

## Global Computer Companies Announce NVIDIA-Powered Enterprise Servers Optimized for Data Science

### Cisco, Dell EMC, Fujitsu, Hewlett Packard Enterprise, Inspur, Lenovo and Sugon Offer NVIDIA T4 Servers Architected to Run CUDA-X AI Accelerated Data Analytics, Machine Learning and Deep Learning

SAN JOSE, Calif., March 18, 2019 GPU Technology Conference -- NVIDIA today announced that mainstream servers optimized to run NVIDIA's data science acceleration software are now available from seven of the world's largest systems manufacturers, including Cisco, Dell EMC, Fujitsu, Hewlett Packard Enterprise (HPE), Inspur, Lenovo and Sugon.

Featuring [NVIDIA® T4 GPUs](#) and fine-tuned to run [NVIDIA CUDA-X AI™](#) acceleration libraries, the servers provide businesses a standard, highly efficient platform for data analytics and a wide range of other enterprise workloads.

Drawing only 70 watts of power and designed to fit into existing data center infrastructures, the T4 GPU can accelerate AI training and inference, machine learning, data analytics and virtual desktops. This unique combination of capabilities has helped create a new class of enterprise servers that, through GPU acceleration, can provide businesses with greater utility and versatility.

"The rapid adoption of T4 on the world's most popular business servers signals the start of a new modern era in enterprise computing -- one in which GPU acceleration has become standard," said Ian Buck, vice president and general manager of Accelerated Computing at NVIDIA. "Now, with a wave of mainstream NVIDIA-powered servers optimized for data science, companies worldwide can deploy accelerated AI at a faster pace across their entire business."

#### New T4 Servers Validated as NGC-Ready

Systems announced today by Cisco, Dell EMC, Fujitsu, HPE, Inspur, Lenovo and Sugon are NVIDIA NGC-Ready validated, a designation reserved for servers with demonstrated ability to excel in a full range of accelerated workloads.

All software tested as part of the NGC-Ready validation process is available from [NVIDIA NGC™](#), a comprehensive repository of GPU-accelerated software, pre-trained AI models, model training for data analytics, machine learning, deep learning and high performance computing accelerated by CUDA-X AI.

Launched in November, the NGC-Ready program features a select set of systems powered by NVIDIA GPUs with Tensor Cores that are ideal for a wide range of AI workloads.

Key NGC-Ready validated T4 servers announced today are:

- [Cisco UCS C240 M5](#)
- [Dell EMC PowerEdge R740/R740xd](#)
- [Fujitsu PRIMERGY RX2540 M5](#)
- [HPE ProLiant DL380 Gen10](#)
- [Inspur NF5280M5](#)
- [Lenovo ThinkSystem SR670](#)
- [Sugon W760-G30](#)

Additionally, several other partners have started the validation process for their T4 servers.

#### New Enterprise Support Program for NGC-Ready Systems

Furthering the value of this program, NVIDIA today launched an enterprise support service exclusively for customers with [NGC-Ready systems](#), including all NGC-Ready T4 systems as well as [previously validated](#) NVLink® and Tesla® V100-based servers and NVIDIA-powered workstations.

[NVIDIA NGC Support Services](#) give customers direct access to NVIDIA technical experts to help ensure their NGC-Ready systems run optimally and maximize system utilization and user productivity.

NVIDIA's support service is available through sellers of NGC-Ready systems, with immediate availability from Cisco for its NGC-Ready validated NVIDIA V100 system, [Cisco UCS C480 ML](#). HPE will offer the service for the [HPE ProLiant DL380 Gen10 server](#) as a validated NGC-Ready NVIDIA T4 server in June. Several other original equipment manufacturers are expected to begin selling the service on their NVIDIA T4 and V100 systems in the second quarter.

#### New T4 Servers Certified with Leading Virtual Desktop and Application Virtualization Solutions

NVIDIA T4-equipped servers have been [certified by leading OEMs](#) for NVIDIA virtual GPU software -- NVIDIA GRID® Virtual PC (vPC) for knowledge workers and NVIDIA Quadro Virtual Data Center Workstation (vDWS) for creative and technical professionals.

With the T4, users will enjoy high-quality virtual desktop experiences on the latest generation of servers, including up to 33 percent more VDI performance than CPU-only with GRID vPC. NVIDIA also announced it has [collaborated with leading OEMs](#) to make VDI easier and more affordable with multi-year discounts available for a limited time only.

#### Broad Industry Support

##### NGC-Ready systems partners

"Cisco is the first OEM to resell NVIDIA NGC support to its customers so they can accelerate their deep learning projects on Cisco's NGC-ready UCS C480 ML M5 server, with eight NVIDIA Tesla V100 Tensor Core GPUs interconnected with NVLink. We are also excited to announce support for the new NVIDIA T4 universal GPU on our UCS C240 M5 and C220 M5 servers, providing our users the flexibility to maximize their infrastructure use for AI/ML and virtualization projects on the same hardware."

-- Siva Sivakumar, senior director of Data Center Solutions, Cisco

"Dell Technologies is focused on helping customers transform their IT while benefiting from advancements such as artificial intelligence. As the world's leading provider of server systems, Dell EMC continues to enhance the PowerEdge server portfolio to help our customers ultimately achieve their goals. Our close collaboration with NVIDIA and historical adoption of the latest GPU accelerators, including the new NVIDIA T4, play a vital role in helping our customers stay ahead of the curve in AI training and inference."

-- Ravi Pendekanti, senior vice president of product management and marketing, Servers & Infrastructure Systems, Dell EMC

"Powered by NVIDIA T4 GPUs and officially validated NGC-Ready, Fujitsu Server PRIMERGY is a cutting-edge server with practicability and versatility. We believe it can satisfy various kinds of workloads desired by our customers today, such as efficient operation of virtual desktop systems, high-speed analysis of big data and deployment of AI."

-- Kenichi Sakai, corporate executive officer, senior vice president, head of Data Center Platform business unit, Fujitsu Ltd

"We continue to deliver advanced performance, resiliency and security across our server portfolio, enabling new experiences and increased business outcomes. By combining HPE ProLiant Gen10 servers with the new NVIDIA T4 GPUs, we are broadening accelerated computing capabilities to the mainstream. We are empowering customers to run a wide range of workloads like machine and deep learning, inferencing, high performance computing and VDI environments, from the edge to cloud."

-- Bill Mannel, vice president and general manager of HPC and AI, Hewlett Packard Enterprise

"Inspur works closely with NVIDIA to keep bringing the most powerful AI computing products to our clients in order to enable the continuous breakthroughs in their business. The new NVIDIA T4 NGC-ready GPU feature server is fine-tuned to run the NVIDIA CUDA-X AI acceleration libraries, providing a comprehensive solution and service support for data scientists, supporting multiple AI workloads while enjoying a high-quality virtual desktop experience."

-- Jun Liu, general manager of AI & HPC, Inspur

"As more enterprises begin to institute artificial intelligence into their business operations, the need for dense servers equipped with versatile GPU accelerators continues to grow. Together with NVIDIA, Lenovo is providing the infrastructure necessary for customers to successfully implement AI initiatives that lead to data-driven insights and competitive advantage. The Lenovo ThinkSystem SR670 supports up to eight NVIDIA T4 GPUs in 2U, and jointly with Lenovo Intelligent Computing Orchestration (LiCO), provides leading economies of scale for customers on their AI journey."

-- Madhu Matta, vice president and general manager of HPC/AI, Lenovo Data Center Group

"Thanks to the excellent performance and universal capabilities of NVIDIA T4, Sugon's latest HPC data center solutions are able to provide AI training/Inference, rendering and 3D design services simply through large-scale deployment of one single kind of GPU servers. One single kind of GPU server, which greatly reduces our deployment and maintenance costs. For this reason, we specifically adjusted our W760 server to give NVIDIA T4 the best performance platforms."

-- Chaoqun Sha, senior vice president, Sugon

#### Ecosystem partners

"NVIDIA and Kinetica have enabled us to do the impossible -- render a high-fidelity 3D view of an oil basin using 100 billion data points at scale. The Kinetica Active Analytics Platform and NVIDIA T4-powered servers bring together active analytics and artificial intelligence to deliver a radically more efficient way to develop oil fields by utilizing subsurface models generated and rendered at high resolution. The combined solution provides revolutionary performance at scale, accelerating the output of our data scientists and geo scientists to run GPU-accelerated models that make spatial and economic predictions faster for the most capital-efficient recovery of resources out of the basin."

-- Sanjay Paranj, CTO at Anadarko Petroleum Corporation

"Active analytics is the backbone of the data-driven Fourth Industrial Revolution. With new NVIDIA T4-powered servers, our customers have a radically more efficient way to GPU-accelerate data workloads to keep pace with extreme velocity from an infinite set of streaming data sources, from wearables to cars, smart devices to connected infrastructure. In short, NVIDIA and Kinetica help companies, cities and societies manage the make-or-break shift from using data as a passive asset to glean insights into using data as an active asset that can help them react immediately."

-- Paul Appleby, CEO at Kinetica

"With the age of AI upon us, mainstream data centers are looking to adopt accelerators such as NVIDIA T4 GPUs to maximize the impact. Having a stable control plane for these new technologies is crucial, a need that Red Hat is helping to address by certifying server configurations that include NVIDIA T4 GPUs on Red Hat Enterprise Linux and Red Hat OpenShift Container Platform. This brings the scale of our OEM partner ecosystem and certification process to artificial intelligence use cases across the globe."

-- Stefanie Chiras, vice president and general manager of Red Hat Enterprise Linux, Red Hat

"NVIDIA and VMware have a long history of working together to empower organizations to maximize their data center resources. Businesses are looking to accelerate both VDI and machine and deep learning workloads, while maximizing their return on investment in GPU-accelerated data center infrastructure. NVIDIA T4-powered servers offer a solution for these organizations, along with increased flexibility, better utilization and protection of their VMware infrastructure investments."

-- Susan Nash, senior vice president of Strategic Corporate Alliances, VMware

#### About NVIDIA

[NVIDIA's](http://nvidianews.nvidia.com/) (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 computer companies using NVIDIA powered servers and the availability of these servers from the world's largest systems manufacturers; the benefits, performance, impact and abilities of NVIDIA-powered enterprise

servers, T4 GPUs, the NGC-Ready program and NVIDIA NGC Support Services; T4 GPU servers being in a new class and providing greater utility and versatility; the rapid adoption of T4 in servers signaling the start of a new era in enterprise computing; access to data center platforms enabling businesses to solve their most pressing challenges and to create opportunities for customers; the companies using NVIDIA-powered servers and those that have completed and are starting the NGC-Ready validation process; NVIDIA NGC Support Services giving customers access to experts and support to run their systems; the timing and expected NCG-Ready systems with support this year; T4 providing users with enjoyable experiences and performance; NVIDIA making VDI easier and more affordable; Cisco servers with NVIDIA GPUs providing users with flexibility to maximize their infrastructure for AI/ML and virtualization projects; Dell EMC enhancing its server portfolio to help customers; NVIDIA customers helping to play a vital role in helping customers stay ahead of the curve; PRIMERGY satisfying various kinds of workloads desired by customers; HPE broadening accelerated computing capabilities to the mainstream and empowering customers to run a wide range of workloads; NVIDIA and Inspur bringing the most powerful AI computing products to clients and enabling the continuous breakthroughs in their business; the need for dense servers equipped with versatile GPU accelerators continuing to grow; NVIDIA and Lenovo providing the infrastructure for customers to implement AI initiatives and economies of scale; mainstream data centers are looking to adopt accelerators and its impact; and addressing the need for a stable control plans and its impact 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 GRID, NVIDIA NGC, NVLink, Quadro, and Tesla 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**

Kristin Bryson

+1 203 241 9190

[kbryson@nvidia.com](mailto:kbryson@nvidia.com)