



3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014

Background

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 XenServer's 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.

XenServer 6.2.0 Service Pack 1 includes the 3D Graphics Pack (vGPU) for XenServer and an updated version of the XenCenter management console.

Available vGPU Types

NVIDIA Grid K1 and K2 cards contain multiple Graphics Processing Units (GPU). 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.

The vGPU types supported by GRID K1 and K2 are defined in the Table 1.

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 K140Q	Workstation	2560x1600	4	16
		GRID K100	VDI	1920x1200	8	32
GRID K2	2	GRID K260Q	Workstation	2560x1600	2	4
		GRID K240Q	Workstation	2560x1600	4	8
		GRID K200	VDI	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 three types of vGPU; GRID K200, GRID 240Q, and GRID 260Q.



3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014

Pre-requisites:

- NVIDIA GRID K1 or K2 cards:
Customers who have installed Hotfix XS62ESP1004 can use additional NVIDIA GRID cards and vGPU types. Hotfix XS62ESP1004 is available from <https://support.citrix.com/article/CTX140417>
 - For the most recent vGPU types refer to <http://www.nvidia.com/object/virtual-gpus.html> .
 - For 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-boards.html>
- A server capable of hosting XenServer and NVIDIA GRID cards.
(Refer to the vGPU Release Notes at www.citrix.com/go/vgpu for details of recommended hardware)
- Citrix XenServer 6.2.0 Service Pack 1 or later
- The NVIDIA GRID vGPU software package for Citrix XenServer, consisting of the GRID Virtual GPU Manager for XenServer, and NVIDIA drivers for Windows 7 32-bit/64-bit
Available from <http://www.nvidia.com/vGPU>

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

- Citrix XenDesktop 7.1 or later full installation

Note: No other versions of Citrix XenServer or XenDesktop are currently supported for use with NVIDIA virtual GPUs.

vGPU Guest Support

The following guests are supported for use with vGPU:

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

Customers who have installed Hotfix XS62ESP1004 can also use the following guests:

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



3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014

Installation Overview:

1. [Install XenServer and Apply Service Pack 1](#)
2. [Installing the NVIDIA GRID vGPU Manager for XenServer](#)
3. [Create a vGPU Enabled VM](#)
4. [Boot the VM and install the NVIDIA drivers](#)



3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014

Install XenServer and Apply Service Pack 1

1. XenServer 6.2.0 is available from <https://www.citrix.com/downloads/xenserver/product-software/xenserver-62.html>: customers will need to install the **XenServer 6.2.0 Base Installation ISO** and **XenCenter 6.2.0 Windows Management Console**.

Refer to the [XenServer Quick Start Guide](http://support.citrix.com/article/CTX137827) for comprehensive details on installation
<http://support.citrix.com/article/CTX137827>
2. Install XenServer 6.2.0 Service Pack 1 available from <http://support.citrix.com/article/CTX139788>
3. Reboot your host.



3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014

Installing the NVIDIA GRID vGPU Manager for XenServer

Customers will need the following two components available from NVIDIA - <http://www.nvidia.com/vGPU> :

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

Note: The driver versions shown above will change as NVIDIA releases future drivers. Citrix supports customers using future releases of these drivers.

The *GRID vGPU Manager* runs in XenServer's Control Domain (dom0). It is provided as an RPM file, which must be copied to XenServer's dom0 and then installed. To do this:

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, 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:



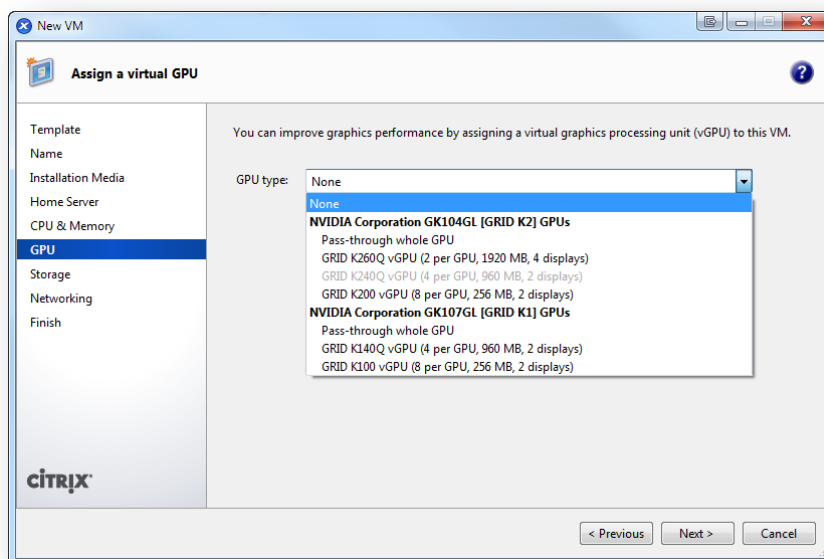
3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014

```
[root@xenserver ~]# nvidia-smi
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   0        | GRID K1     | Off  | N/A |
| N/A  27C      | P0         13W / 31W | 0%      9MB / 4095MB | 0%      Default |
+-----+-----+-----+-----+-----+-----+
|   1   1        | GRID K1     | Off  | N/A |
| N/A  25C      | P0         13W / 31W | 0%      9MB / 4095MB | 0%      Default |
+-----+-----+-----+-----+-----+-----+
|   2   2        | GRID K1     | Off  | N/A |
| N/A  21C      | P0         13W / 31W | 0%      9MB / 4095MB | 0%      Default |
+-----+-----+-----+-----+-----+-----+
|   3   3        | GRID K1     | Off  | N/A |
| N/A  23C      | P0         13W / 31W | 0%      9MB / 4095MB | 0%      Default |
+-----+-----+-----+-----+-----+-----+
|   4   4        | GRID K1     | Off  | N/A |
| N/A  24C      | P0         13W / 31W | 0%      9MB / 4095MB | 0%      Default |
+-----+-----+-----+-----+-----+-----+
|   5   5        | GRID K1     | Off  | N/A |
| N/A  24C      | P0         13W / 31W | 0%      9MB / 4095MB | 0%      Default |
+-----+-----+-----+-----+-----+-----+
|   6   6        | GRID K1     | Off  | N/A |
| N/A  25C      | P0         13W / 31W | 0%      9MB / 4095MB | 0%      Default |
+-----+-----+-----+-----+-----+-----+
|   7   7        | GRID K1     | 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 Enabled VM

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



- From the **GPU Type** drop-down list, select a vGPU type. (Unavailable vGPU types will be greyed-out)
- Click **Next** to configure **Storage** and **Networking**.
- When you are happy with your configuration, click **Create Now**.

Install the NVIDIA drivers

Notes

When viewing the VM console in XenCenter, the VM should initially boot to a standard Windows 7 desktop in VGA mode at 800 x 600 resolution.

Customers can use the standard Windows screen resolution controls to increase the resolution to other standard resolutions. (Control Panel → Display → Screen Resolution)

To fully 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 OS installation completes, reboot the VM.
4. Copy the 32- or 64-bit NVIDIA Windows driver package to the VM, open the zip file, and run **setup.exe**.
5. Accept the License Agreement.
6. Select **Express Installation**:



7. 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.
8. To verify that the NVIDIA driver is running; **right-click** on the desktop and select **NVIDIA Control Panel**.

9. In the NVIDIA Control Panel, select **System Information**.

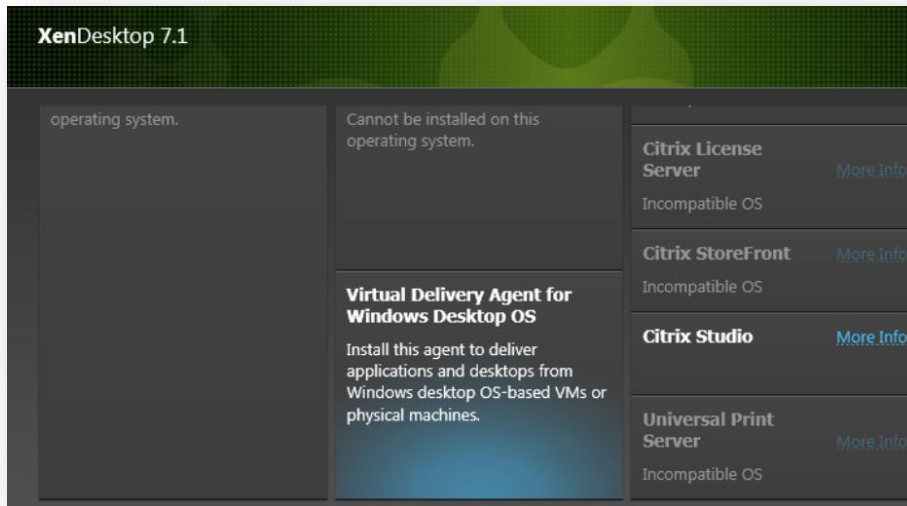
This will show the vGPU Type in use by the VM, its capabilities, and the NVIDIA driver version in use:



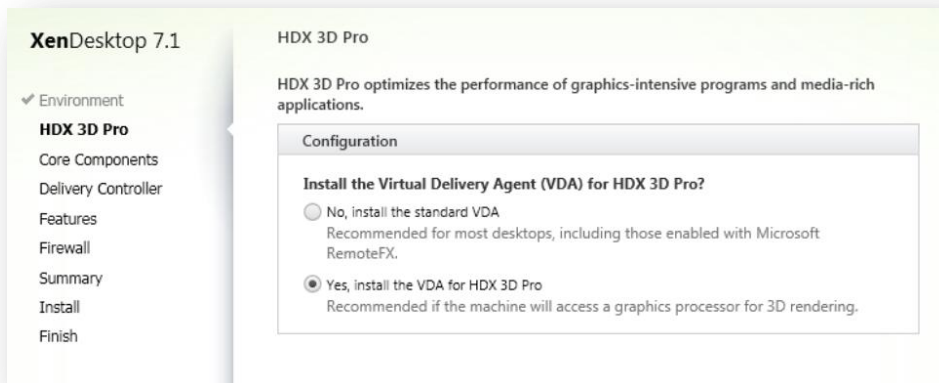
10. Install the **XenServer Tools**.

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

- a) 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**.
 - b) Click **Install XenServer Tools** on the message dialog to go to the VM's console and begin the installation.
 - c) If *Autoplay* is enabled for the VM's CD drive, installation will start automatically. If Autoplay is not enabled, double-click on the CD drive to begin installing the XenServer Tools.
 - d) Follow the on-screen instructions, and reboot the VM when prompted.
11. Mount the **XenDesktop** installation media.
If autorun is not enabled, navigate to and run **AutoSelect.exe** on the installation media.
 12. In the XenDesktop installation wizard, select **Virtual Delivery Agent for Windows Desktop OS**:



13. On the **Configuration** screen, select **Create A Master Image**, then click **Next**.
14. On the **HDX 3D Pro** screen, click **Yes, install the VDA for HDX 3D Pro**, and click **Next**.



15. 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>

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



3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014

Further Information

Customers are advised to review the *GRID VGPU FOR CITRIX XENSERVER* 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.



3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014

About Citrix

Citrix (NASDAQ:CTXS) is the cloud company that enables mobile workstyles empowering people to work and collaborate from anywhere, securely accessing apps and data on any of the latest devices, as easily as they would in their own office. Citrix solutions help IT and service providers build clouds, leveraging virtualization and networking technologies to deliver high-performance, elastic and cost-effective cloud services. With market-leading solutions for mobility, desktop virtualization, cloud networking, cloud platforms, collaboration and data sharing, Citrix helps organizations of all sizes achieve the speed and agility necessary to succeed in a mobile and dynamic world. Citrix products are in use at more than 260,000 organizations and by over 100 million users globally. Annual revenue in 2012 was \$2.59 billion. 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 © 2013 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



3D Graphics Pack: Configuring XenServer to use NVIDIA GRID v2.0

v1.2 April 2014