

NVIDIA Introduces New Breed of High-Performance Workstations for Millions of Data Scientists

Global System Builders Led by HP, Dell, Lenovo Unveil Workstations Integrating NVIDIA Quadro RTX GPUs and NVIDIA CUDA-X AI for Massive Productivity Boost

GPU Technology Conference -- NVIDIA has teamed with the world's leading OEMs and system builders to deliver [powerful new workstations](#) designed to help millions of data scientists, analysts and engineers make better business predictions faster and become more productive.

Purpose-built for data analytics, machine learning and deep learning, the systems provide the extreme computational power and tools required to prepare, process and analyze the massive amounts of data used in fields such as finance, insurance, retail and professional services.

NVIDIA-powered workstations for data science are based on a powerful reference architecture made up of dual, high-end [NVIDIA Quadro RTX™ GPUs](#) and [NVIDIA CUDA-X AI™](#) accelerated data science software, such as [RAPIDS™](#), TensorFlow, PyTorch and Caffe. CUDA-X AI is a collection of libraries that enable modern computing applications to benefit from NVIDIA's GPU-accelerated computing platform.

"Data science is one of the fastest growing fields of computer science and impacts every industry. Enterprises are eager to unlock the value of their business data using machine learning and are hiring at an unprecedented rate data scientists who require powerful workstations architected specifically for their needs," said Jensen Huang, founder and CEO of NVIDIA. "With our partners, we are introducing NVIDIA-powered data science workstations -- made possible by our new Turing Tensor Core GPUs and CUDA-X AI acceleration libraries -- that allow data scientists to develop predictive models that can revolutionize their business."

NVIDIA GPU-Accelerated Data Science Workstation

Data science problems involve data on a massive scale and require large-scale processing capabilities. NVIDIA-powered data science workstations make it easy for scientists to wrangle, prep, train and deploy models quickly and accurately. Features and benefits include:

- Dual, high-end Quadro RTX GPUs -- Powered by the latest NVIDIA Turing™ GPU architecture and designed for enterprise deployment, dual [Quadro RTX™ 8000 and 6000 GPUs](#) deliver up to 260 teraflops of compute performance and 96GB of memory using [NVIDIA NVLink® interconnect technology](#). Quadro RTX-powered data science workstations provide the capacity and bandwidth to handle the largest datasets and compute-intensive workloads as well as the graphics power required for 3D visualization of massive datasets, including VR.
- Data science software stack -- built on the Linux operating system and Docker containers:
 - NVIDIA CUDA-X AI -- A collection of NVIDIA's GPU acceleration libraries to accelerate deep learning, machine learning and data analysis. CUDA-X AI includes cuDNN for accelerating deep learning primitives, cuML for accelerating machine learning algorithms, TensorRT™ for optimizing trained models for inference and over 15 other libraries. Together they work seamlessly with NVIDIA Tensor Core GPUs to accelerate the end-to-end workflows for developing and deploying AI-based applications. CUDA-X AI can be integrated into deep learning frameworks, including TensorFlow, PyTorch and MXNet, and leading cloud platforms, including AWS, Microsoft Azure and Google Cloud.
 - [NVIDIA RAPIDS](#) -- A set of GPU-accelerated libraries analytics for data preparation, traditional machine learning and graph analytics.
 - Anaconda™ Distribution -- With Anaconda, Inc., NVIDIA is providing Anaconda Distribution, an innovative way for data scientists to perform Python/R, data science, AI and machine learning.
- Enterprise ready -- Tested and optimized in conjunction with workstation manufacturers to meet the needs of mission-critical enterprise deployments.
- Optional software support -- Offers peace of mind with NVIDIA-developed software and containers, including deep learning and machine learning frameworks.

By freeing data scientists to work locally, NVIDIA-powered data science workstations are the ideal complement to NVIDIA's data science portfolio.

"The NVIDIA-powered data science workstation enables our data scientists to run end-to-end data processing pipelines on large datasets faster than ever," said Mike Koelemay, chief data scientist at Lockheed Martin Rotary & Mission Systems. "Leveraging RAPIDS to push more of the data processing pipeline to the GPU reduces model development time, which leads to faster deployment and business insights."

Broad Ecosystem Support and Adoption

NVIDIA-powered Data Science Workstations help OEMs and leading data science software providers meet the growing demand for GPU-accelerated data science capabilities and offer powerful new options to customers conducting AI-based exploration.

Read what partners and customers, such as BlazingDB, BOXX, Charter Communications, Datalogue, Dell, Graphistry, H2O.ai, HP, Kinetica, Lenovo, MapR, MIT and OmniSci, among others, [are saying](#).

Availability

NVIDIA-powered systems for data scientists are available immediately from global workstation providers such as Dell, HP and Lenovo and regional system builders, including AMAX, APY, Azken Muga, BOXX, CADNetwork, Carri, Colfax, Delta, EXXACT, Microway, Scan, Sysgen and Thinkmate.

About NVIDIA

[NVIDIA](#)'s (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined modern computer graphics and revolutionized parallel computing. More recently, GPU deep learning ignited modern AI — the next era of computing — with the GPU acting as the brain of computers, robots and self-driving cars that can perceive and understand the world. More information at <http://nvidianews.nvidia.com/>.

Certain statements in this press release including, but not limited to, statements as to: global systems builders and OEMs teaming with NVIDIA and integrating NVIDIA Quadro RTX GPUs and NVIDIA CUDA-X AI; the benefits, impact, performance, abilities and features of NVIDIA GPU-accelerated data science workstations, NVIDIA Quadro RTX GPUs and NVIDIA CUDA-X AI; CUDA-X AI enabling modern computing applications to benefit from NVIDIA's GPU-accelerated computing platform; enterprises being eager to unlock the benefits of machine learning and analytics and hiring data scientists at an unprecedented rate; NVIDIA GPU-accelerated data science workstations providing enterprises with the processing power to build their own machine learning

models and delivering a major boost in AI acceleration enabling data scientists to interpret and manage data, solve complex problems and provide valuable, actionable insights; NVIDIA-powered data science workstations making it easy for scientists to use models quickly and accurately; NVIDIA-powered data science workstations being the ideal complement to NVIDIA's data science portfolio; NVIDIA-powered data science workstation enabling data scientists to run end-to-end data processing pipelines on large datasets faster than ever and leveraging RAPIDS, leading to faster deployment and business insights; NVIDIA-powered data science workstations helping OEMs and data science software providers bring powerful new options to customers; and the availability of NVIDIA-powered systems for data scientists are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

© 2019 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA-X, NVIDIA Turing, NVLink, Quadro RTX, RAPIDS and TensorRT are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability and specifications are subject to change without notice.

Media Contacts

Gail Laguna

+1 408 386 2435

glaguna@nvidia.com