

Disclaimer

- This presentation may contain product features or functionality that are currently under development.
- This overview of new technology represents no commitment from VMware to deliver these features in any generally available product.
- Features are subject to change, and must not be included in contracts, purchase orders, or sales agreements of any kind.
- Technical feasibility and market demand will affect final delivery.
- Pricing and packaging for any new features/functionality/technology discussed or presented, have not been determined.
- This information is confidential.



#NOT
SUPPORTED

VMTN6636U: GPU Enabled Linux VDI


Tony Foster – Sr. Advisor, Technical Marketing, Dell Technologies
VMware vExpert; VMware EUC Champion;
NVIDIA GRID Community Advisor

@wonder_nerd www.wondernerd.net

#NotSupported – VMworld 2012

- The first rule of Not Supported is you don't talk about Not Supported
- The second rule of Not Supported is...

Rules are made to be broken!



Agenda

- Overview
- Host Configuration
- Licensing
- Virtual Machine Configuration
- Testing

NVIDIA GRID “August 2017 Release” scheduled for August 32

Pascal GPUs are #NotSupported for Linux Desktops in VMware Horizon

Assumptions

- Ability to manage a VMware Horizon Environment
- Ability to administer Linux Desktops
- You understand this configuration is **Not Supported** by VMware

Results WILL vary!

Not covered today:

- Integration of Linux Desktops with Directory Services
- VMware Horizon Installation
- Linux VM image optimization
- Basic Linux VM administration tasks

#NOT
SUPPORTED

Overview


vmware
VMTN



#VMworld #VMTN6636U #NotSupported

Hardware Specs

- Testing on Cisco UCS C240 M3 host
 - Dual E5-2640 – 6 Core Procs
 - 64GB of RAM
 - NVIDIA P4 @ 384.37 (Version 5.0 beta)
- VMware vSphere 6.5 (Build 5310538)
- vCenter Server Appliance 6.5.0 (Build 5705665)
- VMware Horizon 7.1.0 (Build 5170113)
 - Basic Environment Only
 - *Sub-optimal*
- Management environment on separate host (Dell R610)
 - vCenter Appliance
 - AD/DNS (Windows 2k8 R2)
 - Jump Box (Windows 2k8 R2)
 - NVIDIA GRID License Server (CentOS7.1 & Windows 2k8 R2)
 - vSphere Connection Server (Windows 2k8 R2)
- Horizon View Client run on Jump box




VM Specs

- CentOS 7.1 (x64)
 - 4 vCPU
 - 8GB vRAM
 - VMware Blast Extreme protocol

GPU

- GRID_p4-4q Profile
- GRID_p4-8q Profile
- Previous Testing
 - GRID_m60-4q vGPU Profile
 - GRID_m60-8q vGPU Profile
 - Passthrough



Host Configuration

vmware
VMTN



#VMworld #VMTN6636U #NotSupported

Installing the Virtual GPU Manager (VIB)

1. Enter maintenance mode on the ESXi host

```
esxcli system maintenanceMode set --enable true
```

2. Stop xorg: `/etc/init.d/xorg stop`

3. Install the Virtual GPU Manager VIB

```
esxcli software vib install -v [full-vib-path] / [vib-name].vib
```

Change ECC Mode

1. Set ECC config for the GPU

1. `nvidia-smi --ecc-config=0`

2. *Note for Maxwell based GPUs it is still necessary to use GPU Mode Switch*

2. Start xorg: `/etc/init.d/xorg start`

3. Reboot ESXi host

4. Verify the VIB version loaded

1. `vmkload_mod -l | grep nvidia`

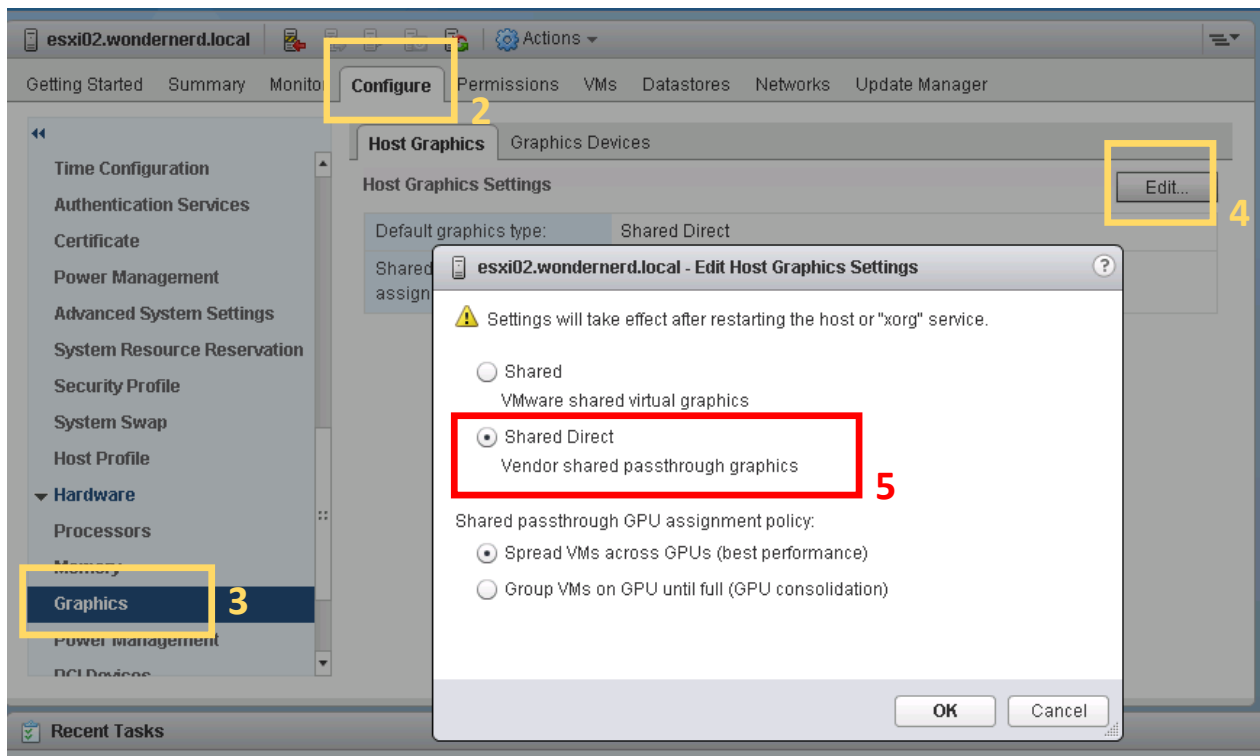
2. `nvidia-smi`

5. Exit Maintenance Mode

`esxcli system maintenanceMode set --enable false`

Set Default Graphics – Shared Direct

1. ESXi Host
2. Configure Tab
3. Graphics option
4. Edit Button
5. Shared Direct



Licensing

vmware
VMTN



#VMworld #VMTN6636U #NotSupported

NVIDIA GRID Licensing

- License Editions

Edition	Features
GRID Virtual Application	<ul style="list-style-type: none">• Virtual GPUs for virtual application computing
NVIDIA GRID vPC (GRID vPC)	<ul style="list-style-type: none">• Virtual GPUs for business desktop computing
NVIDIA® Quadro® Virtual Data Center Workstation (Quadro vDWS)	<ul style="list-style-type: none">• Virtual GPUs for midrange and high-end workstation computing• Workstation graphics on GPU passthrough

- Each edition offers multiple virtual GPU options



Licensing is now enforced across all Operating Systems



vGPU Licensed on Tesla P40

Name	Max Instances	FB Memory	Display Heads	Max X Res	Max Y Res	License
P40-1Q	24	1024MiB	4	4096	2160	GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0
P40-2Q	12	2048MiB	4	4096	2160	GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0
P40-3Q	8	3072MiB	4	4096	2160	GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0
P40-4Q	6	4096MiB	4	4096	2160	GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0
P40-6Q	4	6144MiB	4	4096	2160	GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0
P40-8Q	3	8192MiB	4	4096	2160	GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0
P40-12Q	2	12288MiB	4	4096	2160	GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0
P40-24Q	1	24576MiB	4	4096	2160	GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0
P40-1A	8	1024MiB	1	1280	1024	GRID-Virtual-Apps,3.0
P40-2A	4	2048MiB	1	1280	1024	GRID-Virtual-Apps,3.0
P40-3A	2	3072MiB	1	1280	1024	GRID-Virtual-Apps,3.0
P40-4A	1	4096MiB	1	1280	1024	GRID-Virtual-Apps,3.0
P40-6A	4	6144MiB	1	1280	1024	GRID-Virtual-Apps,3.0
P40-8A	3	8192MiB	1	1280	1024	GRID-Virtual-Apps,3.0
P40-12A	2	12288MiB	1	1280	1024	GRID-Virtual-Apps,3.0
P40-24A	1	24576MiB	1	1280	1024	GRID-Virtual-Apps,3.0
P40-1B	24	1024MiB	4	2560	1600	GRID-Virtual-PC,2.0; GRID-Virtual-WS,2.0; GRID-Virtual-WS-Ext, 2.0

Install GRID License Server

- Windows or Linux
 - Windows 7 / 8 / 8.1 / 10 / Server 2008 R2
 - Red Hat Enterprise Linux 7.1 64-bit
 - CentOS 7.1 64-bit (141MB disk space)
- Static IP Address
- License File Bound to MAC Address
- *Detailed* in “*GRID License Server Release Notes*”
- Supports High Availability

License Server Installation - CentOS

1. Install JRE

```
java -version
```

```
sudo yum install java (if not installed or lower than V.1.7)
```

2. Install Apache Tomcat

```
sudo yum install tomcat tomcat-webapps
```

```
sudo systemctl enable tomcat.service
```

```
sudo systemctl start tomcat.service
```

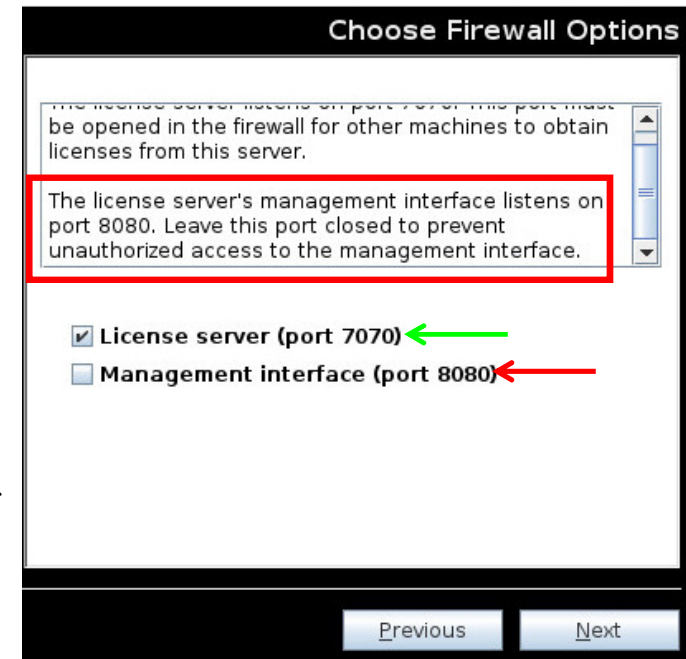
3. Install License Server

```
tar xzf NVIDIA-linux-2015.09-0001.tgz
```

```
sudo ./setup.bin
```

GRID License Server Configuration

1. Respond to questions using “Next” to navigate
2. At the Firewall Options
 - License Server (Port 7070) **Checked**
 - Management Interface (Port 8080) **Unchecked**
3. Click Done to Finish Installation
4. Verify Install by going to `http://localhost:8080/licserver`



Generate License File

1. Login to NVIDIA Enterprise
2. Click **Register License Server** on left (1)
 1. Enter Mac Address (2)
 2. Alias (optional)
 3. Site Name (optional)
3. Click **Map-Add-ons** (3)
 1. Enter quantity of GRID Licenses (4)
 2. Click **Map Add-Ons** button (5)
4. Click **Download License File** (6)
5. Install license by using the “License Management” section (7)
 1. Click the **Choose File** (8) button browse for the `bin` file containing the licenses.
 2. Click the **Upload** button (9)

Check Point

- Physical Cards Installed
- GPU Manager (VIB) Installed
- GRID License Server Installed, Configured, and Licensed
- Linux VM Created for Template

#NOT
SUPPORTED

VM Configuration

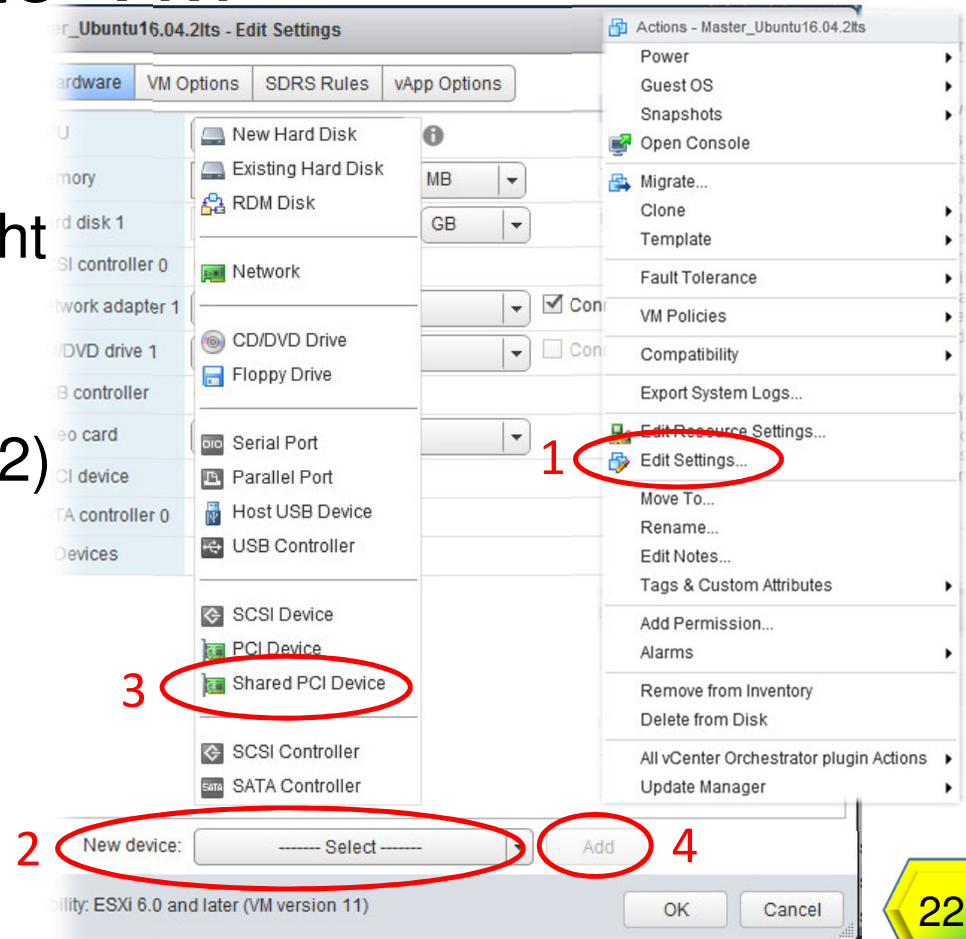
vmware
VMTN



#VMworld #VMTN6636U #NotSupported


Add vGPU to Template VM

1. With Linux VM Shutdown
2. In the vSphere Web Client Right Click VM
 - Select Edit Settings (1)
3. Click **New device** drop down (2)
 - Select Shared PCI Device (3)
4. Click the Add button (4)



Add vGPU to Template VM (Continued)

1. Select desired GPU profile (1)
2. Click Reserve all Memory button (2)
3. Click OK button (3)
4. Power on VM



Configure Linux Template VM

1drnr.me/CfgHs7LD 

- For Ubuntu disable Compiz for improved performance
 - <http://kb.vmware.com/kb/2114809>
- Configure networking to resolve the FQDN of the Connection Server
- For RHEL and CentOS
 - Map Host name to 127.0.0.1 in `/etc/hosts`
 - Verify `virbr0` is disabled
 - `virsh net-destroy default`
 - `virsh net-undefined default`
 - `service libvirtd restart`



Important!

Configure Linux Template VM

- Install gcc
 - `sudo yum group install "Development Tools"`
- Edit the `/etc/nsswitch.conf`
 - `hosts: cache db files dns`
- Configure runlevel for 5
- Disable Nouveau driver
 - Varies based on OS


Install Drivers in the VM

- **GUEST OS Drivers Must Match VIB Version**

1. Copy the VMware Horizon Linux Agent to the VM
2. Copy the Linux GRID Driver Package to the VM
3. **Before attempting to run the driver installer, exit the X server and terminate all OpenGL applications.**
 - *RHEL & CentOS:*
 1. `sudo init 3`
 2. `sudo service gdm stop`
 - *Ubuntu:*
 1. Use **CTRL-ALT-F1** to switch to a console login prompt.
 2. Log in and shut down the display manager: `sudo service lightdm stop`
4. `chmod +x NVIDIA-linux-x86_64-[version]-grid.run`
5. Run the driver installer:
`sudo sh ./NVIDIA-Linux_x86_64-[Version]-grid.run`

Finish Driver Configuration

- Select **Yes** to update the X configuration
- Reboot the VM
- **Switch to Horizon, ssh, or VNC Console – (Black Screen on terminal connection)**
- Install Horizon Linux Agent
 - Unpack the Horizon Linux Agent
 - `sudo sh ./[path]/install_viewagent.sh`
- Run: `nvidia-smi` (to validate that the card is present)
- Reboot



Result of nvidia-smi

- Notice it shows the information about the vGPU added

```
root@basecent:~  
[root@basecent ~]# nvidia-smi  
Wed Aug 23 17:16:12 2017  
-----  
| NVIDIA-SMI 384.37              Driver Version: 384.37      |  
-----  
| GPU   Name                Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |  
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |  
-----  
|    0   GRID P4-4Q             On          | 0000:02:02.0   On      |                     N/A |  
| N/A   N/A   P8     N/A /  N/A |   401MiB /  4095MiB |      0%      Default  |  
-----  
|  
| Processes:                      GPU Memory |  
|   GPU   PID     Type  Process name      Usage   |  
-----  
|    0   16970    G   /usr/bin/gnome-shell  40MiB |  
|    0   22740    G   /usr/bin/Xorg       56MiB |  
-----  
[root@basecent ~]#
```

#NOT
SUPPORTED


Configuring gridd.conf - vGPU

1. As root In Text Editor Open `/etc/nvidia/gridd.conf`
 1. `sudo vi /etc/nvidia/gridd.conf`
 2. Template can be found: `/etc/nvidia/gridd.conf.template`
2. Set the `ServerAddress` to the Address of your GRID License Server
3. Set the `FeatureType` to 1 for **vGPU**
4. Save Changes to the file (`esc : wq` for VI)
5. Restart the `nvidia-gridd` service
 - `sudo service nvidia-gridd restart`



Licensed

The screenshot shows a web browser window at localhost:8080/licserver/. The page title is "Licensed Clients". Below the title, there is a message: "Licensed Clients with features consumed or reserved. Click a Client ID for further details." A table displays one client entry:

Client ID	Client ID Type	Client Type
 005056AD5F5C	ETHERNET	VIRTUAL


Page 1

Copyright (c) 2017 NVIDIA Corporation. All Rights Reserved. 20170501-0001

The left sidebar contains two sections: "License Server" with links for Licensed Clients, Reservations, Licensed Feature Usage, License Management, Configuration, and Login; and "License Client Manager" with links for About and Settings.

Limitations

- Pascal GPUs not supported
- No VMware vMotion
- No VMware DRS
- No Instant Clone Support
- No Snapshotting of running VMs
- Must use VMware Blast Extreme protocol



Automated Pool in Horizon

1. In desktop pools, right click **Add**
2. Select **Automated Desktop Pool**
3. Select pool type
4. Select Clone type
5. Enter pool info
6. Configure Desktop pool settings
 1. Default Display protocol: **VMware Blast**
 2. Allow Users to Choose Protocol: **No**
 3. 3D Render: [**Automatic** | **NVIDIA GRID VGPU** | **Hardware**]
7. Continue as normal

#NOT
SUPPORTED

Add Desktop Pool - CentOS71_8q

Desktop Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Desktop Pool Identification
- Desktop Pool Settings
- Provisioning Settings
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

Desktop Pool Settings

General

State: ▾

Connection Server restrictions: None

Remote Settings

Remote Machine Power Policy: ▾ ?

Automatically logoff after disconnect: ▾

Allow users to reset/restart their machines: ▾

Remote Display Protocol

Default display protocol: ▾

Allow users to choose protocol: ▾

3D Renderer: ▾ ?

Max number of monitors: ▾ ?

Max resolution of any one monitor: ▾ ?

HTML Access: Enabled ?

< Back Next > Cancel

Testing


#NOT
SUPPORTED



#VMworld #VMTN6636U #NotSupported

Testing

- OpenGFX
- Won't run without an OpenGL capable graphics card
- Can highlight other problems



GFXBench – It Works!

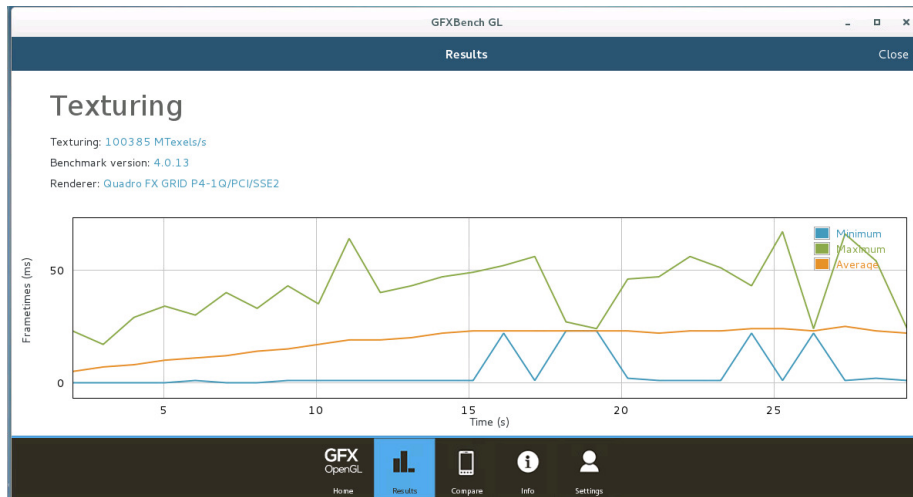
gfxbench.com

More

Single pass results in a non-optimized environment, individual results may vary significantly

Testing with P4-8Q

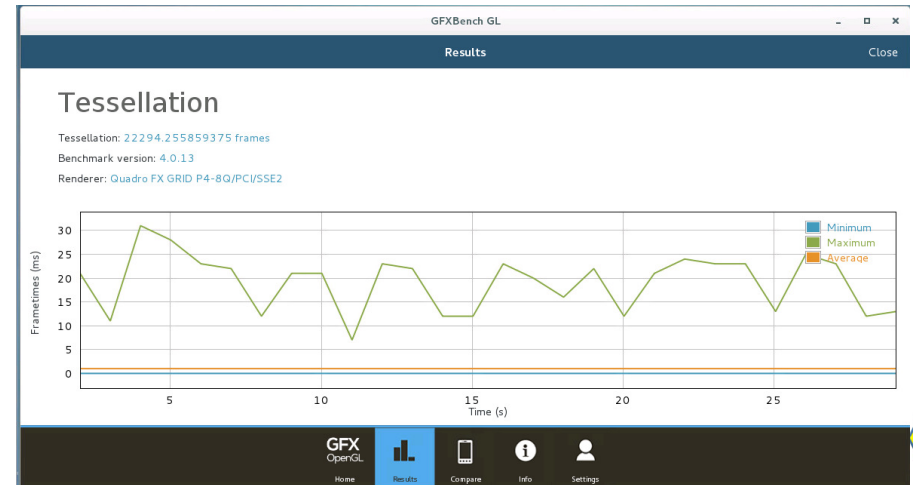
- Lowest FPS: Texturing @ 63.5544 FPS
- Highest FPS: Tessellation @ 713.276 FPS
- Off screen Lowest FPS: 1080p Texturing Offscreen @ 95.8384 FPS
- Off screen Highest FPS: 1080p Tessellation Offscreen @ 1481.04 FPS



Testing with P4-4Q

- Lowest FPS: Texturing @ 64.4044 FPS
- Highest FPS: Tessellation @ 713.862 FPS
- Off screen Lowest FPS: 1080p Texturing Offscreen @ 95.9874 FPS
- Off screen Highest FPS: 1080p Tessellation Offscreen @ 1473.5 FPS

Single pass results in a non-optimized environment, individual results may vary significantly





Thank you for attending this vBrownBag TechTalk – A VMTN Network Member

Presentation available at www.wondernerd.net



#VMworld #VMTN6636U #NotSupported