NVIDIA Announces Instant AI Infrastructure for Enterprises

Built with Enterprise IT Partners and Deployed First at Equinix, NVIDIA AI LaunchPad Includes End-to-End NVIDIA Hardware and Software Stack to Accelerate AI from Hybrid Cloud to Edge

NVIDIA today unveiled NVIDIA AI LaunchPad, a comprehensive program delivered through hybrid-cloud providers that gives enterprises immediate access to NVIDIA-powered infrastructure and software to streamline the entire AI lifecycle.

Digital infrastructure leader Equinix will be the first in the AI LaunchPad program, providing NVIDIA-powered hybrid-cloud solutions globally on Platform Equinix®. The Equinix infrastructure deploys in minutes, allowing enterprises to access an entire spectrum of NVIDIA resources that support virtually every aspect of AI, from data center training and inference to full-scale deployment at the edge.

With AI LaunchPad, Equinix customers can develop advanced AI on NVIDIA DGX™ SuperPODs managed by NVIDIA Base Command™ Platform, an AI development hub. Enterprises will also be able to use Base Command Platform on NVIDIA-Certified Systems™ from leading manufacturers including Dell Technologies and Lenovo. Built on the NVIDIA EGX™ platform, these industry-standard servers are ideal for running the NVIDIA AI Enterprise software suite on VMware vSphere to deploy an AI-ready enterprise platform that scales AI on hybrid clouds.

Additionally, customers can securely deploy and manage AI applications across distributed edge infrastructure at Equinix with the NVIDIA Fleet Command™ managed cloud service — which entered general availability today — connected to NVIDIA-Certified Systems.

“Today’s enterprises are looking for a simple, comprehensive solution that provides instant access to the resources they need to build and deploy AI with ease,” said Manuvir Das, head of Enterprise Computing at NVIDIA. “NVIDIA AI LaunchPad puts AI at the fingertips of enterprises everywhere with fully automated, hybrid-cloud infrastructure and software for every stage of the AI lifecycle.”

Global Ecosystem for AI in Hybrid Clouds
Increasingly, in order to build accurate AI models, applications need to access multiple external data sources that are spread across public clouds and private enterprise data centers. By using AI LaunchPad at an interconnected hub like Equinix, enterprises can have secure, high-speed access via Equinix Fabric™ to an ecosystem of companies that operate their digital infrastructures at Equinix. The global reach of Platform Equinix, which includes more than 220 data centers strategically located in more than 60 metros spanning across five continents, also helps to satisfy real-time performance requirements for digital businesses.

“Many industries use private clouds to keep costs down by having computing resources close to the data, for performance, data privacy, ownership and sovereignty reasons,” said Steve Steinhilber, vice president of Business Development at Equinix. “With NVIDIA AI LaunchPad globally available at Equinix, enterprises will have immediate access to NVIDIA software and NVIDIA-Certified infrastructure in a comprehensive hybrid-cloud solution.”

Turnkey Hybrid Cloud for Industries
Among the key industries served by Equinix is the automotive sector. Continental, one of the world’s largest automotive technology companies, uses NVIDIA DGX™ systems hosted at Equinix data centers to develop autonomous vehicle technologies for manufacturers.

“Developing advanced AI technologies requires substantial accelerated computing infrastructure, and having NVIDIA systems available globally in Equinix data centers ensures that our teams always have access to the resources they need,” said Ismail Dagli, senior vice president of Research and Development in the Advanced Driver Assistance Systems business unit at Continental. “Adding NVIDIA-powered AI inference and edge options to Equinix data centers is significant, as it enables AI to move from development into production.”

Accelerated AI on Mainstream Servers
Key to AI LaunchPad is NVIDIA AI Enterprise, which is an end-to-end, cloud-native suite of AI and data analytics software that is optimized, supported and exclusively certified by NVIDIA to run on VMware vSphere with NVIDIA-Certified Systems.

With AI Enterprise and VMware vSphere running on NVIDIA-Certified Systems, enterprises can easily deploy AI in Equinix data centers with near bare-metal performance across multiple nodes to power large, complex training and machine learning workloads.

“Enterprises are seeking solutions that enable them to easily integrate AI into their existing infrastructure,” said Krish Prasad, senior vice president and general manager of the Cloud Platform business unit at VMware. “VMware vSphere is critical to mainstreaming AI in the enterprise, and now, NVIDIA AI LaunchPad will deliver our AI-ready enterprise platform as a turnkey solution to rapidly deploy, manage and scale AI workloads.”

Enterprise Security and Speed with NVIDIA GPUs and DPUs
AI LaunchPad will leverage systems powered by NVIDIA Ampere architecture GPUs and NVIDIA BlueField® DPUs. This combination provides enterprises the ability to accelerate AI workloads and also take advantage of the security, isolation and performance enhancements provided by DPUs.

Availability
The first AI LaunchPad offerings at Equinix are expected this summer. Additional details are available on the NVIDIA website.

About NVIDIA
NVIDIA’s (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market and has redefined modern computer graphics, high performance computing and artificial intelligence. The company’s pioneering work in accelerated computing and AI is reshaping trillion-dollar industries, such as transportation, healthcare and manufacturing, and fueling the growth of many others. More information at https://nvidianews.nvidia.com/.

Certain statements in this press release including, but not limited to, statements as to: the benefits, impact, performance, features, and availability of our products and services; the AI LaunchPad program deployed at Equinix; enterprises using Base Command Platform to manage and deploy mainstream AI workloads on NVIDIA-Certified Systems from leading manufacturers; the global reach of Platform Equinix; Continental using NVIDIA DGX systems hosted at Equinix data centers to develop autonomous vehicle technologies for manufacturers; AI LaunchPad leveraging systems powered by NVIDIA Ampere architecture GPUs and NVIDIA BlueField DPUs; and the timing of the AI LaunchPad program at Equinix 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.

Shannon McPhee
+1-310-920-9642
smcphee@nvidia.com