



NVIDIA® TESLA®. ONE PLATFORM. UNLIMITED DATA CENTER ACCELERATION.

The Exponential Growth of Computing






Accelerating scientific discovery, visualizing big data for insights, and providing smart services to consumers are everyday challenges for researchers and engineers. Solving these challenges takes increasingly complex and precise simulations, the processing of tremendous amounts of data, or training sophisticated deep learning networks. These workloads also require accelerating data centers to meet the growing demand for exponential computing.

NVIDIA Tesla is the world's leading platform for accelerated data centers, deployed by some of the world's largest

supercomputing centers and enterprises. It combines GPU accelerators, accelerated computing systems, interconnect technologies, development tools, and applications to enable faster scientific discoveries and big data insights.

At the heart of the NVIDIA Tesla platform are the massively parallel GPU accelerators that provide dramatically higher throughput for compute-intensive workloads—without increasing the power budget and physical footprint of data centers.

Choose the Right NVIDIA Tesla Solution for You

WORKLOAD	Mixed-Workload HPC		Hyperscale HPC		Graphics Virtualization	
USED BY	Supercomputing, Academia, Government	Oil and Gas	Artificial Intelligence/Deep Learning		Design and Manufacturing, Architecture Engineering and Constructions, Defense, Higher Education	
OPTIMIZED FOR	Time to Insight	Imaging Accuracy	Training Time	Jobs/Second/Watt	Graphics Accelerated Virtual Desktops and Applications	
WORKLOAD PROFILE	Mixed workloads	Specific applications such as RTM	Deep learning frameworks such as Caffe and TensorFlow	Mixed inference workloads such as image, video, or data processing	Flexible deployments: user experience/ graphics performance vs concurrent user density	
KEY REQUIREMENTS	<ul style="list-style-type: none"> > Performance (Double- and Single-Precision) > Memory Size and Bandwidth > Interconnect Bandwidth 		<ul style="list-style-type: none"> > Performance (Single-Precision) > Memory Size Per GPU > Interconnect Bandwidth 	<ul style="list-style-type: none"> > Power Footprint > Form Factor 	<ul style="list-style-type: none"> > Virtual graphics (vGPU) > Graphics accelerated applications delivered anywhere, on any device > Server form factor: Rack and Blade 	
RECOMMENDED SOLUTION	MIXED WORKLOADS K80 		TRAINING M40 	INFERENCE M4 	RACK FORM FACTOR M60 	BLADE FORM FACTOR M6 

To learn more about the NVIDIA Tesla platform of solutions visit www.nvidia.com/tesla



Penguin Computing

Penguin Computing is one of North America's largest private suppliers of high-performance computing and enterprise scale-out solutions. Penguin pioneers complete compute, storage and networking solutions that deliver the advantages of open technologies while helping our customers successfully navigate the inherent risks in pursuing an open IT strategy. We combine 16 years of Linux experience with system innovation and software technology, services, support and operational excellence to provide our customers with industry leading and best value solutions.

www.penguincomputing.com | +1 (888) 736-4846 | sales@penguincomputing.com

