DELL Technologies

NVIDIA AI Enterprise

2024 Sales Guide

Everything you need to prepare, promote, and sell NVIDIA AI Enterprise

NVIDIA

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

Sales objectives and goals

We want to help enable customers to embark on a successful AI journey through the utilization of Dell Hardware and NVIDIA Hardware, built together with the “Operating System for AI,” which is the AI Enterprise Software Package. This software helps customers build, deploy, support, and scale AI workloads and enable them to move from pilot to production with confidence. As Dell sellers, we should look to quote AI Enterprise licensing to any customer on the path of AI who is confidently engaging with NVIDIA to help provide the best use case examples to ensure customer success on their journey.

What’s in it for me?

- Retire additional quota (ISG Server+) and increase margin on the overall deal.
- Solidify Dell’s role as a trusted advisor for end-to-end AI solutions.
- Maintain Dell’s footprint within the account by providing all the components for AI solutions, becoming the one-stop shop for our customer’s AI needs.
- Uncover additional deal drag from storage, to networking, to servers, and even racks and power.

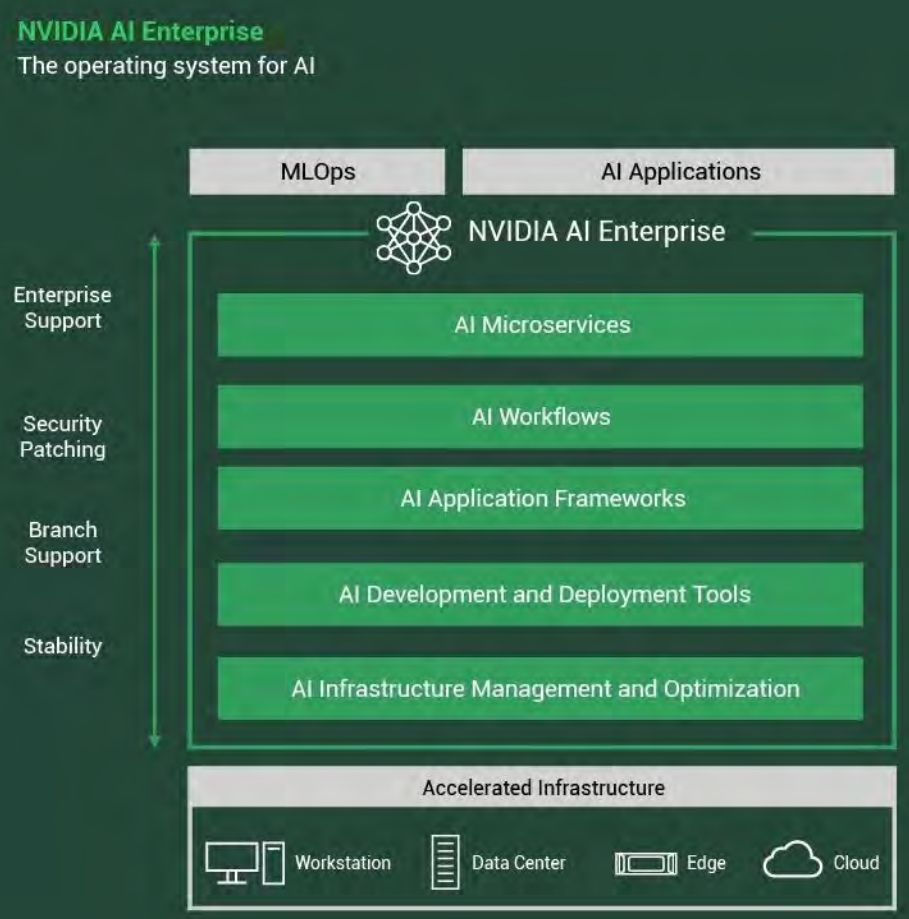

- Cloud-native software platform designed to accelerate data science pipelines and streamline the development and deployment of production-grade AI applications.
- Acts as the “operating system” for enterprise AI, providing optimized model performance, enterprise-grade security, and stability.
- Microservices within NVIDIA AI Enterprise ensure a smooth transition from prototype to production for business running on AI.

What is NVIDIA AI Enterprise?

NVIDIA AI Enterprise is an end-to-end, cloud-native software platform that accelerates data science pipelines and streamlines development and deployment of production-grade co-pilots and other generative AI applications. Easy-to-use microservices provide optimized model performance with enterprise-grade security, support, and stability to ensure a smooth transition from prototype to production for enterprises that run their businesses on AI.

NVIDIA AI Enterprise video >

A 2-minute overview of NVIDIA NIM for scaling generative AI across the enterprise.

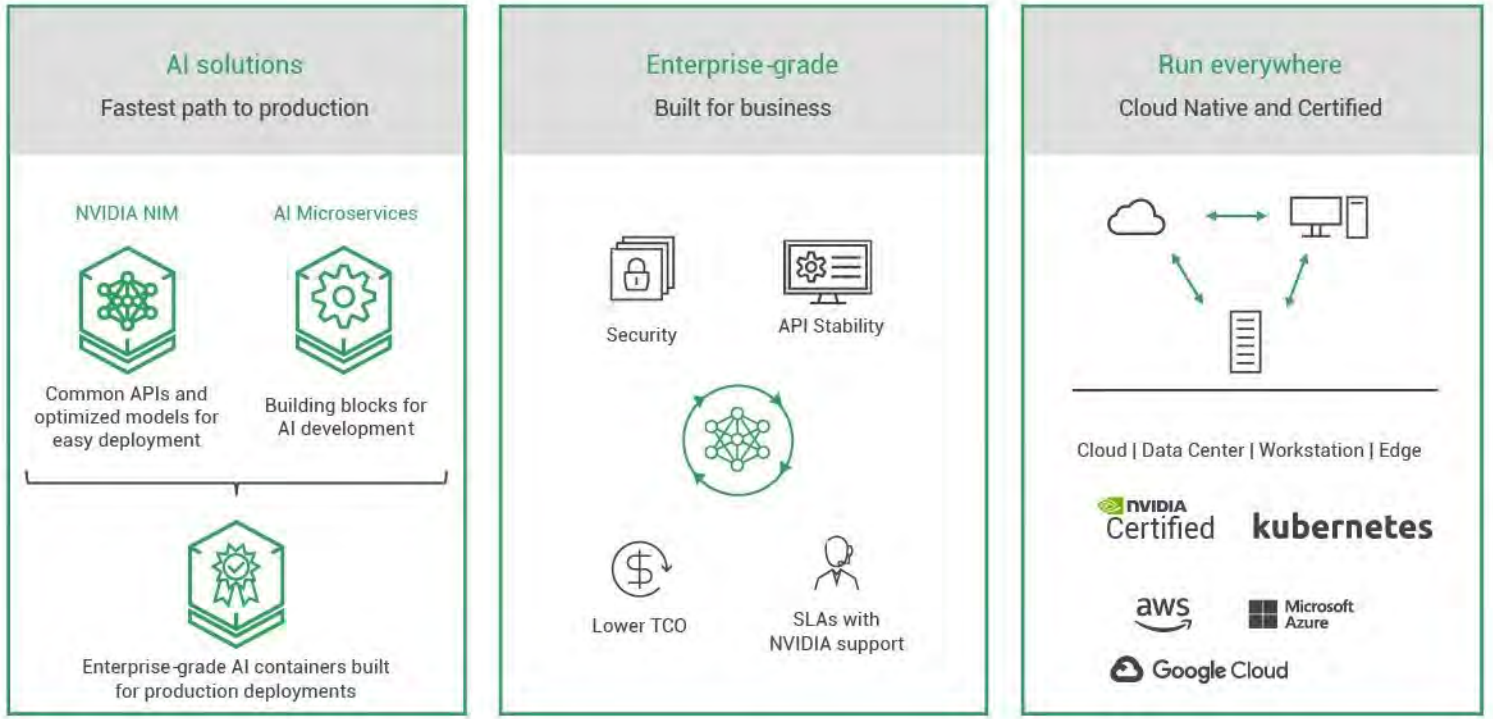

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA.

Challenges enterprises face when developing AI

Complexity
Managing an end-to-end AI stack from disparate products is resource-intensive. Optimizing and tuning the environment also creates roadblocks to deployment and production.

Performance
Fast deployment and high performance are crucial for AI, ML, and data analytics.

Security
The AI software stack's many open-source packages introduce daily risks.

Implementation cost
AI is compute-intensive, requiring thousands of servers, leading to high infrastructure costs and energy use.

Support
AI apps use ever-changing open-source code, which lacks the support, security, and compatibility of enterprise software.

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

NVIDIA AI Enterprise

Key functionality & benefits

NVIDIA AI Enterprise Software Package
Package of software with all AI tools from the NVIDIA NGC library, NVIDIA Inference Microservices (NIM) for easier planning and deployment, Base Command Manager Essentials for cluster management, Pre-Trained Models, AI Workflows, Infrastructure Optimization, and Enterprise-Grade Security with Support for all licensed software.

NVIDIA AI Enterprise Support
Ensure smooth AI Ops with expert guidance, guaranteed response times, case management, and regular updates. Ensure uptime of mission-critical AI apps with a stable and secure environment.

NIM
Scale, adapt, and integrate AI apps by managing AI inference as independent microservices. Bridge the gap between the complex world of AI development and the operational needs of enterprise environments.

NIM Agent Blueprints
NIM Agent Blueprints are reference AI workflows tailored for specific use cases. They include sample applications built with NVIDIA NIM and partner microservices, reference code, customization documentation, and a Helm chart for deployment.

NVIDIA Base Command™ Manager Essentials
Easily maintain AI infrastructure, manage workloads, and monitor resources with cluster management software for AI and HPC.

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

NVIDIA AI Enterprise Software Package

A unified suite for scalable AI development and deployment

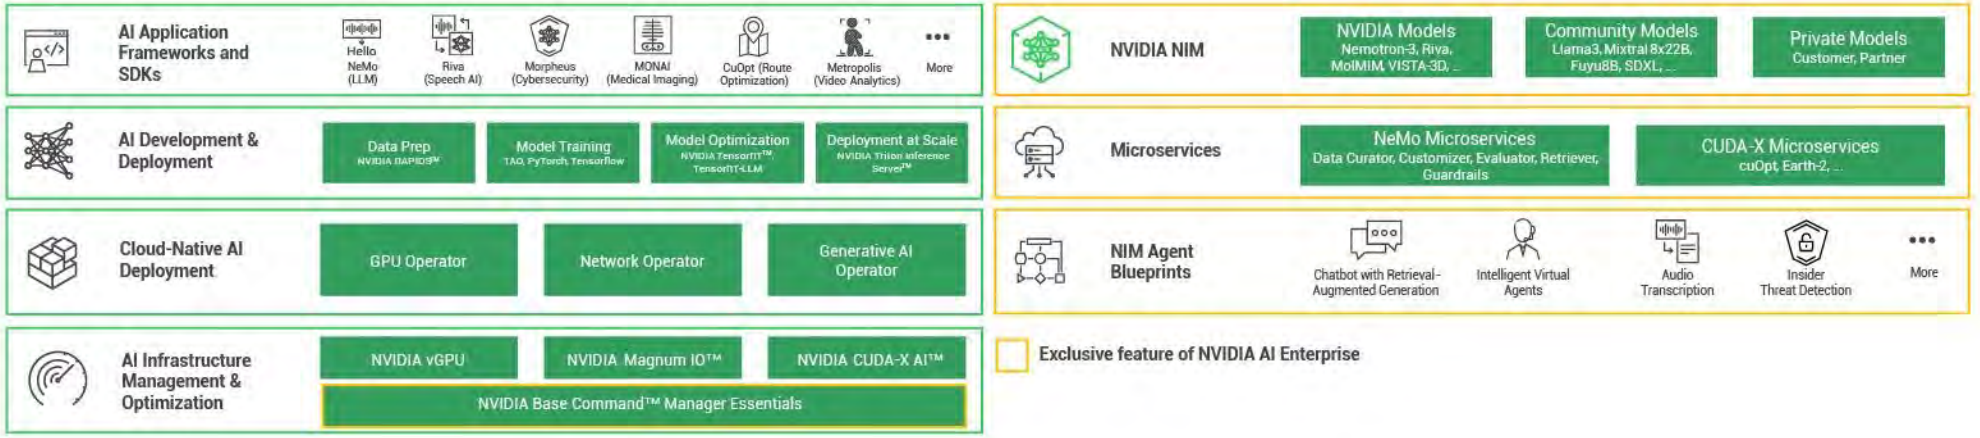

1,000 of open-source AI SDKs, framework, and pre-trained models

- Generative AI
- Digital avatars
- Autonomous vehicles
- Cybersecurity
- Drug discovery
- Speech AI
- Federated learning
- Medical imaging
- Physics AI
- Route optimization
- Robotics
- Video analytics
- Audio / video enhancement
- Sensor data processing

NVIDIA AI Enterprise

- Enterprise support for best-in-class development tools, frameworks, and pretrained models for AI practitioners
- Reliable management and orchestration for IT professionals to ensure performance, high availability, and security

Enterprise Support Levels

Key support offerings to address customer needs

Enterprise Business Standard Support

NVIDIA's foundational support level is Enterprise Business Standard Support. It provides end-to-end enterprise support for designated NVIDIA branded solutions. The Business Standard service level is designed to deliver the best value and optimal use of NVIDIA solutions. In addition, customers may supplement this service with Value Add Services to meet their business needs. If broader and faster support is required, customers may select Enterprise Business Critical Support.

Enterprise Business Critical Support

The Enterprise Business Critical Support service level is NVIDIA's premium support service level. It is designed for mission-critical deployments where a small downtime may cause a significant business impact. The Business Critical Support Service provides 24x7 support and a one-hour response time for Severity Level 1 cases. The Business Critical Support service is available for designated NVIDIA offerings. For complete coverage, Enterprise Business Critical Support needs to be purchased for all NVIDIA offerings being deployed with the solution.

NVIDIA AI Enterprise support levels

Service component	Business standard	Business critical (select products)
Technical support access	Local business hours	24 x 7
Severity 1 initial response times	4 hours	1 hour
Severity 2 initial response times	4 hours	2 hour
Severity 3 initial response times	1 business day	4 hours
Severity 4 initial response times	2 business days	1 business day
Access to Customer Portal	24 x 7	24 x 7
Web	24 x 7	24 x 7
Phone	Local business hours	24 x 7
Product coverage	Hardware / firmware / software / cloud	Hardware / firmware / software / cloud

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

Enterprise Support Levels

Overview of support branches and release cadence

Feature Branch

- Top of tree SW optimization
- Monthly release cadence
- CVE patches and bug fixes in roll forward release

Production Branch

- API Stability
- Monthly CVE patches and bug fixes
- 2 branches/year with 9-month lifetime
- 3-month overlap between 2 PBs

Long-Term Support Branch

- For highly regulated industries Quarterly
- CVE patches/bug fixes
- Up to 3 years support
- 6-month overlap period

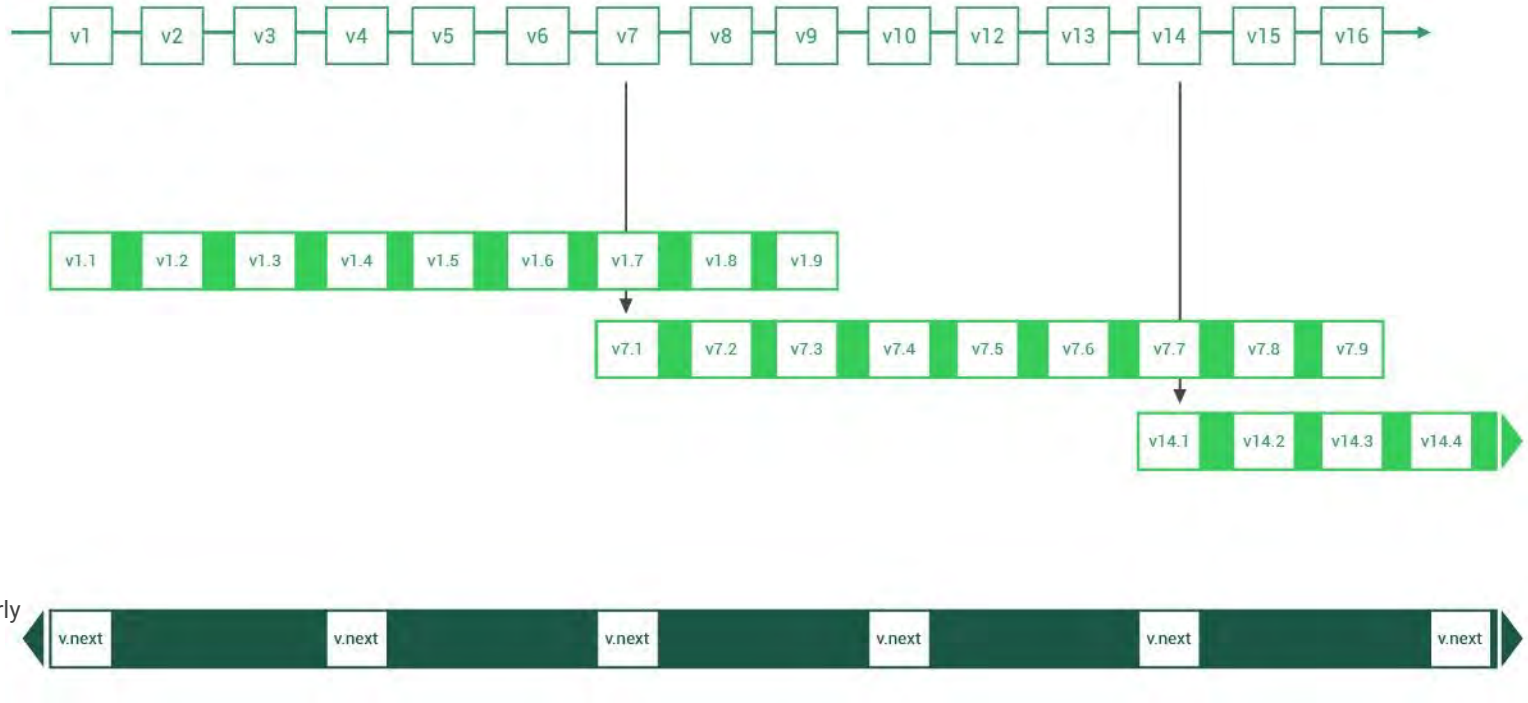

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

NVIDIA Inference Microservices (NIM)

Aligning AI development with enterprise scalability and operations

NVIDIA NIM is designed to bridge the gap between the complex world of AI development and the operational needs of enterprise environments, enabling 10-100X more enterprise application developers to contribute to AI transformations of their companies.

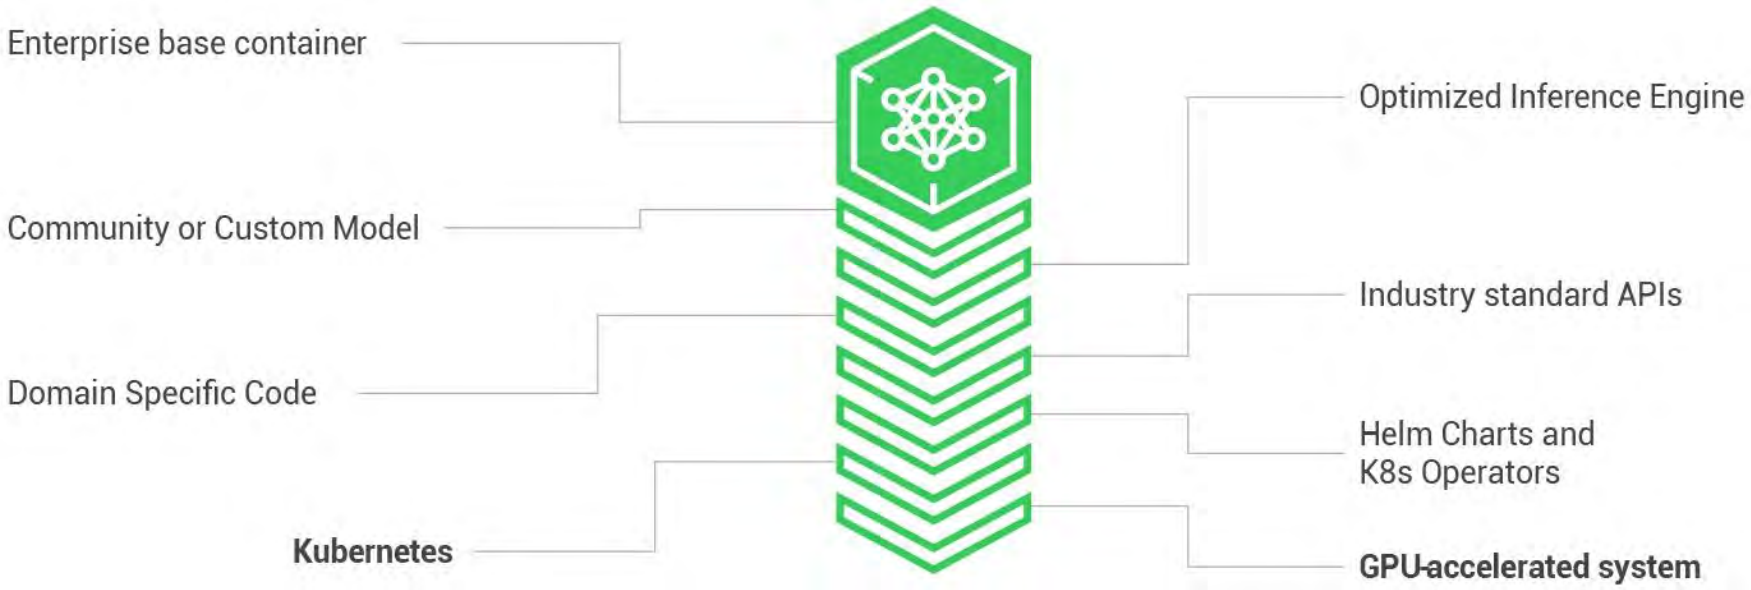

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

NVIDIA Inference Microservices (NIM)

Key benefits for deploying and scaling AI across the enterprise

Deploy anywhere

NIM is built for portability and control, enabling model deployment across various infrastructures, from local workstations to cloud to on-premises data centers. This includes NVIDIA DGX, NVIDIA DGX Cloud, NVIDIA Certified Systems, and NVIDIA RTX workstations and PCs.

Prebuilt containers and Helm charts packaged with optimized models, are rigorously validated and benchmarked across different NVIDIA hardware platforms, cloud service providers, and Kubernetes distributions. This enables support across all NVIDIA-powered environments and ensures that organizations can deploy their generative AI applications anywhere, maintaining full control over their applications and the data they process.

Develop with industry-standard APIs

Developers can access AI models through APIs that adhere to the industry standards for each domain, simplifying the development of AI applications. These APIs are compatible with standard deployment processes within the ecosystem, enabling developers to update their AI applications swiftly—often with as little as three lines of code. This seamless integration and ease of use facilitate rapid deployment and scaling of AI solutions within enterprise environments.

Leverage domain-specific models

NIM also addresses the need for domain-specific solutions and optimized performance through several key features. It packages domain-specific NVIDIA CUDA libraries and specialized code tailored to various domains such as language, speech, video processing, healthcare, and more. This approach makes sure that applications are accurate and relevant to their specific use case.

Run on optimized inference engines

NIM leverages optimized inference engines for each model and hardware setup, providing the best possible latency and throughput on accelerated infrastructure. This reduces the cost of running inference workloads as they scale and improves the end-user experience. In addition to supporting optimized community models, developers can achieve even more accuracy and performance by aligning and fine-tuning models with proprietary data sources that never leave the boundaries of their data center.

Support for enterprise-grade AI

Part of NVIDIA AI Enterprise, NIM is built with an enterprise-grade base container providing a solid foundation for enterprise AI software through feature branches, rigorous validation, enterprise support with service-level agreements, and regular security updates for CVE. The comprehensive support structure and optimization capabilities underscore the role of NIM as a pivotal tool in deploying efficient, scalable, and customized AI applications in production.

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA.

NIM Agent Blueprints

Driving continuous growth of digital intelligence

This is an AI data flywheel. An AI query engine must be able to store and leverage learnings from AI-powered applications and agents in production, automatically increasing the knowledge of the enterprise, creating an AI data flywheel. Data is the raw material, tokens are the new commodity, and NVIDIA is the token generator in the AI factory. This way, every company will produce digital intelligence.

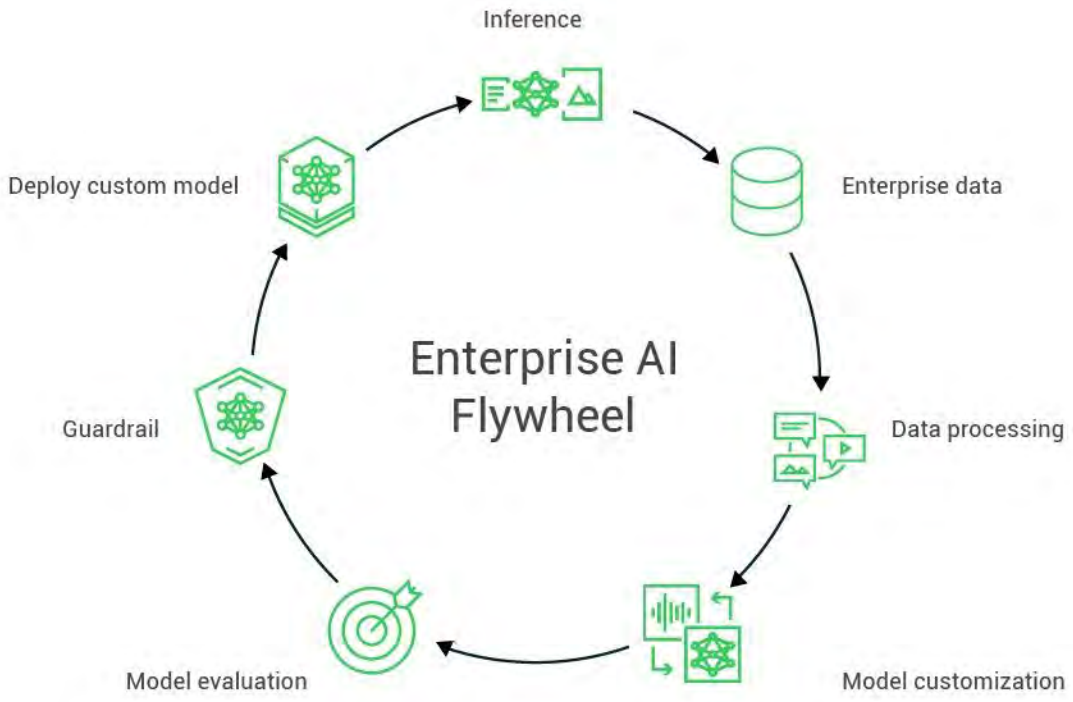

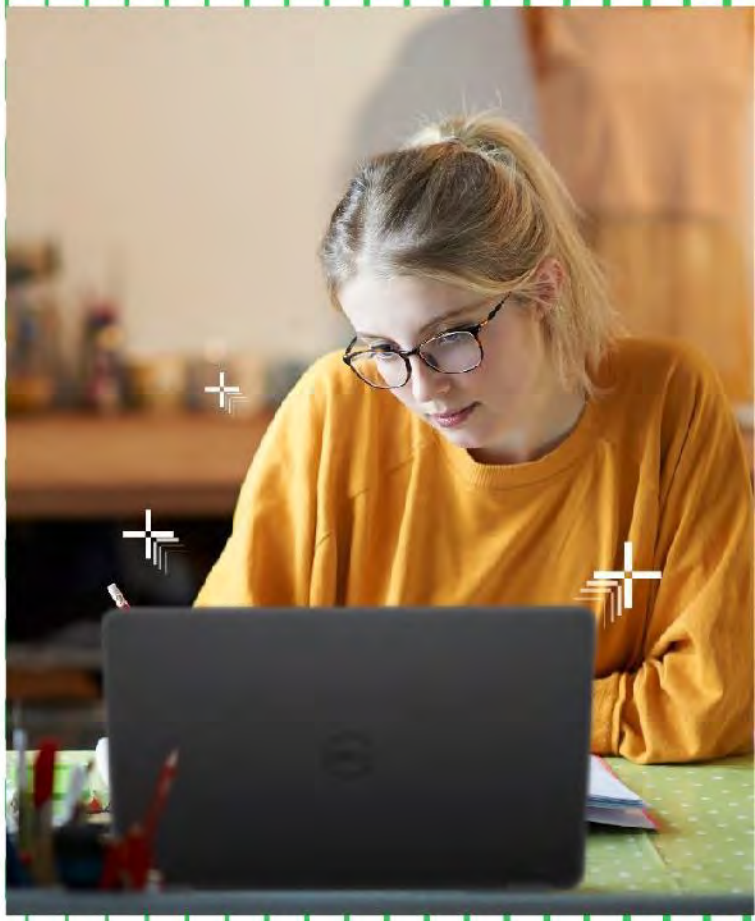

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

NIM Agent Blueprints

Accelerating generative AI with pre-trained, customizable workflows

NVIDIA NIM Agent Blueprints are reference AI workflows, pre-trained for specific use cases, that can be customized by enterprise developers and the NVIDIA partner ecosystem. They include reference code for constructing the workflow using NVIDIA AI Enterprise software, including NVIDIA NIM microservices and NVIDIA NeMo, along with tools and documentation for deploying in your own environment and customizing with your own data. They also include a reference architecture which documents API definitions, how the microservices interoperate, and more.

The collection of NIM Agent Blueprints include AI Virtual Assistants and Digital Humans for Customer Service, Multimodal PDF Data Retrieval, High Throughput Virtual Screening, and Software Security Vulnerability Analysis. More will follow.

NIM Agent Blueprints are foundational building blocks that enable the generative AI-powered enterprise of the future. An AI query engine is the brain that powers AI agents. An army of AI agents will transform enterprises–enabling instant, exponential growth of resources, enhancing productivity and efficiency of enterprise employees in every function. AI agents can quickly search, index, and retrieve the most accurate, relevant, and context-aware insights from petabytes of enterprise data.

You can talk to your AI agents in natural language–and in your local language. They can follow multi-turn conversations. AI agents can have human-like interfaces–digital humans. They are optimized to make the best decisions, taking into account all variables captured across enterprise data sources. These building blocks serve as the basis for building generative AI-powered applications and agents for various use cases across many industries.

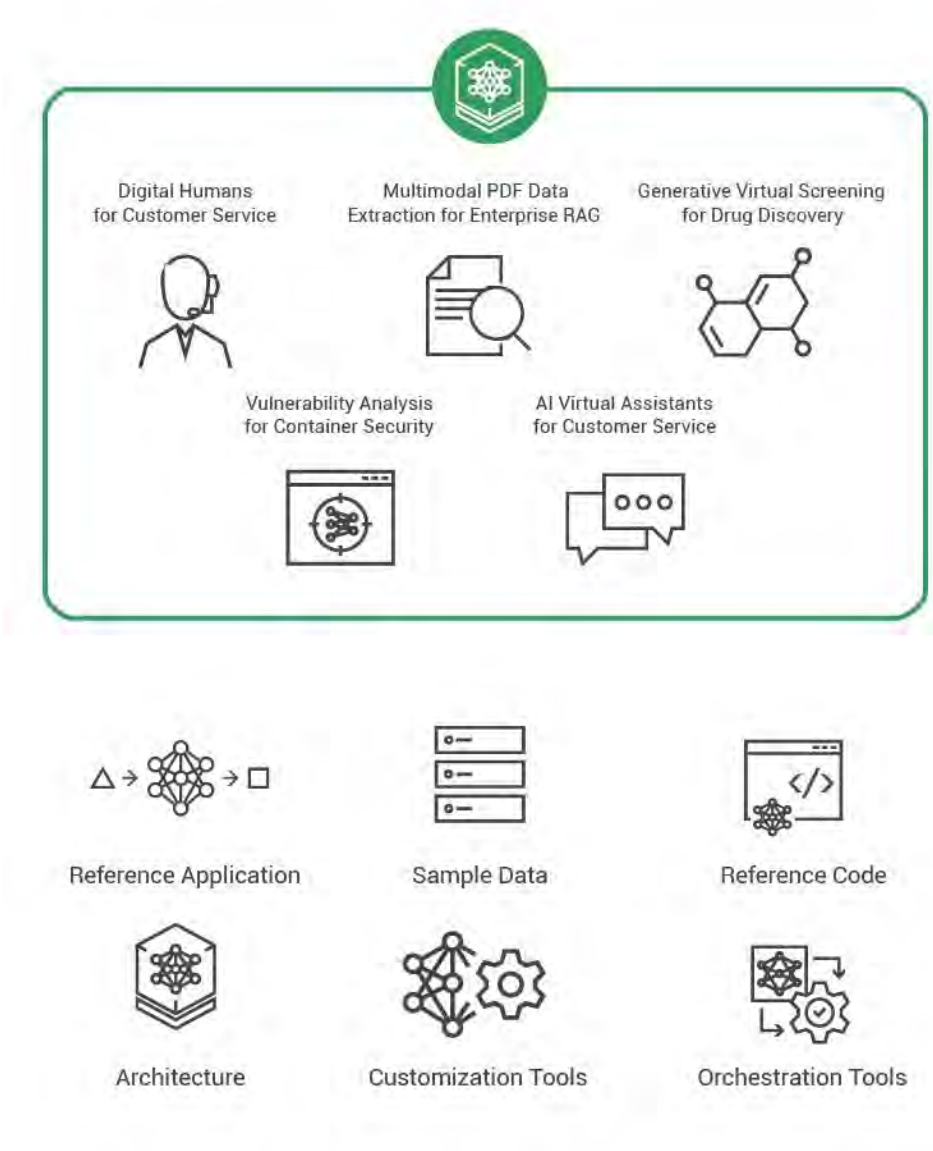

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

NVIDIA base command manager essentials

Key challenges for AI workload infrastructure management

Workload complexity

- Supporting increasingly sophisticated AI pipelines
- Managing and using specialized accelerated computing hardware

Resource optimization

- Optimizing utilization of specialized compute resources
- Gaining insights into cluster usage for informed decision-making

Reliability and scalability

- Operationalizing systems management at scale
- Providing a resilient computing infrastructure for data science

NVIDIA Base Command Manager Essentials streamlines infrastructure provisioning, workload management, and resource monitoring across data center, edge, and hybrid cloud. Built for AI and data science, it facilitates deployment of AI developer and deployment tools—including Kubernetes and Jupyter Notebooks—dynamic scaling, and policy-based resource allocation. It also ensures cluster integrity and reports on cluster usage by project or application, enabling chargeback and accounting.

Automates the complexities of infrastructure management, empowering IT administrators and DevOps teams to focus on running production applications.

Benefits of NVIDIA base command manager essentials

Infrastructure provisioning

- Installs the OS, sets up networking, security, and DNS, and ensures cluster integrity
- Automates server management and updates, preventing server drift

User access and workload management

- Deploys Kubernetes, automates scaling, and enables streamlined Jupyter setup with NVIDIA NGC containers; Run:ai integration is included

Resource monitoring

- Provides comprehensive cluster and job monitoring, GPU metrics, resource allocation, access control, and chargeback options

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

NVIDIA base command manager essentials

Simplifying end-to-end AI infrastructure management

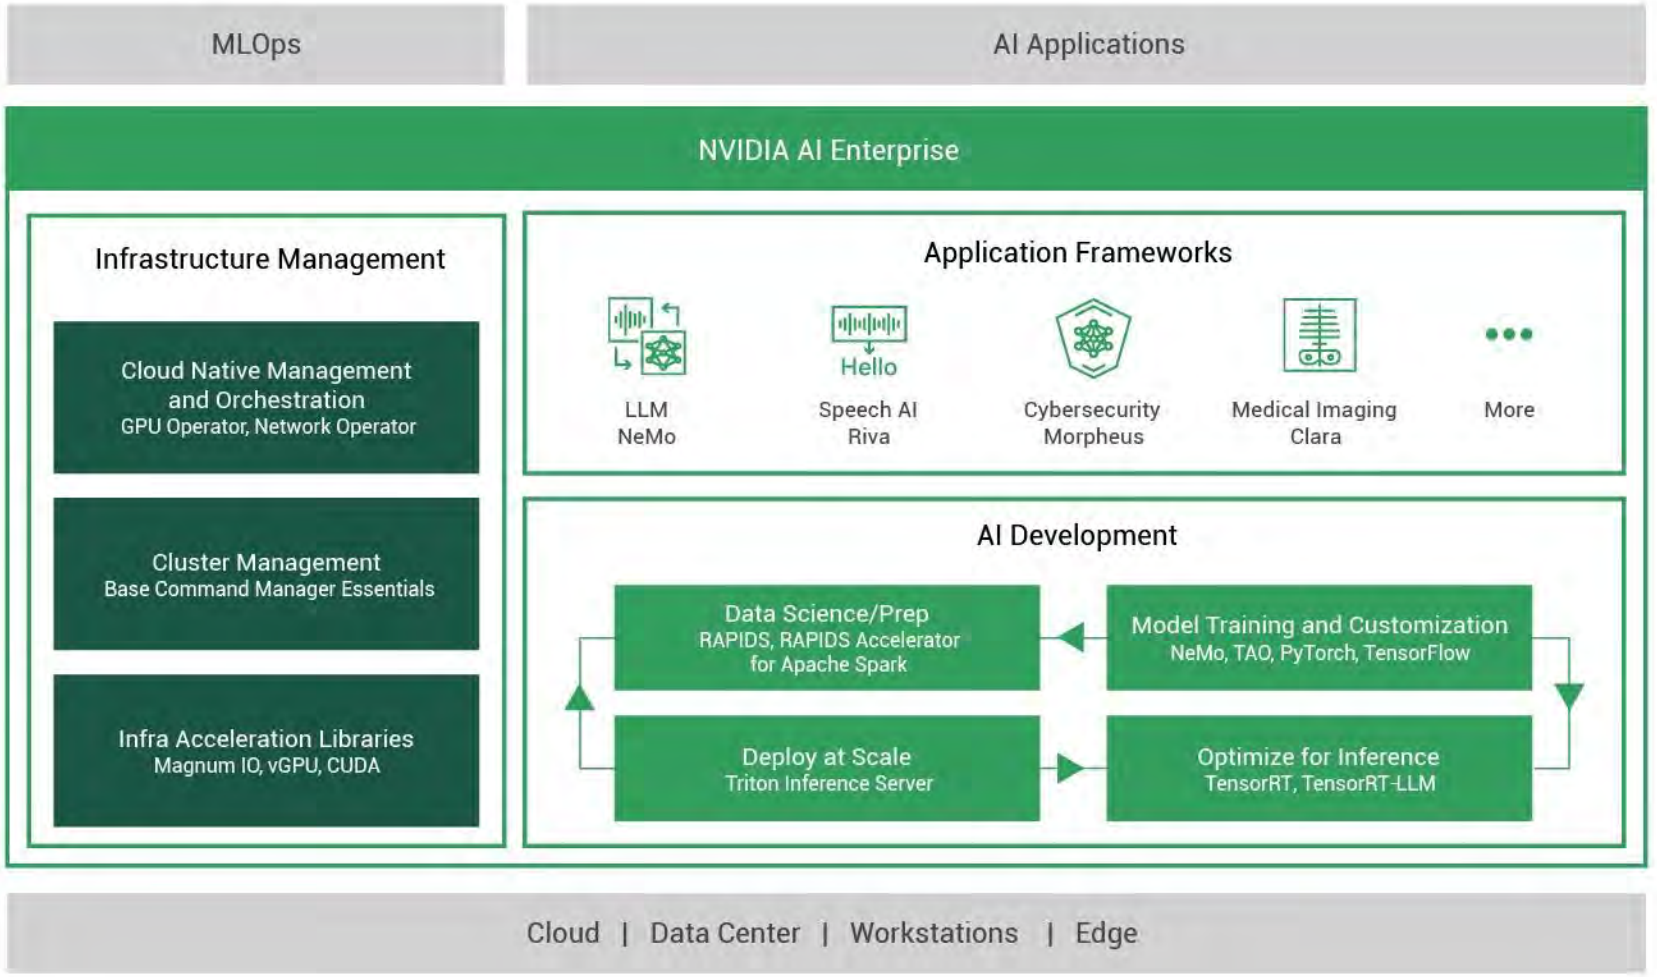

Use cases

AI is already transforming every industry

Demand for fast and easy deployment is greater than ever

Credit card fraud
- 1.1B credit transactions / day

Contact center AI
- 500M calls / day

Meeting transcription
- 15B meeting minutes / day

Public safety
- >1B smart city cameras deployed

Product recommendations
- 300M e-commerce visitors / day

Product recommendations
- $275M inventory loss / day

Medical imaging
- 10M diagnostic scans / day

Industrial inspection
- 94M vision sensors installed by 2025

Use cases

- Industry Use Cases for NVIDIA AI Enterprise >
- NVIDIA AI Enterprise SDK Use Case Mapping >

Enabling enterprise transformation with AI

NVIDIA is leading in all industry AI use cases

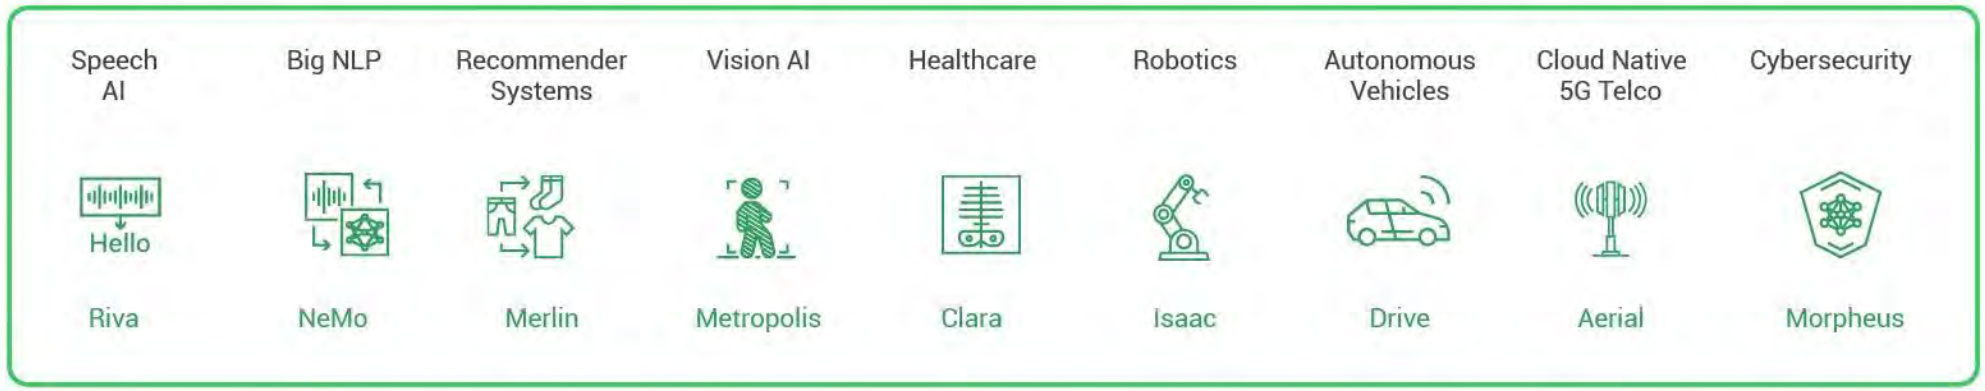

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

Targeting & engaging customers

Understanding your target audience to drive sales success

	AI practitioners	IT administrators	Line of business
Titles	> VP of Innovations > Chief Data Officer > Data Scientists > Developers > Data Engineer	> Infrastructure Experts > CIO, CTO > MLOps Team > VMware Administrators	> VP of Marketing > VP of Customer Support > VP of Loss Prevention > VP of Operations > VP of Product Development
What they do	> Develop and productionize AI models > Integrate AI into applications	> Deploy and maintain IT Infrastructure for AI-enabled application	> Get results as fast as possible > Bring solutions that are efficient and flexible to meet needs and requirements
What they want	> Agile development > Fast roll-out to production > Choice of frameworks/models > Cloud-native, Kubernetes > High performance infrastructure	> High availability/stability > Ensure security > Ease of management, scalability > Upgrade vs Rip and Replace	> Time to business value > Reliability of results > Respond to evolving business needs

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

Targeting & engaging customers

	AI practitioners	IT administrators	Line of business
Prospecting questions	1. What are your top challenges to creating AI applications and services? a. Access to infrastructure b. Access to AI development software tools c. In-house expertise d. Data e. Complexity of using rapidly changing open source tools/models f. Other 2. Would it be helpful if you had a common AI platform that enables you to create AI solutions fasters that are production ready for the cloud, hybrid cloud, or on-prem data center?	1. What are the top challenges in supporting your organization's AI strategy? a. Lack infrastructure b. Insufficient in-house AI expertise c. Data management d. Complexity of managing AI tools, and open source software e. Security needs, enterprise support 2. Would it be helpful if you had a common AI platform that you could easily support AI projects/initiatives across your organization?	1. Does your organization have an AI strategy? 2. Would it be helpful if you had a common AI platform that you can easily support AI projects/initiatives across your organization? 3. Would you be interested in running your AI solutions significantly faster to save money?
Qualifying questions	1. Does your team need high- performance compute resources? 2. Do you use open-source AI tools? 3. Is it important to you that you have access to the latest innovations in AI development? 4. What container orchestration tools do you prefer—Kubernetes, OpenShift, or Tanzu? 5. How do you currently get support when something goes wrong with your model or project?	1. How does your organization provide infrastructure support for AI projects today? 2. Are you concerned about manageability and security of your AI infrastructure? 3. Would you like to have a future- proofed, scalable, and affordable AI platform that could support multiple AI projects? 4. How do you currently stand up your AI infrastructure? (Bare Metal, Virtualization, etc.) 5. If you could manage AI infrastructure in a virtualized environment, how could that impact the way you serve up and deploy AI across the organization?	1. Are multiple business functions/ departments leveraging AI? 2. Would you like to accelerate your organization's adoption of AI?

Overcoming objections

Addressing customer concerns with clear benefits

NVIDIA has software containers for free.

What's the difference between the freeware and your paid software?

Open-source software has played a central role in AI adoption as it allows developers to easily consume complex algorithms developed by a broad community. However, the diverse range of software components and associated interdependencies make maintaining a reliable AI software stack a serious task. The software complexity increases vulnerabilities. These CVEs (Common Vulnerabilities and Exposures) grow with time. Most enterprises don't have the resources to monitor these updates over time. NVIDIA AI Enterprise offers a secure, stable, supported software platform to address these challenges.

The NVIDIA AI Enterprise Production Branch is designed to ensure API stability and security for your AI applications. The Production Branch releases every six months with a 9-month lifetime, including monthly fixes for high-priority and critical software vulnerabilities. Customers can use production branches to get the newest vulnerability fixes without breaking your application. Enterprises that need long term support can also get 3 years with stable APIs.

My ML team is experienced enough to write and support our CUDA application. I don't need support.

Why should I bother looking at NVIDIA AI Enterprise?

AI is a constantly evolving endeavor. When AI is core to conducting business operations, reliability and uptime are critical for building generative AI at the enterprise level. That is why NVIDIA builds NVIDIA AI Enterprise, an enterprise software platform with enterprise-grade security, stability, manageability, and support to maintain AI application uptime and relieve organizations of the burden of maintaining and securing the complex software platform of AI. NVIDIA AI Enterprise support for production deployments comes with the full backing of NVIDIA experts, globally, with SLA response times. Customers with NVIDIA AI Enterprise licenses get prioritized fixes.

We would need to try NVIDIA AI Enterprise before we can commit to buying it.

How can we run a POC?

NVIDIA offers two ways to try NVIDIA AI Enterprise.

1. NVIDIA offers a free 90-day trial license for customers who have existing compatible infrastructure. The 90-day eval licenses give customers access to Production Branches, management software (Base Command Manager Essentials for example), and Enterprise Support. This allows customers to test NVIDIA AI Enterprise on-premises, in their environment, and with their own data sets.

2. NVIDIA LaunchPad – This is an AI development and deployment trial program. It provides hands-on labs for AI practitioners and IT staff. This offers customers a 3-day access to try out an AI workflow and experience the NVIDIA software.

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

What do you get with NVIDIA AI Enterprise Support?

NVIDIA AI Enterprise includes business standard support

Support availability

- Cases accepted via the web portal and email for 24/7 support
- Escalation support during a customer's local business hours (9:00 a.m.-5:00 p.m., Monday-Friday)
- Support provided by NVIDIA experts and engineers for timely resolution
- Enhanced partner interaction on issues

Remote support

- Remote hardware and software support

Software updates and upgrades

- Access to the latest NVIDIA AI Enterprise software suite
- Maintenance release and security fixes
- Priority notifications and distribution

Long-Term Support (LTS)

- 3 years for LTS branches

Portal access benefits

- Full access to a knowledge base via the portal
- Insight into which knowledge base articles, tips, and tutorials are trending and have proven most valuable
- View of current and past issue history
- Salesforce.com features, such as intelligent knowledge base prompts that support enhanced search capabilities

Service terms

- Included with NVIDIA AI Enterprise, whether purchased by itself, with NVIDIA H100 Tensor Core GPUs, or with NVIDIA DGXTM systems

Business-critical support

Increased availability

- All support referenced in Business-Standard Support
- Live NVIDIA agent access 24/7

Faster responses

- One-hour initial response for Severity 1 issues 24/7
- Two-hour initial response for Severity 2 issues 24/7

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

What to expect from NVIDIA AI Enterprise support response times

Understanding response times to manage customer expectations

		Business Standard	Business Critical
Severity Level	**Definition**	**Initial Response Time**	**Initial Response Time**
Catastrophic Severity 1	• The situation has halted your business. • The service or critical functions are unavailable or unusable and no workarounds exist.	4 Business Hours Local Business Hours	1 Hour 24/7
Severe Severity 2	• The situation has impacted your business. • The service or important functions are not working as expected or require workarounds, and no procedural workaround exists.	4 Business Hours Local Business Hours	2 Hours 24/7
Moderate Severity 3	• Non-critical issues that are intermittent or can be addressed using workarounds, but business continues to function.	4 Business Hours Local Business Hours	4 Business Hours Local Business Hours
Minor Severity 4	• Cosmetic or other minor issues that do not cause any significant detrimental effects. • Incremental feature requests.	1 Business Day	1 Business Day

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA.

Licensing models

Understanding licensing terms and options to meet customer needs

- Offered as Subscription Model (One Perpetual options for 5-year term)
- Licensed Per GPU (or per Node if licensing a non-GPU server)
- Sold as 1-, 3-, or 5-year options
- AI Enterprise Licensing Model includes 8x5 Standard Support
- Available license to upgrade to 24x7 Enterprise Support
- Pricing Available for EDUCATION with approval from NVIDIA (must be an educational institution customer
  - Without approval from NVIDIA, orders with EDU SKUs will be rejected)
- ELA SKUs available for opportunities that are $500K+ (based on list price for the software SKU only; upgrade support SKU is not included in this program) for the initial purchase (NVIDIA approval required)
  - ELA SKUs scale from ELA1 at 25% discount from standard list price with minimum of $500K, to ELA2 at 30% discount from standard list price with minimum of $1.25M, to ELA3 at 35% discount from standard list price with minimum of $2M, to ELA4 at 40% discount from standard list price with minimum of $2.75M
  - 1 year subscription license does not qualify for ELA program

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

Dell SKU guide

Understanding Dell bundles and licensing for NVIDIA AI Enterprise

- Offered as subscription model (one perpetual option for 5 year term)
- Licensed per GPU
- Sold as 1, 3, or 5 year options
- Pricing available for EDUCATION with approval from NVIDIA (must be .edu customer)
- AI Enterprise licensing model includes support - renewed per term
- ELA SKUs available for opportunities that are $500K for the initial purchase (NVIDIA approval required)

NVIDIA AI Enterprise — Subscription Bundle (Per GPU)

Dell SKU	Unit price (USD)	NVIDIA part number	Description
AC566091	$4,500.00	731-AI7003+P3CMI12	NVIDIA AI Enterprise subscription per GPU 1 year includes Standard 8x5 Support
AC566092	$13,500.00	731-AI7003+P3CMI36	NVIDIA AI Enterprise subscription per GPU 3 year includes Standard 8x5 Support
AC566093	$18,000.00	731-AI7003+P3CMI60	NVIDIA AI Enterprise subscription per GPU 5 year includes Standard 8x5 Support

The SKU Sheet can be accessed on the NVIDIA AI Enterprise Software TEAMS Page.
Direct link to the SKU Sheet is here.

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

Dell SKU guide

Understanding Dell bundles and licensing for NVIDIA AI Enterprise

NVIDIA AI Enterprise — Perpetual Bundle (Per GPU)

Dell SKU	Unit price (USD)	NVIDIA part number	Description
AC566097	$22,500.00	731-AI7004+P3CMI60	NVIDIA AI Enterprise Perpetual License per GPU 5 year includes Standard 8x5 Support

NVIDIA AI Enterprise — Upgrade to 24x7 Support Services (Per GPU)

Dell SKU	Unit price (USD)	NVIDIA part number	Description
AC566099	$1,100.00	731-AI7003+P3CMI12	Upgrade to 24x7 Enterprise Support Services for NVIDIA AI Enterprise per GPU 1 year
AC566100	$3,000.00	731-AI7003+P3CMI36	Upgrade to 24x7 Enterprise Support Services for NVIDIA AI Enterprise per GPU 3 year
AC566101	$5,000.00	731-AI7003+P3CMI60	Upgrade to 24x7 Enterprise Support Services for NVIDIA AI Enterprise per GPU 5 year

DELL Technologies

Prepare

Engage

One-click asset library

NVIDIA

Assets

Key resources to support your sales efforts

Sales reference assets

- NVIDIA AI Enterprise Software TEAMS page >
- General NVIDIA AI Enterprise documentation >
- NVIDIA AI Enterprise supported platforms >
- NVIDIA AI Enterprise battlecard >

Services and support

- NVIDIA AI Enterprise Support Services Brief data sheet >
- NVIDIA AI Enterprise software & support matrix >

NVIDIA AI Enterprise 90-day trial program

- NVIDIA AI Enterprise trial programs >

Enable customers to evaluate NVIDIA AI Enterprise for free based on their existing infrastructure and time commitment.

Customer-facing sales assets

- NVIDIA AI Enterprise customer deck >

Primary customer deck introducing AI practitioners and IT professionals to NVIDIA AI Enterprise, covers AI market trends, generative AI, data workflows, MLOps, and essential backup and resources.

- NVIDIA NIM customer deck >

This deck highlights NVIDIA NIM's advantages, including its seamless deployment of AI models with industry-standard APIs, domain-specific code, and optimized inference engines. It also covers support for custom models, cost efficiency, and the platform's robust security and performance features.

For assistance with these conversations with the customer, please reach out to the AI Assist team at ai.assist@dell.com.

One-click asset library

Download library

All assets listed below have been compiled into a single, downloadable PDF.

Prepare

Target audience

- Targeting and engaging customers 1-sheet >

Sales enablement assets

- Sales script >
- Overcoming objections script >

Platforms

- General NVIDIA AI Enterprise documentation >
- Dell NVIDIA platform >
- NVIDIA AI Enterprise supported platforms >

Software & support

- NVIDIA AI Enterprise software and support matrix >
- NVIDIA AI Enterprise support included >

Licensing

- Dell SKU guide >
- Trial programs >

Engage

Customer-facing sales assets

- NVIDIA AI Enterprise customer deck >
- NVIDIA NIM customer deck >


DELL Technologies


NVIDIA.