TeamRGE

Remoting Graphics Experts

NVIDIA vGPU – The Full Update

June 2019

Jits Langedijk

Sr. Solution Architect Professional Visualization - NVIDIA

JLangedijk@nvidia.com

@JRLangedijk

/in/JitsLangedijk

JitsLangedijk.com

Workspace IoT

TeamRGE
Remoting Graphics Experts

Citrix Technology Advocates

NVIDIA vGPU – Introduction

4 FACTORS CONTRIBUTING TO VDI MARKET GROWTH

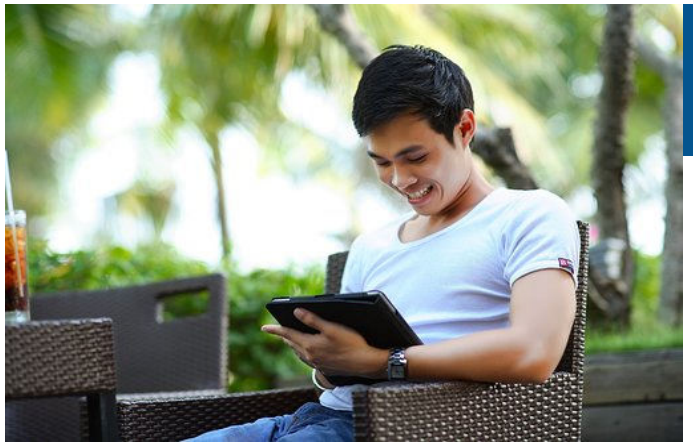

Workplace Flexibility and Business Agility

Collaboration

Security & Risk Management

Reduced CAPEX and OPEX

Source: BusinessWire. Global Virtual Desktop Infrastructure (VDI) Market to Grow Rapidly Through 2021, Due to Its Ability to Offer Improved Security and Flexibility to Enterprises: Technavio.
http://www.businesswire.com/news/home/20161111005754/en/Global-Virtual-Desktop-Infrastructure-VDI-Market-Grow

PERFORMANCE FROM THE DATA CENTER

NVIDIA Virtual GPU technology delivers graphics accelerated virtual desktops and applications

All devices have graphics

Virtual machines also need a GPU

THE NEW DIGITAL WORKER

They Way we Work is Changing

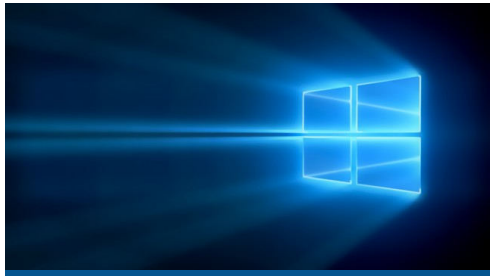

Windows 10

32% increase in CPU requirement over Windows 7[1]

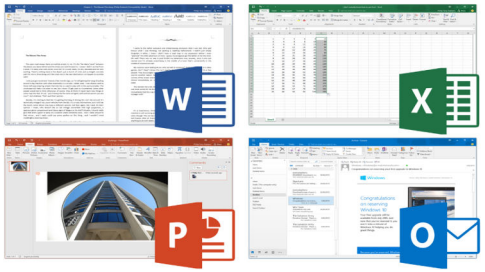

Office 365/Office 2016

50-85% increase in CPU requirements over Windows 7[1]

Web Browsers

Modern browsers are hardware accelerated by default

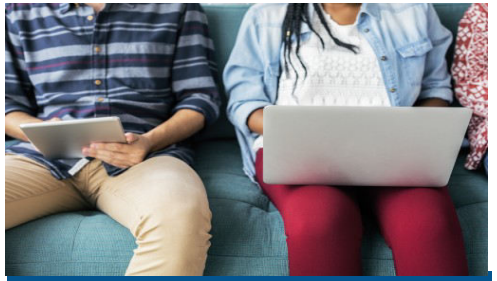

Latest Web Standards

Flash, HTML5, and WebGL are all very taxing to the CPU

PDF Viewers

Adobe® Acrobat® and Microsoft Edge are hardware accelerated by default

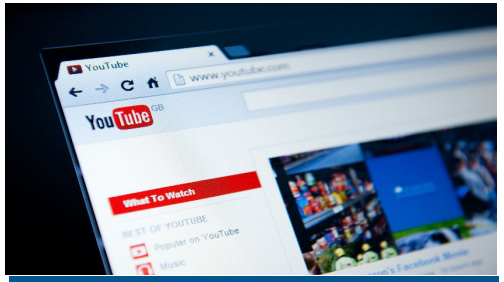

Collaboration and Video

Skype and YouTube are now prevalent across the enterprise

Digital Imaging & Design

Some features in Adobe® Photoshop® won't work without a GPU[2]

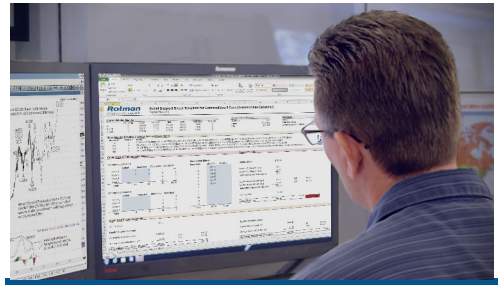

Multi-, High Res Monitors

Multi-monitors is the new normal and 4K is becoming mainstream

HOW IT WORKS

Virtualisation without a GPU

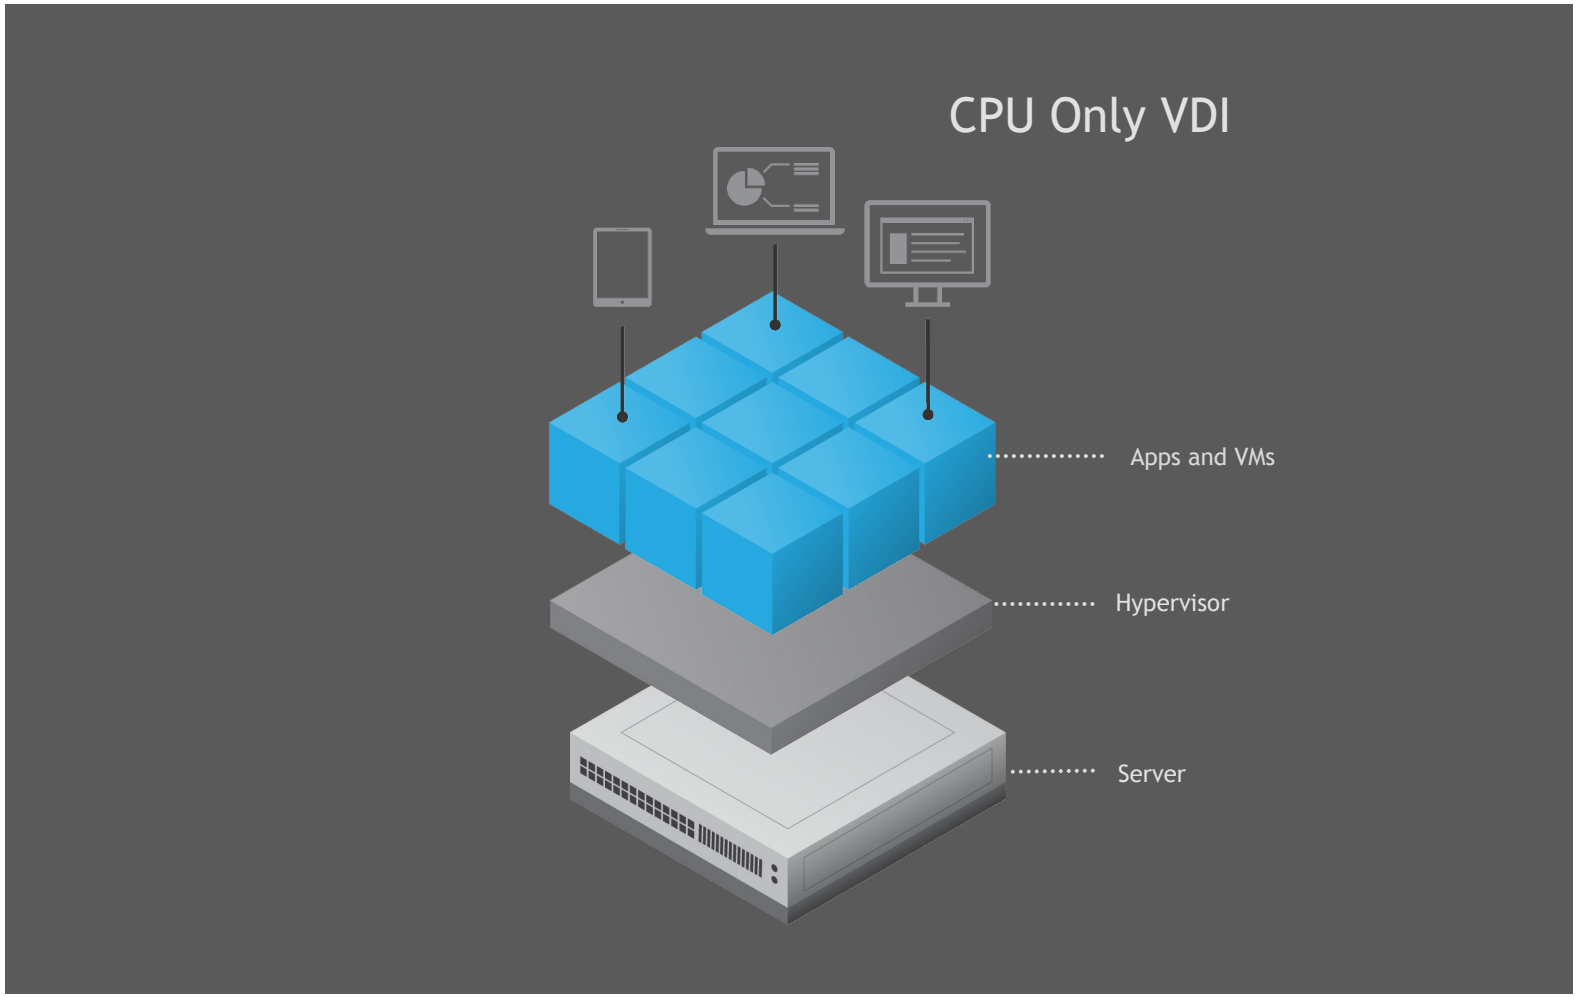

HOW IT WORKS

NVIDIA virtual GPU products deliver a GPU Experience to every Virtual Desktop

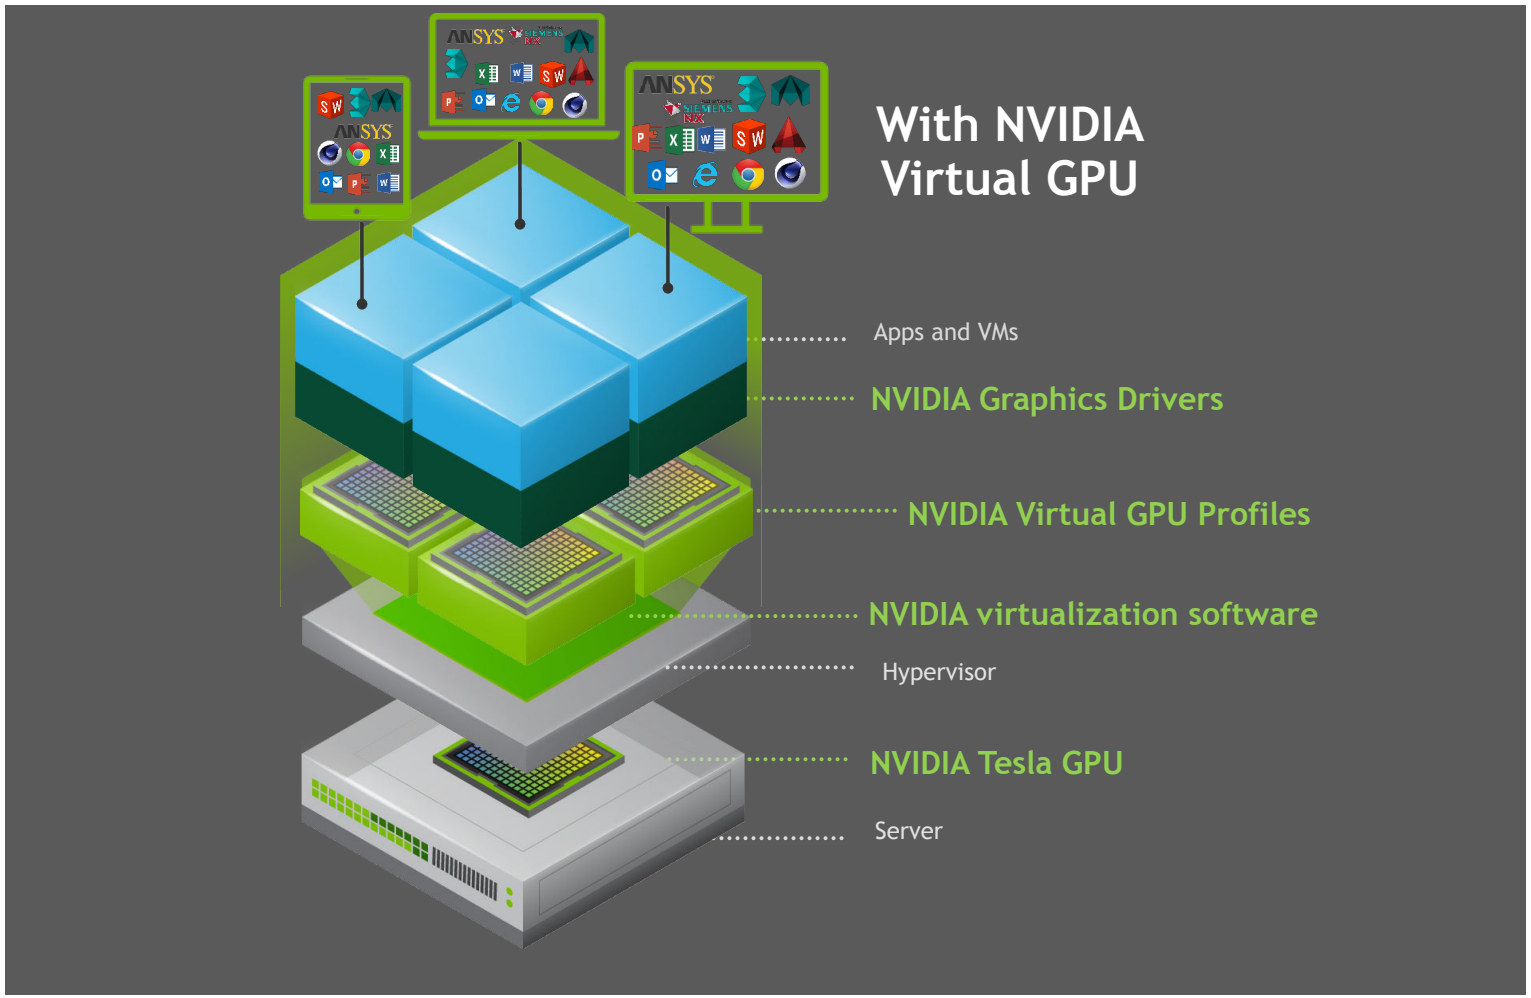

HOW IT WORKS

GPU Pass-Through

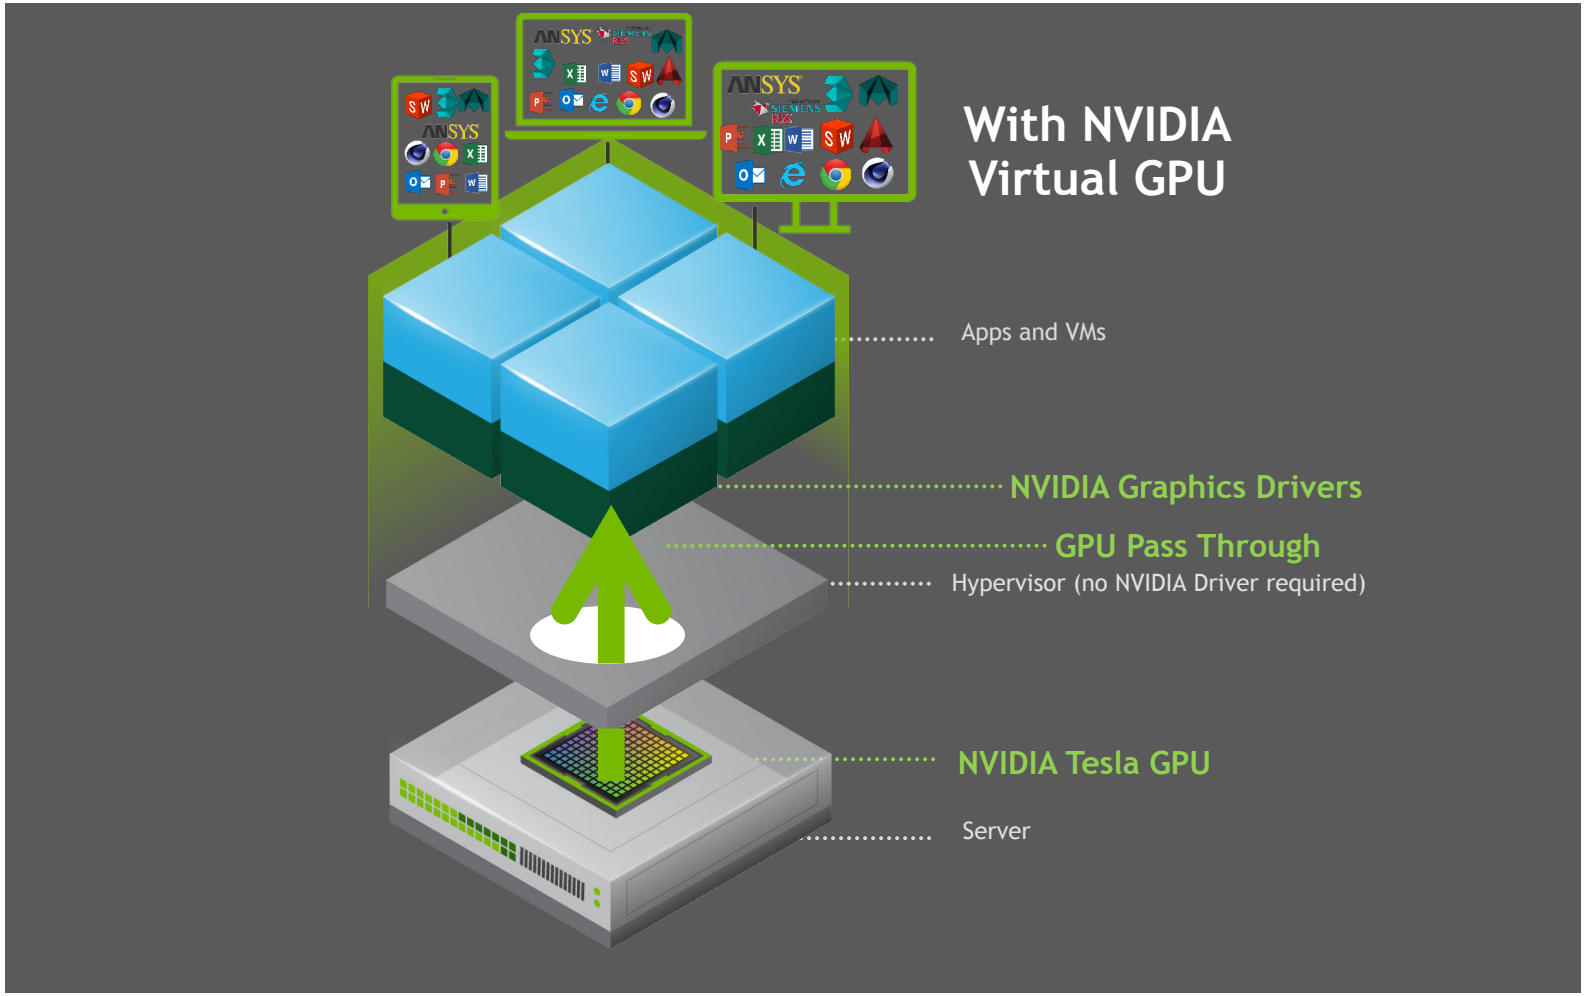

EVOLUTION OF GPU ACCELERATED VDI

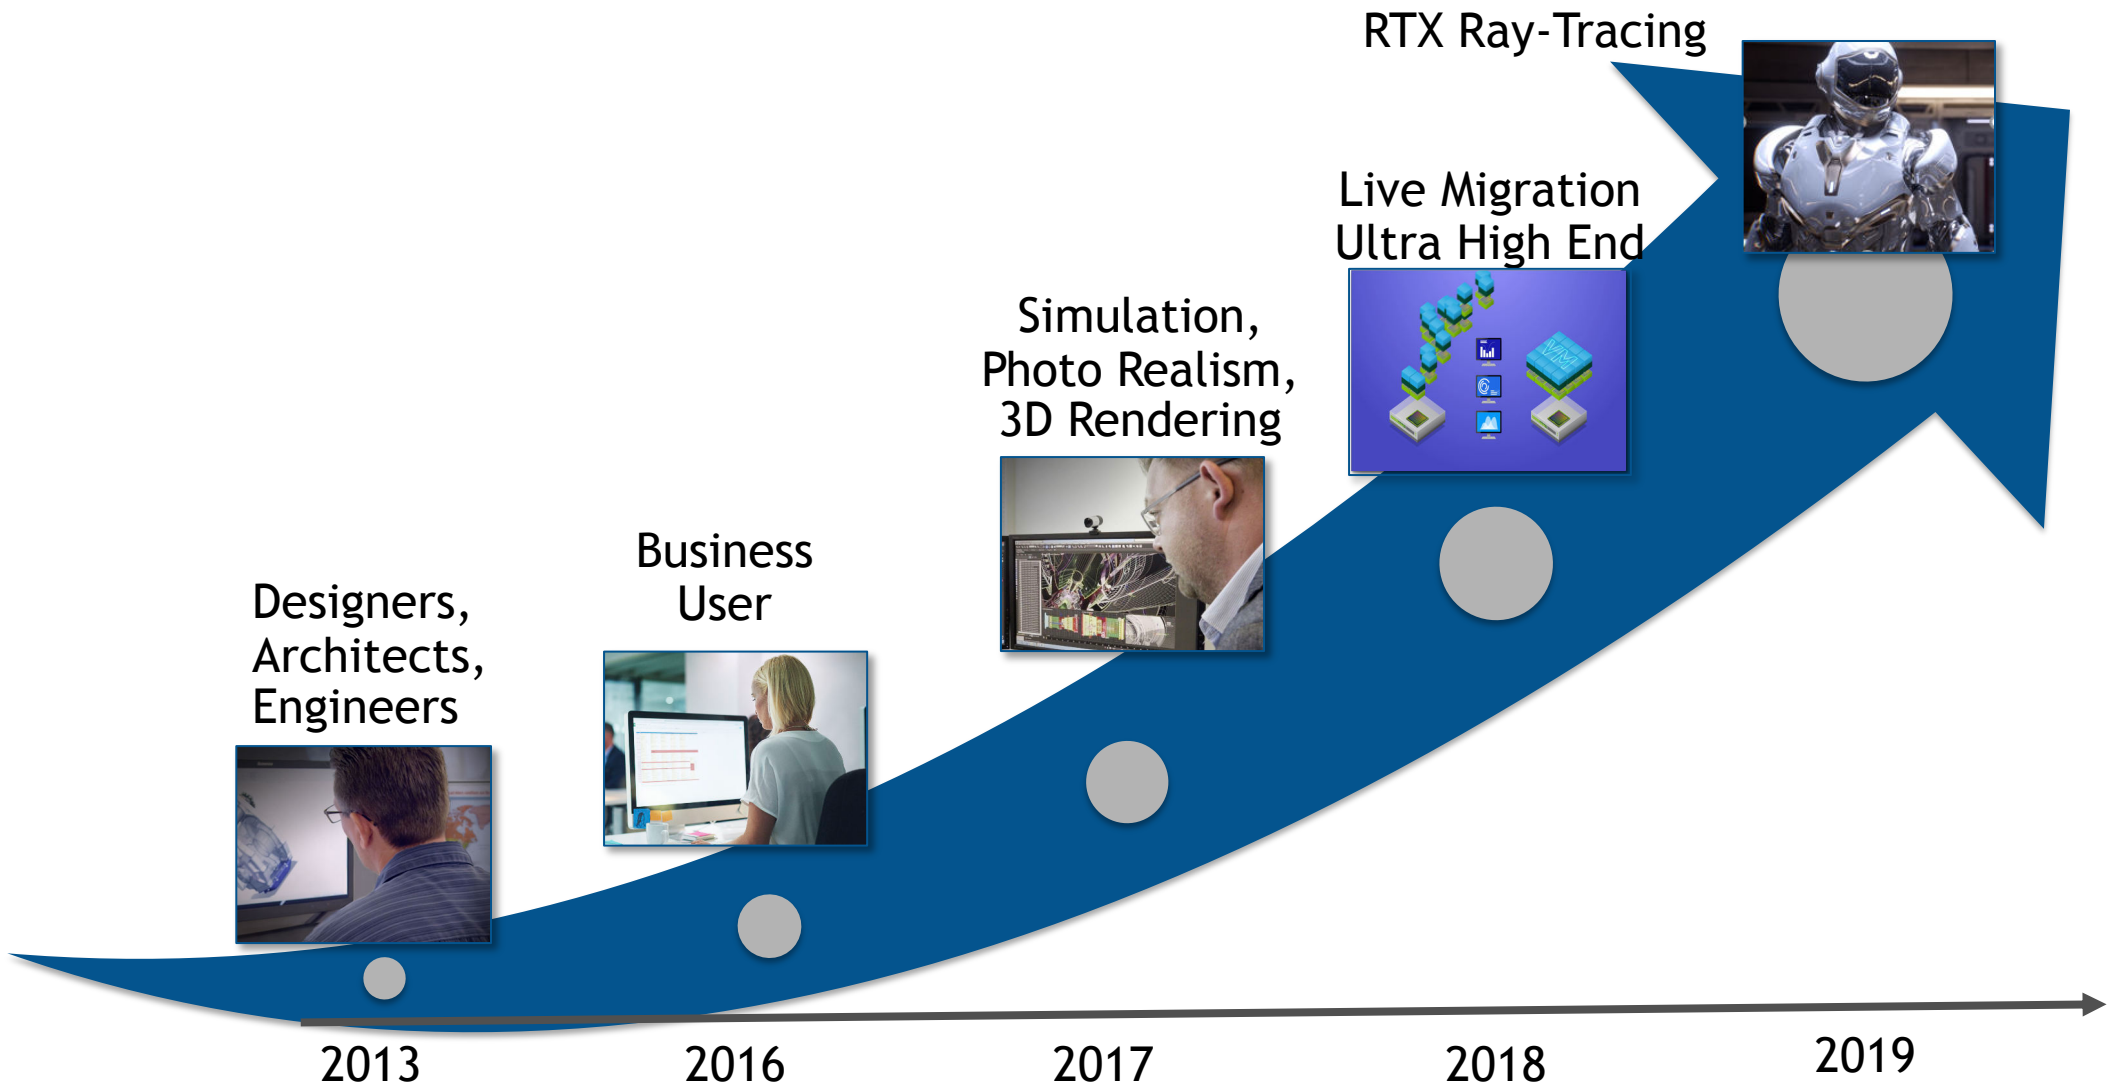

NVIDIA vGPU – Versions

vGPU RELEASE CADENCE

Keep the GPU Manager and VM's Driver within the same major release

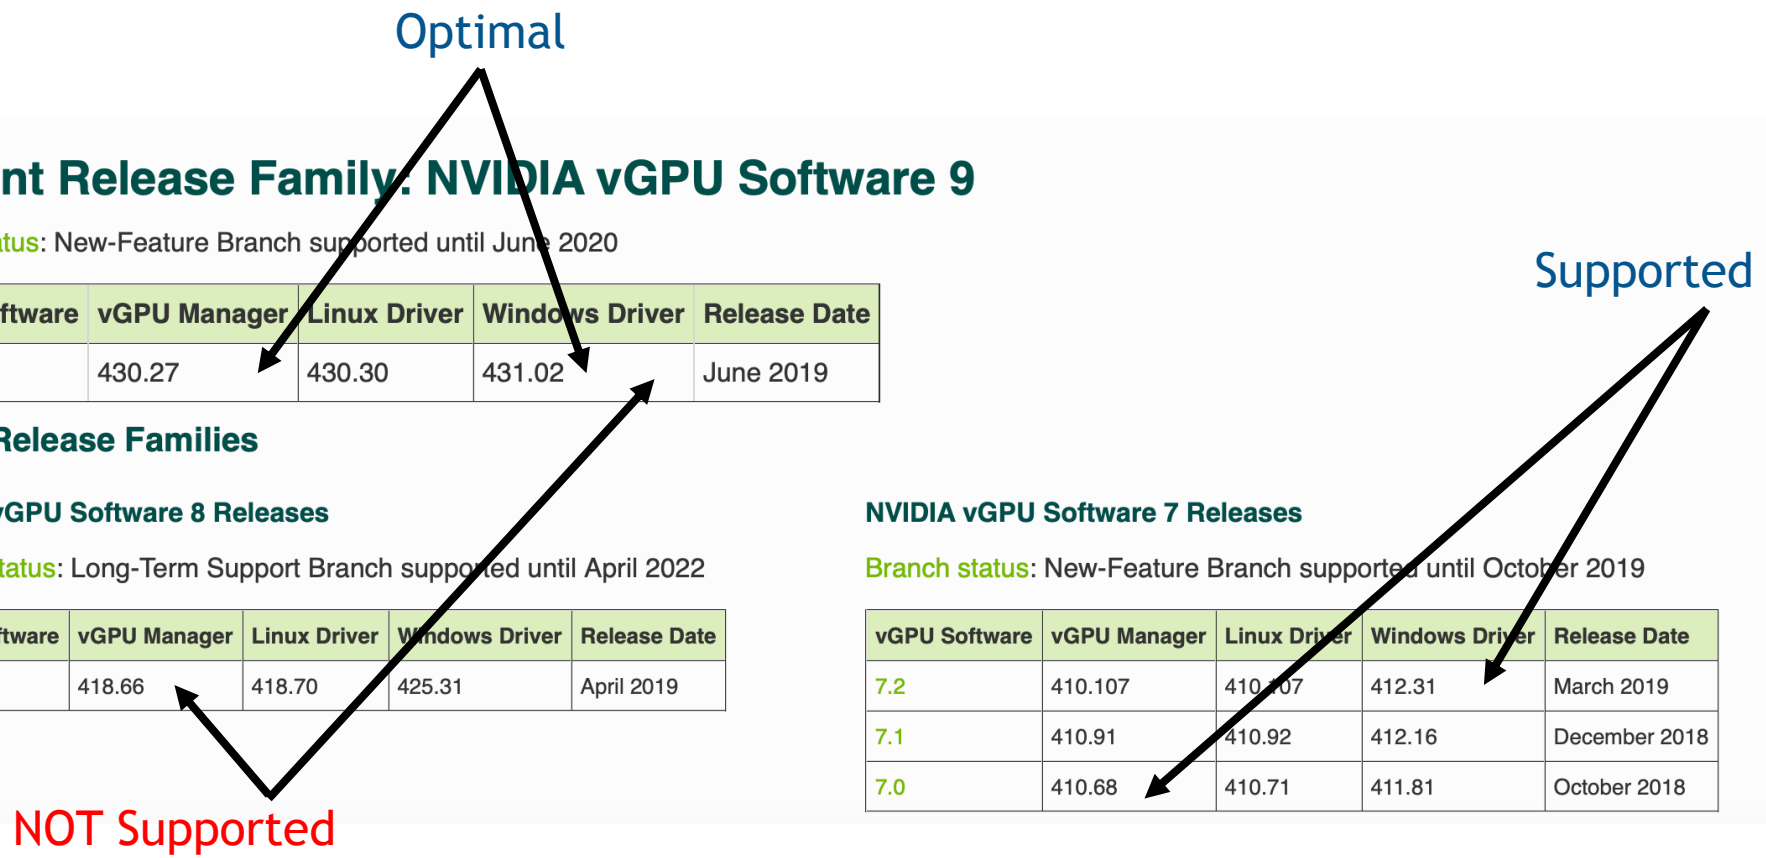

Current Release Family: NVIDIA vGPU Software 9

Branch status: New-Feature Branch supported until June 2020

vGPU Software	vGPU Manager	Linux Driver	Windows Driver	Release Date
9.0	430.27	430.30	431.02	June 2019

Earlier Release Families

NVIDIA vGPU Software 8 Releases

Branch status: Long-Term Support Branch supported until April 2022

vGPU Software	vGPU Manager	Linux Driver	Windows Driver	Release Date
8.0	418.66	418.70	425.31	April 2019

NVIDIA vGPU Software 7 Releases

Branch status: New-Feature Branch supported until October 2019

vGPU Software	vGPU Manager	Linux Driver	Windows Driver	Release Date
7.2	410.107	410.107	412.31	March 2019
7.1	410.91	410.92	412.16	December 2018
7.0	410.68	410.71	411.81	October 2018

https://docs.nvidia.com/grid/

VIRTUAL GPU OCTOBER 2018 (vGPU 7.X)

Unprecedented Performance & Manageability

Multi-vGPU Support
World's Most Powerful
Quadro vDWS

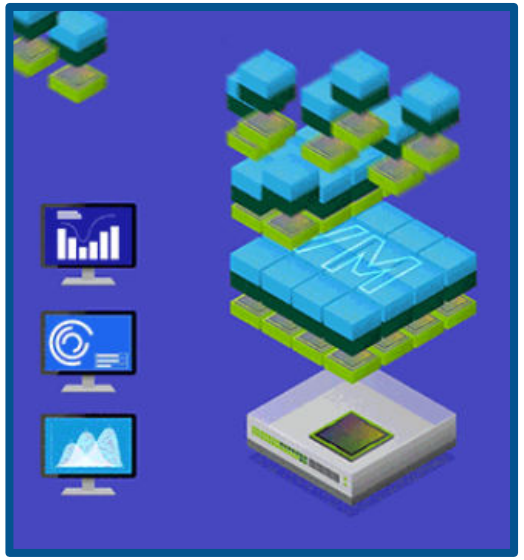

vMotion Support for vGPU
Live Migration of vGPU enabled
VMs Quadro vDWS & GRID

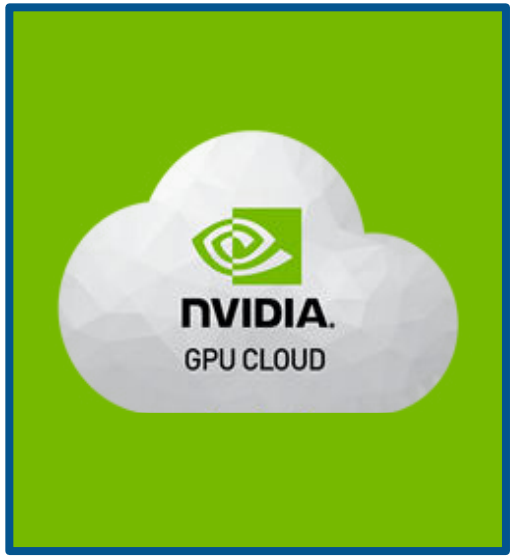

NGC with vGPU
Available with vGPU
Quadro vDWS

MIXED WORKLOADS WITH NVIDIA vGPU

Increase productivity & utilization, decrease costs

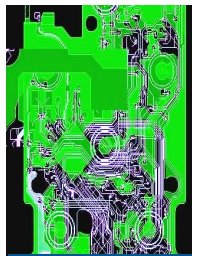

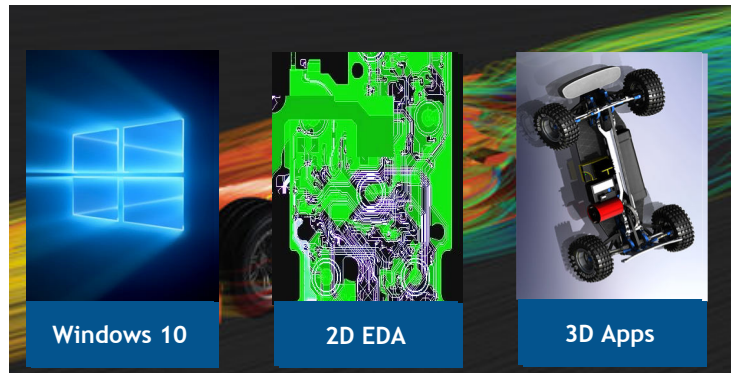

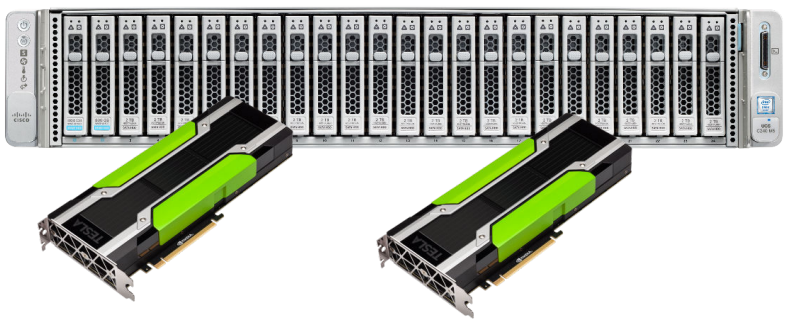

vSphere vMotion with vGPU

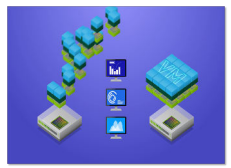

End-to-end GPU insights with vROPS

ENABLING THE AGILE DATA CENTER

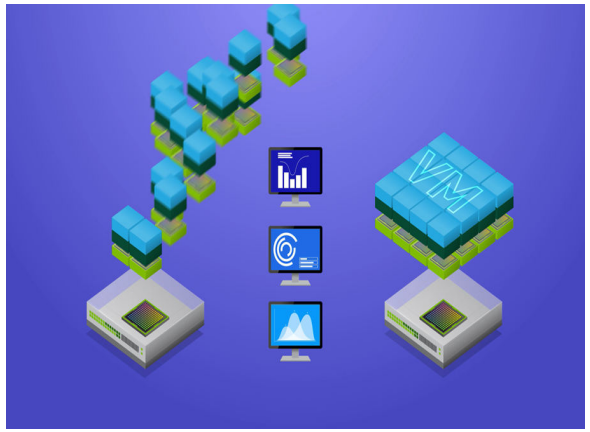

Ensure High Availability

Live migrate GPUs with support for vMotion & XenMotion

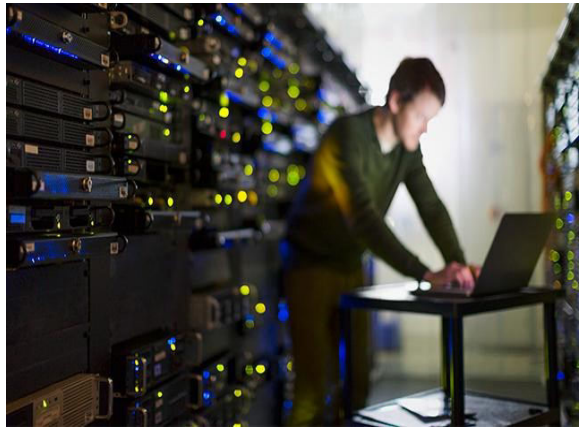

Lower IT Tickets

End-to-end monitoring with vROps integration

Maximize Infrastructure Utilization

Virtualize any workload and run VDI by day, HPC by night

VIRTUAL GPU DECEMBER 2018 (vGPU 7.1)

Unprecedented Performance & Manageability

Nutanix AHV 5.10

VMware Horizon 7.7

Tesla T4 GPU Support

Latest Generation Turing

Quadro vDWS

VIRTUAL GPU MARCH 2019 (vGPU 7.2)

Unprecedented Performance & Manageability

Citrix Virtual Apps and Desktops - 1903

VMware Horizon 7.8

VIRTUAL GPU SPRING 2019 (8.0)

RTX Support for Quadro Virtual Workstations

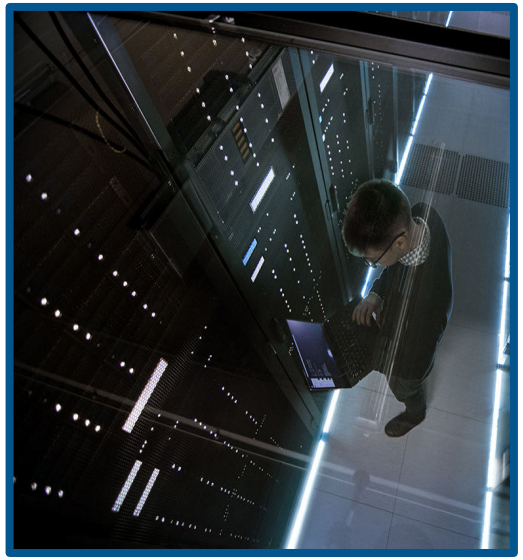

Quadro RTX 6000/8000

Quadro RTX-powered
Virtual Workstations

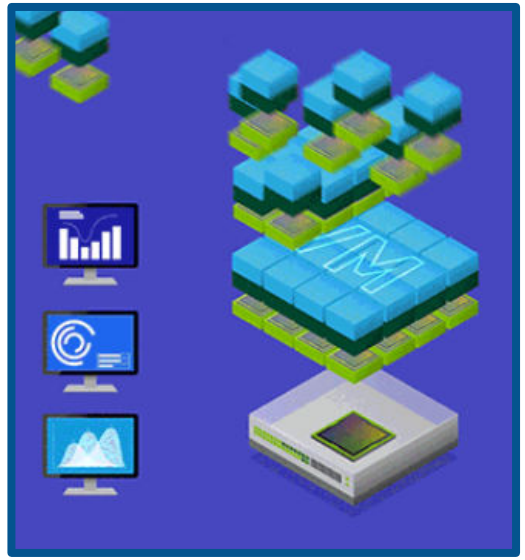

Live Migration for V100

Live Migration of vGPU enabled
VMs with V100

4K Support for 1B Profile

Available with
GRID software

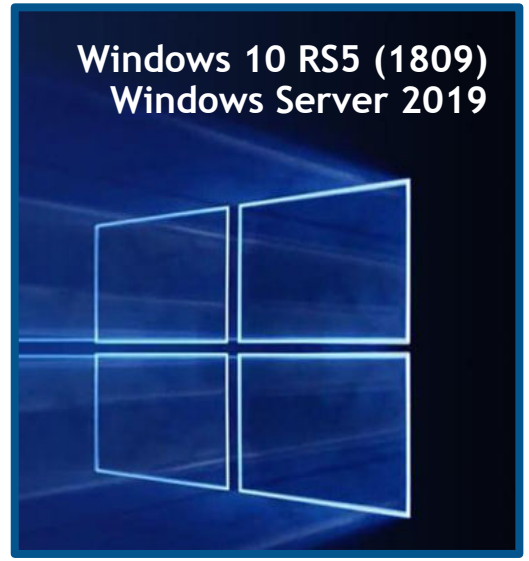

Updated Windows Support

Available with
vGPU software

ACCELERATING CREATIVITY

Most Powerful Rendering Platform

Media & Entertainment
Photorealistic Rendering & Design

Manufacturing
Product Design

Architecture
Design and Rendering

RTX FOR QUADRO VIRTUAL WORKSTATIONS

Value of Virtualization Extended to RTX Server

Design Workstation

Design Workstation

Design Workstation

Render Node

Render Node

Render Node

Render Node

Render Node

VIRTUALIZATION WITH RTX SERVER

INCREASE PRODUCTIVITY

Get up to 5X faster rendering per 1U versus a CPU-only solution and scale up to 4U to maximize rendering speed by 20X.

IMPROVE UTILIZATION

Consolidate multiple physical workstations to one server with multiple virtual workstations, and suspend workstation nodes/spin up as needed

REDUCE COSTS

Buy one virtualized server instead of multiple CPU-only servers, which saves money and is easier to manage.

vGPU SOFTWARE EDITIONS

Product Portfolio

Virtual Apps

For organizations deploying XenApp (CVAD) or RDSH solutions. Designed to deliver Windows applications at full performance

Virtual PC

For users who want a virtual desktop but need great user experience leveraging PC Windows applications, browsers and high definition video

Quadro Virtual Datacenter Workstation

For users who want to be able to use remote professional graphics applications with full performance and application certification

Focus on Compute

vGPU SOFTWARE EDITIONS

End Entitlements

Feature	vApps	vPC	QvDWS	vCS
License Entitlement				
Concurrent user (CCU)	✓	✓	✓	
Per GPU				✓
Capability Entitlement				
Desktop Virtualization		✓	✓	
RDSH App Hosting	✓	✓	✓	
RDSH Desktop Hosting	✓	✓	✓	
Compute Virtualization			✓	✓
Windows Guest OS	✓	✓	✓	
Linux Guest OS		✓	✓	✓
Maximum Displays	1	4	4	1
Maximum Resolution	1280*1024	4096*2160 (4K)	4096*2160 (4K)	4096*2160 (4K)
NVIDIA Quadro Software Features			✓	
CUDA & OpenCL Supported			✓	✓
ECC & Page Retirement			✓	✓
Multi-GPU			✓	✓
NVLINK			✓	✓
GPU Pass-through Supported	✓		✓	✓
Bare Metal Supported	✓		✓	

vGPU SOFTWARE EDITIONS

vGPU Profiles Supported

Feature	vApps	vPC	QvDWS	vCS
vGPU Profiles Supported				
512 MB		✓	✓	
1 GB	✓	✓	✓	
2 GB	✓	✓	✓	
3 GB	✓		✓	
4 GB	✓		✓	✓
6 GB	✓		✓	✓
8 GB	✓		✓	✓
12 Gb	✓		✓	✓
16 GB	✓		✓	✓
24 GB	✓		✓	✓
32 GB	✓		✓	✓
48 GB			✓	

nvidia.com/content/grid/pdf/Virtual-GPU-Packaging-and-Licensing-Guide.pdf

QUADRO vDWS HIGH-END GPUs

Key Features

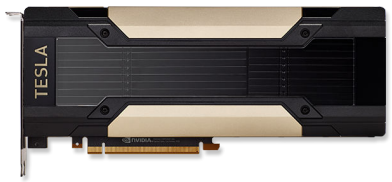

	RTX 6000	RTX 8000	V100
RT Cores	✓	✓	
Tensor Cores	✓	✓	✓
Memory	24 GB GDDR6	48 GB GDDR6	32/16 GB HBM2
FP 32 (single precision)	✓	✓	✓
FP 64 (double precision)			✓

NVIDIA DATA CENTER GPUs
Recommended for Virtualization

	V100	P40	RTX 8000	RTX 6000	T4	M10	P6
GPUs / Board (Architecture)	1 (Volta)	1 (Pascal)	1 (Turing)	1 (Turing)	1 (Turing)	4 (Maxwell)	1 (Pascal)
CUDA Cores	5,120	3,840	4,608	4,608	2,560	2,560 (640 per GPU)	2,048
Tensor Cores	640	---	576	576	320	---	---
RT Cores	---	---	72	72	40	---	---
Memory Size	32 GB/16 GB HBM2	24 GB GDDR5	48 GB GDDR6	24 GB GDDR6	16 GB GDDR6	32 GB GDDR5 (8 GB per GPU)	16 GB GDDR5
vGPU Profiles	1 GB, 2 GB, 4 GB, 8 GB, 16 GB, 32 GB	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 24 GB	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 16 GB, 24 GB, 48 GB	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 24 GB	1 GB, 2 GB, 4 GB, 8 GB, 16 GB	0.5 GB, 1 GB, 2 GB, 4 GB, 8 GB	1 GB, 2 GB, 4 GB, 8 GB, 16 GB
Form Factor	PCIe 3.0 Dual Slot & SXM2 (rack servers)	PCIe 3.0 Dual Slot (rack servers)	PCIe 3.0 Dual Slot	PCIe 3.0 Dual Slot	PCIe 3.0 Single Slot (rack servers)	PCIe 3.0 Dual Slot (rack servers)	MXM (blade servers)
Power	250W/300W	250W	295W	295W	70W	225W	90W
Thermal	passive	passive	active	active	passive	passive	bare board
	PERFORMANCE Optimized					**DENSITY** Optimized	**BLADE** Optimized

WRAP UP

- VDI
- User eXperience; People expect VDI to deliver similar to local PC
- Modern OS and Apps are designed with GPU’s in mind
- vGPU is a true software solution
- Mixed workloads – Graphics; Compute; Rendering
- Enabling virtual Compute usecases

TeamRGE

Remoting Graphics Experts

THANK YOU

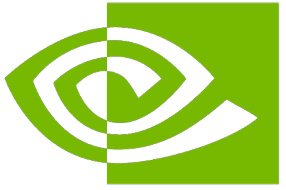

NVIDIA®