

TESLA P100 PERFORMANCE GUIDE

Oil and Gas Applications



OIL AND GAS



Geoscience simulations are key to the discovery of oil and gas and performing geological modeling. Many of the top geoscience applications are accelerated with GPUs today. When running Geoscience applications, a data center with Tesla P100 GPUs can save up to 65% in server acquisition cost.

KEY FEATURES OF THE TESLA PLATFORM AND P100 FOR GEOSCIENCE


- > Servers with P100 replace up to 50 CPU servers for applications such as RTM and SPECFEM 3D
- > Top Oil and Gas applications are GPU-accelerated
- > Up to 10.6 TFLOPS of single precision floating point performance
- > Up to 16 GB of memory capacity with up to 732 GB/s memory bandwidth

View all related applications at:

www.nvidia.com/oil-and-gas-apps

RTM Performance Equivalence

Single GPU Server vs Multiple CPU-Only Servers



CPU Server: Dual Xeon E5-2690 v4 @ 2.6 GHz, GPU Servers: Same as CPU server with NVIDIA® Tesla® P100 for PCIe (12 GB or 16 GB) | NVIDIA CUDA® Version: 8.0.44 | Dataset: TTI R8 3 pass | To arrive at CPU node equivalence, we used measured benchmarks with up to 8 CPU nodes and linear scaling beyond 8 nodes.

RTM

Reverse time migration (RTM) modeling is a critical component in the seismic processing workflow of oil and gas exploration

VERSION

2016

ACCELERATED FEATURES


Batch algorithm

SCALABILITY

Multi-GPU and Multi-Node

SPECFEM 3D Performance Equivalence

Single GPU Server vs Multiple CPU-Only Servers



CPU Server: Dual Xeon E5-2690 v4 @ 2.6 GHz, GPU Servers: Same as CPU server with NVIDIA® Tesla® P100 for PCIe (12 GB or 16 GB) | NVIDIA CUDA® Version: 8.0.44 | Dataset: Globe 112x64, 100 mins | To arrive at CPU node equivalence, we used measured benchmarks with up to 8 CPU nodes and linear scaling beyond 8 nodes.

SPECFEM 3D

Simulates Seismic wave propagation

VERSION

7.0.0

ACCELERATED FEATURES

Wilson-clover fermions, Krylov solvers, and Domain-decomposition

SCALABILITY

Multi-GPU and Multi-Node

MORE INFORMATION

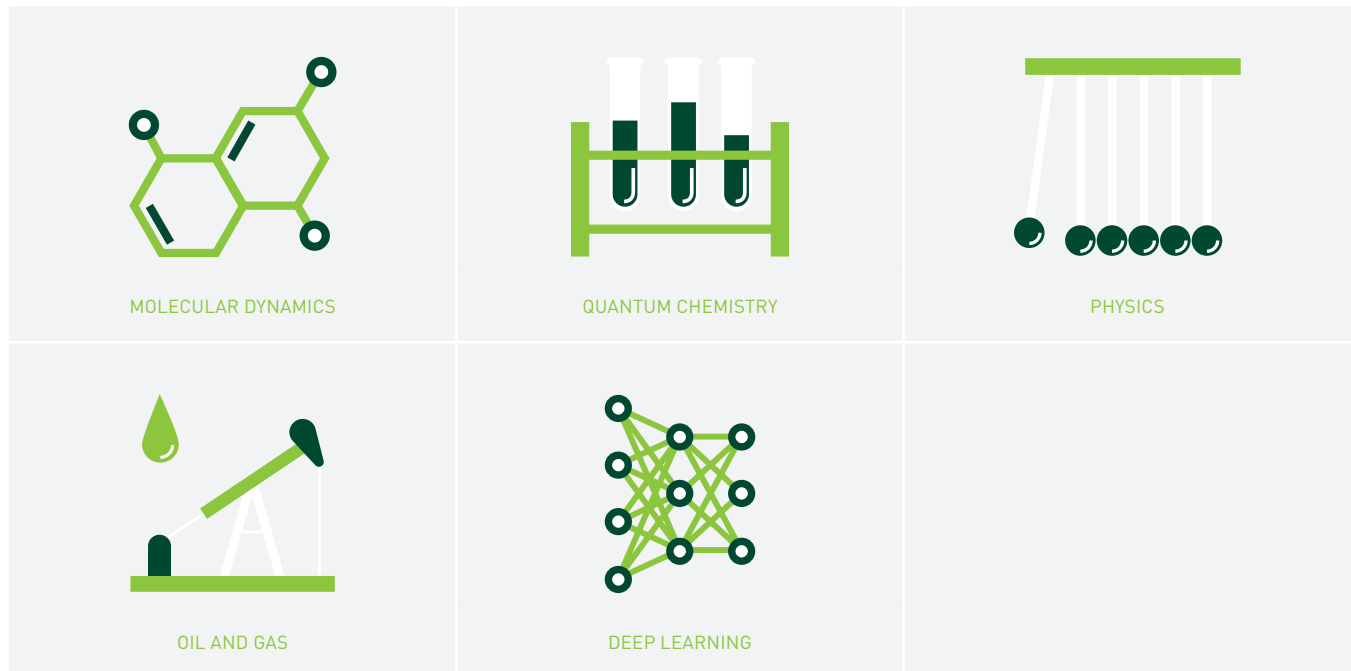
www.nvidia.com/specfem3d-globe

TESLA P100 PERFORMANCE GUIDE

Modern high performance computing (HPC) data centers are key to solving some of the world's most important scientific and engineering challenges. NVIDIA® Tesla® accelerated computing platform powers these modern data centers with the industry-leading applications to accelerate HPC and AI workloads. The Tesla P100 GPU is the engine of the modern data center, delivering breakthrough performance with fewer servers resulting in faster insights and dramatically lower costs.

Every HPC data center can benefit from the Tesla platform. Over 450 HPC applications in a broad range of domains are optimized for GPUs, including all 10 of the top 10 HPC applications and every major deep learning framework.

RESEARCH DOMAINS WITH GPU-ACCELERATED APPLICATIONS INCLUDE:



Over 450 HPC applications and all deep learning frameworks are GPU-accelerated.

- > To get the latest catalog of GPU-accelerated applications visit:
www.nvidia.com/teslaapps
- > To get up and running fast on GPUs with a simple set of instructions for a wide range of accelerated applications visit:
www.nvidia.com/gpu-ready-apps
- > To read the full Tesla P100 Performance Guide visit:
www.nvidia.com/p100-performance-guide

TESLA P100 PRODUCT SPECIFICATIONS



	NVIDIA Tesla P100 for PCIe-Based Servers	NVIDIA Tesla P100 for NVLink-Optimized Servers
Double-Precision Performance	up to 4.7 TFLOPS	up to 5.3 TFLOPS
Single-Precision Performance	up to 9.3 TFLOPS	up to 10.6 TFLOPS
Half-Precision Performance	up to 18.7 TFLOPS	up to 21.2 TFLOPS
NVIDIA NVLink™ Interconnect Bandwidth	-	160 GB/s
PCIe x 16 Interconnect Bandwidth	32 GB/s	32 GB/s
CoWoS HBM2 Stacked Memory Capacity	16 GB or 12 GB	16 GB
CoWoS HBM2 Stacked Memory Bandwidth	732 GB/s or 549 GB/s	732 GB/s

Assumptions and Disclaimers

The percentage of top applications that are GPU-accelerated is from top 50 app list in the i360 report: H PC Application Support for GPU Computing. Calculation of throughput and cost savings assumes a workload profile where applications benchmarked in the domain take equal compute cycles.