NVIDIA GRID™ VGPU™

Steve Harpster SA NALA | October 2013

AGENDA

Overview

Prerequisites

XenServer Installation/Setup

Booting vGPU-enabled VMs

Basic Perf. Tuning

Support Resources

Q&A

Graphics Options in Virtualization

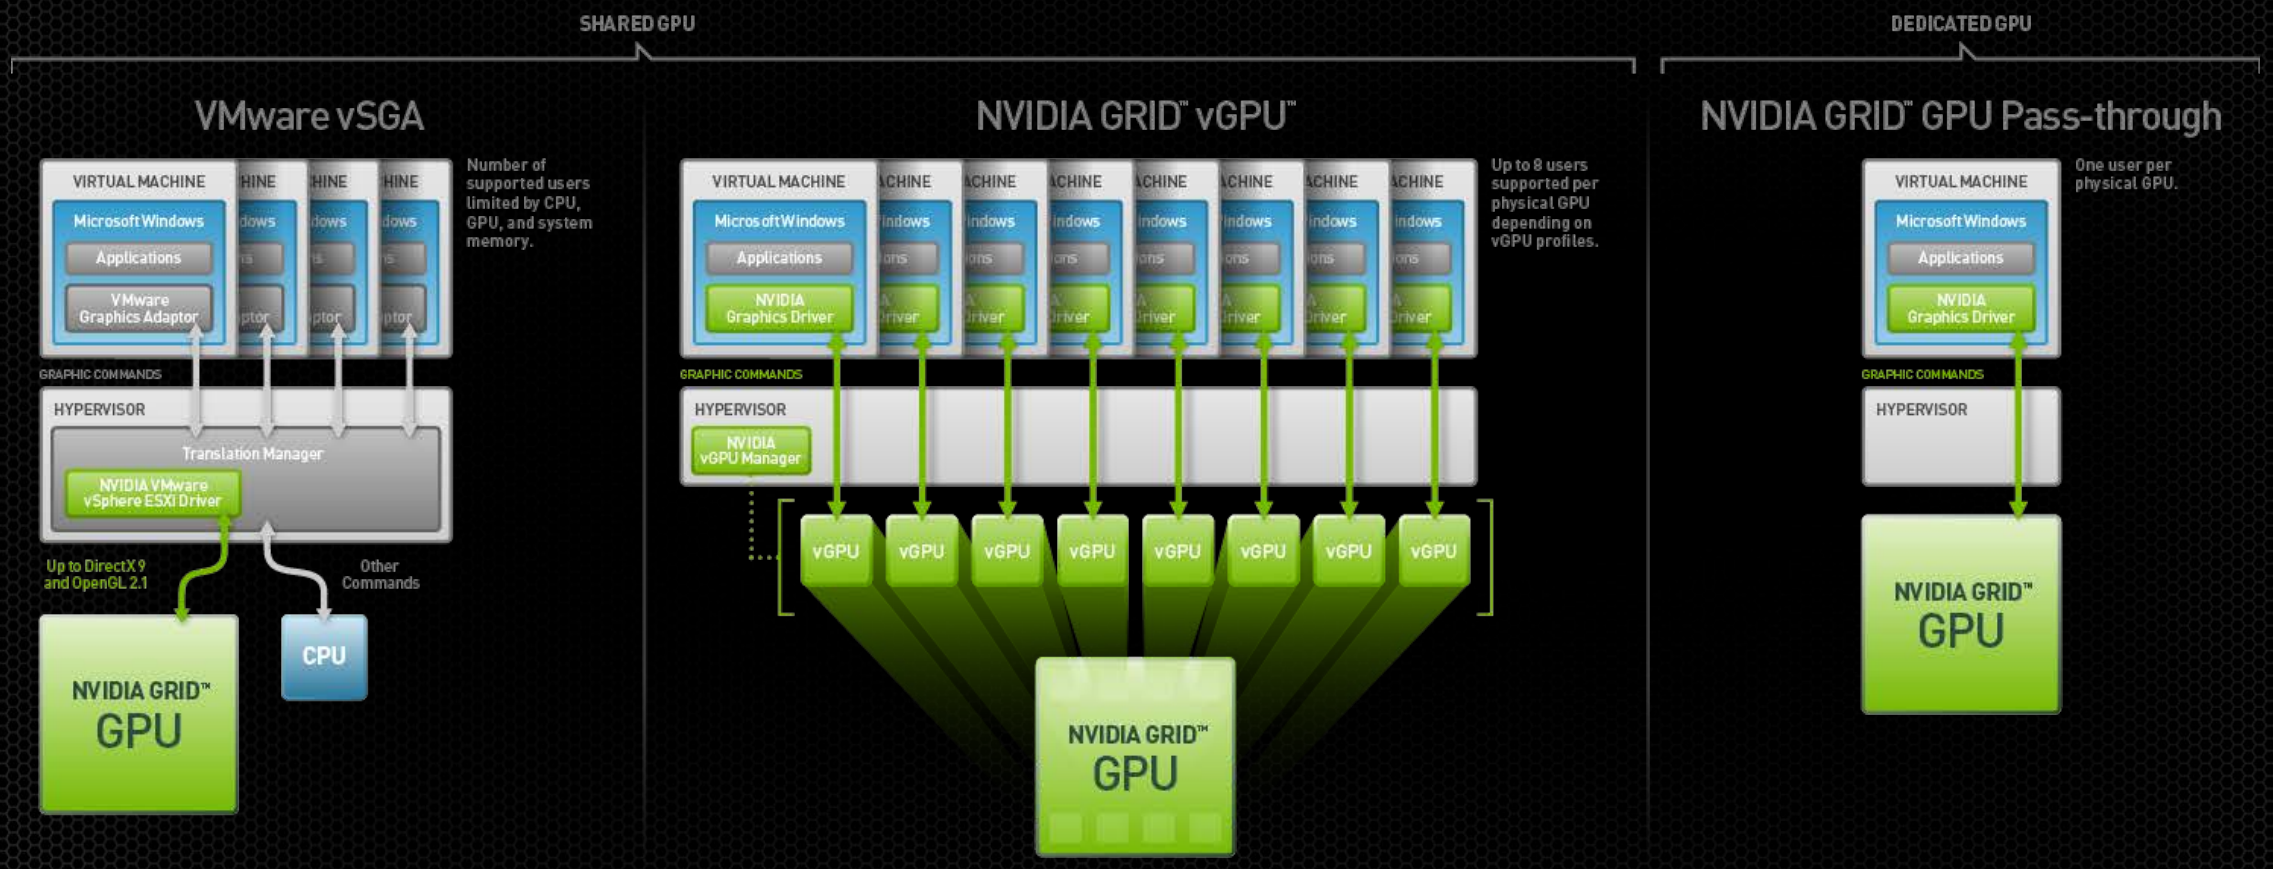

Overview of vGPU on XenServer

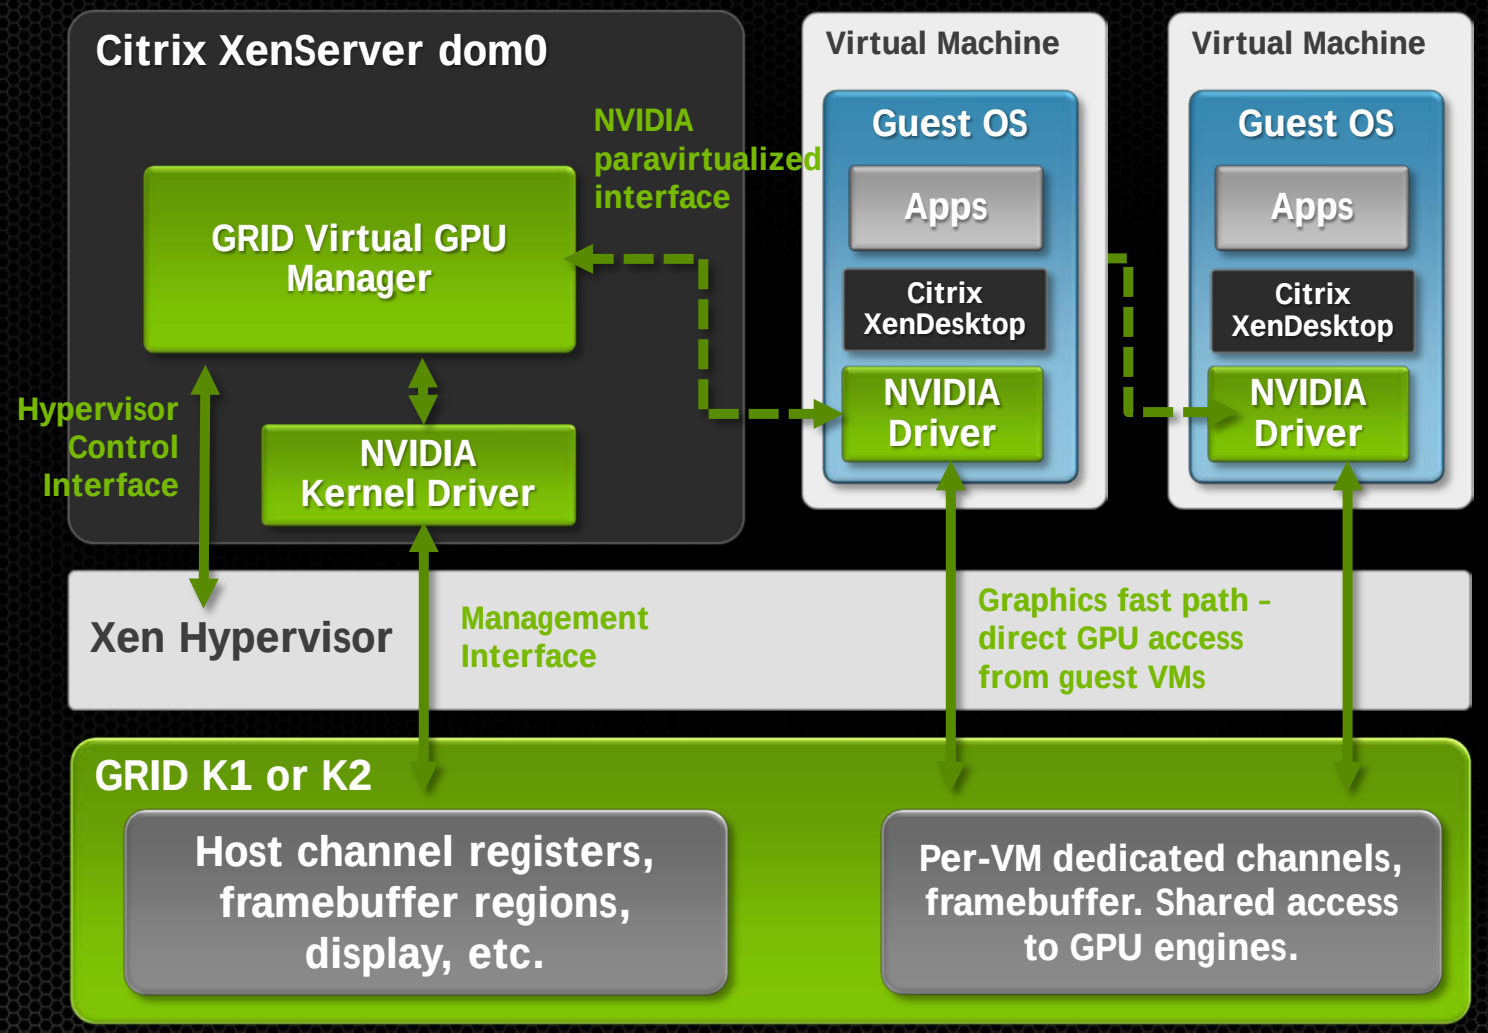

GRID vGPU enables multiple VMs to share a single physical GPU.

VMs run an NVIDIA driver stack and get direct access to the GPU.

- Supports same graphics APIs as our physical GPUs (DX9/10/11, OGL 4.x).

NVIDIA GRID Virtual GPU Manager for XenServer runs in dom0.

AGENDA

Overview

Prerequisites

XenServer Installation/Setup

Booting vGPU-enabled VMs

Basic Perf. Tuning

Support Resources

Q&A

Prerequisites

Server platform with GRID K1/K2 card(s)

NVIDIA GRID vGPU components

- Citrix XenServer 6.2.0 with vGPU support and XenDesktop 7.1 Tech Preview
- vGPU Manager
- vGPU Driver

Miscellaneous tools

NVIDIA GRID™ Enabled OEM Platforms

Cisco UCS C240 M3	Dell PowerEdge R720	Fujitsu Celsius C620 Fujitsu Celsius R930	HP ProLiant WS460c Gen8 HP ProLiant SL250s Gen8	IBM iDataPlex dx360 M4 IBM Flex System

For more information
on GRID enabled servers visit
www.nvidia.com/buygrid

DESIGNER

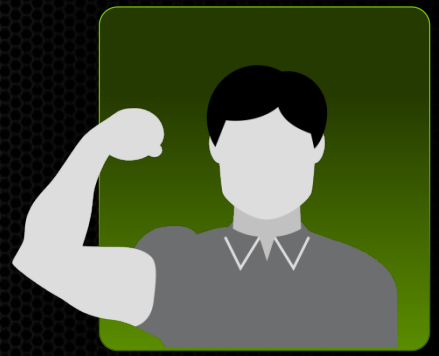

POWER USER

KNOWLEDGE WORKER

NVIDIA GRID K1

NVIDIA GRID K2

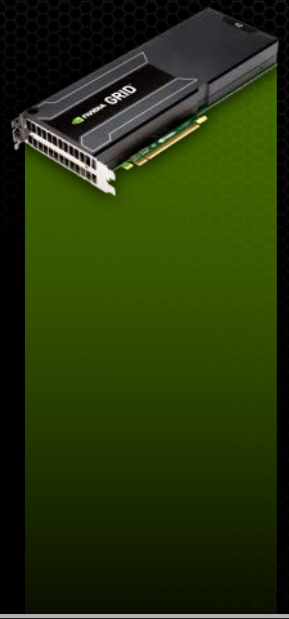

	NVIDIA GRID K1	NVIDIA GRID K2
GPU	4 Kepler GPUs	2 High End Kepler GPUs
CUDA Cores	768 (192/GPU)	3072 (1536/GPU)
Memory Size	16GB DDR3 (4GB/GPU)	8GB GDDR5 (4GB/GPU)
Max Power	130 W	225 W
Equivalent Quadro with Pass-through	Quadro K600 (entry)	Quadro K5000 (high end)

[1] Number of users depends on software solution, workload, and screen resolution

Where do I get the code?

GRID OVERVIEW | GRID BOARDS | DEDICATED GPU | SHARED GPU | PARTNER SOLUTIONS

SHARED VIRTUAL GPU (VGPU) TECHNOLOGY

NVIDIA GRID™ software brings the full benefit of NVIDIA hardware-accelerated graphics to virtualized solutions. This gives you exceptional performance in two primary uses:

- GPU pass-through that provides a virtual machine access to a dedicated GPU
- GRID Virtual GPU (vGPU) technology that provides the most advanced way to share a GPU

NVIDIA GRID VIRTUAL GPU (VGPU) TECHNOLOGY

NVIDIA GRID™ vGPU™ is the industry's most advanced technology for sharing true GPU hardware acceleration between multiple virtual desktops—without compromising the graphics experience. Application features and compatibility are exactly the same as they would be at the desk.

With vGPU technology, the graphics commands of each virtual machine are passed directly to the GPU, without translation by the hypervisor. This allows the GPU hardware to be time-sliced to deliver the ultimate in shared virtualized graphics performance.

Up to 8 users supported per physical GPU depending on vGPU profiles.

CONTACT US

- GRID Certified Servers
- Reseller questions
- General questions
- Sign up for GRID Newsletter

Install NVIDIA GRID vGPU

- Step 1: Configure XenServer + XenDesktop www.citrix.com/go/vgpu
- Step 2: Download vGPU Manager and Guide
- Step 3: Download the Product Series "NVIDIA GRID vGPU" Driver.

VIDEOS

- An overview of NVIDIA GRID™ vGPU™

XenServer and XenDesktop from Citrix

vGPU Manager from NVIDIA

GRID driver from NVIDIA

http://www.nvidia.com/object/virtual-gpus.html

XenServer and XenDesktop download

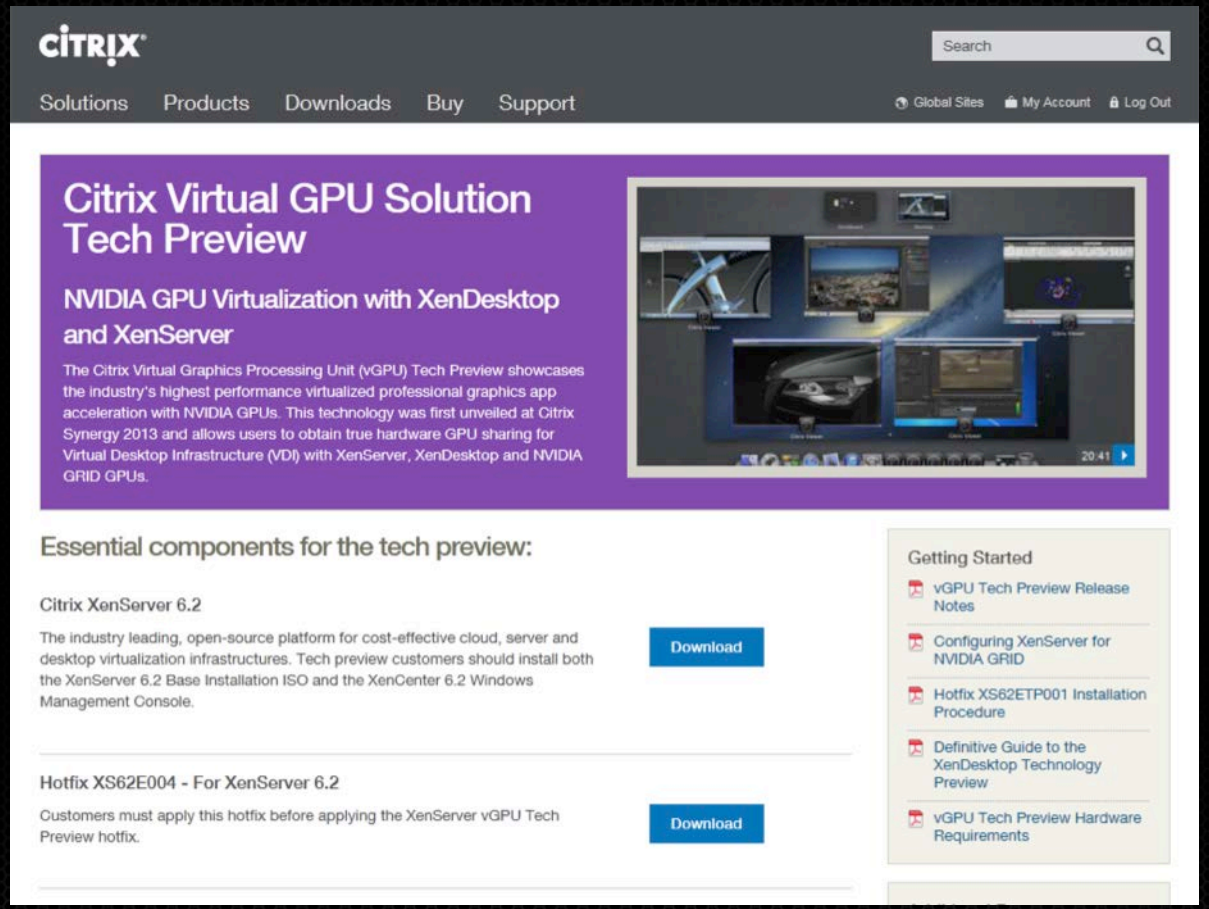

http://www.citrix.com/go/private/vgpu

vGPU Manager and vGPU Driver

GRID vGPU Manager software for Citrix XenServer

- Tech Preview version: 312.53
- NVIDIA-vgx-xenserver-6.2-312.53.i386.rpm
- http://www.nvidia.com/content/cloud-computing/downloads/NVIDIA-GRID-vGPU-Manager-and-Installation-Guide.zip

NVIDIA Windows 7 32/64-bit drivers for guest VMs

- Tech Preview version: 312.56
- http://www.nvidia.com/object/grid-vgpu-win7-64bit-312.56-whql-driver.html

Available on NVIDIA.com

- http://www.nvidia.com/vgpuinstall

Miscellaneous Tools

User Guide and Releases Notes

- Normally packaged with vGPU Manager software
- http://www.nvidia.com/content/cloud-computing/downloads/NVIDIA-GRID-vGPU-Manager-and-Installation-Guide.zip

SSH Client

- Required for login to XenServer's dom0 shell
- Most Linux include ssh
- Putty is a good one for Windows
- http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html

AGENDA

Overview

Prerequisites

XenServer Installation/Setup

Booting vGPU-enabled VMs

Basic Perf. Tuning

Support Resources

Q&A

Installation/Setup

Server Setup

Follow server recommendations for placement of GPUs

- Best to balance GPUs across CPU sockets where possible
- Do NOT add the GPU in a pure PASS THRU mode in XenCenter. They must be un-passed for vGPU to be configured properly

Minimum recommended RAM is 64 GB

- This allows 15 VMs at 4GB RAM / VM, with 4G left over for XenServer
- Less RAM == less VMs: virtualization doesn't offer any silver bullets here

Minimum recommended drive space is 250 GB

- Allows for XenServer install plus up to 32 VMs running identical workloads

SSD highly recommended

- Significantly improves VM boot times, cloning operations, etc

Remote management ("IPMI") highly recommended

- Servers are noisy, best left in a machine room or closet
- XenCenter and remote power/reset/console access means you never need touch the server

XenServer Installation

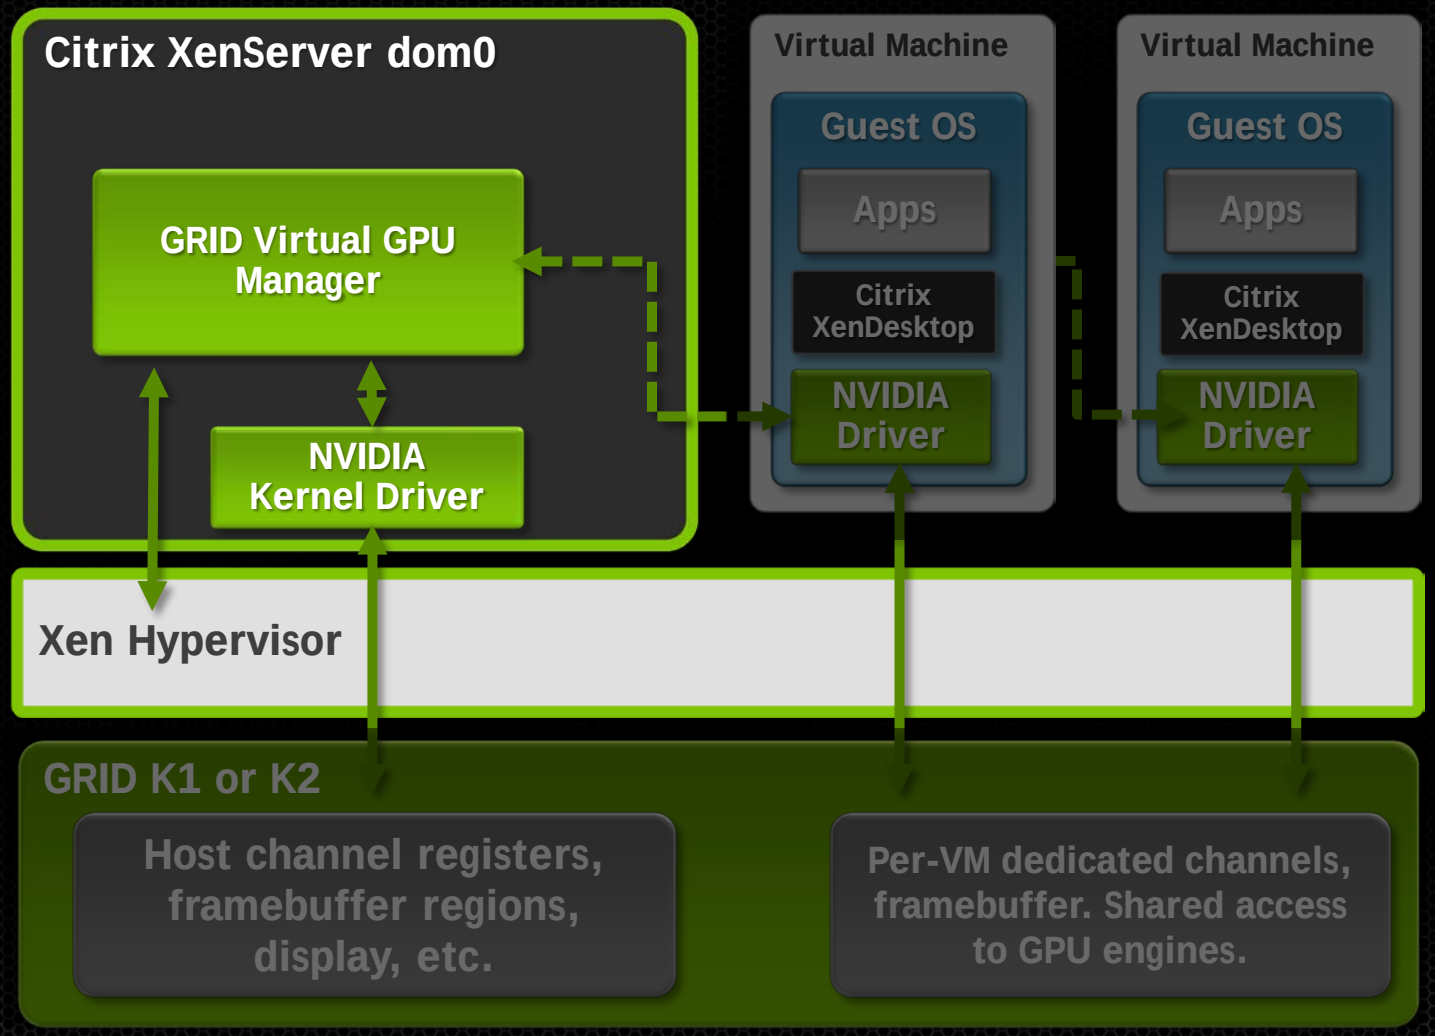

XenServer Installation

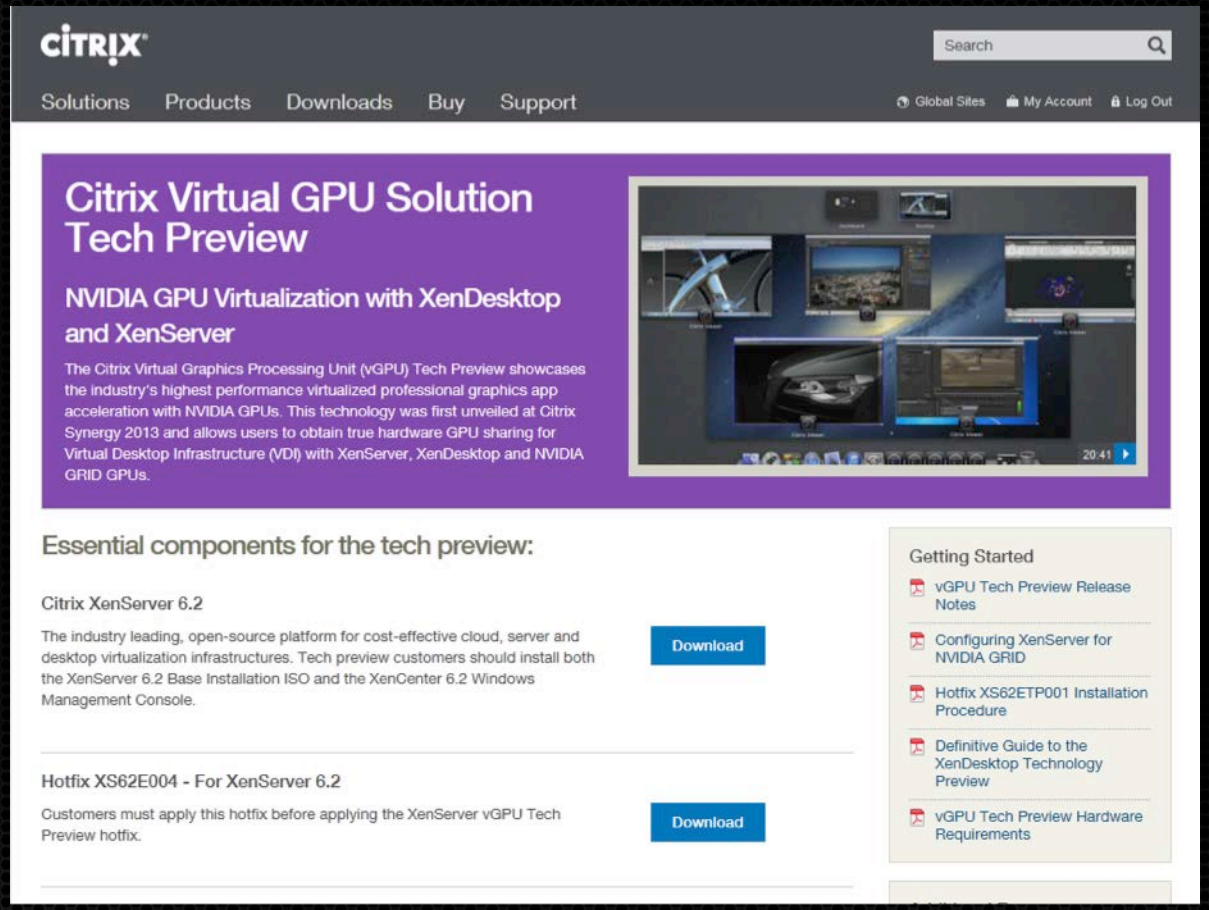

http://www.citrix.com/go/private/vgpu

After you have the SW images

Boot your server from CD or USB, follow the setup instructions

- No supplemental install packs needed
- Don't enable Thin Provisioning for XenDesktop
- Configure a root / management password
- Set up IP address for the server:
  - Static IP address highly recommended
  - With DHCP, you may need ongoing access to server's VGA console to check its IP address
- Installation should complete in a few minutes

XenServer Installation

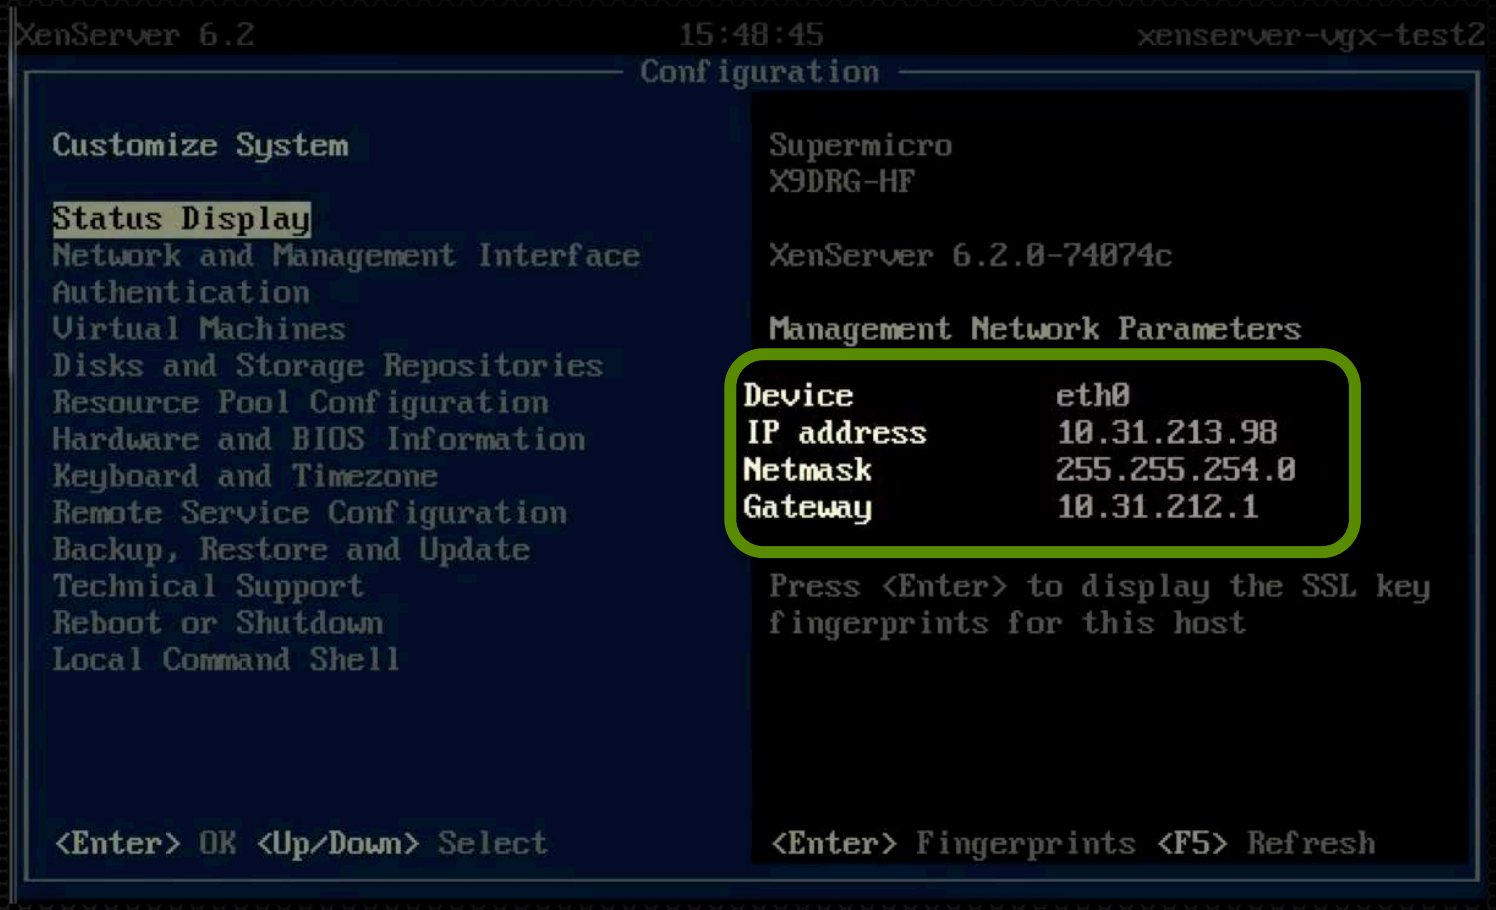

After reboot, server should boot to the XenServer local console

Note the server's IP address, used to connect XenCenter management tool, and to the dom0 shell

XenCenter Installation

XenCenter is a Windows-based remote management tool for XenServer

Install on a PC, using XenCenter.msi from the XenServer bundle

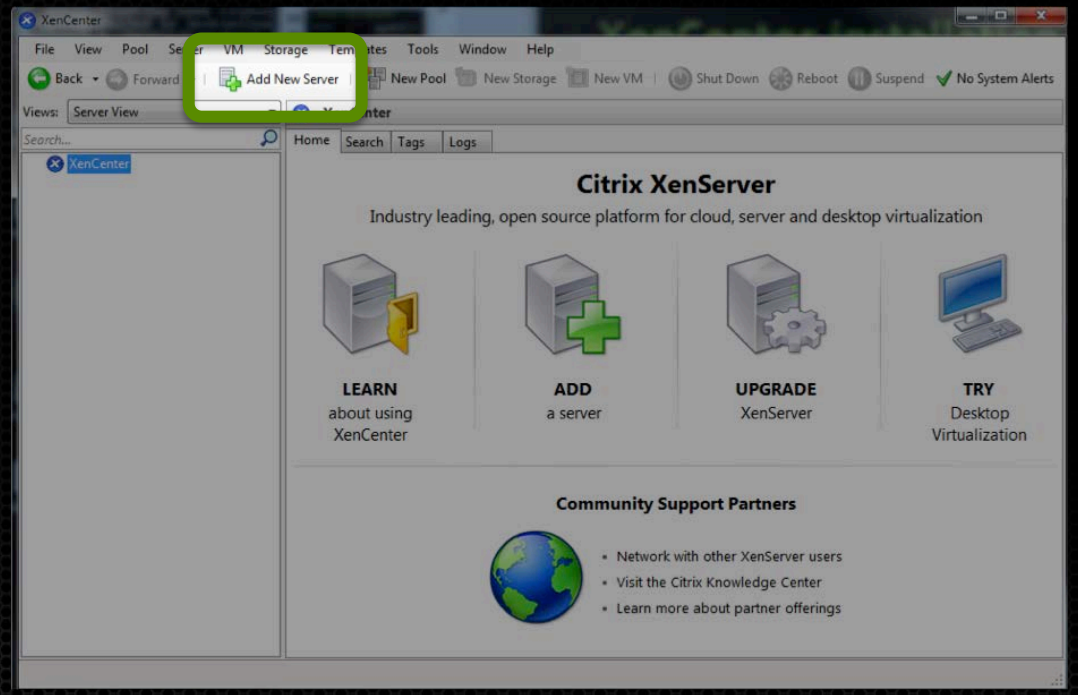

Click "Add New Server"

Enter IP address and root password of your XenServer

Tool should connect and show storage resources on the XenServer

XenServer 6.2 has all features available in the unlicensed version.

Connecting to Xen dom0

Connect to Xen dom0 using putty or some other ssh client

- You can also use the "Console' tab in XenCenter but it's clunky compared to an ssh client

Use the IP address and root password you previously configured for the server

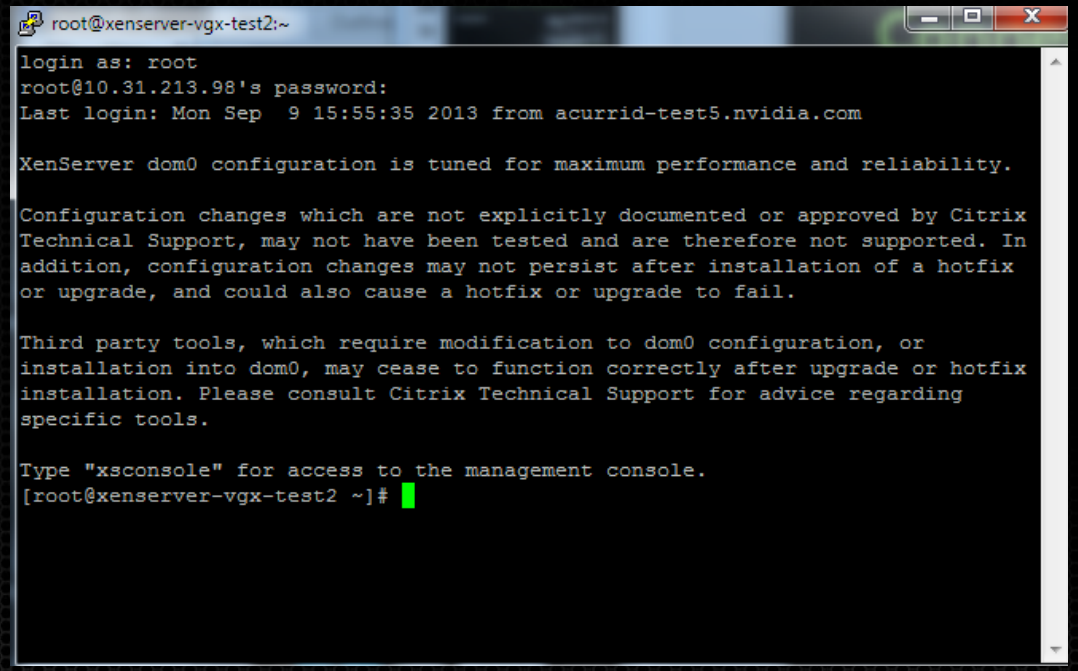

Dom0 runs a paravirtualized Linux OS

- Most all regular Linux commands are present and functional

'xe' and 'xl' commands are primarily used for management of virtual machines

- They support context-sensitive command line completion, using the TAB key

Installing GRID vGPU Manager in dom0

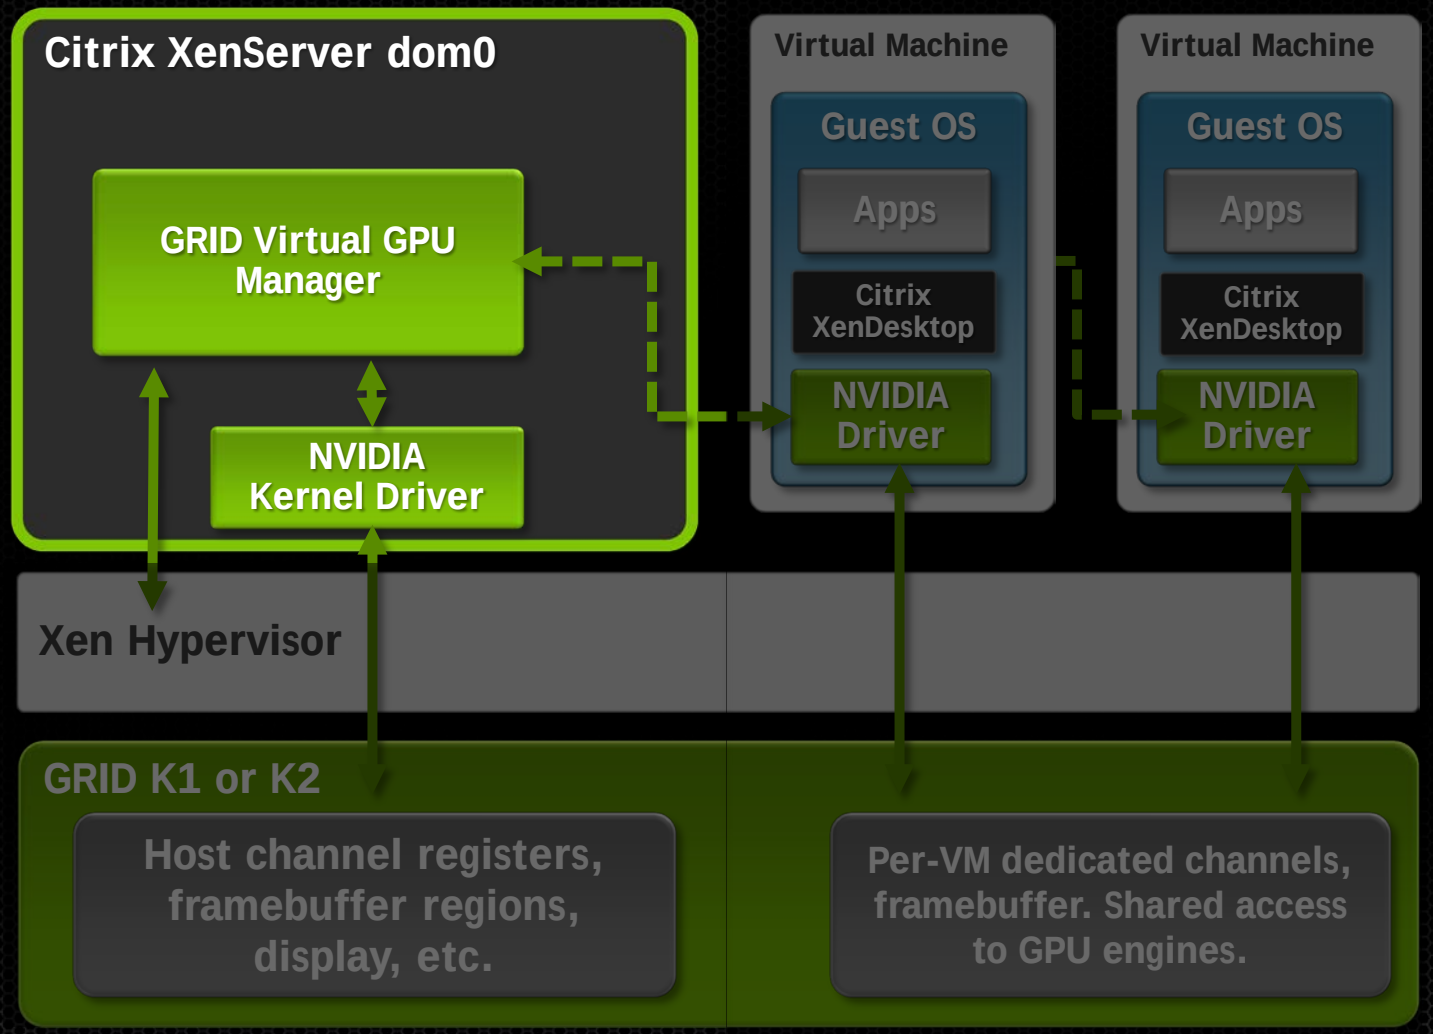

Copy GRID vGPU Manager

Package to dom0

Option 1: Use scp (secure copy – part of the ssh suite)

- On your PC, using the pscp utility that comes with putty
- WinSCP also another good utility with a GUI interface

Option 2: From XenServer dom0 shell, reading the driver package from another system by mounting a share in Linux

Installing GRID vGPU Manager in dom0

Use the rpm command to install the package:

```
[root@xenserver ~]# rpm -iv  NVIDIA-vgx-xenserver-6.2-312.53.i386.rpm
Preparing packages for installation
NVIDIA-vgx-xenserver-6.2-312.53
[root@xenserver ~]#
```

Reboot the server:

```
[root@xenserver ~]# shutdown –r now
Broadcast message from root (pts/1) (Fri Sep 30 14:24:11 2013):
The system is going down for reboot NOW!
```

Reconnect ssh to dom0 after the server has rebooted

- XenCenter should reconnect automatically

Verifying GRID vGPU Manager Installation

```
[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 ~]#
```

Check kernel driver is loaded using lsmod

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

Run nvidia-smi, check it detects GPUs

Creating a VM

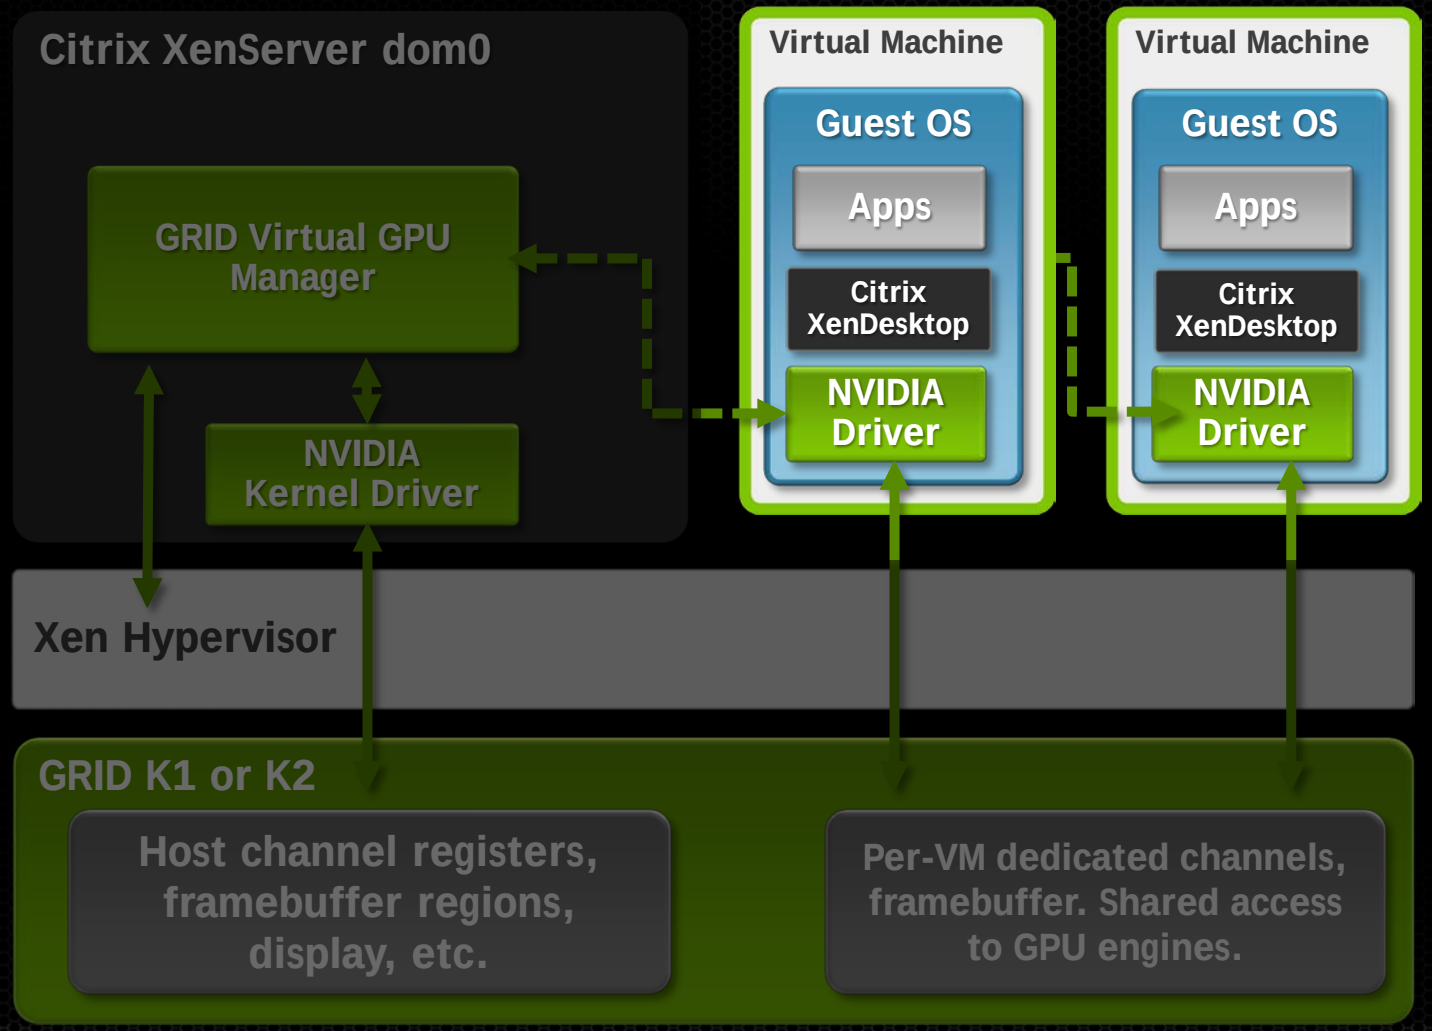

Creating a Windows VM

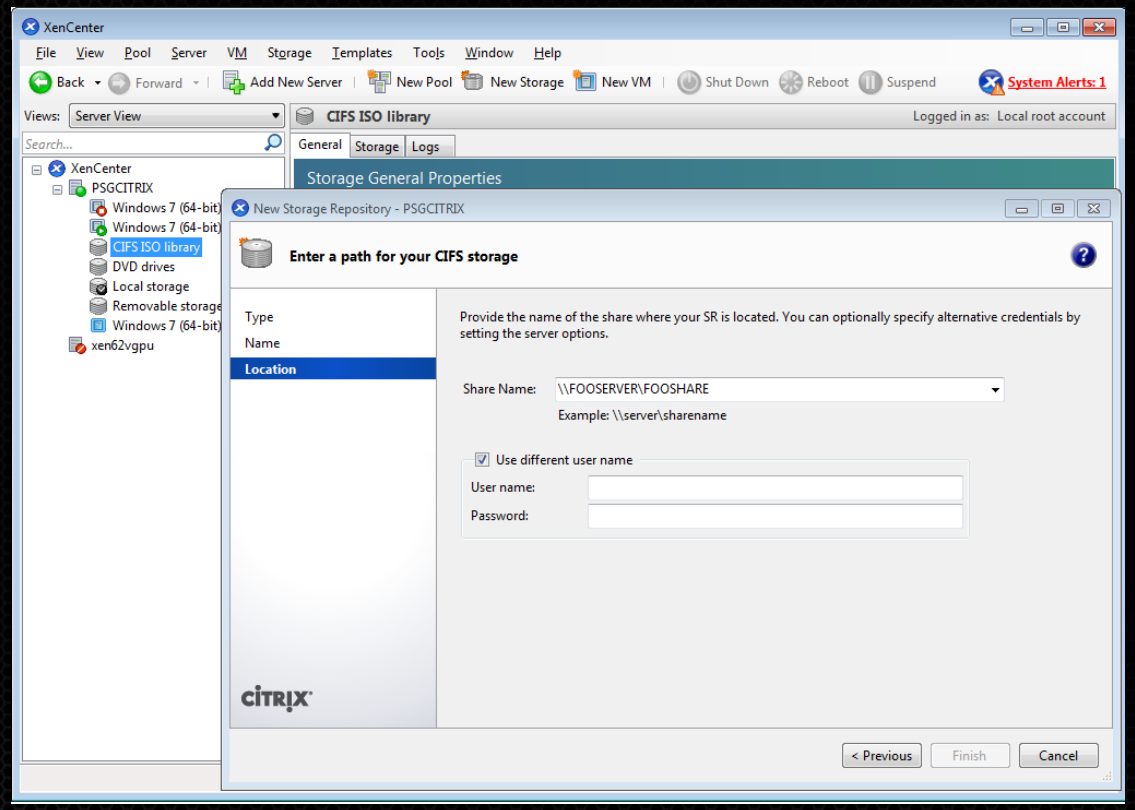

This requires a Windows installation ISO

Can use the Storage tab in XenCenter to mount an ISO share.

Creating a Windows VM

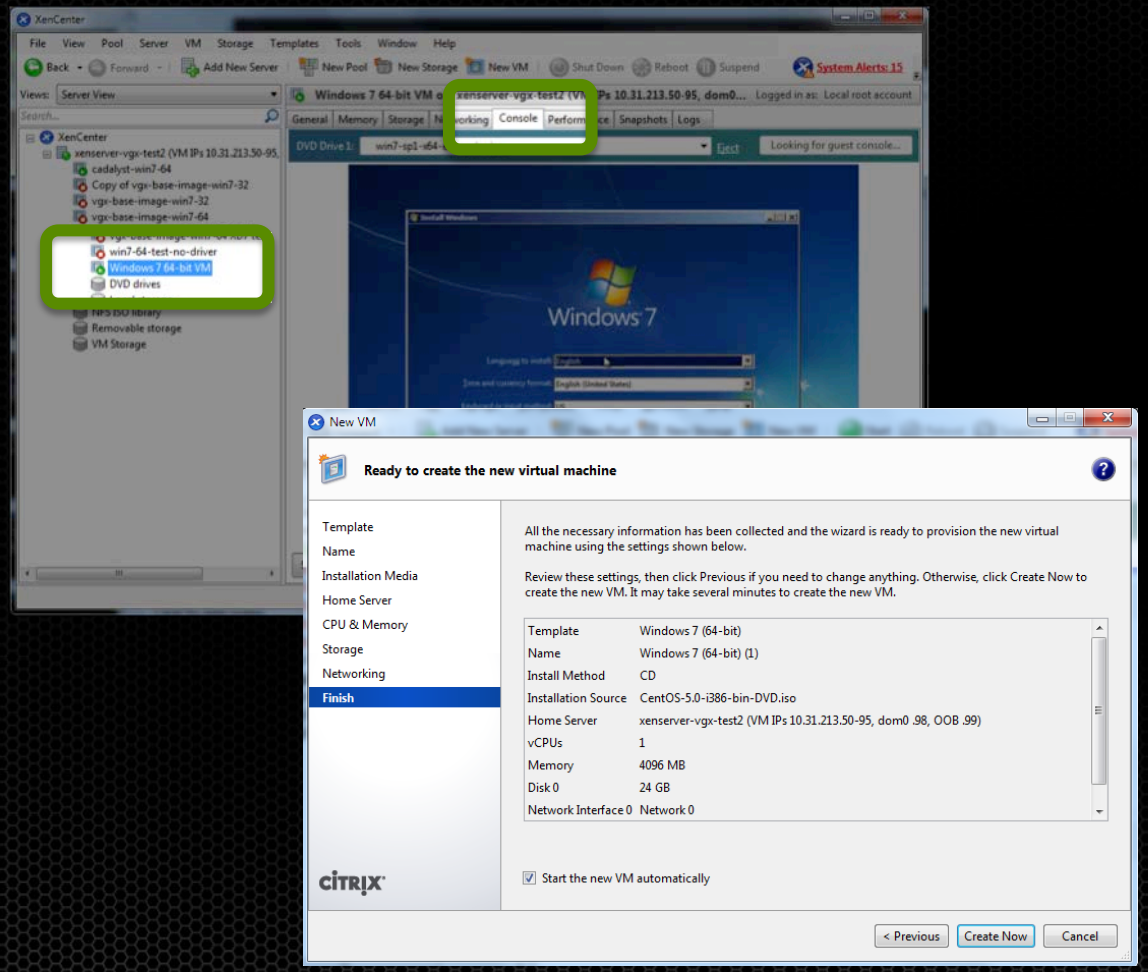

Review the configuration, click “Create Now” to create and start the VM

Once created, VM shows up in the VM list in the left hand pane

Select the VM

Click on console tab to view the VM’s console output

Proceed with Windows installation via the console

Set up networking in the VM – static IP or DHCP

Install XenServer tools

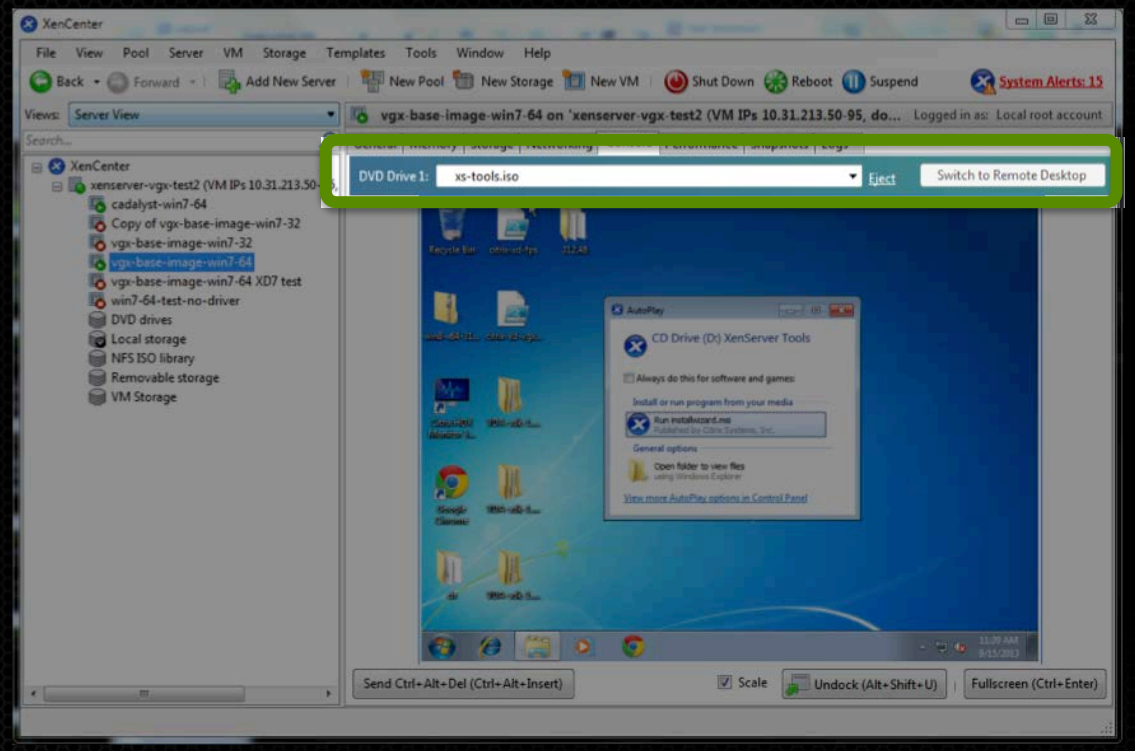

Once Windows installation is complete and VM has rebooted to Windows desktop, install the XenServer tools

Paravirtualized storage and network drivers

Essential for good VM performance

Install by inserting xs-tools.iso into the VM's virtual DVD drive:

- Or right-click on VM, select "Install XenServer tools"
- Or use Tools Menu

After tools installation completes (requires a reboot), shut the VM down

AGENDA

Overview

Prerequisites

XenServer Installation/Setup

Booting vGPU-enabled VMs

Basic Perf. Tuning

Support Resources

Q&A

Creating a vGPU for a Windows VM

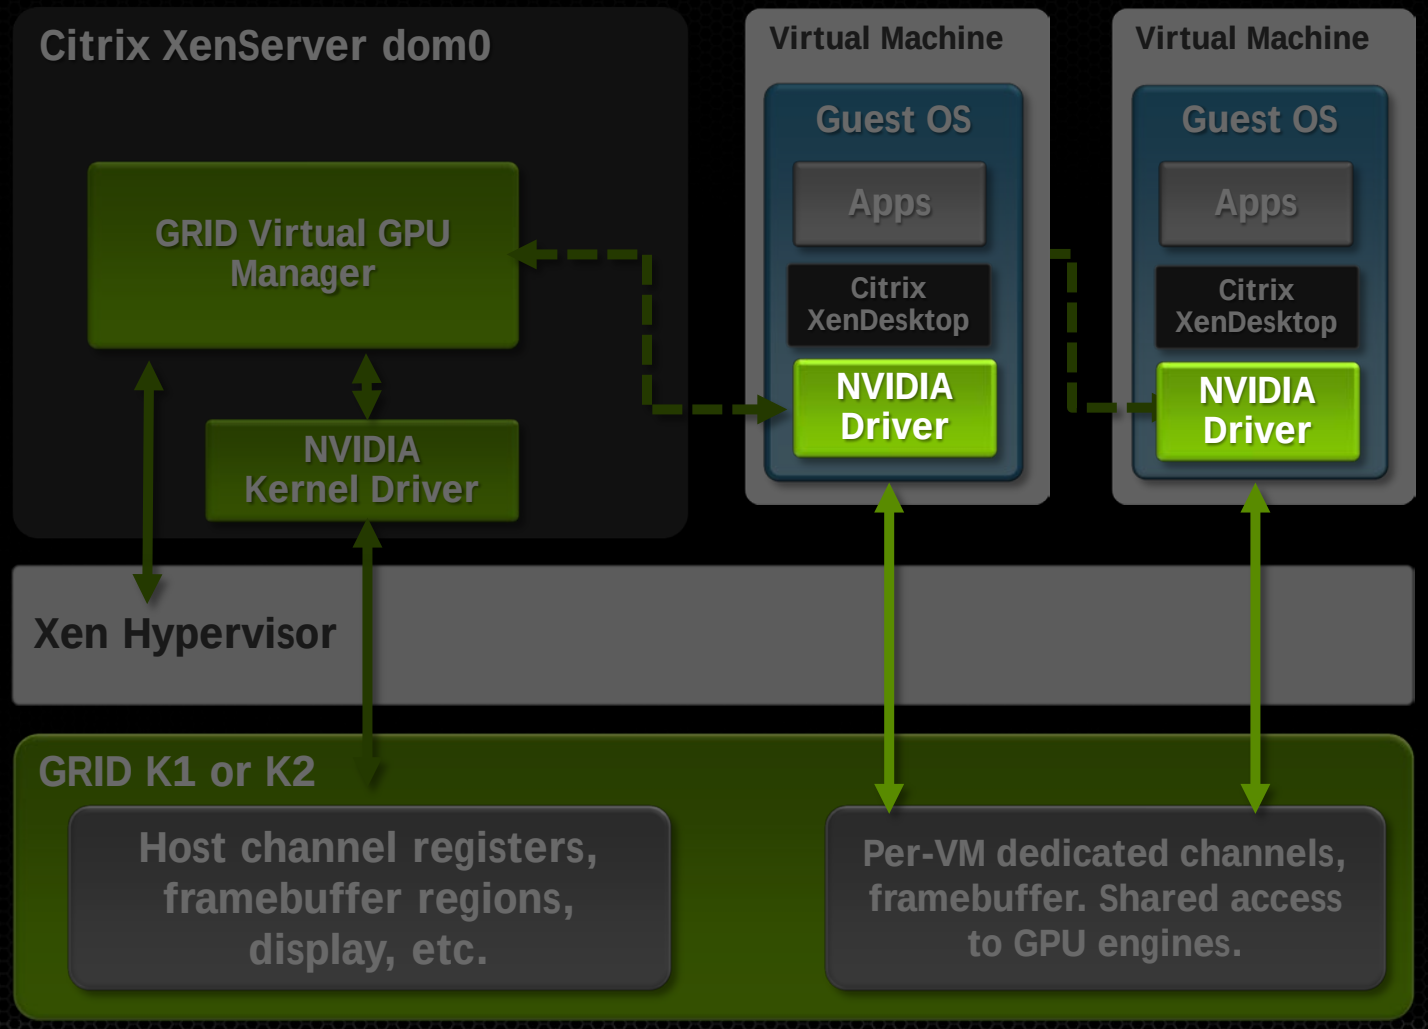

Creating a vGPU – vGPU Types

GRID boards support several different vGPU Types

- e.g. GRID K2 supports K200, K240Q, K260Q vGPUs
- vGPU types differ in amount of framebuffer, virtual display heads, max resolution, and number of simultaneous instances

Each physical GPU on a GRID board can host at most one vGPU type at a time

- e.g. you can't mix a K260Q with K200 vGPUs on the same physical GPU

All Quadro (Q) profiles are being ISV Certified for targeted professional applications.

Card	Physical GPUs	Virtual GPU	Intended Use Case	Framebuffer (Megabytes)	Virtual Display Heads	Max Resolution per Monitor	Maximum vGPUs	
							Per GPU	Per Board
GRID K1	4	GRID K140Q	Power User / Designer	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 / Power User	256	2	1920x1200	8	16

Creating a vGPU – Management Objects

Physical GPU

XenServer automatically creates gpu-group, pgpu, vgpu-type objects for the physical GPUs it discovers on startup

Physical GPUs

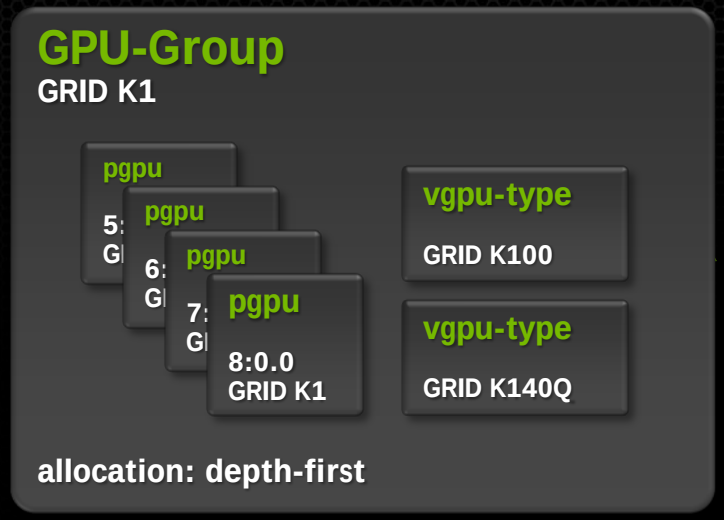

User creates vgpu objects:
- owned by a specific VM
- associated with a gpu-group
- with a specific vgpu-type

At VM boot, XenServer picks an available pgpu in the group to host the vgpu

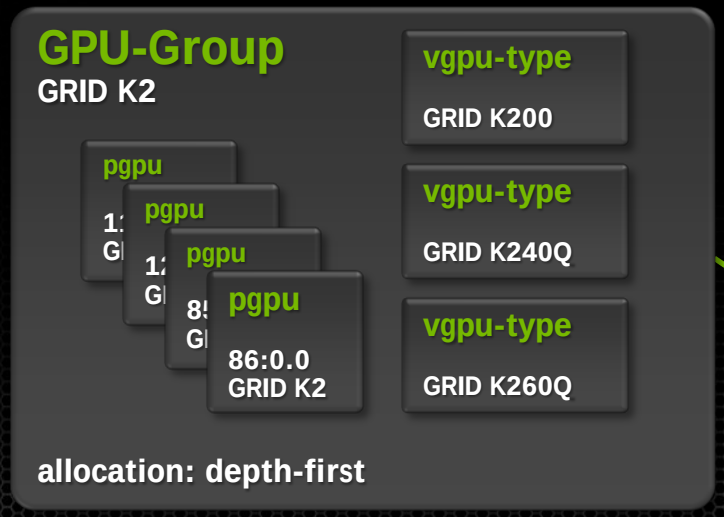

Creating a vGPU – Management Objects

Use ‘xe <obj>-list’ to list management objects

xe gpu-group-list

```
[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):
```

xe pgpu-list

```
[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
```

xe vgpu-type-list

```
[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
```

Creating a vGPU

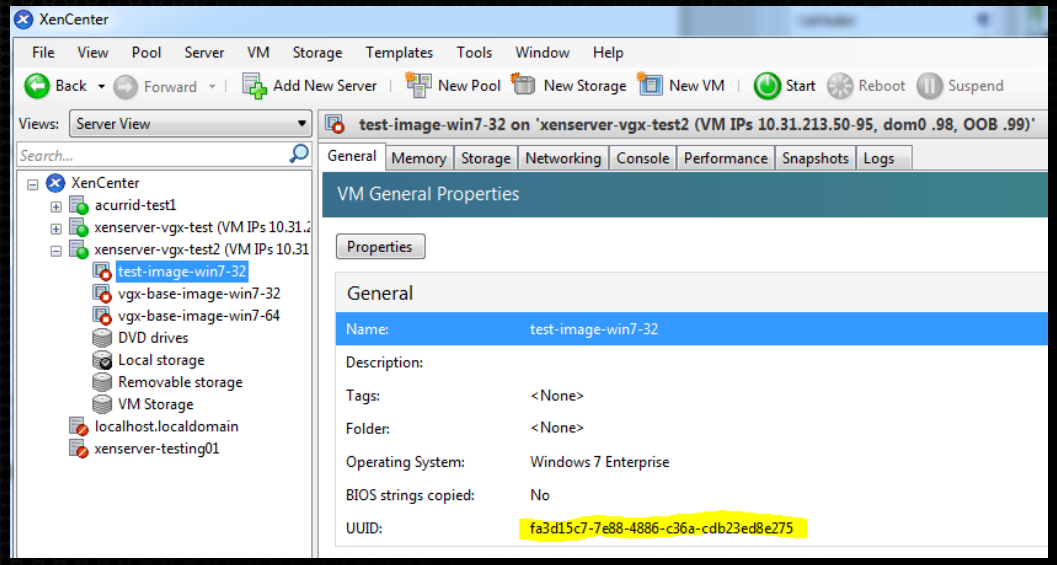

Find VM's UUID using 'xe vm-list'

- Or use XenCenter

Create vGPU using 'xe vgpu-create', specifying:

- UUID of the owning VM
- UUID of the vGPU type (use 'xe vgpu-type-list')
- UUID of the GPU group to allocate the vGPU from (use 'xe gpu-group-list')

```
[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
```

Boot the VM, Install NVIDIA Drivers

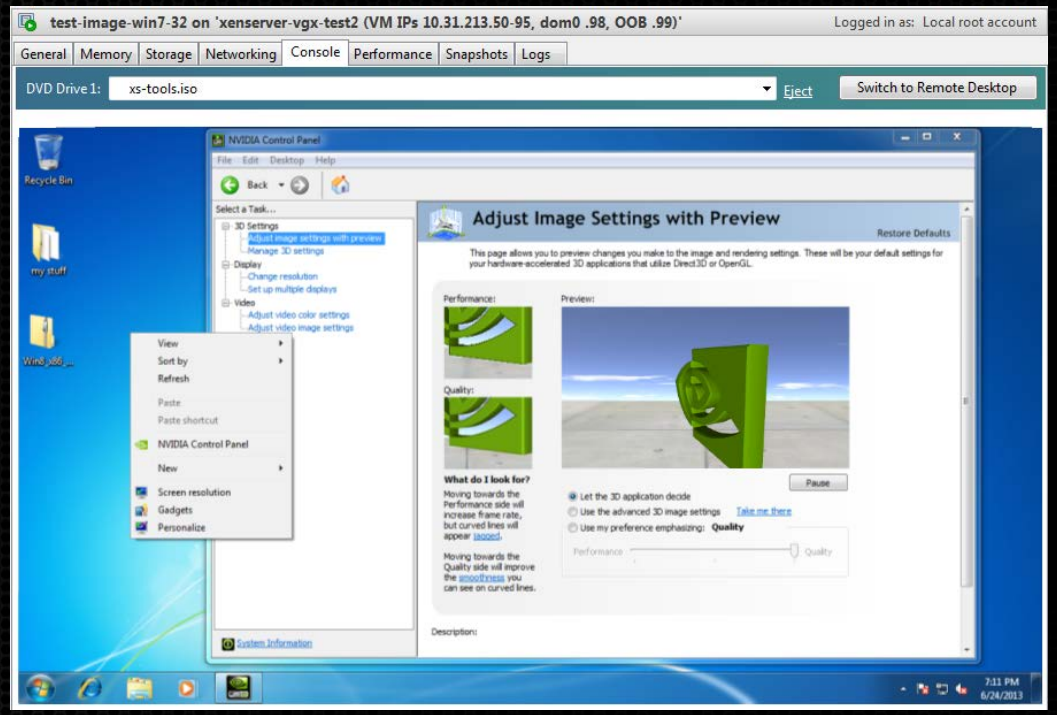

Start the VM via XenCenter or from ssh using
'xe vm-start <vm-uuid>'

VM is now using vGPU as its primary VGA

- Should boot into 800x600 VGA mode

Install NVIDIA Windows 7 32/64-bit drivers

- Tech Preview version: 312.56
  http://www.nvidia.com/object/grid-vgpu-win7-64bit-312.56-whql-driver.html
- Installation requires a reboot

VM's console should boot at max resolution of the vGPU type

- Open NV control panel to verify vGPU is working OK

Cloning VMs

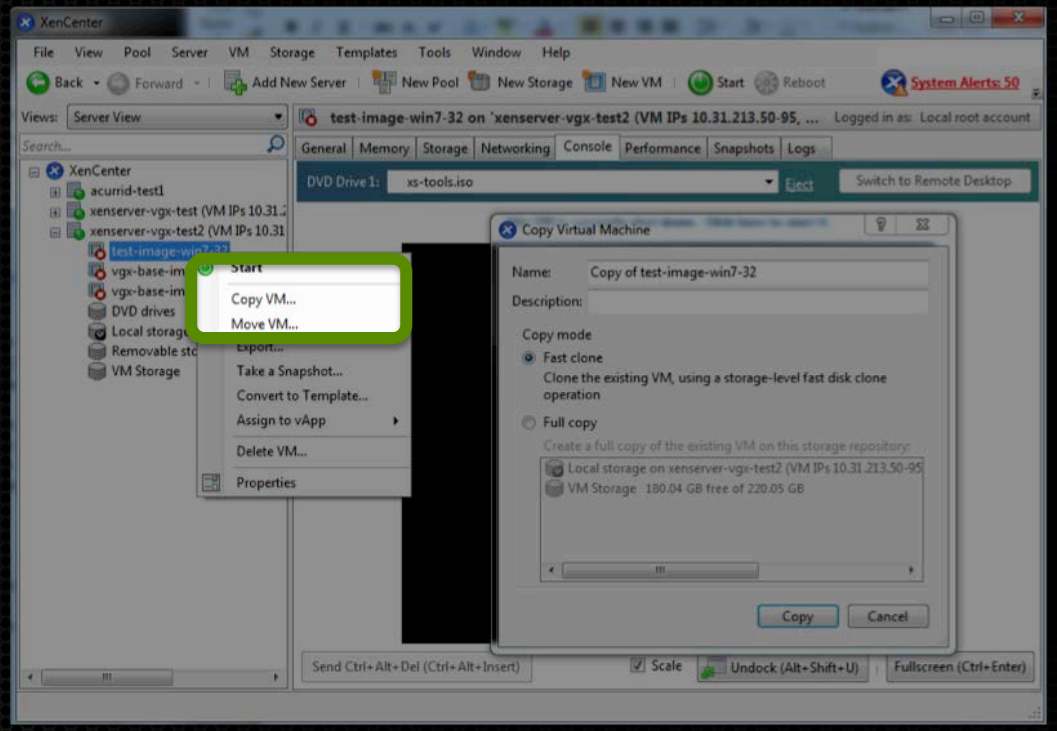

Once a VM is set up and configured the way you want, clone it to create new VMs

"Fast clone" VM uses copy-on-write disk, based off of the original VM's disk

- Host many VMs with little disk space impact

Clones inherit the base VM's vGPU config

- Same vGPU type and GPU group

Cloned VMs are automatically assigned new MAC addresses

- Works well with DHCP

Clone using 'xe vm-clone', or XenCenter GUI "Copy VM…" option

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

AGENDA

Overview

Prerequisites

XenServer Installation/Setup

Booting vGPU-enabled VMs

Basic Perf. Tuning

Support Resources

Q&A

Performance Tuning - Disabling Console VGA

Console VGA is the output path from vGPU that drives the XenCenter console tab

- Requires Reboot of VM
- Not intended for high performance, solely for setup and management
- Imposes a performance penalty on the VM

Once a VM image is set up (including networking)

- Install high-performance remote graphics in the VM, such as XenDesktop
- Or install a VNC server within the VM
- Disable the console VGA output by setting platform:vgpu_vnc_enabled=false

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

Performance tuning – using more than 2 vCPUs

Window client OS (e.g. Win7, Win8) support a maximum of two CPU sockets

By default, XenServer allocates one vCPU per virtual socket

When creating a VM with >2 vCPUs, Windows client will only use two of them

- XenCenter will report all the vCPUs allocated to a VM
- Windows task manager within the VM will only show 2 CPUs

Solution: have XenServer allocate more vCPUs per virtual socket

- Example: configure VM to have 4 vCPUs per socket

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

Performance Tuning - Pinning vCPUs

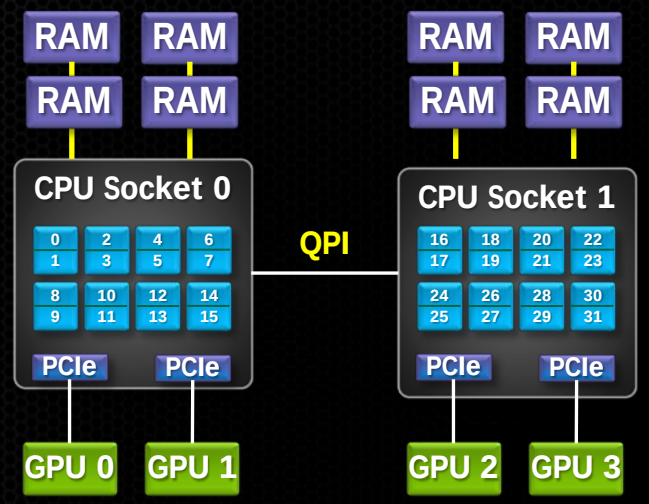

Intel two socket servers have memory and PCIe connected to each socket, with QPI connecting the sockets

QPI crossings cost performance

vGPU performs best when VM's vCPUs are on the same socket as the pGPU hosting the vGPU

To pin a VM's vCPUs to a specific socket, set VCPUs-params:mask

- Setting persists over VM reboots and shutdowns

E.g. 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
```

Performance Tuning
Controlling vGPU Allocation

XenServer creates a VM's vGPU at the time the VM is booted

- Selects an available pGPU in the vGPU's gpu-group based on gpu-group's allocation-algorithm
- Use 'xe gpu-group-param-get <gpu-group-uuid> param-name=allocation-algorithm'
- Use 'xe gpu-group-param-set...' to change it

allocation-algorithm=depth-first

- Attempts to **maximize** number of vGPUs on each pGPU within the gpu-group
- Places new vGPUs on pGPU that can support the new vGPU's vgpu-type, has space remaining, and has **most** number of vGPUs already resident.
- Leads to **higher density** of vGPUs, particularly when different vgpu-types are being used
- May result in **lower performance** because it attempts to maximize sharing of physical GPUs

allocation-algorithm=breadth-first

- Attempts to **minimize** number of vGPUs on each pGPU within the gpu-group
- Places new vGPUs on pGPU that can support the new vGPU's vgpu-type, has space remaining, and has **least** number of vGPUs already resident.
- May lead to **lower density** of vGPUs and may artificially limit the total number of vGPUs that can run, when different vgpu-types are being used
- Generally results in **higher performance** because it attempts to minimize sharing of physical GPUs

Performance Tuning

Allocating vGPU on a Specific pGPU

To precisely control vGPU allocation on pGPUs, create custom gpu-groups, e.g.

- One GPU group for all the identical GPUs connected to the same CPU socket, or
- One GPU group for each individual GPU

Example: creating a GPU group for the individual GPU at bus ID 5:0.0 –

- Find the UUID of the pgpu object for the GPU at 5:0.0 –

```
[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
```

- Ensure no vGPUs are currently executing on the physical GPU
  (This command returns nothing if no vGPUs are resident on the pGPU)

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

Performance Tuning

Allocating vGPU on a Specific pGPU

Create a new GPU group using 'xe gpu-group-create'

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

- Now change the pGPU's gpu-group-uuid to the UUID of the new group we just created

```
[root@xenserver ~]# xe pgpu-param-set uuid=7c1e3cff-1429-0544-df3d-
bf8a086fb70a gpu-group-uuid=585877ef-5a6c-66af-fc56-7bd525bdc2f6
```

Any vGPU created in this new GPU group will always be allocated on GPU 5:0.0

Installing Citrix XenDesktop

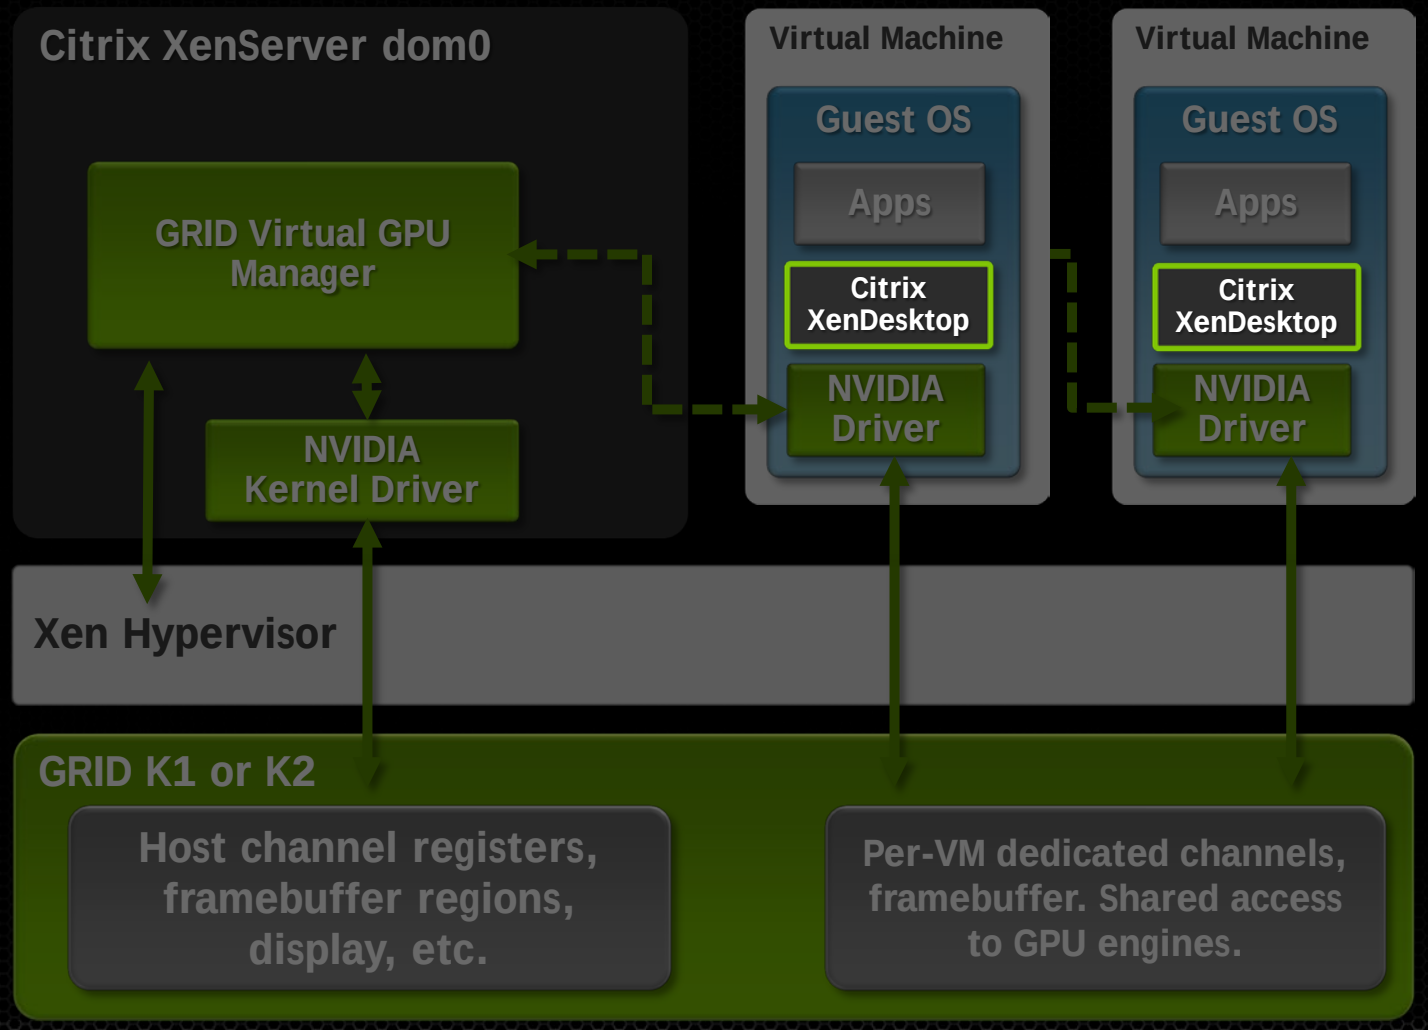

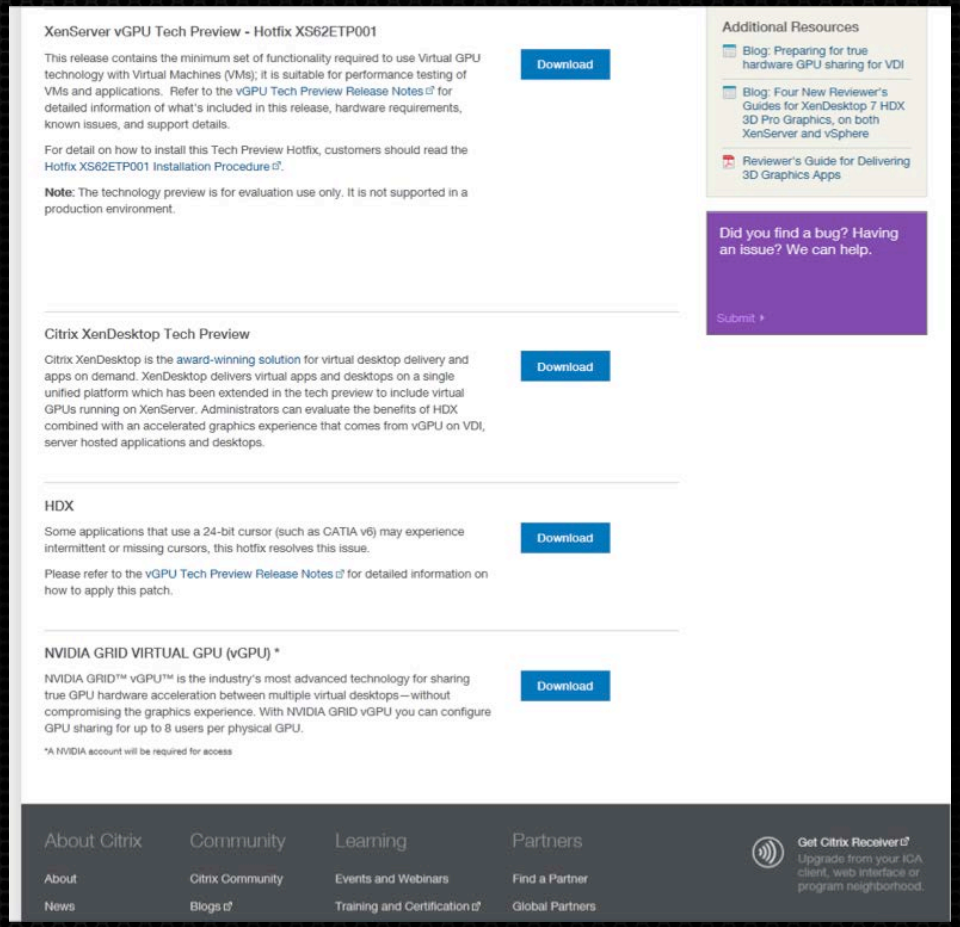

Installing Citrix XenDesktop

Insufficient time/space to cover this in detail.

vGPU requires Citrix XenDesktop Tech Preview.

- https://www.citrix.com/go/private/vgpu

Install the HDX 3D Pro Virtual Desktop Agent in the VM.

- https://www.citrix.com/go/private/vgpu

Install Citrix Receiver from the XenDesktop installation ISO on your client PC.

```
[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  |
+-----------------------------------------------------------------------------+
```

Other Tips
NVIDIA-smi

nvidia-smi run in dom0 shows vGPUs as compute processes

Process memory usage = vGPU FB size

AGENDA

Overview

Prerequisites

XenServer Installation/Setup

Booting vGPU-enabled VMs

Basic Perf. Tuning

Support Resources

Q&A

Citrix Tips

Citrix Optimizations

HDX 3DPro

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

Increasing 2D and 3D Application Scalability and PerformanceConsult the vGPU User Guide and Release Notes

- http://support.citrix.com/proddocs/topic/hd-hdx-parra/hd-hardware-acceleration.html

Improve Graphics Performance for XenApp, XenDesktop, Receiver Technologies

- http://support.citrix.com/article/CTX136671

GRID Remote Workstation Certifications

GRID K2 on Citrix XenDesktop with XenServer

ISV	Application	Certified Driver
Adobe	Premiere Pro	311.35
Auto Desk	AutoCAD 2013	310.90
	Inventor 2013	310.90
	Revit 2013	310.90
	AutoCAD 2014	311.35
	Inventor 2014	311.35
Bunkspeed	Bunkspeed SHOT 2013	311.35
	Bunkspeed PRO 2013	311.35
	Bunkspeed DRIVE 2013	311.35
Chaos	V-Ray	311.35
PTC	Creo (on IBM iDataPlex dx360)	320.00
Synerscope	Synerscope	320.00
TerraSpark	Insight Earth	311.35

http://www.nvidia.com/content/cloud-computing/pdf/GRID_Certifications.pdf

NVIDIA GRID Resources

GRID Website
www.nvidia.com/vdi

Sign up for the monthly GRID VDI Newsletter
http://tinyurl.com/gridinfo

GRID YouTube Channel
http://tinyurl.com/gridvideos

Questions? Email us
GRIDteam@nvidia.com

Follow us on Twitter
@NVIDIAGRID

AGENDA

Overview

Prerequisites

XenServer Installation/Setup

Booting vGPU-enabled VMs

Basic Perf. Tuning

Support Resources

Q&A

Upcoming GTC Express Webinars

November 5 - Accelerating Face-in-the-Crowd Recognition with GPU Technology

November 6 - Bright Cluster Manager: A CUDA-ready Management Solution for GPU-based HPC

November 13 - Innovative Solutions for Energy Exploration and Data Processing

November 20 - How to Improve Performance using the CUDA Memory Model and Features of the Kepler Architecture

November 21 - How to Speed Up Financial Risk Management Cost Efficiently for Intra-day and Pre-deal CVA Calculations

Register at www.gputechconf.com/gtcexpress

GTC 2014 Call for Posters Open

Posters should describe novel or interesting topics in

- Cloud computing
- Science and research
- Professional graphics
- Mobile computing
- Automotive applications
- Game development

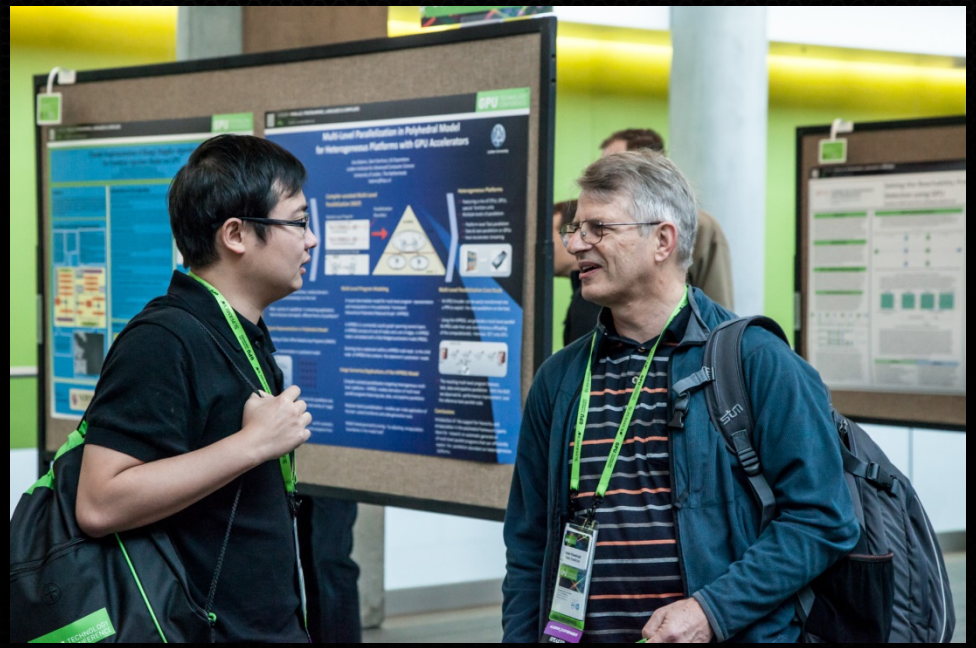

Submit for chance to win Best Poster Award

Submit at www.gputechconf.com