



NVIDIA DGX-1 DEEP LEARNING SYSTEM

The World's First Deep Learning Supercomputer in a Box

Data scientists and artificial intelligence (AI) researchers require accuracy, simplicity, and speed for deep learning success. Faster training and iteration ultimately means faster innovation and time-to-market.


The NVIDIA® DGX-1™ is the world's first purpose-built system optimized for deep learning, with fully integrated hardware and software that can be deployed quickly and easily. Its revolutionary performance significantly accelerates training time, making it the world's first deep learning supercomputer in a box.



SYSTEM SPECIFICATIONS


GPUs	8x Tesla P100
TFLOPS (GPU FP16 / CPU FP32)	170/3
GPU Memory	16 GB per GPU
CPU	Dual 20-core Intel® Xeon® E5-2698 v4 2.2 GHz
NVIDIA CUDA® Cores	28672
System Memory	512 GB 2133 MHz DDR4
Storage	4x 1.92 TB SSD RAID 0
Network	Dual 10 GbE, 4 IB EDR
Software	Ubuntu Server Linux OS DGX-1 Recommended GPU Driver
System Weight	134 lbs
System Dimensions	866 D x 444 W x 131 H (mm)
Packing Dimensions	1180 D x 730 W x 284 H (mm)
Maximum Power Requirements	3200W
Operating Temperature Range	10 - 35 °C

NVIDIA DGX-1 Delivers 56X More Performance



CPU is dual socket Intel Xeon E5-2697 v3. 170 TF is half precision or FP16.

NVIDIA DGX-1 Delivers 75X Faster Training



Note: Caffe benchmark with AlexNet, training 1.28M images with 90 epochs | CPU server uses 2x Xeon E5-2697 v3 CPUs.

Infinite Computing for Infinite Opportunities

NVIDIA DGX-1 is the first system built with groundbreaking Pascal™-powered NVIDIA Tesla® P100 GPU accelerators, interconnected with NVIDIA NVLink™. NVIDIA Pascal architecture was designed as the engine of computers that learn, see, and simulate our world—a world with an infinite appetite for computing.

Supercharge Deep Learning Performance

The NVIDIA DGX-1 software stack includes major deep learning frameworks, the NVIDIA Deep Learning SDK, the DIGITS™ GPU training system, drivers, and CUDA®, for designing the most accurate deep neural networks (DNN). This powerful system includes access to cloud management services for container creation and deployment, system updates, and an application repository. The combination of these software capabilities running on Pascal-powered Tesla GPUs allows applications to run up to 12x faster than any previous GPU-accelerated solutions.

Iterate and Innovate Faster

High-performance training accelerates your productivity, which means faster time-to-insight, and faster time-to-market.

Deploy Quickly and Simply

The turnkey NVIDIA DGX-1 system provides plug-and-play setup and takes you from power-on to deep learning in just minutes.

Maximize Your Investment

NVIDIA DGX-1 support lets you improve productivity and reduce system downtime. Hardware and software support provides access to NVIDIA deep learning expertise and includes cloud management essentials, software upgrades and updates, and priority resolution of your critical issues.

Stay Ahead of the Competition

NVIDIA DGX-1 is engineered with groundbreaking technologies that deliver the fastest solutions for your deep learning training.

For more information on NVIDIA DGX-1, visit www.nvidia.com/dgx1

