NVIDIA TESLA P40
GPU ACCELERATOR

HIGH-PERFORMANCE
VIRTUAL GRAPHICS AND COMPUTE

NVIDIA redefined visual computing by giving designers, engineers, scientists, and graphic artists the power to take on the biggest visualization challenges with immersive, interactive, photorealistic environments. NVIDIA® Quadro® Virtual Data Center Workstation (Quadro vDWS) takes advantage of NVIDIA® Tesla® GPUs to deliver virtual workstations from the data center. Architects, engineers, and designers are now liberated from their desks and can access applications and data anywhere.

The NVIDIA® Tesla® P40 GPU accelerator works with NVIDIA Quadro vDWS software and is the first system to combine an enterprise-grade visual computing platform for simulation, HPC rendering, and design with virtual applications, desktops, and workstations. This gives organizations the freedom to virtualize both complex visualization and compute (CUDA and OpenCL) workloads.

The NVIDIA® Tesla® P40 taps into the industry-leading NVIDIA Pascal™ architecture to deliver up to twice the professional graphics performance of the NVIDIA® Tesla® M60 (Refer to Performance Graph). With 24 GB of framebuffer and 24 NVENC encoder sessions, it supports 24 virtual desktops (1 GB profile) or 12 virtual workstations (2 GB profile), providing the best end-user scalability per GPU. This powerful GPU also supports eight different user profiles, so virtual GPU resources can be efficiently provisioned to meet the needs of the user. And it's available in a wide variety of industry-standard 2U servers.

GPU	1 NVIDIA Pascal GPU
CUDA Cores	3,840
Memory Size	24 GB GDDR5
H.264 1080p30 streams	24
Max vGPU instances	24 (1 GB Profile)
vGPU Profiles	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 24 GB
Form Factor	PCIe 3.0 Dual Slot (rack servers)
Power	250 W
Thermal	Passive

VIRTUALIZE ANY WORKLOAD, ANYWHERE

With NVIDIA virtual GPU software and the NVIDIA Tesla P40, organizations can now virtualize high-end applications with large, complex datasets for rendering and simulations, as well as virtualizing modern business applications. Resource allocation ensures that users have the right GPU acceleration for the task at hand.

NVIDIA software shares the power of Tesla P40 GPUs across multiple virtual workstations, desktops, and apps. This means you can deliver an immersive user experience for everyone from office workers to mobile professionals to designers through virtual workspaces with improved management, security, and productivity.

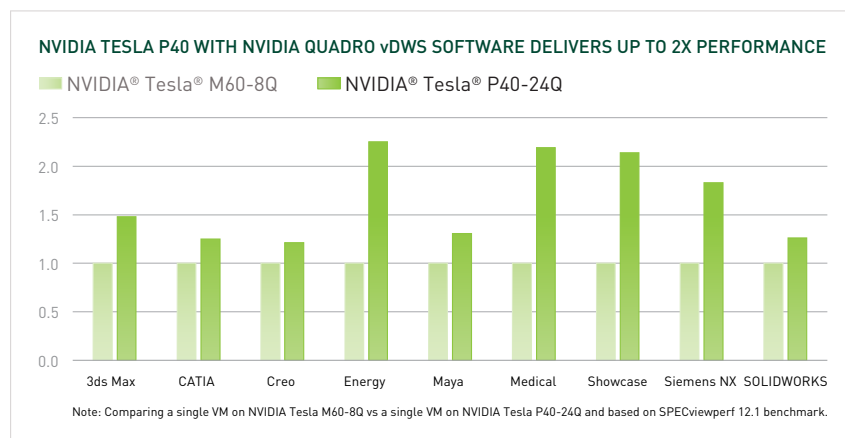## KEY BENEFITS

Exceptional User Experience

Get the ultimate user experience for any workload or vGPU profile. NVIDIA Quadro vDWS software with Tesla P40 GPU supports compute workloads (CUDA and OpenCL) for every vGPU, enabling professional and design engineering workflows at peak performance. The Tesla P40 delivers up to 2X the graphics performance compared to the M60 (Refer to Performance Graph). Users can count on consistent performance with the new resource scheduler, which provides deterministic QoS and eliminates the problem of a "noisy neighbor."

Optimal Management and Monitoring

Management tools give you vGPU visibility into the host or guest level, with application level monitoring capabilities. This lets IT intelligently design, manage, and support their end user's experience. End-to-end management and monitoring also deliver real-time insight into GPU performance. And integration with VMware vRealize Operations (vROps), Citrix Director and XenCenter put flexibility and control in the palm of your hand.

Flexible GPU Infrastructure

Support up to 50% more users per Pascal GPU relative to a single Maxwell GPU, for scaling high performance virtual graphics and compute. More granular user profiles give you more precise provisioning of vGPU resources, and larger profile sizes - up to 3X larger GPU framebuffer than the M60 – for supporting your most demanding users. The P40 provides utilization and flexibility to your NVIDIA Quadro vDWS solution helping you drive down overall TCO.

NVIDIA TESLA P40 WITH NVIDIA QUADRO vDWS SOFTWARE DELIVERS UP TO 2X PERFORMANCE

Legend: ■ NVIDIA® Tesla® M60-8Q   ■ NVIDIA® Tesla® P40-24Q

Categories: 3ds Max, CATIA, Creo, Energy, Maya, Medical, Showcase, Siemens NX, SOLIDWORKS

Note: Comparing a single VM on NVIDIA Tesla M60-8Q vs a single VM on NVIDIA Tesla P40-24Q and based on SPECviewperf 12.1 benchmark.

To learn more about NVIDIA Virtual GPU Software with Tesla GPUs visit **www.nvidia.com/grid**

NVIDIA.