



Configuring XenServer v6.5.0 Service Pack 1 for Graphics

May 2015

Contents

Introduction.....	2
GPU Pass-Through	2
Guest Support and Constraints.....	2
Windows Guests.....	2
HVM Linux Guests.....	3
NVIDIA GRID vGPU.....	3
Available NVIDIA GRID vGPU Types.....	4
NVIDIA GRID System Requirements	5
Preparation Overview:.....	5
Installation on XenServer 6.5.0 Service Pack 1	7
Installing the NVIDIA GRID vGPU Manager for XenServer.....	8
Create a vGPU or GPU Pass-Through Enabled VM	10
To Enable Intel GPU Pass-through.....	11
In Guest Driver Installation.....	12
Install the NVIDIA drivers.....	12
Install the AMD Drivers.....	14
Install the Intel Drivers	16
Install the XenServer Tools and Configure VDA for HDX 3D Pro.....	18

Introduction

Citrix XenServer is leading the way in the virtual delivery of 3D professional graphics applications and workstations. Its offerings include GPU Pass-through (for NVIDIA, AMD and Intel GPUs) as well as NVIDIA GRID vGPU, the first virtualization solution to support hardware-based GPU sharing.

GPU Pass-Through

Unlike the rest of the physical system components, which are shared and represented as multiple virtual instances to multiple clients by the hypervisor, the pass-through GPU is not abstracted at all, but remains one physical device. Each hosted virtual machine gets its own dedicated GPU, eliminating the software abstraction and the performance penalty that goes with it.

XenServer allows you to assign a physical GPU (in the XenServer host) to a Windows or HVM Linux VM running on the same host. This GPU Pass-Through feature is intended for graphics power users, such as CAD designers, who require high performance graphics capabilities.

The following table summarizes GPU and vGPU support for guests at XenServer 6.5 Service Pack 1:

	GPU for Windows	GPU For HVM Linux	vGPU for Windows
AMD			
Intel			
NVIDIA			

Guest Support and Constraints

Windows Guests

Customers can use AMD or NVIDIA GPUs for GPU Pass-through. In addition, XenServer 6.5 Service Pack 1 introduces support for select Intel Integrated GPUs.

The following guests are currently supported for use with GPU pass-through and vGPU:

- Windows 7 (32-bit/64-bit)
- Windows Server 2008 R2 SP1
- Windows 8 (32-bit/64-bit)
- Windows 8.1 (32-bit/64-bit)
- Windows Server 2012
- Windows Server 2012 R2

HVM Linux Guests

At XenServer 6.5 Service Pack 1, the following HVM Linux guests are supported to use GPU Pass-Through with NVIDIA GPUs.

- RHEL 7
- CentOS 7
- Oracle Linux 7
- Scientific Linux 7
- Ubuntu 14.04

Note

- XenServer supports only one GPU per VM.
- vGPU and GPU Pass-Through are **not** compatible with High Availability, XenMotion, Storage XenMotion or VM Suspend. However, VMs, using GPU Pass-Through or vGPU, can still be started on any host that has the appropriate resources.
- XenServer automatically detects and groups together identical physical GPUs across hosts in the same pool. Once assigned to a group of GPUs, a VM may be started on any host in the pool that has an available GPU in the group.

NVIDIA GRID vGPU

NVIDIA GRID vGPU enables multiple Virtual Machines (VM) to have simultaneous, direct access to a single physical GPU, using the same NVIDIA graphics drivers that are deployed on non-virtualized Operating Systems. Under the control of NVIDIA's GRID Virtual GPU Manager, which runs in the XenServer Control Domain (dom0), GRID physical GPUs are capable of supporting multiple virtual GPU devices (vGPUs) that can be assigned directly to VMs.

Guest VMs use GRID virtual GPUs in the same manner as a physical GPU that has been passed through by the hypervisor; an NVIDIA driver loaded in the guest VM provides direct access to the GPU for performance critical fast paths, and a paravirtualized interface to the GRID Virtual GPU Manager. NVIDIA GRID is enabled by default in XenServer v6.5.0.

Licensing Note

vGPU is available for XenServer Enterprise edition customers or those who have access to XenServer through their XenDesktop/XenApp entitlement. To learn more about XenServer editions, and to find out how to upgrade, visit the Citrix website [here](#). For detailed information on Licensing, refer to [CTX141511 - XenServer 6.5 Licensing FAQ](#).

Available NVIDIA GRID vGPU Types

NVIDIA GRID cards can contain multiple Graphics Processing Units (GPU).

For example, GRID K1 cards contain four GK107GL GPUs, and GRID K2 cards contain two GK104GL GPUs. Each physical GPU (pGPU) can host several different *types* of virtual GPU (vGPU). vGPU types have a fixed amount of framebuffer, number of supported display heads and maximum resolutions, and are targeted at different classes of workload.

Some of the vGPU types supported by GRID K1 and K2 are defined in Table 1 below:

Card	No. of Physical GPUs	Virtual GPU Types	Intended Use Case	Max Resolution	No. of vGPUs per GPU	No. of vGPUs per Card
GRID K1	4	GRID K180Q	Designer/Power User	2560x1600	2	8
		GRID K160Q	Designer/Power User	2560x1600	1	4
		GRID K140Q	Power User	2560x1600	4	16
		GRID K120Q	Power User	2560x1600	8	32
		GRID K100	Knowledge Worker	1920x1200	8	32
GRID K2	2	GRID K280Q	Designer/Power User	2560x1600	1	2
		GRID K260Q	Designer/Power User	2560x1600	2	4
		GRID K240Q	Designer/Power User	2560x1600	4	8
		GRID K220Q	Designer/Power User	2560x1600	8	16
		GRID K200	Knowledge Worker	1920x1200	8	16

Table 1 Virtual GPU Types

Note

At any given time, vGPUs hosted on a physical GPU **must all be of the same type**. However, there is no corresponding restriction between physical GPUs on the same card. This restriction is automatic and may cause unexpected capacity planning issues.

For example, a GRID K2 card, has two physical GPUs, and can support four types of vGPU; GRID K200, GRID 220Q, GRID 240Q, and GRID 260Q.

NVIDIA GRID System Requirements

- NVIDIA GRID card:
For the most recent vGPU types, refer to <http://www.nvidia.com/object/virtual-gpus.html>
For a list of the most recently supported NVIDIA cards refer to:
 - XenServer Hardware Compatibility List: <http://hcl.vmd.citrix.com/vGPUDeviceList.aspx> in conjunction with
 - NVIDIA product information: <http://www.nvidia.com/object/grid-technology.html>
- XenServer v6.5.0 Enterprise edition (or access to XenServer through a XenDesktop/XenApp entitlement)
- A server capable of hosting XenServer and NVIDIA GRID cards.
The NVIDIA GRID vGPU software package for XenServer, consisting of the GRID Virtual GPU Manager for XenServer, and NVIDIA drivers, refer to: <http://> For a list of the most recently supported NVIDIA cards refer to:
- XenServer Hardware Compatibility List:
<http://hcl.vmd.citrix.com/vGPUDeviceList.aspx>www.nvidia.com/vGPU

To run XenDesktop with VMs running NVIDIA vGPU, you will also need:

- XenDesktop 7.1 or later, full installation

Notes:

- No other versions of XenServer or XenDesktop are currently supported for use with NVIDIA virtual GPUs.
Note: Customers running XenServer 6.2 should refer to the guide [Configuring XenServer 6.2 to use NVIDIA GRID](#)
- Customers are advised to review the GRID vGPU FOR CITRIX XENSERVEN User Guide (Ref: DU-06920-001) available from NVIDIA - <http://www.nvidia.com/vGPU>
Customers will need to register with NVIDIA to gain access to these components for additional performance optimization steps.

Preparation Overview:

1. [Install XenServer](#)
2. [Install the NVIDIA Virtual GPU Manager for XenServer](#)
3. [Reboot the XenServer host](#)



Configuring XenServer v6.5.0 Service Pack 1 for Graphics

May 2015



Configuring XenServer v6.5.0 Service Pack 1 for Graphics

May 2015

Installation on XenServer 6.5.0 Service Pack 1

XenServer 6.5.0 is available for download from:

<https://www.citrix.com/downloads/xenserver/product-software/xenserver-65.html>.

Customers should install:

- **XenServer 6.5.0 Base Installation ISO**
- **XenServer 6.5 Service Pack 1**
- **XenCenter 6.5.0 Service Pack 1 Windows Management Console**

Refer to the *XenServer 6.5.0 Installation Guide* for comprehensive details on installation:

<http://support.citrix.com/article/CTX141501>

Licensing Note:

vGPU is available for XenServer Enterprise edition customers or those who have access to XenServer through their XenDesktop/XenApp entitlement. To learn more about XenServer v6.5 editions, and to find out how to upgrade, visit the Citrix website [here](#).

Installing the NVIDIA GRID vGPU Manager for XenServer

Customers should install NVIDIA GRID vGPU Software available from [NVIDIA](#).

The NVIDIA GRID software consists of:

1. GRID vGPU Manager
(for example: NVIDIA-vgx-xenserver-6.2-331.30.i386.rpm)
2. Windows Display Driver
(for example: 332.07_grid_win7_64bit_english.exe)

The *GRID vGPU Manager* runs in the XenServer Control Domain (dom0). It is provided as an RPM file, which must be copied to dom0 and then installed. Please also refer to the User Guide included in the NVIDIA GRID vGPU Software for more detailed installation steps and specifics.

Note: The RPM names and versions below are examples and will most likely be different in your environment.

1. Use the **rpm** command to install the package:

```
[root@xenserver ~]# rpm -iv NVIDIA-vgx-xenserver-6.2-331.30.i386.rpm
Preparing packages for installation...
NVIDIA-vgx-xenserver-6.2-331.30
[root@xenserver ~]
```

2. Reboot the XenServer host:

```
[root@xenserver ~]# shutdown -r now
Broadcast message from root (pts/1) Wed Sept 25 13:05:31 2013):
[root@xenserver
```

3. After the XenServer host has rebooted, verify that the GRID package has installed and loaded correctly by checking for the NVIDIA kernel driver in the list of kernel loaded modules:

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

Note: If at this stage the NVIDIA module (*nvidia*) is not loaded correctly and you are using XenServer v6.2.0, refer to [CTX139834](#) for troubleshooting.

4. Verify that the NVIDIA kernel driver can successfully communicate with the GRID physical GPUs in your host by running the **nvidia-smi** command, which produces a listing of the GPUs in your platform similar to:

```
root@xenserver ~]# nvidia-smi
```




Configuring XenServer v6.5.0 Service Pack 1 for Graphics

May 2015

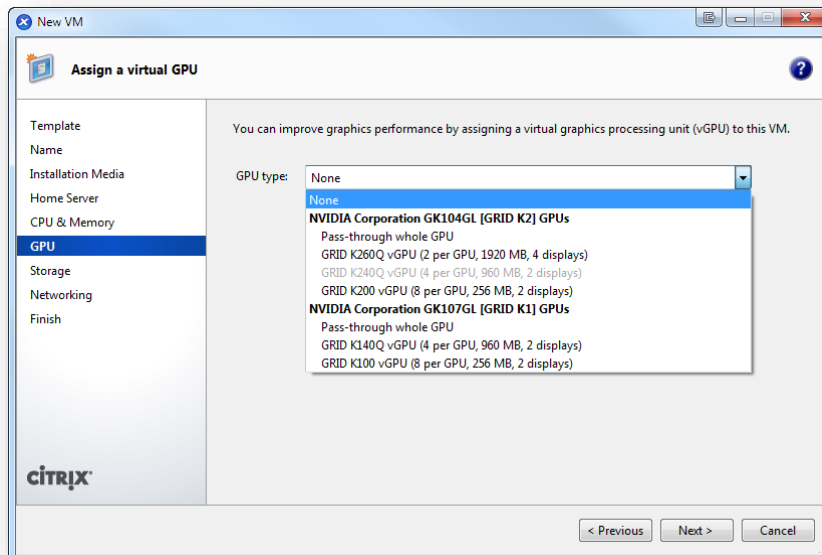
```
Wed Sep 25 13:05:31 2013
-----+
| NVIDIA-SMI 4.312.36 Driver Version: 312.36 |
-----+
| GPU  Name      | Bus-Id      Disp. | Volatile Uncorr. ECC |
| Fan  Temp      | Perf      Pwr:Usage/Cap| Memory-Usage  | GPU-Util  Compute M. |
-----+-----+-----+-----+-----+-----+
|      0      | GRID K1    | 0000:04:00.0 Off | N/A |
| N/A  27C      | P0        13W / 31W | 0%   9MB / 4095MB | 0%   Default |
-----+-----+-----+-----+-----+
|      1      | GRID K1    | 0000:05:00.0 Off | N/A |
| N/A  25C      | P0        13W / 31W | 0%   9MB / 4095MB | 0%   Default |
-----+-----+-----+-----+-----+
|      2      | GRID K1    | 0000:06:00.0 Off | N/A |
| N/A  21C      | P0        13W / 31W | 0%   9MB / 4095MB | 0%   Default |
-----+-----+-----+-----+-----+
|      3      | GRID K1    | 0000:07:00.0 Off | N/A |
| N/A  23C      | P0        13W / 31W | 0%   9MB / 4095MB | 0%   Default |
-----+-----+-----+-----+-----+
|      4      | GRID K1    | 0000:86:00.0 Off | N/A |
| N/A  24C      | P0        13W / 31W | 0%   9MB / 4095MB | 0%   Default |
-----+-----+-----+-----+-----+
|      5      | GRID K1    | 0000:87:00.0 Off | N/A |
| N/A  24C      | P0        13W / 31W | 0%   9MB / 4095MB | 0%   Default |
-----+-----+-----+-----+-----+
|      6      | GRID K1    | 0000:88:00.0 Off | N/A |
| N/A  25C      | P0        13W / 31W | 0%   9MB / 4095MB | 0%   Default |
-----+-----+-----+-----+-----+
|      7      | GRID K1    | 0000:89:00.0 Off | N/A |
| N/A  25C      | P0        12W / 31W | 0%   9MB / 4095MB | 0%   Default |
-----+-----+-----+-----+-----+

+-----+
| Compute processes: GPU Memory |
| GPU  PID      Process name  Usage  |
-----+-----+-----+-----+
|      No running compute processes found  |
-----+
[root@xenserver ~]#
```

Create a vGPU or GPU Pass-Through Enabled VM

Note: Customers using Intel GPUs should first refer to [To Enable Intel GPU Pass-through](#) for additional configuration, and then follow the steps below.

1. Using XenCenter, create a VM.
To do this, on the toolbar, select **New VM** and step through the **New VM** wizard.
2. Select the **Installation Media, Home Server, CPU & Memory**.
3. GPU-enabled servers will now show a **GPU** configuration page:



hg heads

4. From the **GPU Type** drop-down list, select either **Pass-through whole**, or a **vGPU type**. (Unavailable vGPU types will be greyed-out)
5. Click **Next** to configure **Storage** and **Networking**.
6. When you complete your configuration, click **Create Now**.

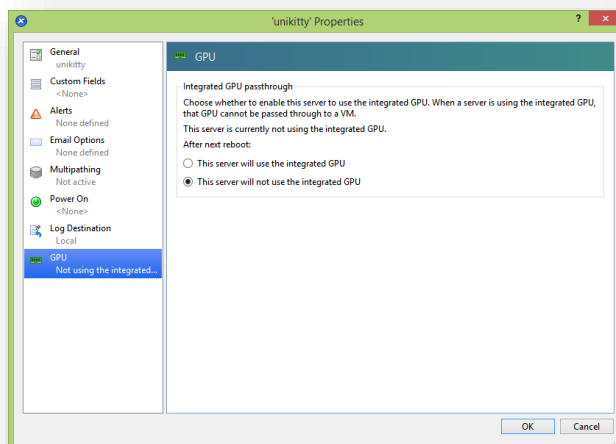
To Enable Intel GPU Pass-through

XenServer 6.5 Service Pack 1 supports the GPU Pass-Through feature for Windows 7 and Windows 8 (32-/64-bit) VMs using an Intel integrated GPU device. For more information on supported hardware, refer to the [XenServer Hardware Compatibility List](#).

When using Intel GPU on Intel servers, the XenServer host's control domain (dom0) will have access to the integrated GPU device. In such cases, the GPU will not be available for pass-through. Customers who wish to use the Intel GPU Pass-Through feature on Intel servers should *disable* the connection between dom0 and the GPU before passing through the GPU to the VM.

To do this:

1. On the **Resources** pane, select the XenServer host.
2. On the **General** tab, click **Properties**, and in the left pane, click on **GPU**
3. In the **Integrated GPU passthrough** section, click **This server will not use the integrated GPU**:



4. This disables the connection between dom0 and the Intel integrated GPU device
4. Click **OK**
5. Reboot the XenServer host for the changes to take effect.

The Intel GPU will now be visible on the GPU type drop-down list during new VM creation, and on the VM's Properties tab.

Note: The XenServer hosts's external console output (for example, VGA, HDMI, DP) will not be available after disabling the connection between dom0 and the GPU.

In Guest Driver Installation


Note

When viewing the VM console in XenCenter, the VM will typically boot to the desktop in VGA mode with a 800 x 600 resolution. The standard Windows screen resolution controls can be used to increase the resolution to other standard resolutions. (Control Panel → Display → Screen Resolution)

Install the NVIDIA drivers


To enable vGPU operation (as for a physical NVIDIA GPU) the NVIDIA driver must be installed into the VM.

1. Start the VM. In the **Resources** pane, **right-click** on the VM, and click **Start**. During this boot process XenServer dynamically allocates a vGPU to the VM.
2. Follow the Windows operating system installation screens.
3. Once the operating system installation completes, **reboot** the VM.
4. Install the appropriate driver for the GPU inside the guest. The following example shows the specific case for in guest installation of the NVIDIA GRID drivers.
 - a. Copy the 32- or 64-bit NVIDIA Windows driver package to the VM, open the zip file, and run **setup.exe**.
 - b. Accept the License Agreement.
 - c. Select **Express Installation**:



- d. Once the driver installation has completed, you may be prompted to reboot the VM. Select **Restart Now** to reboot the VM immediately, alternatively, exit the installer package, and reboot the VM when ready.

- When the VM starts it will boot to a Windows desktop.
- e. To verify that the NVIDIA driver is running, **right-click** on the desktop and select **NVIDIA Control Panel**.
 - f. In the NVIDIA Control Panel, select **System Information**.
This will show the GPU Type in use by the VM, its capabilities, and the NVIDIA driver version in use:

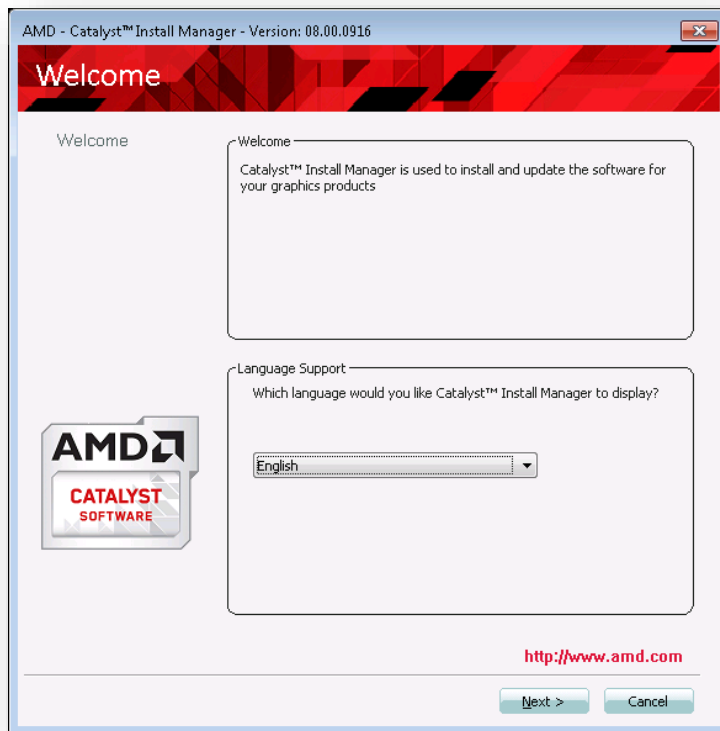


This completes the process for setting up a single VM to use GPU. The VM is now ready to run the full range of DirectX and OpenGL graphics applications supported by the GPU.

Install the AMD Drivers

To enable GPU operation the AMD drivers must be installed into the VM.

1. Start the VM. In the **Resources** pane, **right-click** on the VM, and click **Start**. During this boot process XenServer dynamically allocates a GPU to the VM.
2. Follow the Windows operating system installation screens.
3. Once the operating system installation completes, **reboot** the VM.
4. **Copy** the 32 or 64-bit AMD Windows drivers (AMD Catalyst Install Manager) to the VM.
5. Run the AMD Catalyst Install Manager; select your **Destination Folder** and then click **Install**.



6. Select your **Language**, and on the Welcome Screen select **Express**, and then click **Next**.

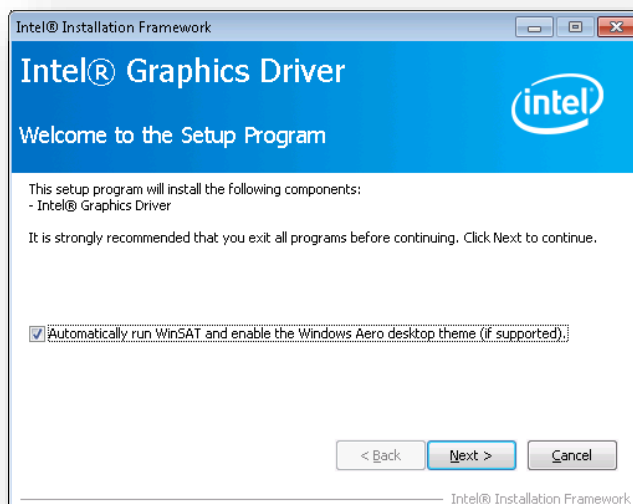


7. Accept the AMD End User License Agreement; this starts the installation.
8. When installation has completed, click **Finish**.
9. To complete the installation, **reboot** your VM.
10. When the VM restarts, check that graphics are working correctly; open the Windows **Device Manager**, expand **Display adapters**, and ensure that the AMD Graphics Adapter does not have any warning symbols.

Install the Intel Drivers

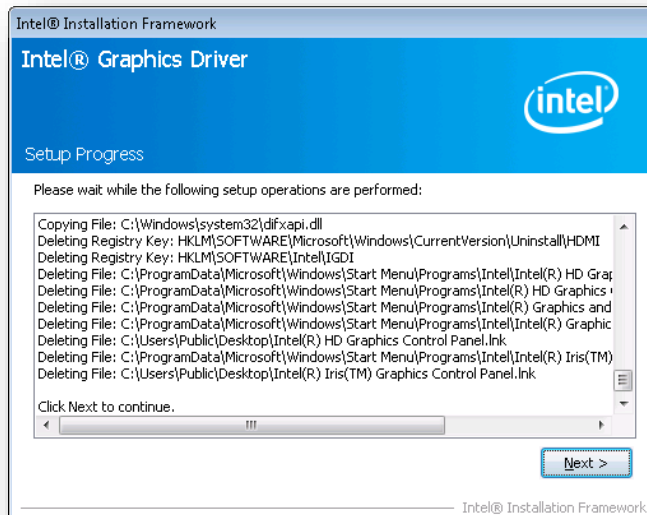
To enable GPU operation the Intel drivers must be installed into the VM.

1. Start the VM. In the **Resources** pane, **right-click** on the VM, and click **Start**. During this boot process XenServer dynamically allocates a GPU to the VM.
2. Follow the Windows operating system installation screens.
3. Once the operating system installation completes, **reboot** the VM.
4. Copy the 32 or 64-bit Intel Windows driver (Intel Graphics Driver) to the VM.
5. Run the **Intel Graphics Driver** setup program
6. Select **Automatically run WinSAT...**, and then click **Next**.



7. To accept the License Agreement, click **Yes**, and on the Readme File Information screen, click **Next**

8. Wait while setup operations complete. When you are prompted, click **Next**.



9. In order to complete the installation, you will be prompted to restart the VM. Select **Yes, I want to restart this computer now**, and click **Finish**.
10. When the VM restarts, check that graphics are working correctly; open the Windows **Device Manager**, expand **Display adapters**, and ensure that the Intel Graphics Adapter does not have any warning symbols.

Install the XenServer Tools and Configure VDA for HDX 3D Pro

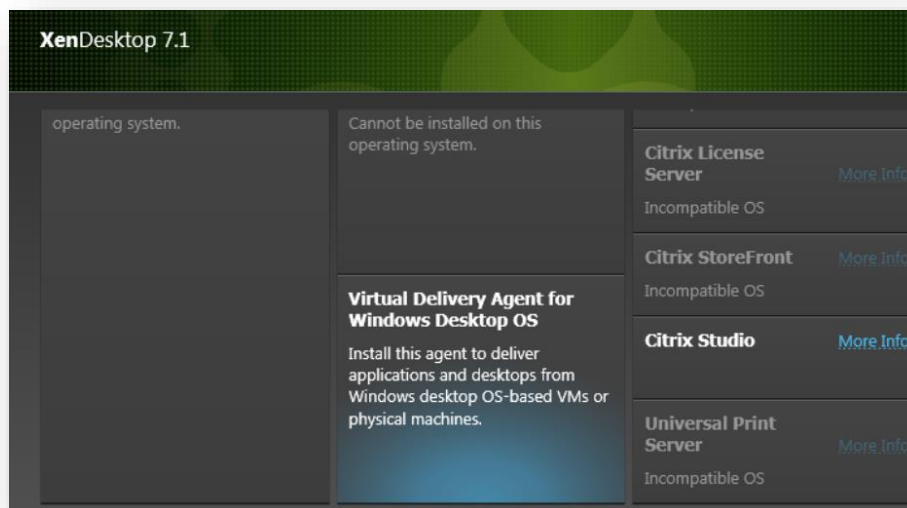
1. Install the **XenServer Tools**.

Without the optimized networking and storage drivers provided by the XenServer Tools, remote graphics applications running on GRID vGPU will **not** deliver maximum performance.

- Select the VM in the **Resources** pane, right-click, and then click **Install XenServer Tools** on the shortcut menu. Alternatively, on the **VM** menu, click **Install XenServer Tools**.
- Click **Install XenServer Tools** on the message dialog to go to the VM's console and begin the installation.
- If *Autoplay* is enabled for the VM's CD drive, installation will be started automatically after a few moments.
If Autoplay is not enabled, double-click on the CD drive to begin installing the XenServer Tools.
- Follow the on-screen instructions, and reboot the VM when prompted.

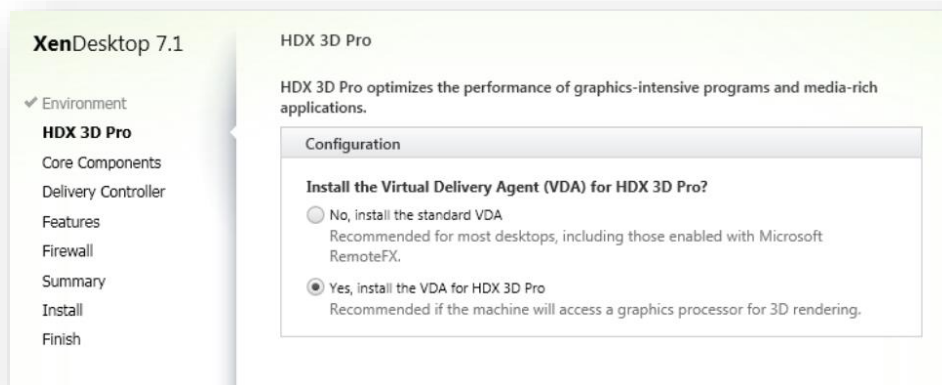
2. Mount the **XenDesktop** installation media. If autorun is not enabled, navigate to and run **AutoSelect.exe** on the installation media.

3. In the XenDesktop installation wizard, select **Virtual Delivery Agent for Windows Desktop OS**:



4. On the **Configuration** screen, select **Create A Master Image**, then click **Next**.

5. On the **HDX 3D Pro** screen, click **Yes, install the VDA for HDX 3D Pro**, and click **Next**.



6. Work through the remainder of the Installation wizard.

For detailed installation information for the XenDesktop VDA, refer to

<http://support.citrix.com/proddocs/topic/xendesktop-7/hd-3d-install.html>



Configuring XenServer v6.5.0 Service Pack 1 for Graphics

May 2015

About Citrix

Citrix (NASDAQ:CTXS) is a leader in mobile workspaces, providing virtualization, mobility management, networking and cloud services to enable new ways to work better. Citrix solutions power business mobility through secure, personal workspaces that provide people with instant access to apps, desktops, data and communications on any device, over any network and cloud. Citrix solutions are in use at more than 330,000 organizations and by over 100 million users globally. Learn more at www.citrix.com.

The copyright in this report and all other works of authorship and all developments made, conceived, created, discovered, invented or reduced to practice in the performance of work during this engagement are and shall remain the sole and absolute property of Citrix, subject to a worldwide, non-exclusive license to you for your internal distribution and use as intended hereunder. No license to Citrix products is granted herein. Citrix products must be licensed separately. Citrix warrants that the services have been performed in a professional and workman-like manner using generally accepted industry standards and practices. Your exclusive remedy for breach of this warranty shall be timely re-performance of the work by Citrix such that the warranty is met. THE WARRANTY ABOVE IS EXCLUSIVE AND IS IN LIEU OF ALL OTHER WARRANTIES, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE WITH RESPECT TO THE SERVICES OR PRODUCTS PROVIDED UNDER THIS AGREEMENT, THE PERFORMANCE OF MATERIALS OR PROCESSES DEVELOPED OR PROVIDED UNDER THIS AGREEMENT, OR AS TO THE RESULTS WHICH MAY BE OBTAINED THEREFROM, AND ALL IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR AGAINST INFRINGEMENT. Citrix's liability to you with respect to any services rendered shall be limited to the amount actually paid by you. IN NO EVENT SHALL EITHER PARTY BE LIABLE TO THE OTHER PARTY HEREUNDER FOR ANY INCIDENTAL, CONSEQUENTIAL, INDIRECT OR PUNITIVE DAMAGES (INCLUDING BUT NOT LIMITED TO LOST PROFITS) REGARDLESS OF WHETHER SUCH LIABILITY IS BASED ON BREACH OF CONTRACT, TORT, OR STRICT LIABILITY. Disputes regarding this engagement shall be governed by the internal laws of the State of Florida.

LINKS TO THIRD PARTY SITES. These release notes may contain links to web sites controlled by parties other than Citrix. Citrix is not responsible for and does not endorse or accept any responsibility for the contents or use of these third party web sites. Citrix is providing these links to you only as a convenience, and the inclusion of any link does not imply endorsement by Citrix of the linked web site. It is your responsibility to take precautions to ensure that whatever you select for your use is free of viruses or other items of a destructive nature.

Copyright © 2015 Citrix Systems, Inc. All rights reserved.

Citrix and Xen are registered trademarks. XenServer and XenCenter are trademarks of Citrix Systems, Inc. in the United States and other countries.

All other product names, company names, marks, logos, and symbols are trademarks of their respective owners.

851 West Cypress Creek Road

Fort Lauderdale, FL 33099

954-267-3000

www.citrix.com