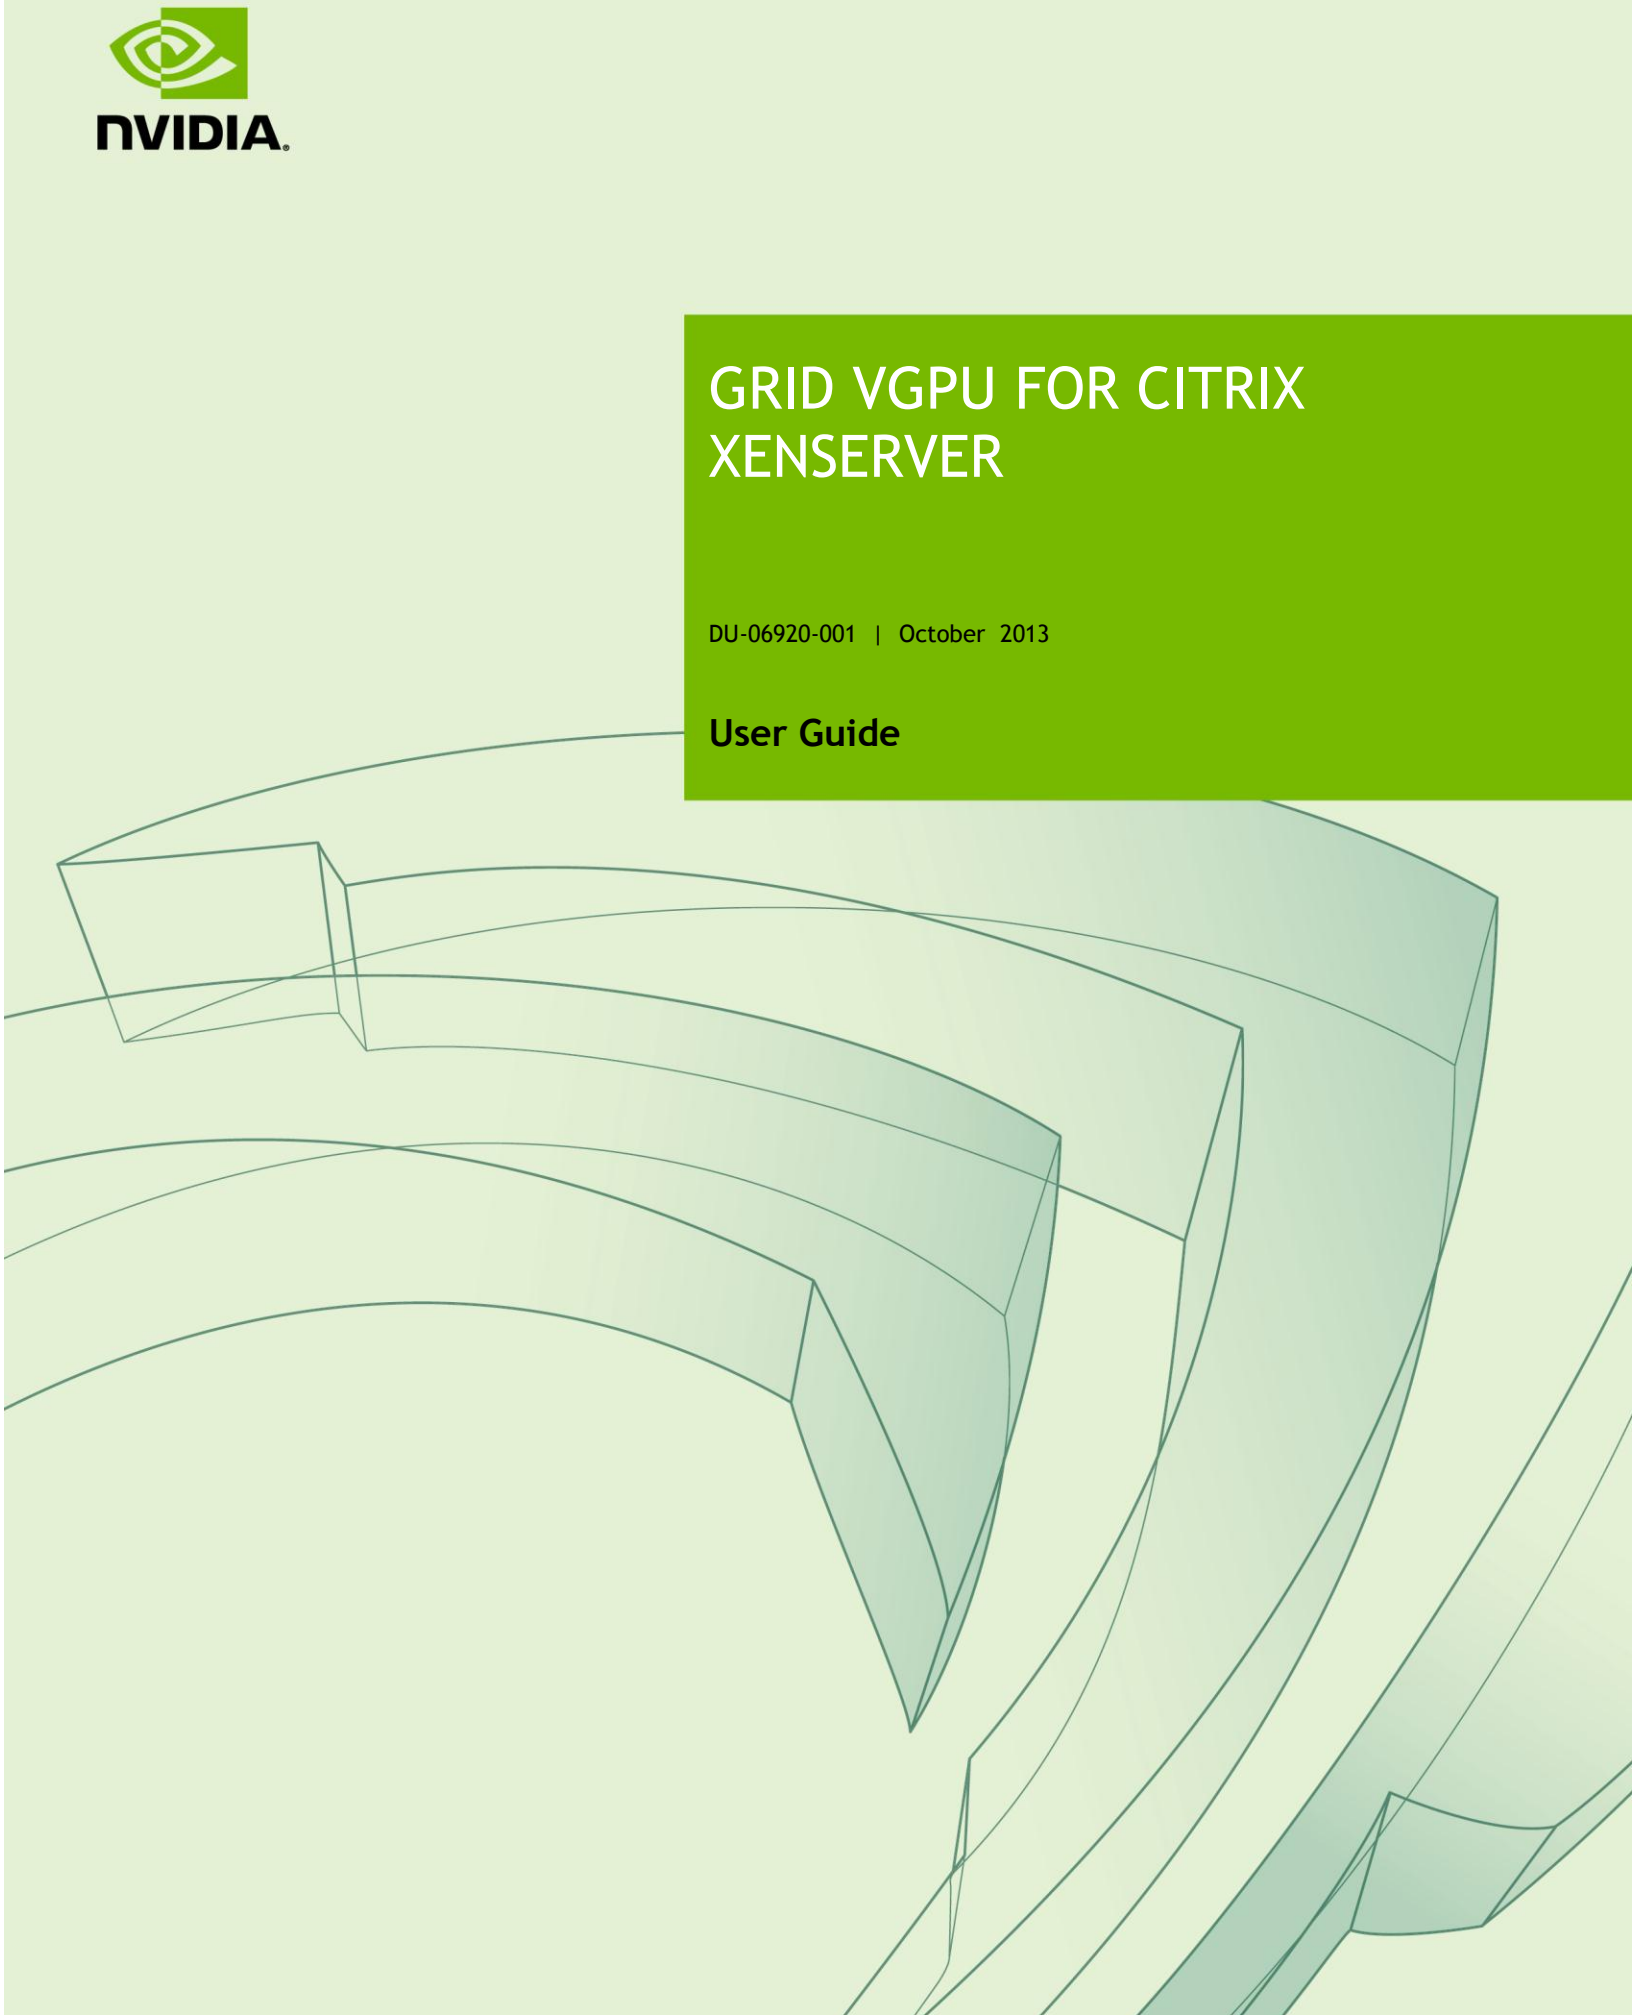

GRID VGPU FOR CITRIX XENSERVER

DU-06920-001 | October 2013

User Guide

DOCUMENT CHANGE HISTORY

DU-06920-001

Version	Date	Authors	Description of Change
0.3	7/1/2013	AC	Initial release for vGPU private beta
0.9	9/1/2013	AC	Updated for vGPU Tech Preview.

LIST OF FIGURES

LIST OF TABLES

Chapter 1. INTRODUCTION

NVIDIA GRID™ vGPU™ enables 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. By doing this, GRID vGPU provides VMs with unparalleled graphics performance and application compatibility, together with the cost-effectiveness and scalability brought about by sharing a GPU among multiple workloads.

This chapter introduces the architecture and features of vGPU. Chapter 2 provides a step-by-step guide to getting started with vGPU on Citrix XenServer. Chapter 3 covers performance optimization and management, and Chapter 4 provides guidance on troubleshooting.

1.1 ARCHITECTURE

GRID vGPU's high-level architecture is illustrated in Figure 1. Under the control of NVIDIA's GRID Virtual GPU Manager running in XenServer dom0, GRID physical GPUs are capable of supporting multiple virtual GPU devices (vGPUs) that can be assigned directly to guest 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 is used for non-performant management operations.

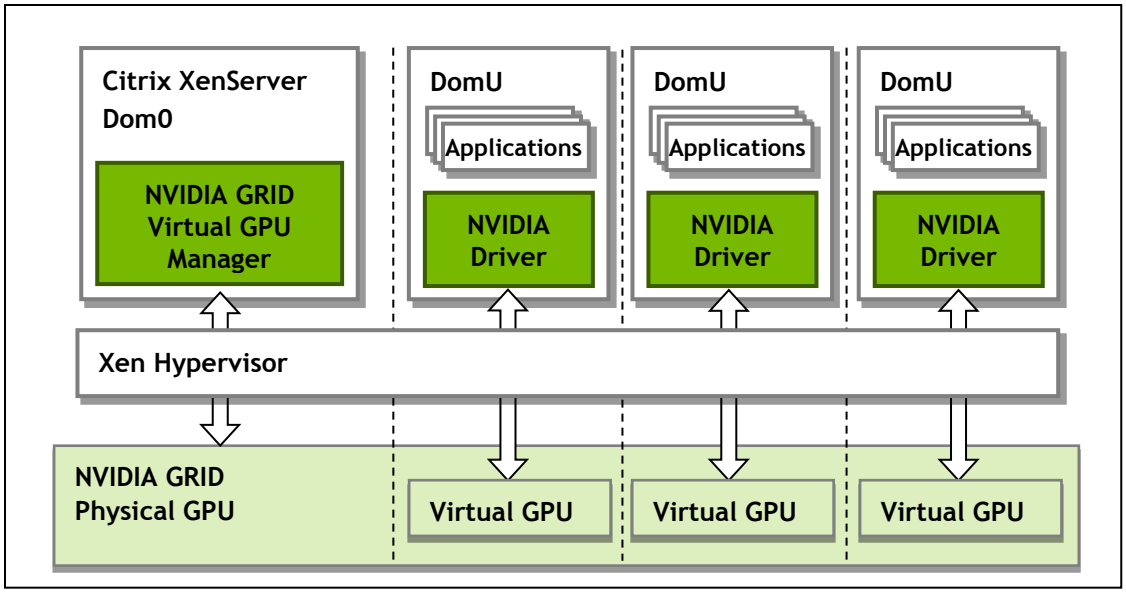

Figure 1 GRID vGPU System Architecture

GRID vGPUs are analogous to conventional GPUs, having a fixed amount of GPU framebuffer, and one or more virtual display outputs or "heads". The vGPU's framebuffer is allocated out of the physical GPU's framebuffer at the time the vGPU is created, and the vGPU retains exclusive use of that framebuffer until it is destroyed.

All vGPUs resident on a physical GPU share access to the GPU's engines including the graphics (3D), video decode, and video encode engines.

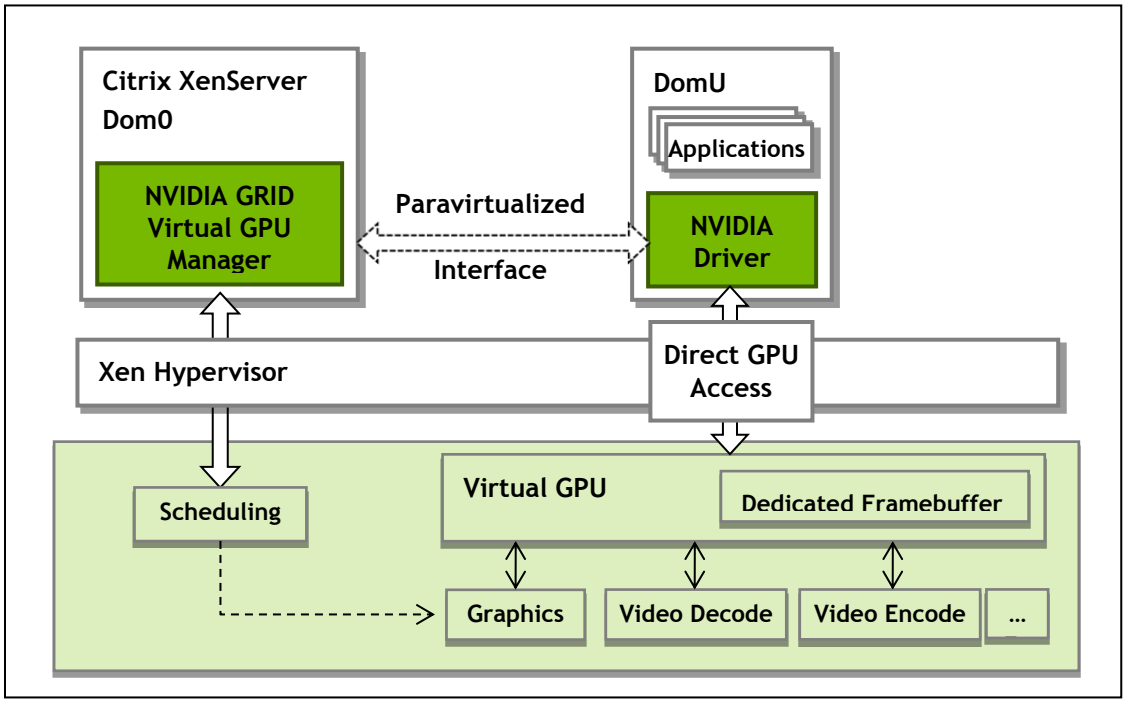

Figure 2 GRID vGPU Internal Architecture

1.2 SUPPORTED GPUS

GRID vGPU is supported on NVIDIA GRID K1 and GRID K2 GPUs. Refer to the release notes for a list of recommended server platforms to use with GRID K1 and K2.

1.2.1 Virtual GPU types

GRID K1 and K2 each implement multiple physical GPUs; GRID K2 has 2 GPUs onboard, and GRID K1 has 4 GPUs.

Each physical GPU can support several different types of virtual GPU. Virtual GPU types have a fixed amount of framebuffer, number of supported display heads and maximum resolutions, and are targeted at different classes of workload

The virtual GPU types supported by GRID K1 and K2 are defined in Table 1.

Card	Physical GPUs	Virtual GPU	Intended Use Case	Frame Buffer (Megabytes)	Virtual Display Heads	Max Resolution	Maximum vGPUs	
							Per GPU	Per Board
GRID K1	4	GRID K140Q	Power User	1024	2	2560x1600	4	16
		GRID K100	Knowledge Worker	256	2	1920x1200	8	32
GRID K2	2	GRID K260Q	Power User, Designer	2048	4	2560x1600	2	4
		GRID K240Q	Power User, Designer	1024	2	2560x1600	4	8
		GRID K200	Knowledge Worker	256	2	1920x1200	8	16

Table 1 Virtual GPU types

Due to their differing resource requirements, the maximum number of vGPUs that can be created simultaneously on a physical GPU varies according to the vGPU type. For example, a GRID K2 physical GPU can support up to 4 K240Q vGPUs on each of its two physical GPUs, for a total of 8 vGPUs, but only 2 K260Qs vGPUs, for a total of 4 vGPUs.

1.2.2 Homogeneous virtual GPUs

This release of GRID vGPU supports homogeneous virtual GPUs: at any given time, the virtual GPUs resident on a single physical GPU must be all of the same type. However, this restriction doesn't extend across physical GPUs on the same card. Each physical GPU on a K1 or K2 may host different types of virtual GPU at the same time.

For example, a GRID K2 card has two physical GPUs, and can support three types of virtual GPU; GRID K200, GRID K240Q, and GRID K260Q. Figure 3 shows some example virtual GPU configurations on K2:

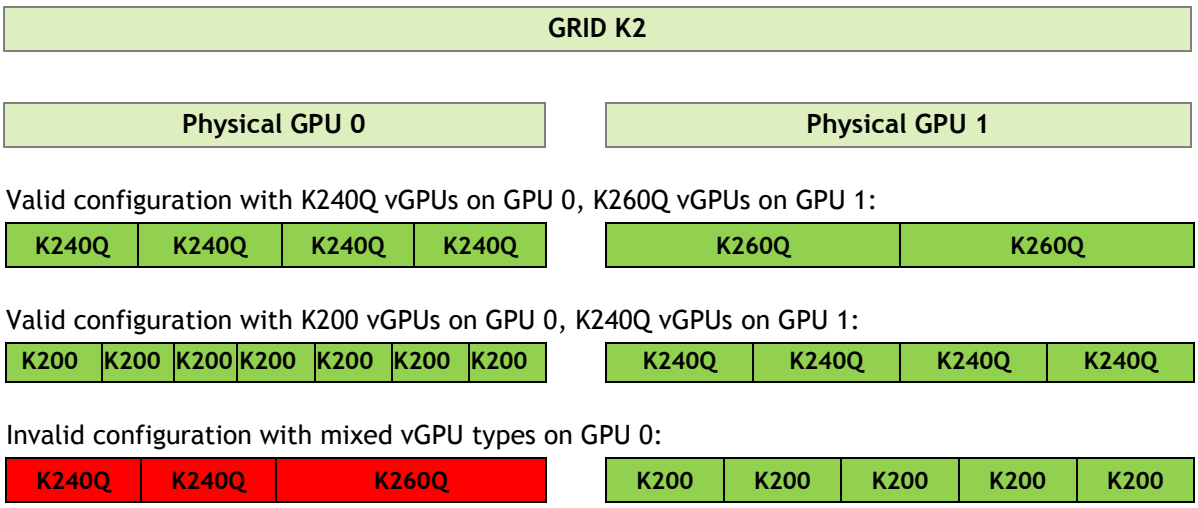

Figure 3 Example vGPU configurations on GRID K2

1.3 GUEST OS

This release of GRID vGPU includes support for Windows 7 32- and 64-bit guest VMs.

Note: At the time of writing, support for Windows Server 2008 R2, Windows Server 2012, and Windows 8 VMs is experimental. Windows Vista, Windows XP, and Linux VMs are not supported.

1.4 FEATURES

GRID vGPU 1.0 includes support for:

- Full DirectX 9/10/11, Direct2D, and DirectX Video Acceleration (DXVA)
- OpenGL 4.3.
- NVIDIA GRID SDK (remote graphics acceleration).

The following are not supported by the vGPU 1.0 release:

- CUDA, OpenCL

Chapter 2. GETTING STARTED

This chapter provides a step-by-step guide to booting a VM on XenServer with NVIDIA Virtual GPU, and assumes familiarity with the XenServer skills covered in Appendix A.

2.1 PREREQUISITES

Before proceeding, ensure that you have these prerequisites:

- NVIDIA GRID K1 or K2 cards.
- A server platform capable of hosting XenServer and the NVIDIA GRID cards. Refer to the release notes for a list of recommended servers.
- The NVIDIA GRID vGPU software package for Citrix XenServer, consisting of the GRID Virtual GPU Manager for XenServer, and NVIDIA drivers for Windows, 32- and 64-bit.
- The vGPU-enabled Tech Preview build of Citrix XenServer 6.2, obtainable from Citrix.
- An installed Windows 7 VM (32- or 64-bit) to be enabled with vGPU.

To run Citrix XenDesktop with virtual machines running NVIDIA Virtual GPU, you will also need:

- The vGPU-enabled Tech Preview build of Citrix XenDesktop, obtainable from Citrix.

Note: No other versions of Citrix XenServer and XenDesktop are currently supported for use with NVIDIA Virtual GPU.

Review the release rotes and known issues for GRID Virtual GPU before proceeding with installation.

2.2 INSTALLING CITRIX XENSERVER

Install Citrix XenServer and any applicable patches, following Citrix's installation instructions.

2.3 INSTALLING THE NVIDIA VIRTUAL GPU MANAGER FOR XENSERVER

The NVIDIA Virtual GPU Manager runs in XenServer's dom0. It is provided as an RPM file, which must be copied to XenServer's dom0 and then installed.

CAUTION: It is not possible to upgrade from the Private Beta version of GRID Virtual GPU Manager to the Tech Preview version. If you have the Private Beta package installed on XenServer, first uninstall it -

```
[root@xenserver ~]# rpm -qa | grep NVIDIA
NVIDIA-vgx-312.38-xenserver-6-2
[root@xenserver ~]# rpm -ev NVIDIA-vgx-312.38-xenserver-6-2
[root@xenserver ~]#
```

- then proceed with installation of the Tech Preview package.

2.3.1 Package installation

Use the `rpm` command to install the package:

```
[root@xenserver ~]# rpm -iv  NVIDIA-vgx-xenserver-6.2-312.47.i386.rpm
Preparing packages for installation...
NVIDIA-vgx-xenserver-6.2-312.47
[root@xenserver ~]#
```

Reboot the XenServer platform:

```
[root@xenserver ~]# shutdown -r now

Broadcast message from root (pts/1) (Fri Aug 30 14:24:11 2013):

The system is going down for reboot NOW!
[root@xenserver ~]#
```

2.3.2 Update installation

If an existing GRID Virtual GPU Manager is already installed on the system and you wish to upgrade, follow these steps:

- Shut down any VMs that are using GRID vGPU.

- Install the new package using the `-U` option to the `rpm` command, to upgrade from the previously installed package:

```
[root@xenserver ~]# rpm -Uv NVIDIA-vgx-xenserver-6.2-312.47.i386.rpm
Preparing packages for installation...
NVIDIA-vgx-xenserver-6.2-312.47
[root@xenserver ~]#
```

Note: You can query the version of the current GRID package using the `rpm -q` command:

```
[root@xenserver ~]# rpm -q NVIDIA-vgx-xenserver
NVIDIA-vgx-xenserver-6.2-312.47
[root@xenserver ~]#
```

If an existing NVIDIA GRID package is already installed and you don't select the upgrade (-U) option when installing a newer GRID package, the rpm command will return many conflict errors.

```
Preparing packages for installation...

        file /usr/bin/nvidia-smi from install of NVIDIA-vgx-
xenserver-6.2-312.47.i386 conflicts with file from package
NVIDIA-vgx-xenserver-6.2-312.46.i386

        file /usr/lib/libnvidia-ml.so from install of NVIDIA-
vgx-xenserver-6.2-312.47.i386 conflicts with file from package
NVIDIA-vgx-xenserver-6.2-312.46.i386

        ...
```

Reboot the XenServer platform:

```
[root@xenserver ~]# shutdown -r now

Broadcast message from root (pts/1) (Fri Aug 30 14:24:11 2013):

The system is going down for reboot NOW!
[root@xenserver ~]#
```

2.3.3 Verifying installation

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

```
[root@xenserver ~]# lsmod | grep nvidia
nvidia               8512994  0
i2c_core               20294  2 nvidia,i2c_i801
[root@xenserver ~]#
```

Verify that the NVIDIA kernel driver can successfully communicate with the GRID physical GPUs in your system by running the `nvidia-smi` command, which should produce a listing of the GPUs in your platform:

```
[root@xenserver ~]# nvidia-smi
Fri Jun 21 14:31:55 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 ~]#
```

The `nvidia-smi` command is described in more detail in section 3.7.

If `nvidia-smi` fails to run or doesn't produce the expected output for all the NVIDIA GPUs in your system, see Chapter 4 for troubleshooting steps.

2.4 CONFIGURING A VM WITH VIRTUAL GPU

This release of virtual GPU for XenServer supports configuration of virtual GPUs using the `xe` command line tool that is run in a XenServer dom0 shell.

> **CAUTION:** The Citrix XenCenter GUI does not currently support GRID vGPU configuration. All vGPU and passthrough GPU configuration should be done using the `xe` command line tool.

2.4.1 XenServer management objects for GPUs

Management of GPUs using `xe` is abstracted via four management objects; physical GPUs, GPU groups, vGPU types, and vGPUs.

2.4.1.1 pgpu

A `pgpu` object represents a physical GPU, such as one of the GPUs present on a GRID K1 or K2 card. XenServer automatically creates `pgpu` objects at startup to represent each physical GPU present on the platform.

To list the physical GPU objects present on a platform, use `xe pgpu-list`. For example, this platform contains a single GRID K2 card with two physical GPUs:

```
[root@xenserver ~]# xe pgpu-list
uuid ( RO)              : 7c1e3cff-1429-0544-df3d-bf8a086fb70a
       vendor-name ( RO): NVIDIA Corporation
       device-name ( RO): GK104GL [GRID K2]
    gpu-group-uuid ( RW): be825ba2-01d7-8d51-9780-f82cfaa64924

uuid ( RO)              : d07fa627-7dc9-f625-92be-ce5d2655e830
       vendor-name ( RO): NVIDIA Corporation
       device-name ( RO): GK104GL [GRID K2]
    gpu-group-uuid ( RW): be825ba2-01d7-8d51-9780-f82cfaa64924
[root@xenserver ~]#
```

To see detailed information about a `pgpu`, use `xe pgpu-param-list`:

```
[root@xenserver ~]# xe pgpu-param-list uuid=d07fa627-7dc9-f625-92be-
ce5d2655e830
uuid ( RO)                    : d07fa627-7dc9-f625-92be-ce5d2655e830
             vendor-name ( RO): NVIDIA Corporation
             device-name ( RO): GK104GL [GRID K2]
          gpu-group-uuid ( RW): be825ba2-01d7-8d51-9780-f82cfaa64924
    gpu-group-name-label ( RO): Group of NVIDIA Corporation GK104GL [GRID K2]
GPUs
               host-uuid ( RO): 6f6209a6-0f11-4c51-b12d-2bce361e9639
         host-name-label ( RO): xenserver (VM IPs 10.31.213.50-95, dom0 .98,
OOB .99)
                  pci-id ( RO): 0000:08:00.0
            dependencies (SRO):
            other-config (MRW):
    supported-VGPU-types ( RO): fa50b0f0-9705-6c59-689e-ea62a3d35237; 1a312df9-
5397-bd44-c447-c6da804d2fe7; d1fb00dd-02e6-e7df-ccd5-1944965ece55; 3f318889-
7508-c9fd-7134-003d4d05ae56
      enabled-VGPU-types (SRW): fa50b0f0-9705-6c59-689e-ea62a3d35237; 1a312df9-
5397-bd44-c447-c6da804d2fe7; d1fb00dd-02e6-e7df-ccd5-1944965ece55; 3f318889-
7508-c9fd-7134-003d4d05ae56
          resident-VGPUs ( RO): bddadca3-5b84-b9a6-091e-e64794efe6f5

[root@xenserver ~]#
```

2.4.1.2 vgpu-type

A `vgpu-type` represents a type of virtual GPU, such as GRID K100, K140Q, K200, etc. An additional, *passthrough* vGPU type is defined to represent a physical GPU that is directly assignable to a single guest VM.

XenServer automatically creates `vgpu-type` objects at startup to represent each virtual type supported by the physical GPUs present on the platform.

To list the `vgpu-type` objects present on a platform, use `xe vgpu-type-list`. For example, this platform contains multiple GRID K2 cards, therefore the vGPU types reported are solely those supported by GRID K2:

```
[root@xenserver ~]# xe vgpu-type-list
uuid ( RO)                : 3f318889-7508-c9fd-7134-003d4d05ae56
         vendor-name ( RO): NVIDIA Corporation
          model-name ( RO): GRID K240Q
    framebuffer-size ( RO): 1006632960

uuid ( RO)                : fa50b0f0-9705-6c59-689e-ea62a3d35237
         vendor-name ( RO):
          model-name ( RO): passthrough
    framebuffer-size ( RO): 0

uuid ( RO)                : 1a312df9-5397-bd44-c447-c6da804d2fe7
```

```
         vendor-name ( RO): NVIDIA Corporation
          model-name ( RO): GRID K200
    framebuffer-size ( RO): 268435456

uuid ( RO)                : d1fb00dd-02e6-e7df-ccd5-1944965ece55
         vendor-name ( RO): NVIDIA Corporation
          model-name ( RO): GRID K260Q
    framebuffer-size ( RO): 2013265920

 [root@xenserver ~]#
```

To see detailed information about a `vgpu-type`, use `xe vgpu-type-param-list`:

```
[root@xenserver ~]# xe vgpu-type-param-list uuid=3f318889-7508-c9fd-7134-
003d4d05ae56
uuid ( RO)                  : 3f318889-7508-c9fd-7134-003d4d05ae56
           vendor-name ( RO): NVIDIA Corporation
            model-name ( RO): GRID K240Q
      framebuffer-size ( RO): 1006632960
              max-heads ( RO): 2
    supported-on-PGPUs ( RO): a4a4df34-4e5f-de9f-82d6-2134d9e339dc; 84c76e93-
555c-5ffa-e9a9-0d6fcb9ff48d; d07fa627-7dc9-f625-92be-ce5d2655e830; 7c1e3cff-
1429-0544-df3d-bf8a086fb70a
      enabled-on-PGPUs ( RO): a4a4df34-4e5f-de9f-82d6-2134d9e339dc; 84c76e93-
555c-5ffa-e9a9-0d6fcb9ff48d; d07fa627-7dc9-f625-92be-ce5d2655e830; 7c1e3cff-
1429-0544-df3d-bf8a086fb70a
            VGPU-uuids ( RO): 8481cb68-66e5-25e6-a0c0-bd691df682b3; b73cbd30-
096f-8a9a-523e-a800062f4ca7

[root@xenserver ~]#
```

2.4.1.3 gpu-group

A `gpu-group` is a collection of physical GPUs, all of the same type. XenServer automatically creates `gpu-group` objects at startup to represent the distinct types of physical GPU present on the platform.

To list the `gpu-group` objects present on a platform, use `xe gpu-group-list`. For example, a system with a single GRID K2 card contains a single GPU group of type GRID K2:

```
[root@xenserver ~]# xe gpu-group-list
uuid ( RO)                : be825ba2-01d7-8d51-9780-f82cfaa64924
          name-label ( RW): Group of NVIDIA Corporation GK104GL [GRID K2] GPUs
    name-description ( RW):

[root@xenserver ~]#
```

To see detailed information about a `gpu-group`, use `xe gpu-group-param-list`:

```
[root@xenserver ~]# xe gpu-group-param-list uuid=be825ba2-01d7-8d51-9780-
f82cfaa64924
uuid ( RO)                    : be825ba2-01d7-8d51-9780-f82cfaa64924
             name-label ( RW): Group of NVIDIA Corporation GK104GL [GRID K2]
GPUs
        name-description ( RW):
              VGPU-uuids (SRO): 6493ff45-d895-764c-58d8-96f1bc0307aa; 8481cb68-
66e5-25e6-a0c0-bd691df682b3; b73cbd30-096f-8a9a-523e-a800062f4ca7
              PGPU-uuids (SRO): a4a4df34-4e5f-de9f-82d6-2134d9e339dc; 84c76e93-
555c-5ffa-e9a9-0d6fcb9ff48d; d07fa627-7dc9-f625-92be-ce5d2655e830; 7c1e3cff-
1429-0544-df3d-bf8a086fb70a
           other-config (MRW):
      enabled-VGPU-types ( RO): d1fb00dd-02e6-e7df-ccd5-1944965ece55; 1a312df9-
5397-bd44-c447-c6da804d2fe7; fa50b0f0-9705-6c59-689e-ea62a3d35237; 3f318889-
7508-c9fd-7134-003d4d05ae56
    supported-VGPU-types ( RO): d1fb00dd-02e6-e7df-ccd5-1944965ece55; 1a312df9-
5397-bd44-c447-c6da804d2fe7; fa50b0f0-9705-6c59-689e-ea62a3d35237; 3f318889-
7508-c9fd-7134-003d4d05ae56
    allocation-algorithm ( RW): depth-first

[root@xenserver ~]
```

2.4.1.4 vgpu

A `vgpu` object represents a virtual GPU. Unlike the other GPU management objects, vGPUs are not created automatically by XenServer.

2.4.2 Creating a vGPU

Use `xe vgpu-create` to create a `vgpu` object, specifying the type of vGPU required, the GPU group it will be allocated from, and the VM it is associated with:

```
[root@xenserver ~]# xe vgpu-create vm-uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2
gpu-group-uuid=be825ba2-01d7-8d51-9780-f82cfaa64924 vgpu-type-uuid=3f318889-
7508-c9fd-7134-003d4d05ae56
b73cbd30-096f-8a9a-523e-a800062f4ca7
[root@xenserver ~]#
```

Creating the vGPU object for a VM does not immediately cause a virtual GPU to be created on a physical GPU. Instead, the vGPU is created whenever its associated VM is started. For more details on how vGPUs are created at VM startup, see section 3.3.

Note: the owning VM must be in the powered-off state in order for the `vgpu-create` command to succeed.

A `vgpu` object's owning VM, associated GPU group, and vGPU type are fixed at creation and cannot be subsequently changed. To change the type of vGPU allocated to a VM, delete the existing `vgpu` object and create another one (see section 3.53.5).

2.5 BOOTING THE VM AND INSTALLING DRIVERS

Once you have configured a VM with a vGPU, start the VM, either from XenCenter or by using `xe vm-start` in a dom0 shell.

Viewing the VM's console in XenCenter, the VM should boot to a standard Windows desktop in VGA mode at 800x600 resolution. The Windows screen resolution control panel may be used to increase the resolution to other standard resolutions, but to fully enable vGPU operation, as for a physical NVIDIA GPU, the NVIDIA driver must be installed.

- Copy the 32- or 64-bit NVIDIA Windows driver package to the guest VM, open the zipped driver package and run `setup.exe`:

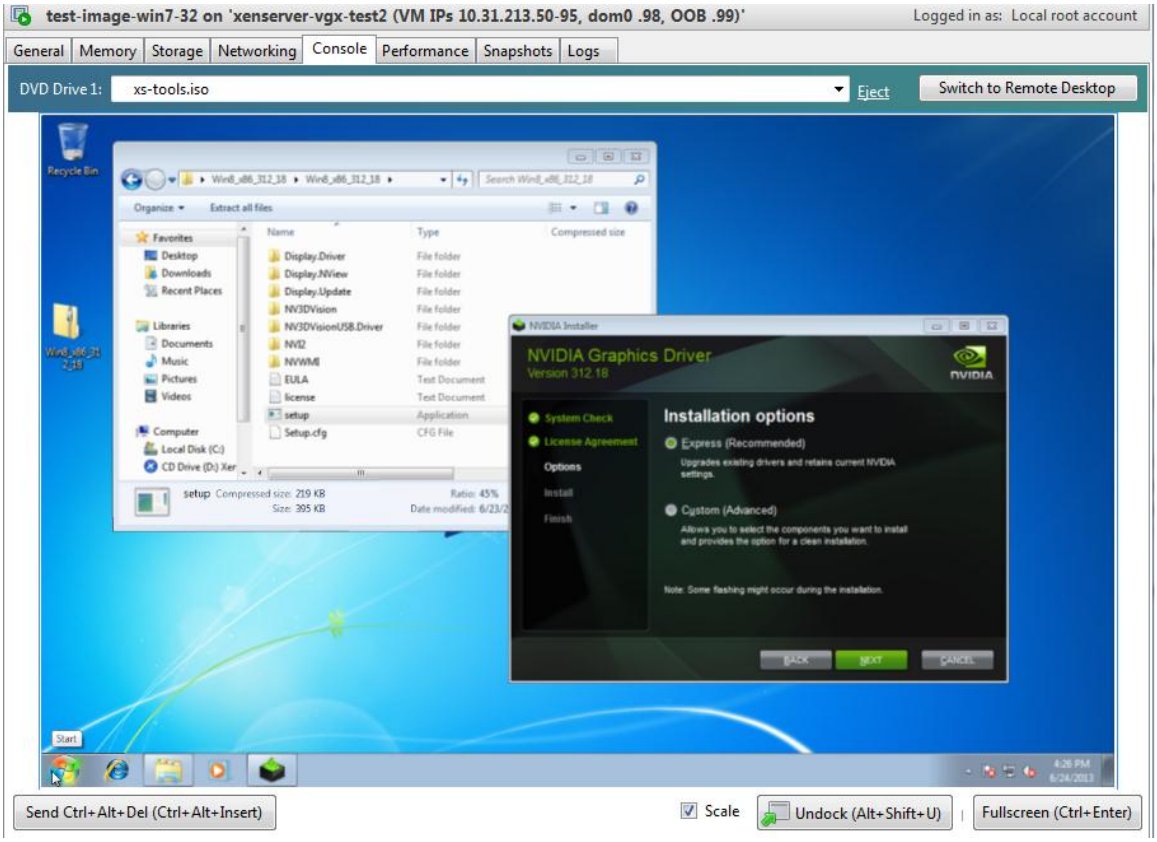

Figure 4 NVIDIA driver installation in the guest VM

- Click through the license agreement
- Select Express Installation
- Once driver installation completes, the installer may prompt you to restart the platform. Select Restart Now to reboot the VM, or exit the installer and reboot the VM when ready.

Once the VM restarts, it will boot to a Windows desktop. Verify that the NVIDIA driver is running by right-clicking on the desktop. The NVIDIA Control Panel will be listed in the menu; select it to open the control panel. Selecting "System Information" in the NVIDIA control panel will report the Virtual GPU that the VM is using, its capabilities, and the NVIDIA driver version that is loaded.

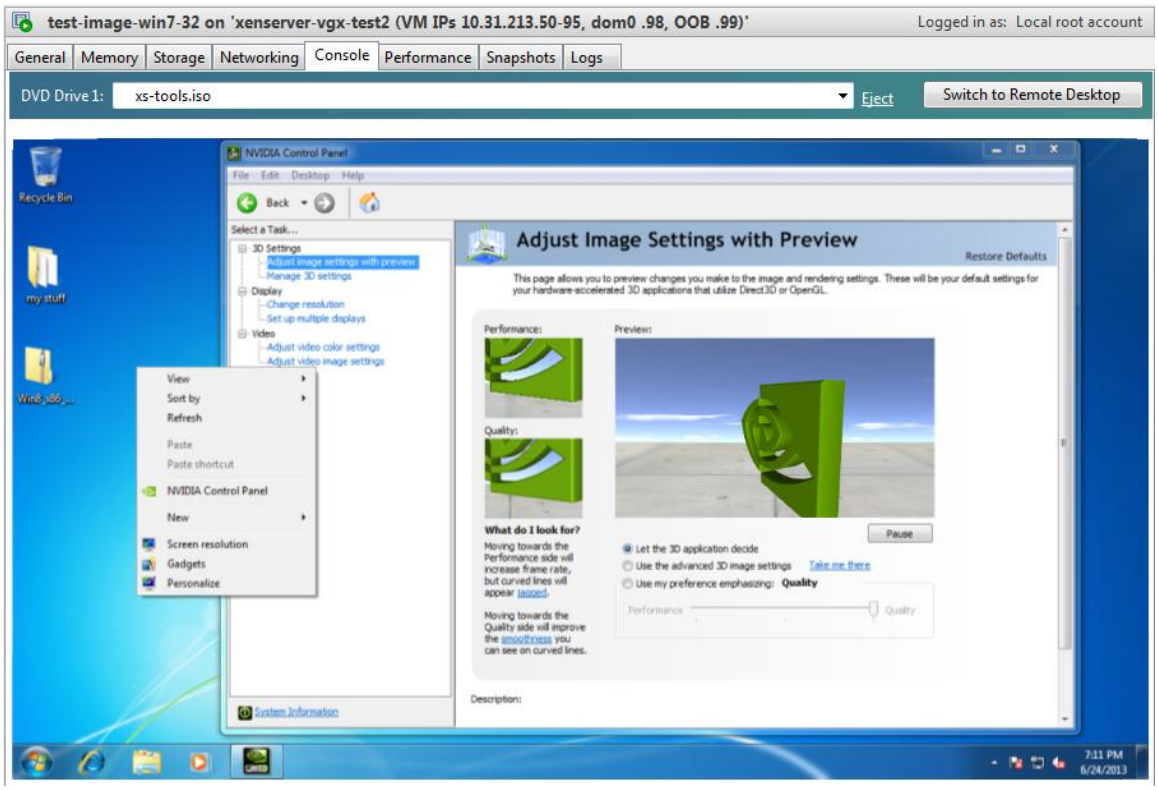

Figure 5 Verifying NVIDIA driver operation using NVIDIA Control Panel

This completes the process of setting up a single VM to use GRID vGPU. The VM is now capable of running the full range of DirectX and OpenGL graphics applications, but in order to deliver the full performance and capabilities of vGPU, additional configuration steps are required. These are reviewed in Chapter 3.

Chapter 3. MANAGEMENT AND PERFORMANCE OPTIMIZATION

This chapter describes vGPU management and techniques to optimize the performance of VMs running with GRID vGPU.

3.1 INSTALL XENSERVER TOOLS

To get maximum performance out of a VM running on Citrix XenServer, regardless of whether you are using GRID vGPU, you must install Citrix XenServer tools within the VM. Without the optimized networking and storage drivers that the XenServer tools provide, remote graphics applications running on GRID vGPU will not deliver maximum performance.

3.2 USING REMOTE GRAPHICS

GRID vGPU implements a console VGA interface that permits the VM's graphics output to be viewed via XenCenter's console tab. This feature allows the desktop of a vGPU-enabled VM to be visible in XenCenter before any NVIDIA graphics driver is loaded in the virtual machine, but it is intended solely as a management convenience; it only supports output of vGPU's primary display and isn't designed or optimized to deliver high frame rates.

To deliver high frames from multiple heads on vGPU, we recommend installation of a high-performance remote graphics stack such as Citrix XenDesktop® with HDX 3D Pro remote graphics and, once this is done, disable vGPU's console VGA.

CAUTION: Using Windows Remote Desktop (RDP) to access Windows 7 / Windows Server 2008 VMs running GRID vGPU will cause the NVIDIA driver in the VM to be unloaded. GPU-accelerated DirectX, OpenGL, and the NVIDIA control panel will be unavailable whenever RDP is active. Installing a VNC server in the VM will allow for basic, low-performance remote access while leaving the NVIDIA driver loaded and vGPU active, but for high performance remote accesses, use an accelerated stack such as XenDesktop.

3.2.1 Disabling console VGA

Once you have installed an alternate means of accessing a VM (such as XenDesktop or a VNC server), we recommend disabling GRID vGPU's console VGA feature. Do this by specifying `vgpu_vnc_enabled=false` in the VM's `platform` parameter:

```
[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2
platform:vgpu_vnc_enabled=false
[root@xenserver ~]#
```

The new console VGA setting takes effect the next time the VM is started or rebooted.

CAUTION: If you disable console VGA before you have installed/enabled an alternate mechanism to access the VM (such as XenDesktop), you will not be able to interact with the VM once it has booted. The XenCenter console will display the Windows boot splash screen, and nothing else.

You can recover console VGA access by removing the `vgpu_vnc_enabled` key from the `platform` parameter or by setting it to `true`:

```
[root@xenserver ~]# xe vm-param-set uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2
platform:vgpu_vnc_enabled=true
```

3.3 CONTROLLING VGPU ALLOCATION

Creation of a `vgpu` object for a VM does not immediately cause a virtual GPU to be created; rather, the virtual GPU is created at the time the VM is next booted or rebooted, using the following steps:

- The GPU group that the `vgpu` object is associated with is checked for a physical GPU that can host a vGPU of the required type (i.e. the `vgpu` object's associated `vgpu-type`). Because vGPU types cannot be mixed on a single physical GPU, the new vGPU can only be created on a physical GPU that has no vGPUs resident on it, or only vGPUs of the same type, and less than the limit of vGPUs of that type that the physical GPU can support.
- If no such physical GPUs exist in the group, the `vgpu` creation fails and the VM startup is aborted.

- Otherwise, if more than one such physical GPU exists in the group, a physical GPU is selected according to the GPU group's *allocation policy*, as described here:

3.3.1 GPU allocation policy

XenServer creates GPU groups with a default allocation policy of *depth-first*. The depth-allocation policy attempts to maximize the number of vGPUs running on each physical GPU within the group, by placing newly-created vGPUs on the physical GPU that can support the new vGPU and that has the most number of vGPUs already resident. This policy generally leads to higher density of vGPUs, particularly when different types of vGPUs are being run, but may result in lower performance because it attempts to maximize sharing of physical GPUs.

Conversely, a *breadth-first* allocation policy attempts to minimize the number of vGPUs running on each physical GPU within the group, by placing newly-created vGPUs on the physical GPU that can support the new vGPU and that has the least number of vGPUs already resident. This policy generally leads to higher performance because it attempts to minimize sharing of physical GPUs, but in doing so it may artificially limit the total number of vGPUs that can run.

The allocation policy of a GPU group is stored in the `allocation-algorithm` parameter of the `gpu-group` object, and can be changed using `gpu-group-param-set`:

```
[root@xenserver ~]# xe gpu-group-param-get uuid=be825ba2-01d7-8d51-9780-
f82cfaa64924 param-name=allocation-algorithm
depth-first
[root@xenserver ~]# xe gpu-group-param-set uuid=be825ba2-01d7-8d51-9780-
f82cfaa64924 allocation-algorithm=breadth-first
[root@xenserver ~]#
```

3.3.2 Determining the physical GPU that a virtual GPU is resident on

The `vgpu` object's `resident-on` parameter returns the UUID of the `pgpu` object for the physical GPU the vGPU is resident on:

```
[root@xenserver ~]# xe vgpu-param-get uuid=8481cb68-66e5-25e6-a0c0-bd691df682b3
param-name=resident-on
a4a4df34-4e5f-de9f-82d6-2134d9e339dc

[root@xenserver ~]# xe pgpu-param-list uuid=a4a4df34-4e5f-de9f-82d6-
2134d9e339dc
uuid ( RO)                    : a4a4df34-4e5f-de9f-82d6-2134d9e339dc
             vendor-name ( RO): NVIDIA Corporation
             device-name ( RO): GK104GL [GRID K2]
          gpu-group-uuid ( RW): be825ba2-01d7-8d51-9780-f82cfaa64924
```

```
    gpu-group-name-label ( RO): Group of NVIDIA Corporation GK104GL [GRID K2]
GPUs
                host-uuid ( RO): 6f6209a6-0f11-4c51-b12d-2bce361e9639
          host-name-label ( RO): xenserver (VM IPs 10.31.213.50-95, dom0 .98,
OOB .99)
                   pci-id ( RO): 0000:09:00.0
             dependencies (SRO):
             other-config (MRW):
    supported-VGPU-types ( RO): fa50b0f0-9705-6c59-689e-ea62a3d35237; 1a312df9-
5397-bd44-c447-c6da804d2fe7; d1fb00dd-02e6-e7df-ccd5-1944965ece55; 3f318889-
7508-c9fd-7134-003d4d05ae56
      enabled-VGPU-types (SRW): fa50b0f0-9705-6c59-689e-ea62a3d35237; 1a312df9-
5397-bd44-c447-c6da804d2fe7; d1fb00dd-02e6-e7df-ccd5-1944965ece55; 3f318889-
7508-c9fd-7134-003d4d05ae56
          resident-VGPUs ( RO): 8481cb68-66e5-25e6-a0c0-bd691df682b3

[root@xenserver ~]#
```

Note: If the vGPU is not currently running, the `resident-on` parameter is not instantiated for the vGPU, and the `vgpu-param-get` operation returns `<not in database>`.

3.3.3 Creating vGPUs on specific physical GPUs

To precisely control allocation of vGPUs on specific physical GPUs, create a separate GPU group for each physical GPU you wish to allocate vGPUs on. When creating a virtual GPU, create it on the GPU group containing the physical GPU you want it to be allocated on.

For example, to create a new GPU for the physical GPU at PCI bus ID `0000:05:0.0`, start by creating the new GPU group with an appropriate name:

```
[root@xenserver ~]# xe gpu-group-create name-label="GRID K2 5:0.0"
585877ef-5a6c-66af-fc56-7bd525bdc2f6
[root@xenserver ~]#
```

Next, find the UUID of the physical GPU at `0000:05:0.0` that you wish to assign to the new GPU group

```
[root@xenserver ~]# xe pgpu-list pci-id=0000:05:00.0
uuid ( RO)              : 7c1e3cff-1429-0544-df3d-bf8a086fb70a
       vendor-name ( RO): NVIDIA Corporation
       device-name ( RO): GK104GL [GRID K2]
```

```
    gpu-group-uuid ( RW): be825ba2-01d7-8d51-9780-f82cfaa64924

[root@xenserver ~]
```

Note: the pci-id parameter passed to the pgpu-list command must be in the exact format shown, with the PCI domain fully specified (e.g. `0000`) and the PCI bus and devices numbers each being two digits (e.g. `05:00.0`).

Ensure that no vGPUs are currently operating on the physical GPU by checking the `resident-VGPUs` parameter:

```
[root@xenserver ~]# xe pgpu-param-get uuid=7c1e3cff-1429-0544-df3d-bf8a086fb70a
param-name=resident-VGPUs

[root@xenserver ~]#
```

If any vGPUs are listed, shut down the VMs associated with them.

Now change the `gpu-group-uuid` parameter of the physical GPU to the UUID of the newly-created GPU group:

```
[root@xenserver ~]# xe pgpu-param-set uuid=d07fa627-7dc9-f625-92be-ce5d2655e830
gpu-group-uuid=585877ef-5a6c-66af-fc56-7bd525bdc2f6
[root@xenserver ~]#
```

Any `vgpu` object now created that specifies this GPU group UUID will always have its vGPUs created on the GPU at PCI bus ID `0000:05:0.0`.

Note: you can add more than one physical GPU to a manually-created GPU group - for example, to represent all the GPUs attached to the same CPU socket in a multi-socket server platform - but as for automatically-created GPU groups, all the physical GPUs in the group must of the same type.

3.4 CLONING VGPU-ENABLED VMS

XenServer's fast-clone or copying feature can be used to rapidly create new VMs from a "golden" base VM image that has been configured with GRID vGPU, the NVIDIA

driver, applications, and remote graphics software. Cloning/copying can be initiated via XenCenter (see Figure 6) or from the dom0 shell:

```
[root@xenserver ~]# xe vm-clone new-name-label="new-vm" vm="base-vm-name"
7f7035cb-388d-1537-1465-1857fb6498e7
[root@xenserver ~]#
```

When a VM is cloned, any vGPU configuration associated with the base VM is copied to the cloned VM. Starting the cloned VM will create a vGPU instance of the same type as the original VM.

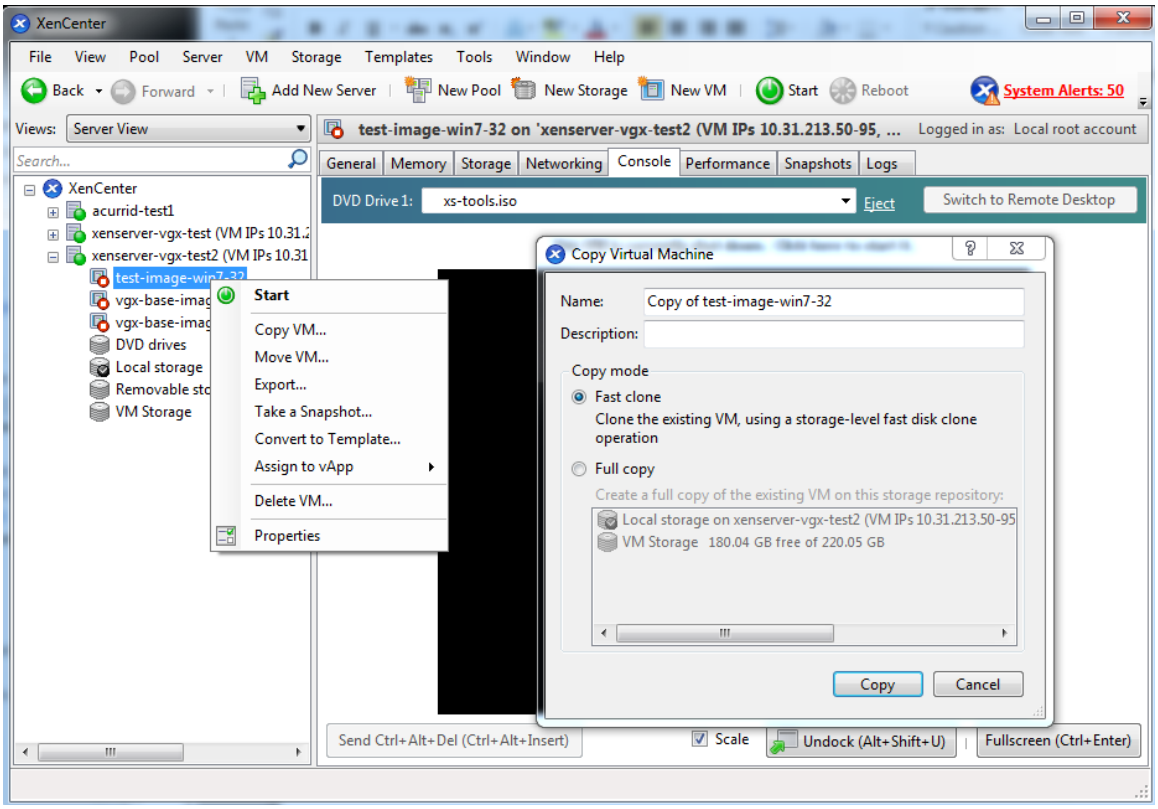

Figure 6 Cloning a VM using XenCenter

3.5 REMOVING A VGPU CONFIGURATION FROM A VM

To remove a virtual GPU assignment from a VM, such that it no longer uses a virtual GPU, use `vgpu-destroy` to delete the virtual GPU object.

```
[root@xenserver ~]# xe vgpu-list vm-uuid=e71afda4-53f4-3a1b-6c92-a364a7f619c2
uuid ( RO)              : c1c7c43d-4c99-af76-5051-119f1c2b4188
```

```
           vm-uuid ( RO): e71afda4-53f4-3a1b-6c92-a364a7f619c2
    gpu-group-uuid ( RO): d53526a9-3656-5c88-890b-5b24144c3d96

[root@xenserver ~]# xe vgpu-destroy uuid=c1c7c43d-4c99-af76-5051-119f1c2b4188
[root@xenserver ~]#
```

Note: the VM must be in the powered-off state for the `vgpu-destroy` operation to succeed.

3.6 ALLOCATION STRATEGIES

3.6.1 NUMA considerations

Server platforms typically implement multiple CPU sockets, with system memory and PCI Express expansion slots local to each CPU socket, as illustrated in Figure 7.

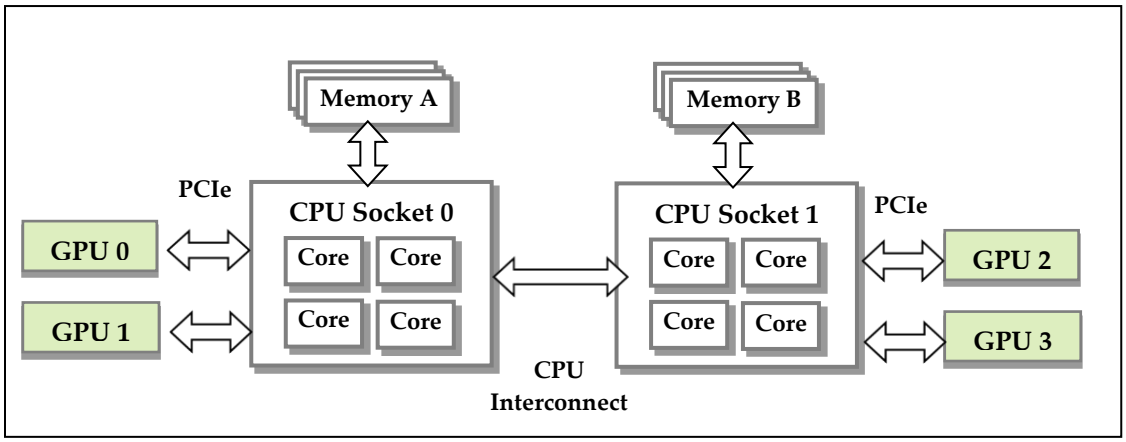

Figure 7 A NUMA server platform

These platforms are typically configured to operate in Non-Uniform Memory Access (NUMA) mode; physical memory is arranged sequentially in the address space, with all the memory attached to each socket appearing in a single contiguous block of addresses. The cost of accessing a range of memory from a CPU or GPU varies; memory attached to the same socket as the CPU or GPU is accessible at lower latency than memory on another CPU socket, because accesses to remote memory must additionally traverse the interconnect between CPU sockets.

To obtain best performance on a NUMA platform, we recommend pinning VM vCPU cores to physical cores on the same CPU socket to which the physical GPU hosting the VM's vGPU is attached. For example, using Figure 7 as a reference, a VM with a vGPU

allocated on physical GPU 0 or 1 should have its vCPUs pinned to CPU cores on CPU socket 0. Similarly, a VM with a vGPU allocated on physical GPU 2 or 3 should have its vCPUs pinned to CPU cores on socket 1.

See Appendix A.5 for guidance on pinning vCPUs, and A.7 for guidance on determining which CPU socket a GPU is connected to. Section 3.3.3 describes how to precisely control which physical GPU is used to host a vGPU, by create GPU groups for specific physical GPUs.

3.6.2 Maximizing performance

To maximize performance as the number of vGPU-enabled VMs on the platform increases, we recommend adopting a *breadth-first* allocation: allocate new VMs on the least-loaded CPU socket, and allocate the VM's vGPU on an available, least-loaded, physical GPU connected via that socket.

XenServer's creates GPU groups with a default allocation policy of *depth-first*. See section 3.3.1 for details on switching the allocation policy to breadth-first.

Note: Due to vGPU's requirement that only one type of vGPU can run on a physical GPU at any given time, not all physical GPUs may be available to host the vGPU type required by the new VM.

3.7 USING NVIDIA-SMI

NVIDIA System Management Interface, `nvidia-smi`, is a command line tool that reports management information for NVIDIA physical GPUs present in the system. `nvidia-smi` is run from the dom0 shell, and when invoked without additional arguments, it provides a summary of all GPUs in the system, along with PCI bus IDs, power state, temperature, current memory usage, and so on.

In this release of GRID vGPU, `nvidia-smi` provides basic reporting of vGPU instances running on physical GPUs; each vGPU instance is reported in the "Compute processes" section, together with its physical GPU index and the amount of framebuffer memory assigned to it. In the example that follows, five vGPUs are running; one on physical GPU 0, and four on physical GPU 1:

```
[root@xenserver ~]# nvidia-smi
Tue Jun 25 18:33:21 2013
+------------------------------------------------------+
| NVIDIA-SMI 4.312.37   Driver Version: 312.37         |
|-------------------------------+----------------------+----------------------+
| 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    P8     8W /  31W |   7%  270MB / 4095MB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   1  GRID K1                  | 0000:05:00.0     Off |                  N/A |
| N/A   26C    P8     8W /  31W |  26% 1048MB / 4095MB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   2  GRID K1                  | 0000:06:00.0     Off |                  N/A |
| N/A   22C    P0    13W /  31W |   0%    9MB / 4095MB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   3  GRID K1                  | 0000:07:00.0     Off |                  N/A |
| N/A   25C    P0    13W /  31W |   0%    9MB / 4095MB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   4  GRID K1                  | 0000:86:00.0     Off |                  N/A |
| N/A   27C    P0    14W /  31W |   0%    9MB / 4095MB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   5  GRID K1                  | 0000:87:00.0     Off |                  N/A |
| N/A   27C    P0    13W /  31W |   0%    9MB / 4095MB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   6  GRID K1                  | 0000:88:00.0     Off |                  N/A |
| N/A   29C    P0    13W /  31W |   0%    9MB / 4095MB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   7  GRID K1                  | 0000:89:00.0     Off |                  N/A |
| N/A   28C    P0    12W /  31W |   0%    9MB / 4095MB |      0%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Compute processes:                                               GPU Memory |
|  GPU       PID  Process name                                     Usage      |
|=============================================================================|
|    0     10300  /usr/lib/xen/bin/vgpu                             256MB  |
|    1     10350  /usr/lib/xen/bin/vgpu                             256MB  |
|    1     10321  /usr/lib/xen/bin/vgpu                             256MB  |
|    1     11512  /usr/lib/xen/bin/vgpu                             256MB  |
|    1     10210  /usr/lib/xen/bin/vgpu                             256MB  |
+-----------------------------------------------------------------------------+
[root@xenserver ~]#
```

3.8 USING GPU PASS-THROUGH

GPU pass-through is used to directly assign an entire physical GPU to one VM, bypassing the GRID Virtual GPU Manager. GPU pass-through can be used in a server platform alongside GRID vGPU, with some restrictions:

- A GRID physical GPU can host GRID vGPUs, or be used for pass-through, but cannot do both at the same time.
- When GRID vGPUs are active, the `other-config:pci` parameter setting must not be used to manage pass-through, otherwise a GPU may erroneously be assigned to a vGPU while already being used for passthrough. Use only the `xe gpu-group` interface to manage passthrough.

3.8.1 Configuring a VM to use GPU pass-through

To configure a VM for GPU pass-through, create a `vgpu` object with the "passthrough" vGPU type:

```
[root@xenserver ~]# xe vgpu-type-list model-name="passthrough"
uuid ( RO)                : fa50b0f0-9705-6c59-689e-ea62a3d35237
         vendor-name ( RO):
          model-name ( RO): passthrough
    framebuffer-size ( RO): 0

[root@xenserver ~]# xe vgpu-create vm-uuid=753e77a9-e10d-7679-f674-65c078abb2eb
vgpu-type-uuid=fa50b0f0-9705-6c59-689e-ea62a3d35237 gpu-group-uuid=585877ef-
5a6c-66af-fc56-7bd525bdc2f6
6aa530ec-8f27-86bd-b8e4-fe4fde8f08f9
[root@xenserver ~]#
```

3.8.2 Removing a GPU pass-through configuration from a VM

To remove a pass-through GPU assignment from a VM, use `vgpu-destroy` to delete the VM's virtual GPU object, as described in section 3.5.

Chapter 4. TROUBLESHOOTING

This chapter describes basic troubleshooting steps and how to collect debug information when filing a bug report.

4.1 KNOWN ISSUES

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

4.2 TROUBLESHOOTING STEPS

If a vGPU-enabled VM fails to start, or doesn't display any output when it does start, follow these steps to narrow down the probable cause.

4.2.1 Verify the NVIDIA kernel driver is loaded

Use lsmod to verify that the kernel driver is loaded:

```
[root@xenserver-vgx-test2 ~]# lsmod|grep nvidia
nvidia               8525302  84
i2c_core               20294  2 nvidia,i2c_i801
[root@xenserver-vgx-test2 ~]#
```

If the nvidia driver is not listed in the output, check `dmesg` for any load-time errors reported by the driver (see section 4.2.3). Also use the 'rpm' command to verify that the NVIDIA GPU Manager package is correctly installed (see section 2.3).

If the `nvidia` kernel driver appears to be correctly loaded, verify that is active on the physical GPU that a VM's virtual GPU is configured to use. First, confirm the PCI bus ID

of the physical GPU the VM's vGPU is using, following the steps described in section 3.3.2. Then use `lspci -k` to verify that the `nvidia` driver is loaded on the physical GPU:

```
[root@xenserver ~]# lspci -k -s 0:4:0.0
04:00.0 VGA compatible controller: NVIDIA Corporation GK107GL [GRID K1] (rev
a1)
        Subsystem: NVIDIA Corporation Device 1012
        Kernel driver in use: nvidia
        Kernel modules: nvidia
[root@xenserver ~]#
```

If the "kernel driver in use" is `pciback` rather than `nvidia`, the GPU is being used for passthrough and cannot be simultaneously used to support vGPU. Use a different physical GPU, or shutdown the VM using the physical GPU in passthrough mode, and unload `pciback`.

Once unloaded, or if no driver is listed as in use, load the NVIDIA kernel driver:

```
[root@xenserver ~] modprobe nvidia
```

4.2.2 Verify that nvidia-smi works

If the NVIDIA kernel driver is correctly loaded on the physical GPU, run `nvidia-smi` and verify that all physical GPUs are listed in the output. For details on expected output, see section 3.7.

If `nvidia-smi` fails to report the expected output, check `dmesg` and `/var/log/messages` for NVIDIA kernel driver messages.

4.2.3 dmesg output

Information and debug messages from the NVIDIA kernel driver are logged in dmesg, prefixed with "`NVRM`" or '`nvidia`':

```
[root@xenserver ~]# dmesg | grep -E "NVRM|nvidia"
[   22.054928] nvidia: module license 'NVIDIA' taints kernel.
[   22.390414] NVRM: loading
[   22.829226] nvidia 0000:04:00.0: enabling device (0000 -> 0003)
[   22.829236] nvidia 0000:04:00.0: PCI INT A -> GSI 32 (level, low) -> IRQ 32
[   22.829240] NVRM: This PCI I/O region assigned to your NVIDIA device is
invalid:
[   22.829241] NVRM: BAR0 is 0M @ 0x0 (PCI:0000:00:04.0)
[   22.829243] NVRM: The system BIOS may have misconfigured your GPU.
```

4.2.4 /var/log/messages

Information and debug messages from the GRID Virtual GPU Manager are written to `/var/log/messages`, prefixed with '`vmiop`':

```
[root@xenserver ~]# grep vmiop /var/log/messages
Jul  5 17:03:42 xenserver fe: vgpu-4[24578]: vmiop_log: notice: vmiop-env:
guest_max_gpfn:0xfffff
Jul  5 17:03:42 xenserver fe: vgpu-4[24578]: vmiop_log: notice: pluginconfig:
/usr/share/nvidia/vgx/grid_k200.conf,disable_vnc=1,gpu-pci-id=0:4:0.0
Jul  5 17:03:42 xenserver fe: vgpu-4[24578]: vmiop_log: notice: Loading
Plugin0: libnvidia-vgx
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: vgpu_type : vdi
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: Framebuffer:
0x10000000
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: Virtual Device
Id: 0x118D:0x101D
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: ######## Host
NVIDIA Driver Information: ########
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: Driver Version:
312.38
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: Host Version:
rel/gpu_drv/r312/r312_00-3781
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: Host Title:
Private r312_00 rel/gpu_drv/r312/r312_00-3781 unknown
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: Changelist:
16364079
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: ########
libnvidia-vgx.so Information: ########
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: Title: Private
r312_00 rel/gpu_drv/r312/r312_00-3783 unknown
Jul  5 17:03:43 xenserver fe: vgpu-4[24578]: vmiop_log: notice: ChangeList:
16364079
```

4.3 FILING A BUG REPORT

When filing a bug report with NVIDIA, capture relevant configuration data from the platform exhibiting the bug, using `nvidia-bug-report.sh`, or the XenServer server status report.

4.3.1 nvidia-bug-report.sh

Run `nvidia-bug-report.sh` from the dom0 shell to capture debug information from your XenServer installation into a gzip'd log file on the server:

```
[root@xenserver ~]# nvidia-bug-report.sh
```

```
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 Linux 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.

For Xen open source/XCP users, if you are reporting a domain issue,
please run: nvidia-bug-report.sh --domain-name <"domain_name">

Please include the 'nvidia-bug-report.log.gz' log file when reporting
your bug via the NVIDIA Linux forum (see devtalk.nvidia.com)
or by sending email to 'linux-bugs@nvidia.com'.

Running nvidia-bug-report.sh...

If the bug report script hangs after this point consider running with
--safe-mode command line argument.

 complete.

[root@xenserver ~]#
```

4.3.2 XenServer status report

From XenCenter, select the Tools menu, Server Status Report, then select the XenServer instance from which you wish to collect a status report. Select the data to include in the report, check “NVIDIA-logs” to include GRID vGPU debug information, then generate the report.

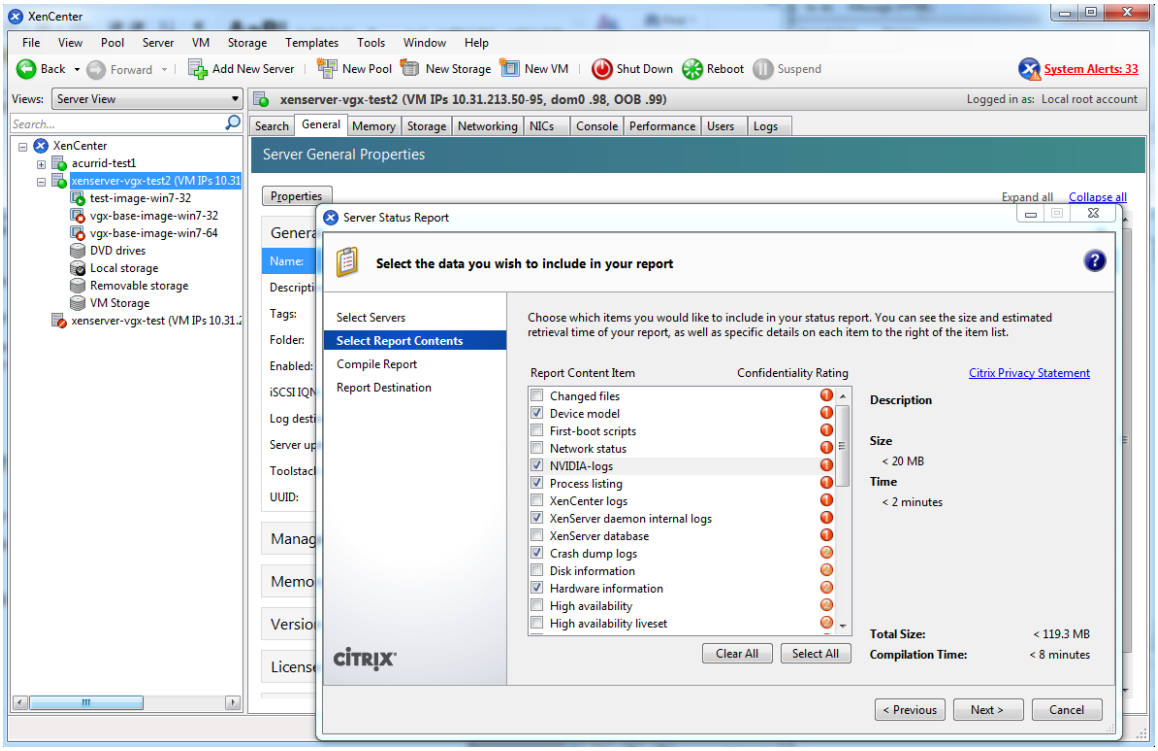

Figure 8 Including NVIDIA logs in server status report

APPENDIX A. XenServer Basics

This appendix outlines basic operations on XenServer that are needed in order to install and configure GRID vGPU, and optimize XenServer operation with vGPU.

A.1. Opening a dom0 shell

Most configuration commands must be run in a command shell on XenServer's dom0. There are two ways to open a shell on XenServer's dom0; using the console window in XenCenter, or using a standalone secure shell (ssh) client:

A.1.1 Accessing the dom0 shell via XenCenter

To access the dom0 shell via XenCenter, in the left-hand pane click on the XenServer host you wish to connect to. Then click on the Console tab to open the XenServer's console, and press enter to start a shell prompt:

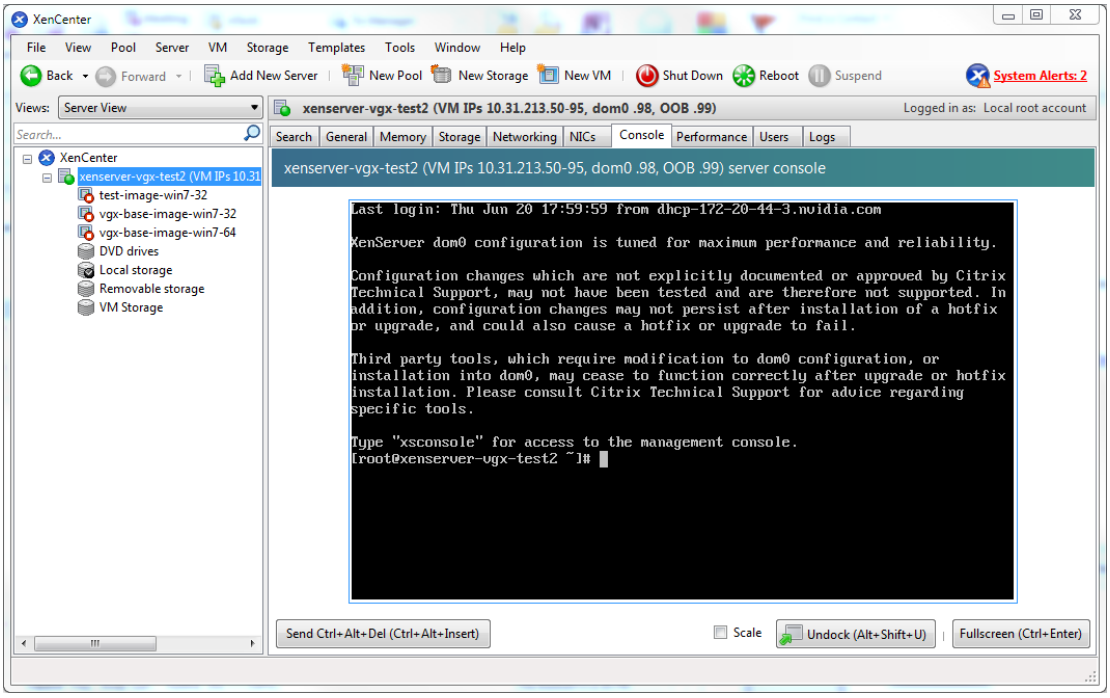

Figure 9 Connecting to the dom0 shell via XenCenter

A.1.2 Accessing the dom0 shell using ssh

To access the dom0 shell via an ssh client, you will need an ssh client suite such as *putty* on Windows, or the ssh client from OpenSSH on Linux.

Connect your ssh client to the management IP address of the XenServer, and log in as the root user.

A.2. Copying files to dom0

Files can be easily copied to/from XenServer dom0 using an scp client or using a network-mounted filesystem.

A.2.1 Copying files using scp

scp is a secure copy program that is part of the ssh suite of applications. scp is implemented in dom0 and can be used to copy from a remote ssh-enabled server:

```
[root@xenserver ~]# scp root@10.31.213.96:/tmp/somefile .
The authenticity of host '10.31.213.96 (10.31.213.96)' can't be established.
RSA key fingerprint is 26:2d:9b:b9:bf:6c:81:70:36:76:13:02:c1:82:3d:3c.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.31.213.96' (RSA) to the list of known hosts.
```

```
root@10.31.213.96's password:
somefile                                      100%  532     0.5KB/s   00:00
[root@xenserver ~]#
```

Alternatively, scp can be used to copy files from a remote system to XenServer. Using the `pscp` program from the putty suite on Windows:

```
C:\Users\nvidia>pscp somefile root@10.31.213.98:/tmp
root@10.31.213.98's password:
somefile             | 80 kB |  80.1 kB/s | ETA: 00:00:00 | 100%

C:\Users\nvidia>
```

A.2.2 Copying files via an CIFS-mounted filesystem

Files can be copied to/from a CIFS/SMB file share by mounting the share from dom0.

The following example assumes that the fileshare is part of an Active Directory domain called `domain.com`, and user `myuser` has permissions to access the share. To mount a network share `\\myserver.domain.com\myshare` at `/mnt/myshare` on dom0,

```
[root@xenserver ~]# mkdir /mnt/myshare
[root@xenserver ~]# mount -t cifs -o username=myuser,workgroup=domain.com
//myserver.domain.com/myshare /mnt/myshare
Password:
[root@xenserver ~]#
```

When prompted for a password, enter the password for `myuser` in the `domain.com` domain. After completion, files can be copied to/from the fileshare by copying to/from /mnt/myshare:

```
[root@xenserver ~]# cp /mnt/myshare/NVIDIA-vgx-312.36-xenserver-6-2.i386.rpm .
[root@xenserver ~]#
```

A.3. Determining a VM's UUID

To determine a virtual machine's UUID, use the `xe vm-list` command in a dom0 shell, or XenCenter:

A.3.1 Using xe vm-list

To list all VMs and their associated UUIDs, use `xe vm-list`:

```
[root@xenserver ~]# xe vm-list
uuid ( RO)           : 6b5585f6-bd74-2e3e-0e11-03b9281c3ade
     name-label ( RW): vgx-base-image-win7-64
    power-state ( RO): halted


uuid ( RO)           : fa3d15c7-7e88-4886-c36a-cdb23ed8e275
     name-label ( RW): test-image-win7-32
    power-state ( RO): halted


uuid ( RO)           : 501bb598-a9b3-4afc-9143-ff85635d5dc3
     name-label ( RW): Control domain on host: xenserver
    power-state ( RO): running


uuid ( RO)           : 8495adf7-be9d-eee1-327f-02e4f40714fc
     name-label ( RW): vgx-base-image-win7-32
    power-state ( RO): halted
```

To find the UUID of a specific named VM, use the `name-label` parameter to `xe vm-list`:

```
[root@xenserver ~]# xe vm-list name-label=test-image-win7-32
uuid ( RO)           : fa3d15c7-7e88-4886-c36a-cdb23ed8e275
     name-label ( RW): test-image-win7-32
    power-state ( RO): halted
```

A.3.2 Using XenCenter

In the left-hand pane click on the VM, then click on the General tab . The UUID is listed in the VM's General Properties.

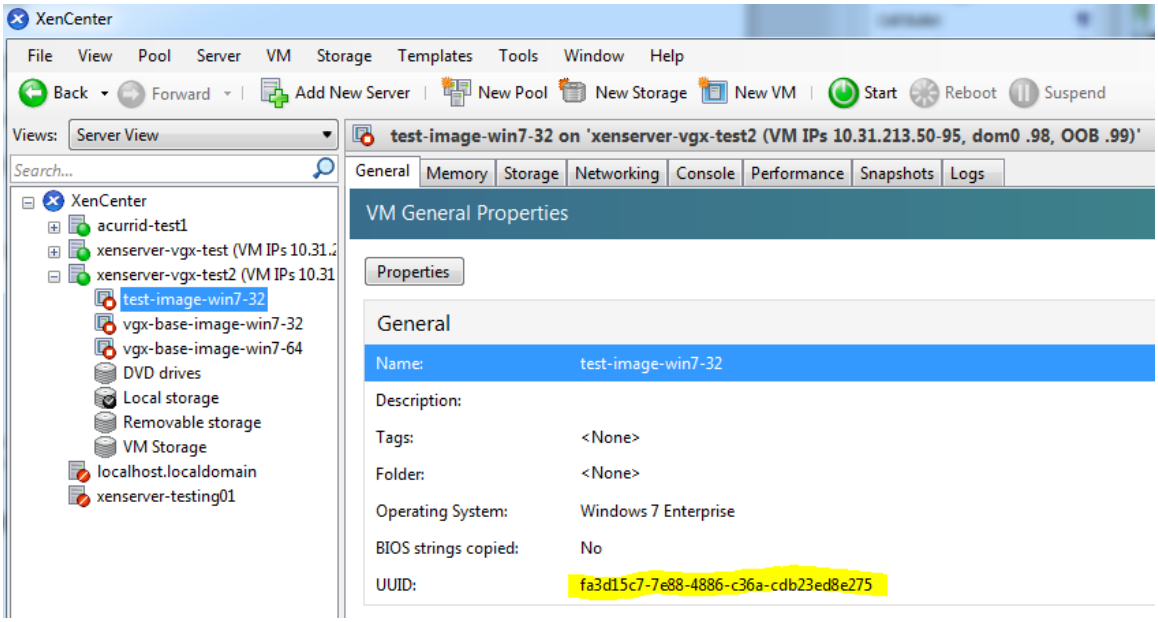

Figure 10 Using XenCenter to determine a VM's UUID

A.4. Using more than two vCPUs with Windows client VMs

Window client operating systems support a maximum of two CPU sockets. When allocating vCPUs to virtual sockets within a guest VM, XenServer defaults to allocating one vCPU per socket; any more than two vCPUs allocated to the VM won't be recognized by the Windows client OS.

To fix this, set `platform:cores-per-socket` to the number of vCPUs allocated to the VM:

```
[root@xenserver ~]# xe vm-param-set uuid=<vm-uuid> platform:cores-per-socket=4
VCPUs-max=4 VCPUs-at-startup=4
```

A.5. Pinning VMs to a specific CPU socket and cores

Use `xe host-cpu-info` to determine the number of CPU sockets and logical CPU cores in the server platform. In this example the server implements 32 logical CPU cores across two sockets:

```
[root@xenserver ~]# xe host-cpu-info
cpu_count                 : 32
             socket_count: 2
```

```
                  vendor: GenuineIntel
                   speed: 2600.064
               modelname: Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz
                  family: 6
                   model: 45
                stepping: 7
                   flags: fpu de tsc msr pae mce cx8 apic sep mtrr mca cmov
pat clflush acpi mmx fxsr sse sse2 ss ht nx constant_tsc nonstop_tsc aperfmperf
pni pclmulqdq vmx est ssse3 sse4_1 sse4_2 x2apic popcnt aes hypervisor ida arat
tpr_shadow vnmi flexpriority ept vpid
                features: 17bee3ff-bfebfbff-00000001-2c100800
    features_after_reboot: 17bee3ff-bfebfbff-00000001-2c100800
        physical_features: 17bee3ff-bfebfbff-00000001-2c100800
                maskable: full
```

To pin a VM's vCPUs to a specific socket, set `VCPUs-params:mask`. This setting persists over VM reboots and shutdowns. In a dual socket platform with 32 total cores, cores 0-15 are on socket 0, and cores 16-31 are on socket 1. To restrict a VM to only run on socket 0:

```
[root@xenserver ~]# xe vm-param-set uuid=<vm-uuid> VCPUs-
params:mask=0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15
```

Similarly, to restrict a VM to only run on socket 1:

```
[root@xenserver ~]# xe vm-param-set uuid=<vm-uuid> VCPUs-
params:mask=16,17,18,19,20,21,22,23,24,25,26,27,28,29,30,31
```

To pin vCPUs to specific cores within a socket, specify the cores directly:

```
[root@xenserver ~]# xe vm-param-set uuid=<vm-uuid> VCPUs-
params:mask=16,17,18,19
```

Use `xl vcpu-list` to list the current assignment of vCPUs to physical CPUs:

```
[root@xenserver ~]# xl vcpu-list
Name                                ID  VCPU   CPU State   Time(s) CPU Affinity
Domain-0                             0     0   25   -b-    9188.4  any cpu
Domain-0                             0     1   19   r--    8908.4  any cpu
Domain-0                             0     2   30   -b-    6815.1  any cpu
Domain-0                             0     3   17   -b-    4881.4  any cpu
Domain-0                             0     4   22   -b-    4956.9  any cpu
Domain-0                             0     5   20   -b-    4319.2  any cpu
Domain-0                             0     6   28   -b-    5720.0  any cpu
Domain-0                             0     7   26   -b-    5736.0  any cpu
```

```
test-image-win7-32                           34     0     9   -b-       17.0  4-15
test-image-win7-32                           34     1     4   -b-       13.7  4-15
```

A.6. Changing dom0 vCPUs and pinning

The default number of vCPUs assigned to dom0 is 8. To change this number, modify the dom0_max_vcpus parameter in the Xen bootline. For example:

```
[root@xenserver ~]# /opt/xensource/libexec/xen-cmdline --set-xen
dom0_max_vcpus=4
```

Note: After applying this setting, you must reboot the system for it to take effect. Use `shutdown -r now` to reboot the server, or reboot it from XenCenter.

By default, dom0's vCPUs are unpinned, and able to run on any physical CPU in the system. To pin dom0 vCPUs to specific physical CPUs, use `xl vcpu-pin`. For example, to pin dom0's vCPU 0 to physical CPU 18, use:

```
[root@xenserver ~]# xl vcpu-pin Domain-0 0 18
```

CPU pinnings applied this way take effect immediately but do not persist over reboots. To make settings persistent, add `xl vcpu-pin` commands into `/etc/rc.local`, for example:

```
xl vcpu-pin 0 0 0-15
xl vcpu-pin 0 1 0-15
xl vcpu-pin 0 2 0-15
xl vcpu-pin 0 3 0-15
xl vcpu-pin 0 4 16-31
xl vcpu-pin 0 5 16-31
xl vcpu-pin 0 6 16-31
xl vcpu-pin 0 7 16-31
```

A.7. Determining GPU locality

As noted in section 3.6.1, current multi-socket, servers typically implement PCIe expansion slots local to each CPU socket and it is advantageous to pin VMs to the same socket that their associated physical GPU is connected to.

For current Intel platforms, CPU socket 0 typically has its PCIe root ports located on bus 0, so any GPU below a root port located on bus 0 is connected to socket 0. CPU socket 1 has its root ports on a higher bus number, typically bus 0x20 or bus 0x80 depending on the specific server platform.

Notice

HDMI

OpenCL

Trademarks

Copyright