

# Unleash Greater Compute Power for the Modern Data Center

The NVIDIA Tesla P100 is the most advanced data center accelerator ever built, leveraging the groundbreaking NVIDIA Pascal™ GPU architecture to deliver the world's fastest compute node. It's powered by four innovative technologies with huge jumps in performance for HPC and deep learning workloads.

With the NVIDIA Tesla P100 your data center can keep up with ever-growing user demands and deliver lightning-fast performance in a broad range of HPC applications.


Silicon Mechanics is an open technology integrator of Servers, High-Performance Computing, Software Defined Storage, Cloud and Virtualization solutions.

For over 15 years, Silicon Mechanics has enabled organizations to deploy purpose-built compute and storage solutions with our customer focused approach. Offering customers deep technical experience, along with our defined build methodology, we partner with customers to architect, build, deploy and support flexible solutions from a large network of technology partners. Founded in 2001 and recognized as one of the fastest growing companies in the Seattle metropolitan technology corridor, Silicon Mechanics is empowering innovative organizations as they transform the world through open technology.



## NVIDIA® Tesla® P100

Infinite Compute Power For the Modern Data Center



The NVIDIA® Tesla® P100 is the world's most advanced datacenter accelerator ever built, a brand new GPU architecture to deliver the world's fastest compute node. Powered by four ground-breaking technologies with discontinuous jumps in performance, Tesla P100 enables lightning-fast nodes to deliver the highest absolute performance for HPC and deep learning workloads with infinite computing needs.

The Tesla P100 also features NVIDIA NVLink™ technology that enables superior strong-scaling performance for HPC and hyperscale applications. Up to eight Tesla P100 GPUs interconnected in a single node can deliver the performance of racks of commodity CPU servers.

### SPECIFICATIONS

GPU Architecture	NVIDIA Pascal
NVIDIA CUDA® Cores	3584
Double-Precision Performance	5.3 TeraFLOPS
Single-Precision Performance	10.6 TeraFLOPS
Half-Precision Performance	21.2 TeraFLOPS
GPU Memory	16 GB CoWoS HBM2
Memory Bandwidth	720 GB/s
Interconnect	NVIDIA NVLink
Max Power Consumption	300 W
ECC	Native support with no capacity or performance overhead
Thermal Solution	Passive
Form Factor	SXM2
Compute APIs	NVIDIA CUDA, DirectCompute, OpenCL™, OpenACC

TeraFLOPS measurements with NVIDIA GPU Boost™ technology

## NVIDIA® Tesla® P100 PCIe

World's Most Advanced Data Center Accelerator of PCIe-based Servers



The NVIDIA Tesla P100 accelerators are the world's most advanced data center GPUs ever built, designed to boost throughput and save money for HPC and hyperscale data centers. Powered by the brand new NVIDIA Pascal™ architecture, Tesla P100 for PCIe-based servers enables a single node to replace up to half-rack of commodity CPU nodes by delivering lightning-fast performance in a broad range of HPC applications.

### SPECIFICATIONS

GPU Architecture	NVIDIA Pascal
NVIDIA CUDA® Cores	3584
Double-Precision Performance	4.7 TeraFLOPS
Single-Precision Performance	9.3 TeraFLOPS
Half-Precision Performance	18.7 TeraFLOPS
GPU Memory	16GB CoWoS HBM2 at 720 GB/s or 12GB CoWoS HBM2 at 540 GB/s
System Interface	PCIe Gen3
Max Power Consumption	250 W
ECC	Yes
Thermal Solution	Passive
Form Factor	PCIe Full Height/Length
Compute APIs	CUDA, DirectCompute, OpenCL™, OpenACC

TeraFLOPS measurements with NVIDIA GPU Boost™ technology



Expert included.