



NVIDIA Unveils AI Enterprise Software Suite to Help Every Industry Unlock the Power of AI

Runs on VMware vSphere; Optimized, Certified and Supported by NVIDIA; Hundreds of Thousands of Customers in World's Largest Industries Can Now Adopt NVIDIA AI Enterprise at Scale

NVIDIA today announced [NVIDIA® AI Enterprise](#), a comprehensive software suite of enterprise-grade AI tools and frameworks optimized, certified and supported by NVIDIA, exclusively with [VMware vSphere 7 Update 2](#), separately announced today.

Through a first-of-its-kind industry collaboration to develop an [AI-Ready Enterprise platform](#), NVIDIA teamed with VMware to virtualize AI workloads on VMware vSphere with NVIDIA AI Enterprise. The offering gives enterprises the software required to develop a broad range of AI solutions, such as advanced diagnostics in healthcare, smart factories for manufacturing, and fraud detection in financial services.

With the NVIDIA AI Enterprise software suite, IT professionals at the hundreds of thousands of enterprises that use vSphere for compute virtualization can now support AI with the same tools they use to manage large-scale data centers and hybrid cloud environments. The NVIDIA software suite provides scale-out, multi-node, AI application performance on vSphere that is indistinguishable from bare-metal servers.

“Until now, the world has run AI on bare-metal servers,” said Justin Boitano, vice president and general manager of Enterprise and Edge Computing at NVIDIA. “NVIDIA AI Enterprise enables customers to reduce AI model development time from 80 weeks to just eight weeks, and allows them to deploy and manage advanced AI applications on VMware vSphere with the same scale-out, [record-breaking NVIDIA accelerated computing performance](#) that’s possible on bare metal.”

“Every enterprise is exploring how to modernize their infrastructure to meet the demands of AI applications,” said Lee Caswell, vice president of marketing for the Cloud Platform Business Unit at VMware. “With NVIDIA AI Enterprise and vSphere 7 Update 2, VMware customers are now able to fast-track AI in their virtualized data centers and easily deploy certified, AI-ready infrastructure for their modern apps.”

Flexible AI Infrastructure for the Hybrid Cloud

NVIDIA AI Enterprise provides compatibility for a broad set of accelerated CUDA® applications, AI frameworks, pre-trained models and software development kits running in the hybrid cloud. Optimizations enable workloads to scale across multiple nodes to support large deep learning training models with full GPU virtualization.

To power the applications in NVIDIA AI Enterprise, VMware vSphere 7 Update 2 is now certified for [NVIDIA A100 Tensor Core GPUs](#) on [NVIDIA-Certified Systems™](#), which include high-volume servers from Dell Technologies, HPE, Lenovo and Supermicro. With this certification comes direct customer support from NVIDIA for vSphere customers that acquire a license for NVIDIA AI Enterprise.

NVIDIA has also certified vSphere as the only compute virtualization software to provide hypervisor support for live migration with [NVIDIA Multi-Instance GPU](#) technology, which allows each A100 GPU to be partitioned into up to seven instances at the hardware level to maximize efficiency for workloads of all sizes.

Additionally, select NVIDIA ConnectX adapters are now certified for VMware vSAN over RDMA (Remote Direct Memory Access), which offloads CPU communication tasks to boost application performance and improve infrastructure returns on investment.

NVIDIA AI Enterprise is available as a perpetual license at \$3,595 per CPU socket. Enterprise Business Standard Support for NVIDIA AI Enterprise is \$899 annually per license. Customers can [apply for early access](#) to NVIDIA AI Enterprise as they plan their upgrades to VMware vSphere 7 Update 2.

For more information on bringing AI to VMware hybrid clouds, read the [NVIDIA blog](#) and the [VMware vSphere 7 Update 2 release blog](#). Free registration is now open for [GTC 2021](#), where NVIDIA and VMware will present on bringing AI to every enterprise.

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: NVIDIA helping industries unlock the power of AI; the benefits, performance, features and abilities of NVIDIA AI Enterprise and VMware vSAN and what they enable; NVIDIA AI Enterprise letting customers adopt AI at scale; NVIDIA AI Enterprise enabling enterprises to develop AI solutions and support AI with the same tools used to manage large-scale data centers; NVIDIA AI Enterprise enables customers to reduce AI model development time and to deploy and manage AI applications with the same levels of scale-out; enterprises exploring how to modernize their infrastructures; VMware customers being able to fast-track AI in their data centers and deploy AI ready infrastructure; the customers using NVIDIA AI Enterprise; AI playing an important role in data-driven services; NVIDIA AI Enterprise's ability to scale across nodes and its effects; VMware vSphere 7's certifications and its impacts; and the price and availability of NVIDIA AI Enterprise and its support 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.

© 2021 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA and NVIDIA-Certified Systems are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. VMware vSphere and VMware vSAN are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and other jurisdictions. 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.

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