

AI-Ready Enterprise Platform

Unleashing AI for every enterprise with an end-to-end platform from NVIDIA and VMware

KEY CHALLENGES FOR ENTERPRISE AI WORKLOADS

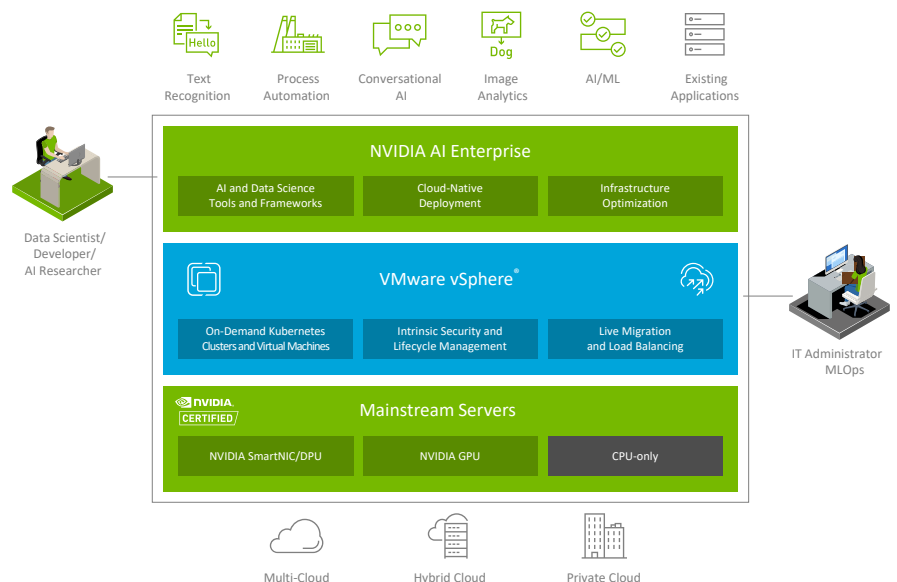
- **Risk** – It’s difficult to pull together an end-to-end AI solution from disparate products and integrate them with existing infrastructure.
- **Performance** – High performance is critical for AI, machine learning, and data analytics workloads. Fast time to deployment is also critical.
- **Scaling** – Going from proof of concept to enterprise deployments requires effective scaling—through efficient use of resources to ensure manageability and availability—and infrastructure cost management.

Artificial intelligence (AI) is transforming every industry, whether it’s by improving customer relationships in financial services, streamlining manufacturer supply chains, or helping doctors deliver better outcomes for patients. While most organizations know they need to invest in AI to secure their future, they struggle with finding the strategy and platform that can enable success.

Unlike traditional enterprise applications, AI apps are a relatively recent development for many IT departments. They’re anchored in rapidly evolving, open-source, bleeding-edge code and lack proven approaches that meet the rigors of scaled production settings in enterprises. In fact, according to Gartner, “On the average 54% of AI projects make it from pilot to production”¹.

AI-Ready Platform from NVIDIA and VMware

VMware and NVIDIA have partnered together to unlock the power of AI for every business by delivering an end-to-end enterprise platform optimized for AI workloads. This integrated platform delivers best-in-class AI software, the NVIDIA AI Enterprise Suite, optimized and certified for the enterprise workload platform, VMware vSphere®. Running on NVIDIA-Certified Systems™, industry-leading accelerated servers, this platform accelerates the speed at which developers can build AI and high-performance data analytics, enables organizations to scale modern workloads on the same vSphere infrastructure they’ve already invested in, and delivers enterprise-class manageability, security and availability.



BENEFITS FOR ENTERPRISES

- **Deploy with confidence** – Leverage the jointly certified NVIDIA and VMware solution to deploy and manage AI workloads with optimized software.
- **Boost AI performance** – Achieve near bare-metal performance for AI training and inference.
- **Scale without compromise** – Enable large, more complex AI training and data analytics with new scale-out capabilities.

BENEFITS FOR ENTERPRISES

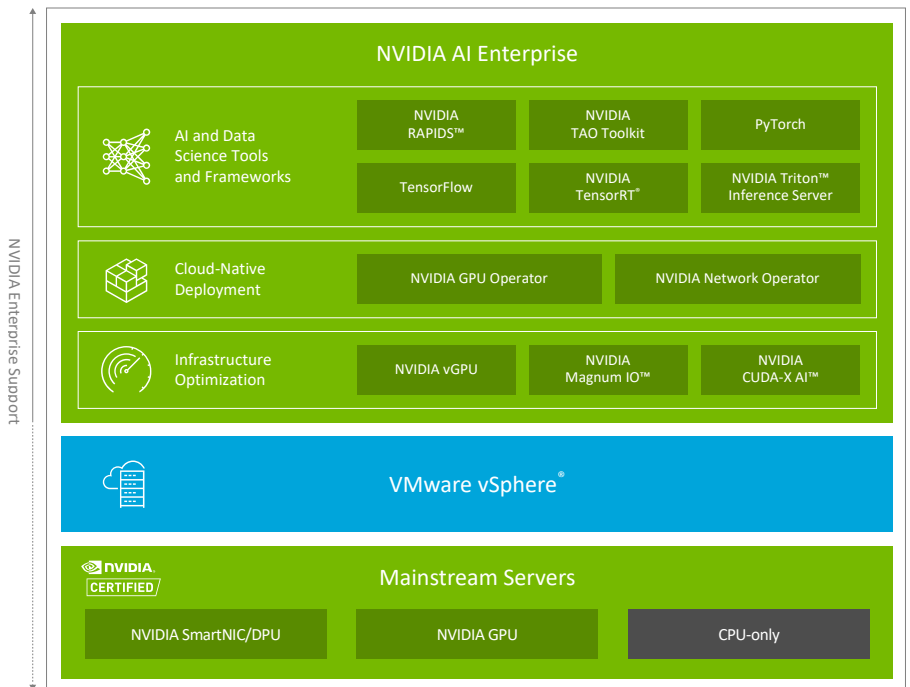
- Avoid AI silos and simplify management by leveraging virtualization to fold AI deployments into existing enterprise infrastructure.
- Get near bare-metal performance for AI and machine learning training and inference, with the added benefits of enterprise-grade security, manageability, and efficiency.
- Enable high availability for AI workloads while simplifying infrastructure maintenance such as consolidation, expansion and upgrades.
- Simplify management at scale across AI infrastructure with automatic initial placement of workloads for optimal resource consumption and deterministic costs.

Key Enabling Technologies

The solution is comprised of the following building blocks.

NVIDIA AI Enterprise Suite

NVIDIA AI Enterprise is an end-to-end, cloud native suite of AI and data science applications and frameworks optimized and exclusively certified by NVIDIA to run on vSphere with NVIDIA-Certified Systems. It includes key enabling technologies and software from NVIDIA for rapid deployment, management and scaling of AI workloads in the modern hybrid cloud. NVIDIA AI Enterprise is licensed and supported by NVIDIA.



vSphere

vSphere, the enterprise workload platform runs over 70 million workloads across hundreds of thousands of organizations worldwide. vSphere transforms bare-metal servers (including CPU- and GPU-based resources) into centrally managed AI and machine learning infrastructure pools that can quickly provision virtual machines and containers on demand.

vSphere optimizes performance, increases availability, tightens security, streamlines maintenance, and reduces costs to create an agile, efficient, resilient and intrinsically secure infrastructure platform that supports existing workloads and next-gen applications such as AI.

vSphere has been enhanced to optimally run the NVIDIA AI Enterprise software suite leveraging NVIDIA’s latest GPUs.

NVIDIA-Certified Systems

The NVIDIA AI Enterprise Suite is certified to run on NVIDIA-Certified Systems, which include the following:

- NVIDIA Ampere architecture-based GPUs such as the NVIDIA A100 Tensor Core GPU. The Tensor Core technology included in the Ampere architecture has brought dramatic speedups to AI operations, bringing down training times from weeks to hours and providing massive acceleration to inference.
- NVIDIA® Mellanox® ConnectX® SmartNICs and the NVIDIA BlueField® data processing unit (DPU) provide a host of software-defined hardware engines for accelerating networking and security. These enable the best of both worlds: best-in-class AI training and inference performance, with all the necessary levels of enterprise data privacy, integrity and reliability.

Get Started

Learn more at vmware.com/aiml

To learn more about the VMware and NVIDIA partnership, visit

www.vmware.com/partners/global-alliances/nvidia and
www.nvidia.com/en-us/data-center/vmware/

or contact Sales at

www.vmware.com/company/contact_sales and
www.nvidia.com/en-us/contact/sales/.

1. Gartner®, "Press Release: Gartner Survey Reveals 80% of Executives Think Automation Can Be Applied to Any Business Decision", August 22nd 2022. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved