



NVIDIA DGX STATION PERSONAL AI SUPERCOMPUTER

The Personal Supercomputer for Leading-Edge AI Development

Your data science team depends on computing performance to gain insights, and innovate faster through the power of deep learning and data analytics. Until now, AI supercomputing was confined to the data center, limiting the experimentation needed to develop and test deep neural networks prior to training at scale. Now there's a solution, offering the power to experiment with deep learning while bringing AI supercomputing performance within arm's reach.

Groundbreaking AI, at Your Desk

Now you can get the computing capacity of 400 CPU's, in a workstation that conveniently fits under your desk, drawing less than 1/20th the power. NVIDIA® DGX Station™ delivers incredible deep learning and analytics performance, designed for the office and whisper quiet with only 1/10th the noise of other workstations. Data scientists and AI researchers can instantly boost their productivity with a workstation that includes access to optimized deep learning software and runs popular analytics software.

Get Started in Deep Learning, Faster

DGX Station breaks through the limitations of building your own deep learning platform. You could spend a month or longer, procuring, integrating, and testing hardware and software. Then additional expertise and effort are needed to optimize frameworks, libraries, and drivers. That's valuable time and money spent on systems integration and software engineering that could be spent training and experimenting.

NVIDIA DGX Station is designed to kickstart your AI initiative, with a streamlined plug-in and power-up experience that can have you training deep neural networks in just one day.



SYSTEM SPECIFICATIONS

GPUs	4X Tesla V100
TFLOPS (GPU FP16)	480
GPU Memory	64 GB total system
NVIDIA Tensor Cores	2,560
NVIDIA CUDA® Cores	20,480
CPU	Intel Xeon E5-2698 v4 2.2 GHz (20-Core)
System Memory	256 GB LRDIMM DDR4
Storage	Data: 3X 1.92 TB SSD RAID 0 OS: 1X 1.92 TB SSD
Network	Dual 10 Gb LAN
Display	3X DisplayPort, 4K resolution
Acoustics	< 35 dB
System Weight	88 lbs / 40 kg
System Dimensions	518 D x 256 W x 639 H (mm)
Maximum Power Requirements	1,500 W
Operating Temperature Range	10-30 °C
Software	Ubuntu Desktop Linux OS DGX Recommended GPU Driver CUDA Toolkit

Productivity That Takes You from Desk to Data Center

Deep learning platforms require software engineering expertise to keep today's frameworks optimized for maximum performance, with time spent waiting on stable versions of open source software. This means hundreds of thousands of dollars in lost productivity, dwarfing the initial hardware cost.

NVIDIA DGX Station includes the same software stack found in all DGX solutions. This innovative, integrated system includes access to popular deep learning frameworks, updated monthly, each optimized by NVIDIA engineers for maximized performance. It also includes access to NVIDIA DIGITS™ deep learning training application, third-party accelerated solutions, the the NVIDIA Deep Learning SDK (e.g. cuDNN, cuBLAS, NCCL), CUDA® Toolkit, and NVIDIA drivers.

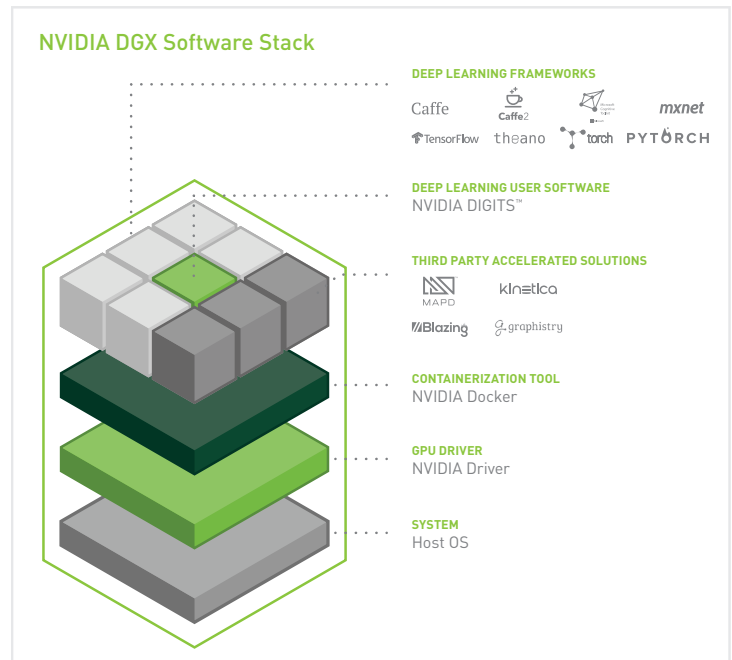
Built on container technology powered by NVIDIA Docker, this unified deep learning software stack simplifies workflow, saving you days in re-compilation time when you need to scale your work and deploy your models in the data center or cloud. The same workload running on DGX Station can be effortlessly migrated to a DGX-1 or the cloud, without modification.

Supercomputing Performance, at Your Desk


DGX Station brings the incredible performance of an AI supercomputer in a workstation form factor that takes advantage of innovative engineering and a water-cooled system that runs whisper-quiet.

The NVIDIA DGX Station packs 480 TeraFLOPS of performance, with the first and only workstation built on four NVIDIA Tesla® V100 accelerators, including innovations like next generation NVLink™ and new Tensor Core architecture. This groundbreaking solution offers:

- > 47X the performance for deep learning training, compared with CPU-based servers
- > 100X in speed-up on large data set analysis, compared with a 20 node Spark server cluster
- > 5X increase in bandwidth compared to PCIe with NVIDIA NVLink technology
- > maximized versatility with deep learning training and over 30,000 images/second inferencing



NVIDIA DGX Station Delivers 47X Faster Training



DGX Station performance projected based on DGX-1 (with Tesla V100) Workload: ResNet50, 90 epochs to solution | CPU Server: Dual Xeon E5-2699 v4, 2.6GHz. Projections subject to change.

Investment Protection

With DGX Station, you get enterprise grade support with access to NVIDIA deep learning expertise, a library of expert training, software upgrades and updates, and priority resolution of your critical issues—all in one place.



Microway, Inc.

Incorporated in 1982, Microway designs state-of-the-art, high-end Linux clusters, servers, and workstations for NVIDIA GPUs applied for HPC, HPDA, and Deep Learning. Classified as a small business, woman owned and operated, Microway's GSA Schedule is GS-35F-0431N

www.microway.com | +1 (508) 746-7341 | wespeakhpc@microway.com

For more information, visit microway.com/preconfiguredsystems/nvidia-dgx-station or nvidia.com/dgx-station

© 2017 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, and Iray are trademarks and/or registered trademarks of NVIDIA Corporation. All company and product names are trademarks or registered trademarks of the respective owners with which they are associated. Features, pricing, availability, and specifications are all subject to change without notice. JUN17

