



NVIDIA DGX STATION MAXIMIZE YOUR DATA SCIENCE TEAM PRODUCTIVITY

Your data science team depends on computing performance to gain insights, and innovate faster through the power of AI and deep learning. 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. Designed for your data science team, NVIDIA® DGX Station™ is the world's fastest workstation for leading-edge AI development. This fully integrated and optimized system enables your team to get started faster, and effortlessly experiment with the power of a data center in your office.

World Class Computing Performance in the Hands of Your Team

Your real work is innovation and discovery. DGX Station is the only workstation with four NVIDIA® Tesla® V100 Tensor Core GPUs, integrated with a fully-connected four-way NVLink™ architecture. With 500 TFLOPS of supercomputing performance, your entire data science team can experience over 2X the training performance of today's fastest workstations. This groundbreaking solution offers:

- > 72X 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

Experiment faster, iterate more frequently, with effortless productivity.



SYSTEM SPECIFICATIONS

GPUs	4X Tesla V100
TFLOPS (Mixed precision)	500
GPU Memory	128 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 RDIMM DDR4
Storage	Data: 3X 1.92 TB SSD RAID 0 OS: 1X 1.92 TB SSD
Network	Dual 10GBASE-T (RJ45)
Display	3X DisplayPort, 4K resolution
Additional Ports	2x eSATA, 2x USB 3.1, 4x USB 3.0
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

Get the Fastest Start in Data Science and AI Research

Spend less time and money on configuration, and more time on data science. DGX Station™ can save you hundreds of thousands of dollars in engineering hours and lost productivity waiting for stable versions of open source code. Powered by the NVIDIA DGX software stack, DGX Station allows you to start innovating within one hour.

This integrated hardware and software solution allows your data science team to easily access a comprehensive catalog of NVIDIA optimized GPU-accelerated containers that offer the fastest possible performance for AI and data science workloads. It also includes access to NVIDIA DIGITS™ deep learning frameworks, HPC containers, third-party accelerated solutions, the NVIDIA Deep Learning SDK (e.g. cuDNN, cuBLAS, NCCL), CUDA® toolkit, RAPIDS open source libraries, and NVIDIA drivers. Built on container technology and powered by NVIDIA Container Runtime for 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 an NVIDIA DGX-1™, NVIDIA DGX-2™, or the cloud, without modification.

With GPU-aware Kubernetes from NVIDIA, your data science team can benefit from industry-leading orchestration tools to better schedule AI resources and workloads. Data scientists can run compute workloads 24/7 by scheduling and queuing jobs, running multiple simultaneous jobs, and easily monitoring GPU health. Eliminate any idle usage of GPUs, drive down the cost per training run, and maximize the productivity and return on investment for your data science team. Enjoy productive experimentation and spend more time focused on insight.

No Data Center? No Problem.


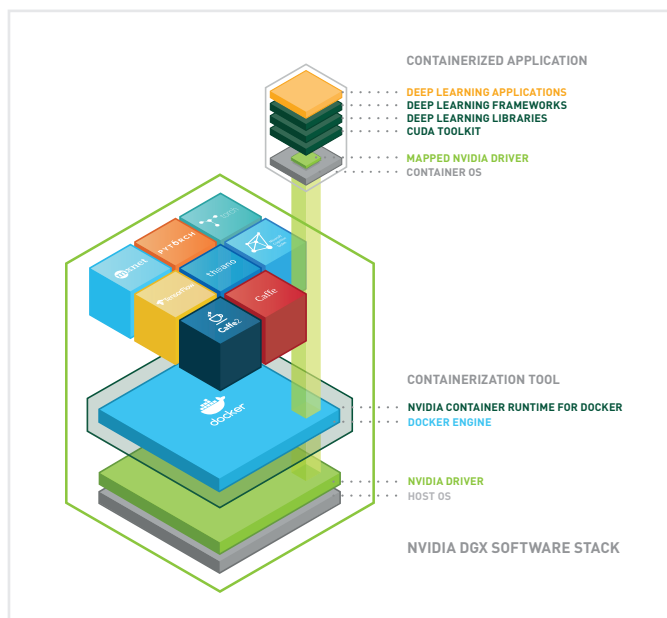
Need to deliver AI breakthroughs but lack the data center to develop them? With NVIDIA DGX Station, you do not need to worry about the challenges and complexity of building an enterprise-scale infrastructure to access compute capacity.

NVIDIA DGX Station has the computing capacity of racks of servers in an office-friendly package. DGX Station is the fastest deep learning system for your data science and AI research teams. Designed for the office and whisper-quiet, you can access water-cooled supercomputing performance at your fingertips. You can instantly boost your productivity with a workstation that includes access to NVIDIA optimized AI and HPC software.

Innovate on your terms instead of waiting for compute cycles in the data center.

Access to AI Expertise and Enterprise Grade Support

Today's AI relies on open source and sometimes unstable code that can disrupt the productivity of your data science team. DGX Station with NVIDIA enterprise support provides your workgroup with rapid response, timely issue resolution, and world-class AI expertise, ensuring your data science team can get the most out of your AI investment. Now your developers can focus on innovation, instead of wasting time troubleshooting software and hardware stacks.



For more information, visit www.nvidia.com/dgx-station

© 2019 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, DGX Station, Tesla, NVLink, DIGITS, CUDA, DGX-1, and DGX-2 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. JAN19

