SOLUTION OVERVIEW

NVIDIA Elite Partner

XENON High Performance Computing

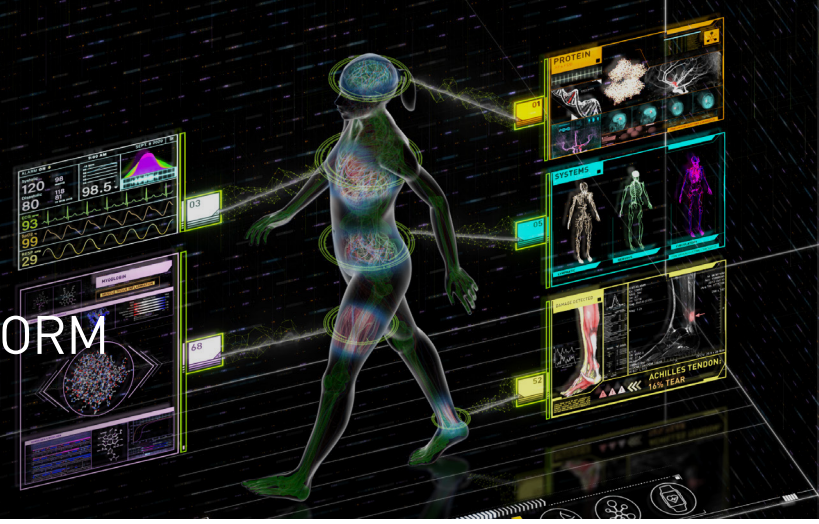

AI-READY ENTERPRISE PLATFORM FOR HEALTHCARE

Unleashing AI for Every Healthcare Enterprise with an End-to-End Platform from NVIDIA

Artificial intelligence is transforming healthcare in many ways, from helping doctors deliver better outcomes for patients to helping organizations run more efficiently. 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 and a rapid return on investment.

Outside of AI, hospitals and healthcare providers continue to face infrastructure challenges. As each medical specialty continues to advance, so do their needs for compute, bandwidth, and storage. Digital breast tomography and whole slide imaging are two examples where the amount of storage and processing required exert tremendous pressure on hospitals.

Unlike traditional enterprise applications, AI applications are a relatively recent development that challenge many IT departments. They're anchored in rapidly evolving, open-source code and lack proven approaches that meet the rigors of scaled production settings. In fact, just 53 percent of AI projects make it from pilot to production. According to Gartner, the complex integration of AI solutions with existing infrastructure remains among the top three barriers to implementation.[1]

A hospital data center is built to provide resilient, always-up services to doctors, nurses, and technicians, as well as clinical and operational teams.

Here are the three sets of systems that will challenge the healthcare enterprise:

- **Existing core hospital clinical and non-clinical applications:** This includes electronic medical record (EMR) systems, enterprise and radiology imaging, lab systems, pharmacy systems, email and communication systems, networking and security systems, and more.
- **New AI applications:** This includes analytics and AI software that support clinical and business processing, including the triage of incoming patient exams, identification and classification of clinical conditions, resource utilization planning, and alerts for service-level agreement (SLA) compliance.
- **Emerging AI infrastructure:** This includes software and hardware that support the creation, tuning, and deployment of AI applications in hospitals.

Healthcare organizations require infrastructure that can encompass not only their core critical applications today but also be AI-ready in the future. Ultimately, they need a foundational platform that supports the caregivers and the patients they serve.

KEY CHALLENGES FOR AI WORKLOADS IN HEALTHCARE

- **Risk:** Pulling together an end-to-end AI solution from disparate products—and integrating them with existing infrastructure—is difficult.
- **Performance:** Fast time to deployment and high performance are critical for AI, machine learning, and data analytics workloads.
- **Scaling:** Going from proof of concept to enterprise deployment requires effective scaling through efficient use of resources to ensure manageability, availability, and infrastructure cost management.

BENEFITS FOR HEALTHCARE ORGANIZATIONS

- **Simplifies AI development** and deployment with included AI frameworks and containers, enabling insights to be gathered faster and business value to be delivered sooner
- **Certified to deploy anywhere,** including on popular data center platforms from VMware and Red Hat, on mainstream NVIDIA-Certified Systems, and in the public cloud
- **Supported by NVIDIA** AI experts to keep AI projects on track, with included priority notifications, long-term support, and customized support-upgrade options

[1]Gartner Identifies the Top Strategic Technology Trends for 2021," Gartner, October 2020

AI-Ready Platform from NVIDIA

The NVIDIA AI Enterprise software suite enables hospitals and healthcare organizations to harness the power of AI, even if they don't have AI expertise today. Optimized to streamline AI development and deployment, NVIDIA AI Enterprise includes proven, open-source containers and frameworks, certified to run on common data center platforms from VMware and Red Hat, mainstream NVIDIA-Certified Systems™ configured with GPUs or CPU-only, and on the public cloud. Since support is included, hospitals benefit from the transparency of open source and the assurance that they can get help from the global NVIDIA Enterprise Support team to keep their AI projects on track. With the NVIDIA AI Enterprise software suite, AI is accessible to organizations of any size, and provides the compute power, tools, and support that organizations need to focus on creating AI business value.

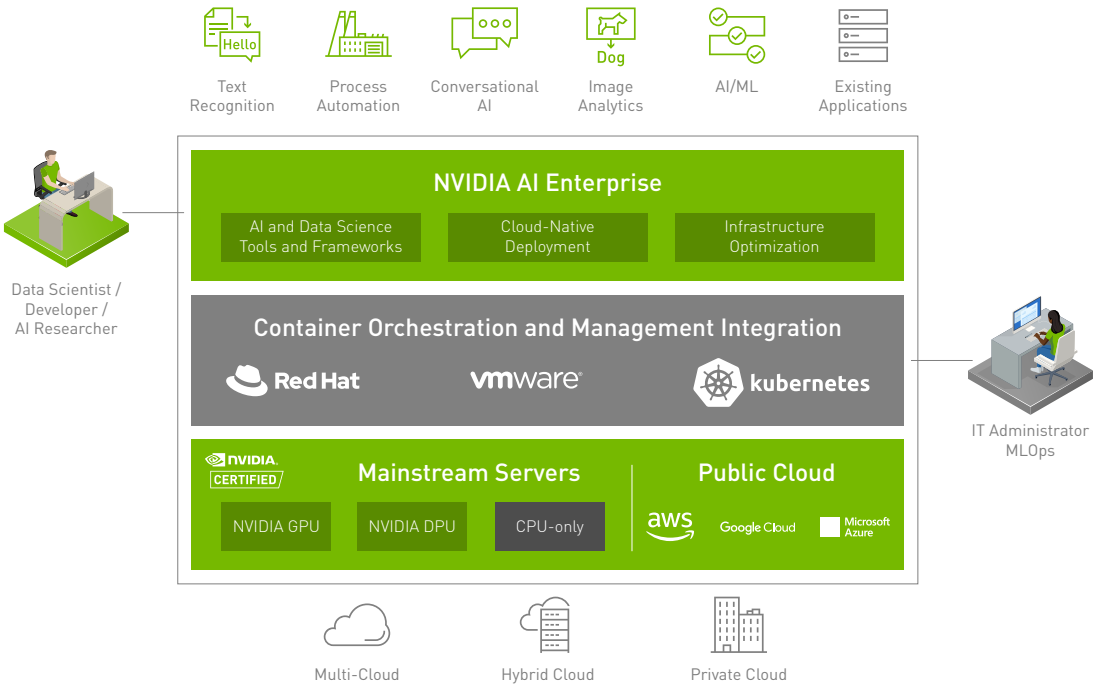

Figure 1. With the AI-Ready Platform, end users can access the software they need to build successful AI projects, and IT admins can support the projects using familiar tools.

Solution Highlights

Consolidate AI Applications for Your Healthcare Enterprise

NVIDIA powers the core hospital and AI applications and frameworks, which include three major components:

- Clinical and business healthcare applications: Core healthcare applications include EMR systems, picture archiving and communication systems (PACS), and advanced visualization solutions.
- Healthcare data science research: Solutions like NVIDIA Clara™ and MONAI are used to create and deploy in-house AI applications. NVIDIA Clara is a healthcare application framework for AI-powered imaging, genomics, and the development and deployment of smart sensors. Clara Imaging is an application framework for medical imaging and pathology. Clara Guardian is a smart hospital application framework that consists of the NVIDIA DeepStream toolkit for video analytics, the NVIDIA Riva SDK for conversational AI, and NVIDIA Triton™ Inference Server for the deployment of AI models at scale.
- Emerging healthcare AI applications: Thousands of healthcare AI applications speed triaging, diagnostic, and operational insights.

Create and Deliver AI with NVIDIA AI Enterprise

NVIDIA AI Enterprise is an end-to-end, cloud-native suite of AI and data analytics software, optimized so every healthcare organization can succeed with AI. It's also certified to deploy anywhere—from the enterprise data center to the public cloud—and includes global enterprise support and training.

NVIDIA AI frameworks and containers enable performance-optimized data science, training, and inference. They also simplify the building, sharing, and deployment of AI applications, so enterprises can gather insights faster and deliver business value sooner. For example, NVIDIA RAPIDS™ makes it possible for organizations to streamline the data science process with up to 70X faster performance, while improving cost-effectiveness by up to 20X. Additionally, NVIDIA TAO Toolkit boosts AI development by up to 10X, the NVIDIA Triton™ Inference Server allows organizations to improve their total cost of ownership (TCO) by 97X, and the NVIDIA A100 Tensor Core GPU enables them to achieve 100X lower latency when compared to CPU.

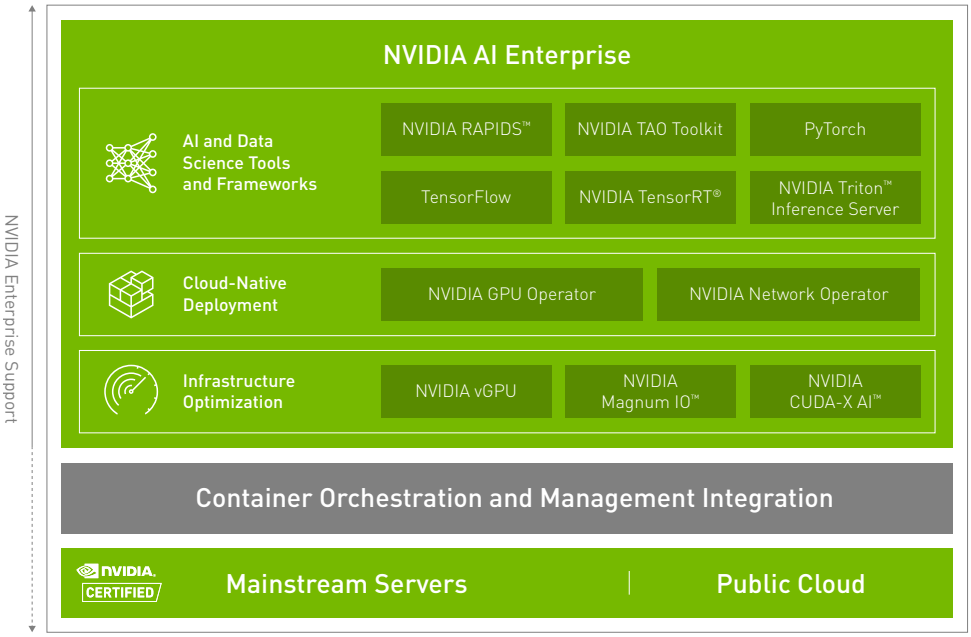

Figure 2. The NVIDIA AI Enterprise software suite includes the applications, frameworks, and tools used by AI researchers, data scientists, and developers, as well as tools for cloud-native deployments and infrastructure optimizations.

Deploy Anywhere

NVIDIA AI Enterprise can be deployed anywhere, including in a virtualized enterprise data center with VMware vSphere with Tanzu, on mainstream NVIDIA-Certified Servers, with Red Hat's OpenShift container platform, in the public cloud, and even on CPU-only servers. Tested and certified by NVIDIA with broadly adopted container, management, and server platforms, hospitals have the flexibility to deploy NVIDIA AI Enterprise with confidence, in just about any data center environment—on premises, in a colocation facility, or in a private cloud. Organizations with a hybrid cloud strategy also have the flexibility to run NVIDIA AI Enterprise on GPU-accelerated public cloud instances—on Amazon Web Services (AWS), Azure, and Google—with full NVIDIA support.

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 contain Tensor Core technology that delivers dramatic speedups to AI operations, reduces training times from weeks to hours, and provides massive inference acceleration.
- NVIDIA® ConnectX® smart network interface cards (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.
- NVIDIA converged accelerators combine the powerful performance of the NVIDIA Ampere architecture with the enhanced security and latency-reduction capabilities of the NVIDIA BlueField-2 DPU. Enterprises can use converged accelerators to create faster, more efficient, and secure AI systems in data centers and at the edge.

NVIDIA Enterprise Support

With NVIDIA AI Enterprise, hospitals benefit from the transparency of open source with the assurance of full enterprise support and platform certification. They can extend their team to include NVIDIA experts, get support ticket prioritization and coordinated support across the full solution and partner products until resolution, control upgrades and maintenance schedules with long-term support (LTS) options, and access the latest customer training and knowledge base resources.

NVIDIA AI Enterprise support includes:

- Full enterprise-grade assistance for every deployment option, including bare metal, virtualized, containerized, GPU and CPU, and public cloud
- Access to NVIDIA AI experts during local business hours for guidance on configuration and performance, including access to engineering
- Priority notifications related to the latest security fixes and maintenance releases
- Long-term support for up to three years for designated software branches
- Customized support-upgrade options, including a designated technical account manager (TAM) and business critical support for 24x7 live agent access.

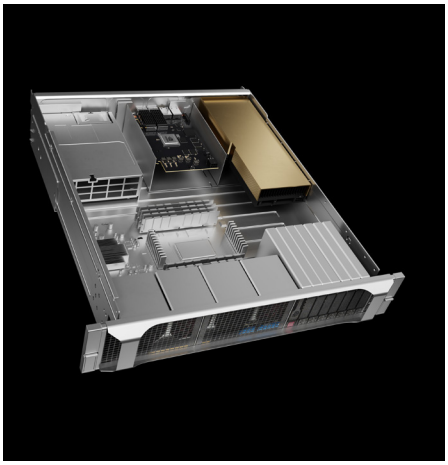

Figure 3. NVIDIA-Certified Systems bring powerful speedups to AI training and inference.

NVIDIA-CERTIFIED SYSTEMS

- Confidently deploy scalable hardware and software solutions that securely and optimally run accelerated workloads.
- Learn more about accelerated servers at **nvidia.com/certified-systems**

NVIDIA CERTIFIED

NVIDIA AI Enterprise Trial Programs

NVIDIA offers the following trial programs that help customers evaluate products for free, depending on their existing infrastructure.

- **NVIDIA LaunchPad**
  NVIDIA LaunchPad provides organizations around the world with immediate, short-term access to the NVIDIA AI Enterprise software suite running on private accelerated computing infrastructure, including a set of hands-on labs for AI practitioners and IT staff. With NVIDIA LaunchPad, enterprises can speed the development and deployment of modern, data-driven applications and quickly test and prototype the entire AI workflow on the same complete stack for purchase and deployment.
- **Evaluation Software**
  The NVIDIA AI Enterprise evaluation software is available for customers who have existing NVIDIA-Certified Systems and is ideal for starting a proof-of-concept (POC) project for deployment at scale.

Building Smart Hospitals with the AI-Ready Platform

The AI-Enterprise Ready Platform from NVIDIA offers the end-to-end hardware and software that hospitals need to digitally transform. Developers, data scientists, and researchers can access the resources they need to build and deploy AI applications efficiently, and IT administrators can confidently deliver uncompromised support using the tools and infrastructure they know. This comprehensive solution makes smart hospitals a reality and, ultimately, helps deliver a better experience to clinicians and the patients they serve.

Ready to Get Started?

To learn more about the AI solutions for healthcare, visit:
https://xenon.com.au/category/case-studies/health-life-science/
https://xenon.com.au/machine-learning-artificial-intelligence-solutions/
nvidia.com/ai-enterprise-suite
nvidia.com/try-ai
nvidia.com/certified-systems

Visit XENON

Contact XENON today!
www.xenon.com.au | info@xenon.com.au | 1300 888 030

NVIDIA Elite Partner

XENON® High Performance Computing