.adobe.com/welcome
  - Dassault Systèmes User Groups: https://www.3ds.com/support/users-communities/

13.2 Filing a Bug Report

When filing a bug or requesting support assistance, it is critical to include information about the environment, so that the technical staff that can help you resolve the issue. NVIDIA includes the

nvidia-bug-report.sh script within the. vib installation package to collect and package this critical information. The script collects the following information:

- Citrix version
- X.Org log and configuration
- PCI information
- CPU information
- GPU information
- **esxcfg** information for PLX devices
- **esxcfg** information for GPU devices
- VIB information
- NVRM messages from vmkernel.log
- System **dmesg** output
- Which virtual machines have vGPU or vSGA configured
- NSMI output

When running this script:

- You may specify the output location for the bug report using either the -o or –output switch followed by the output file name. If you do not specify an output directory, the script will write the bug report to the current directory.
- If you do not specify a file name, the script will use the default name nvidia-bug-report.log.gz.
- If the selected directory already contains a bug report file, then the script will change the name of that existing report file to nvidia-bug-report.log.old.gz before generating a new nvidia-bug-report.log.gz file.

To collect a bug report, issue the command:

```
$ nvidia-bug-report.sh
```

The system displays the following message during the collection process:

```
nvidia-bug-report.sh will now collect information about your system and
create the file 'nvidia-bug-report.log.gz' in the current directory. It may
take several seconds to run. In some cases, it may hang trying to capture
data generated dynamically by the vSphere kernel and/or the NVIDIA kernel
module. While the bug report log file will be incomplete if this happens, it
may still contain enough data to diagnose your problem.
```

Be sure to include the **nvidia-bug-report.log.gz log** file when reporting problems to NVIDIA.

Appendix A. About This Document

A.1 Related Documentation

Refer to the NVIDIA Virtual GPU (vGPU) resources page http://www.nvidia.com/gridresources for additional information about NVIDIA vGPU technology, including:

- NVIDIA Virtual GPU Technology
- *https://www.nvidia.com/en-us/design-visualization/technologies/virtual-gpu/*
- Purchasing Guide for NVIDIA vGPU Solutions
- https://www.nvidia.com/en-us/design-visualization/buy-grid/
- NVIDIA GPU Datasheets
- *http://www.nvidia.com/object/grid-enterprise-resources.html#datasheet*
- Application Deployment Guides and Solution Overviews
- *http://www.nvidia.com/object/grid-enterprise-resources.html#guides*
- Customer Success Stories
- *http://www.nvidia.com/object/grid-enterprise-resources.html#case*
- White Papers
- *http://www.nvidia.com/object/grid-enterprise-resources.html#whitepapers*
- Videos
  http://www.nvidia.com/object/grid-enterprise-resources.html#videos

A.2 Support Contact Information

For technical support you should reach out to your local Citrix and NVIDIA vGPU teams for guidance.

For support when architecting your solution, your Citrix and NVIDIA vGPU teams are available to assist. Please be sure you are in touch with them and keep them up to date with your progress. If you do not know your correct account management teams, please reach out to the appropriate email below:

- NVIDIA vGPU team: gridteam@nvidia.com

The NVIDIA vGPU resources page includes additional contact methods to help you get the answers you need as soon as possible.

Appendix B. Installing & Licensing NVIDIA Driver in Linux Virtual Desktop

B.1 Installing NVIDIA Driver in Linux Virtual Desktop

- Installation in a VM: After you create a Linux VM on the hypervisor and boot the VM, install the NVIDIA vGPU software display driver in the VM to fully enable GPU operation.
- Installation on bare metal: When the physical host is booted before the NVIDIA vGPU software display driver is installed, the vesa Xorg driver starts the X server. If a primary display device is connected to the host, use the device to access the desktop. Otherwise, use secure shell (SSH) to log in to the host from a remote host. If the Nouveau driver for NVIDIA graphics cards is present, disable it before installing the NVIDIA vGPU software display driver.

Installation of the NVIDIA vGPU software display driver for Linux requires:

- Compiler toolchain
- Kernel headers

1. Copy the NVIDIA vGPU software Linux driver package, for example NVIDIA-Linux_x86_64-390.75-grid.run, to the guest VM or physical host where you are installing the driver.

2. Before attempting to run the driver installer, exit the X server and terminate all OpenGL applications.

    a) On Red Hat Enterprise Linux and CentOS systems, exit the X server by transitioning to runlevel 3:

    ```
    [nvidia@localhost ~]$ sudo init 3
    ```

    b) On Ubuntu platforms, do the following:

    i) Use **CTRL**-**ALT**-**F1** to switch to a console login prompt.

    ii) Log in and shut down the display manager:

    ```
    [nvidia@localhost ~]$ sudo service lightdm stop
    ```

3. From a console shell, run the driver installer as the root user.

```
sudo sh ./ NVIDIA-Linux_x86_64-352.47-grid.run
```

In some instances, the installer may fail to detect the installed kernel headers and sources. In this situation, re-run the installer, specifying the kernel source path with the --kernel-source-path option:

```
sudo sh ./ NVIDIA-Linux_x86_64-352.47-grid.run \
–kernel-source-path=/usr/src/kernels/3.10.0-229.11.1.el7.x86_64
```

4. When prompted, accept the option to update the X configuration file (xorg.conf).

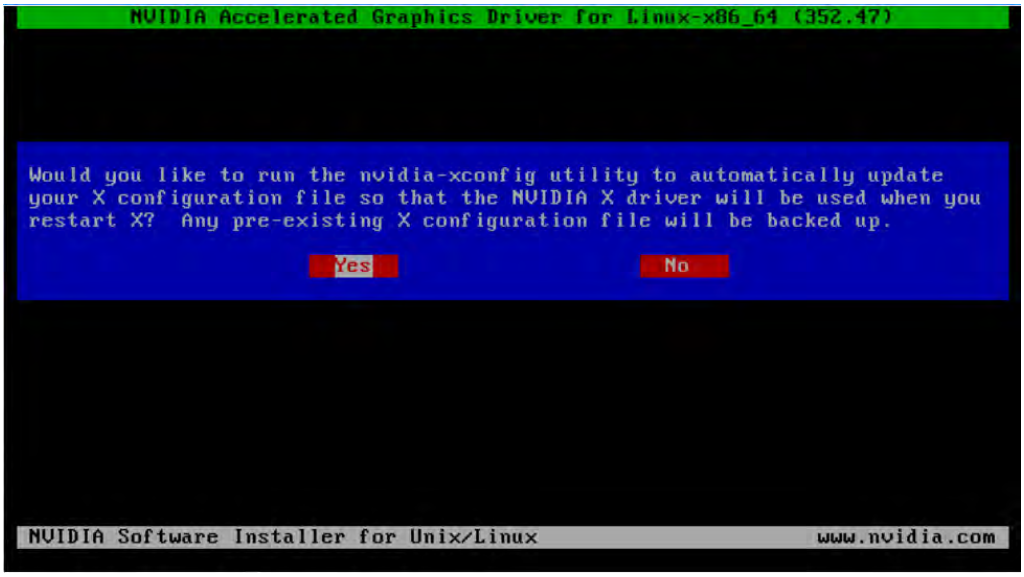

5. Once installation has completed, select **OK** to exit the installer.

6. Verify that the NVIDIA driver is operational.

   a) Reboot the system and log in.

   b) Run nvidia-settings.

```
[nvidia@localhost ~]$ nvidia-settings
```

The NVIDIA X Server Settings dialog box opens to show that the NVIDIA driver is operational.

7. Installation in a VM: After you install the NVIDIA vGPU software display driver, you can license any NVIDIA vGPU software licensed products that you are using. For instructions, refer to Virtual GPU Client Licensing User Guide.

B.2 Licensing NVIDIA vGPU on Linux

1. Start NVIDIA X Server Settings by using the method for launching applications provided by your Linux distribution. For example, on Ubuntu Desktop, open the Dash, search for NVIDIA X Server Settings, and click the **NVIDIA X Server Settings** icon.
2. In the NVIDIA X Server Settings window that opens, click **Manage NVIDIA License**. The License Edition section of the NVIDIA X Server Settings window shows that NVIDIA vGPU is currently unlicensed.
3. In the **Primary Server** field, enter the address of your primary NVIDIA vGPU software License Server. The address can be a fully qualified domain name such as gridlicense1.example.com, or an IP address such as 10.31.20.45. If you have only one license server configured, enter its address in this field.
4. Leave the **Port Number** field under the **Primary Server** field unset. The port defaults to 7070, which is the default port number used by NVIDIA vGPU software License Server.
5. In the **Secondary Server** field, enter the address of your secondary NVIDIA vGPU software License Server. If you have only one license server configured, leave this field unset. The address can be a fully qualified domain name such as gridlicense2.example.com, or an IP address such as 10.31.20.46.
6. Leave the **Port Number** field under the **Secondary Server** field unset. The port defaults to 7070, which is the default port number used by NVIDIA vGPU software License Server.

7. Click **Apply** to assign the settings. The system requests the appropriate license for the current vGPU from the configured license server.
8. The vGPU within the VM should now exhibit full frame rate, resolution, and display output capabilities. The VM is now capable of running the full range of DirectX and OpenGL graphics applications.
9. If the system fails to obtain a license, see Virtual GPU Client Licensing User Guide for guidance on troubleshooting.

Appendix C. GPU Resource Allocation

C.1 vGPU Assignment

This chapter covers how VM's vGPUs will be allocated across the available physical GPUs in a host.

- vGPU Placement Policy
- Integrated GPU-pass-through
- Which vGPU Profiles are allowed on physical GPUs

1. Select the host server in the left resources pane. Navigate to the GPU tab. All physical GPU and vGPU resources available will display.

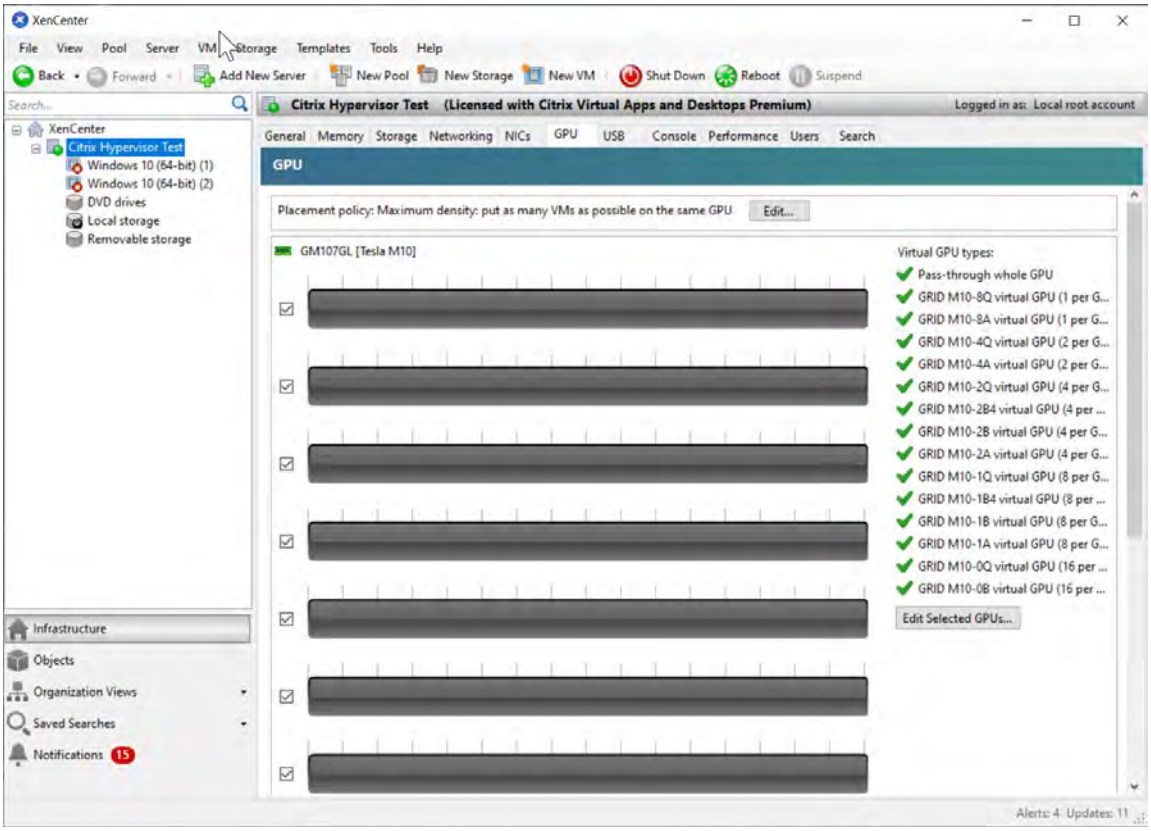

2. Select Edit under placement policy to modify GPU settings.

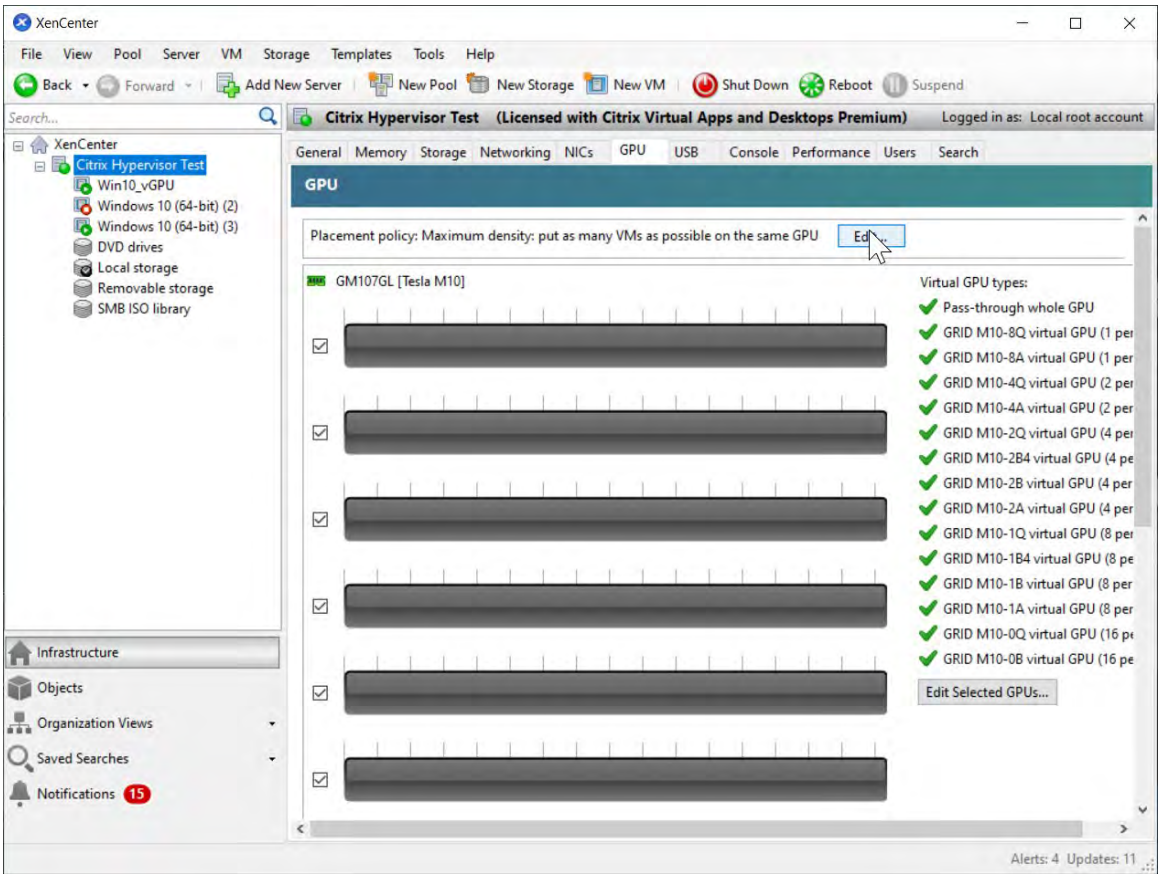

3. The options of maximum density or maximum performance can be viewed here for assigning the VMs to GPUs. The option for selecting integrated GPU will also display. Select the appropriate options for your deployment strategy.

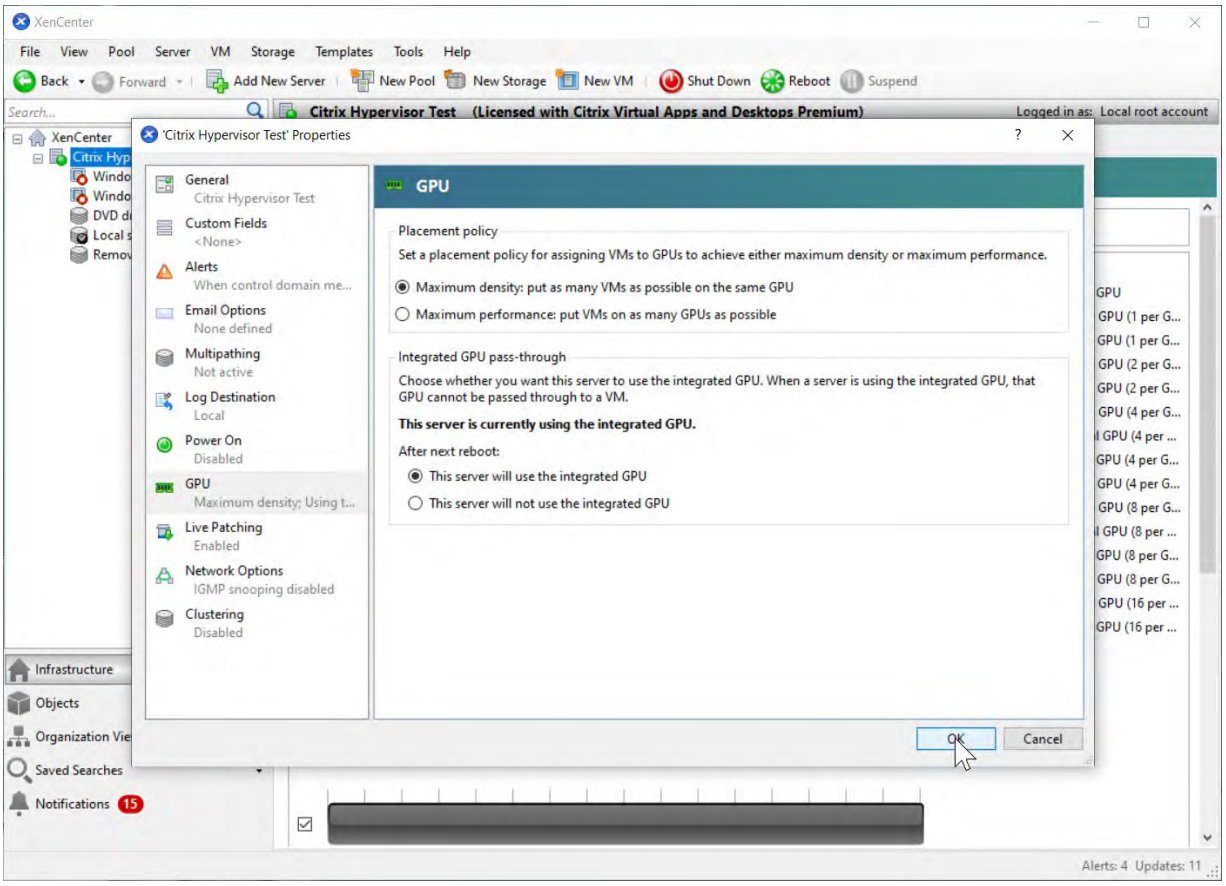

4. In the right panel, select **Edit Selected GPUs** to see the resources allocated to each GPU as well as manage which vGPU profiles are allowed for use on the host server.

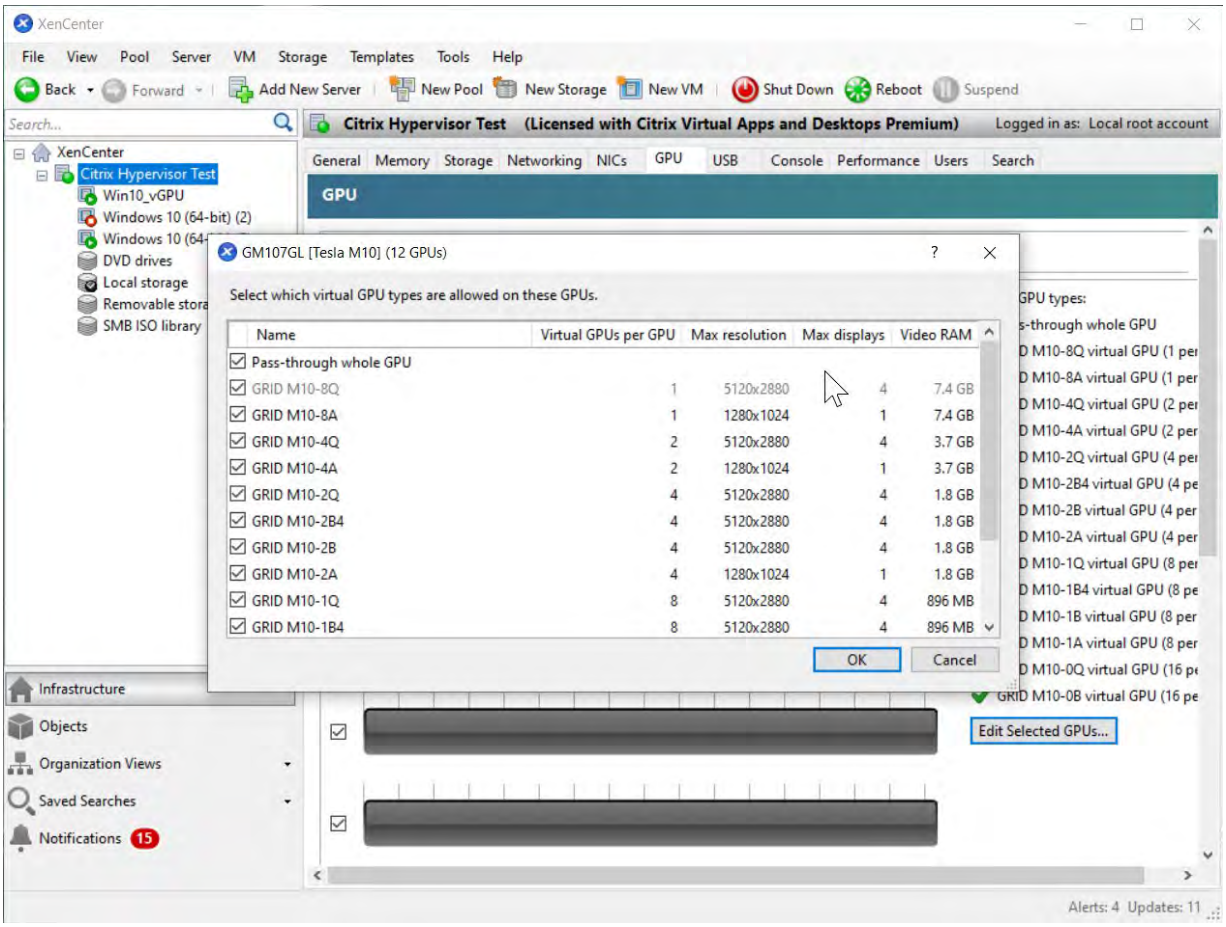

How vGPUs should be allocated will be dependent on your deployment needs. Consult Citrix and your NVIDIA representative to determine the best configuration for your environment.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, CUDA, NVIDIA OptiX, NVIDIA RTX, NVIDIA Turing, Quadro, Quadro RTX, and TensorRT trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2020 NVIDIA Corporation. All rights reserved.

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051

http://www.nvidia.com