NVIDIA Software Product Guide
Product Guide

Enterprise adoption of AI is rapidly accelerating, but many organizations face challenges moving from pilot projects to full-scale production. Lenovo, in partnership with NVIDIA, delivers a comprehensive portfolio of software solutions designed to overcome these hurdles.

NVIDIA's GPU software stack consists of:

- **NVIDIA vGPU** – Enables scalable GPU virtualization for VDI and simulation.
- **NVIDIA AI Enterprise** – Offers a secure, production-grade AI platform.
- **Run:ai** – Orchestrates GPU workloads for maximum efficiency.
- **NVIDIA Omniverse Enterprise** – Powers collaborative 3D simulation and digital twin environments.
- **NVIDIA AI Workbench** (included as part of NVIDIA AI Enterprise) – Enables developers to harness the power of Lenovo's Data science Workstations for AI development

While these solutions are optimized to run on Lenovo's AI-ready infrastructure, the software is fully supported by NVIDIA Enterprise Support and Services Together, Lenovo and NVIDIA provide a path from experimentation to enterprise-grade AI, ensuring better ROI, reduced deployment times, and scalable innovation.

Did you know?

NVIDIA AI Enterprise software, supported on Lenovo servers, provides an end-to-end, cloud-native suite of AI and data analytics frameworks optimized for enterprise use. Recognized as a leader in GPU-accelerated computing, NVIDIA powers more than 40,000 companies worldwide with its AI platforms and has been adopted by over 90% of Fortune 500 companies to drive innovation. With more than 4 million NVIDIA AI Enterprise licenses distributed globally, organizations count on this software to deliver enterprise-grade reliability, security, and performance for their most demanding AI workloads.

Lenovo and NVIDIA

As AI transforms industries, from healthcare and finance to retail and manufacturing, enterprises need a reliable infrastructure and software stack to develop, deploy, and scale AI solutions.

Lenovo and NVIDIA address this demand by offering a tightly integrated hardware-software ecosystem that:

- Ensures consistent GPU access for AI teams.
- Improves infrastructure visibility and utilization.
- Accelerates deployment of GenAI, RAG, and digital twin solutions.
- Delivers enterprise-grade support and governance.

This product guide explores how NVIDIA software solutions run seamlessly on Lenovo platforms to unlock business value.

NVIDIA vGPU

NVIDIA Virtual GPU (vGPU) technology enables the virtualization of physical GPUs so they can be shared among multiple virtual machines (VMs) or containers. Unlike traditional GPU passthrough, which assigns an entire GPU to a single user or workload, vGPU allows multiple users to access a fraction of the GPU's power, enabling better resource efficiency and workload density.

When deployed on Lenovo ThinkSystem servers, vGPU provides a high-performance, scalable platform for powering virtual desktop infrastructure (VDI), simulation workloads, 3D rendering, and AI-powered design applications.

The key capabilities of NVIDIA Virtual GPU offerings include the following:

- **GPU Partitioning** – Allocates dedicated slices of GPU memory and compute cores per VM or container.
- **Support for AI Workloads** – Enables virtualized access to AI acceleration for model training and inference.
- **Application Certification** – Supports a broad ecosystem of ISVs including AutoDesk, Siemens, Adobe, and Dassault Systèmes.
- **Security & Isolation** – Keeps user data and workloads securely isolated, critical for regulated environments.
- **Remote Work Enablement** – Provides high-performance desktops to remote users with full GPU acceleration.

Topics in this section:

Choosing the Right NVIDIA vGPU Software License

NVIDIA vGPU software comes in three licensing tiers, each aligned to user profiles, workloads, and performance needs. Selecting the correct license ensures optimal performance and cost-efficiency.

Key recommendations:

- **Use vApps** if you are delivering apps to multiple users with shared OS instances.
- **Choose vPC** when individual, persistent desktops are needed with lightweight GPU acceleration.
- **Deploy vWS** for power users in industries like architecture, automotive, and M&E.

The following table shows the details of each recommendation, the target user, and ideal workloads for each license type.

Table 1. Choosing the Right NVIDIA vGPU Software License

License Type	Description	Recommended GPU	Target User	Ideal Workloads
NVIDIA Virtual Applications (vApps)	Application streaming via Remote Desktop Session Host (RDSH).	A2, A16	Knowledge users (shared systems)	Multi-user OS sessions, remote app access
NVIDIA Virtual PC (vPC)	Full virtual desktop infrastructure (VDI) for office apps, browsers, and multimedia.	A16, L4	Mainstream knowledge workers	PC-like VDI experience, office productivity
NVIDIA RTX Virtual Workstation (vWS)	High-end GPU acceleration for professional graphics and compute workloads.	L40, L40S, RTX Pro 6000 Blackwell SE	Pro Viz and technical professionals	CAD, CAE, simulation, rendering, modeling

Lenovo part numbers for vGPU

All tiers are licensed per Concurrent User (CCU) and can be centrally managed.

The following table lists the licensing options

Table 2. Part numbers

License	Description
Subscription (1–5 Years)	Includes access to vGPU drivers, software stack, and 24x7 support.
Perpetual + SUMS	Permanent license with a mandatory 5 years Sums contract
EDU & Inc Variants	Discounted licensing models for education and non-profits.

Perpetual licenses: Perpetual licenses cannot be sold as standalone products, you must add a 5 years SUMS contract.

The following license types are offered:

- **Perpetual license**

  A non-expiring, permanent software license that can be used on a perpetual basis without the need to renew. For each perpetual license, customers are also required to purchase a 5-year SUMS support contract. Without this contract, the perpetual license cannot be ordered.

- **Annual subscription**

  A software license that is active for a fixed period as defined by the terms of the subscription license, typically yearly. The subscription includes Support, Upgrade and Maintenance (SUMS) for the duration of the license term.

- **Concurrent User (CCU)**

  A method of counting licenses based on active user VMs. If the VM is active and the NVIDIA vGPU software is running, then this counts as one CCU. A vGPU CCU is independent of the connection to the VM.

The following table lists the ordering part numbers and feature codes.

Table 3. NVIDIA vGPU Software

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
NVIDIA vApps			
7S020004WW	B1MQ	711-VAP002+P3CMI12	NVIDIA vApps Subscription License 1 Year, 1 CCU
7S020005WW	B1MR	711-VAP002+P3CMI36	NVIDIA vApps Subscription License 3 Years, 1 CCU
7S02003DWW	S832	711-VAP002+P3CMI48	NVIDIA vApps Subscription License 4 Years, 1 CCU
7S02003EWW	S833	711-VAP002+P3CMI60	NVIDIA vApps Subscription License 5 Years, 1 CCU
7S020046WW	SDHB	711-VAP001+P3CMI00	NVIDIA vApps Perpetual License, 1 CCU
7S020003WW	B1MP	711-VAP001+P3CMI00	NVIDIA vApps SUMS ONLY 5Yr, 1 CCU (required for perpetual license)
NVIDIA vPC			
7S02000AWW	B1MW	711-VPC022+P3CMI12	NVIDIA vPC Subscription License 1 Year, 1 CCU
7S02000BWW	B1MX	711-VPC022+P3CMI36	NVIDIA vPC Subscription License 3 Years, 1 CCU
7S02003FWW	S834	711-VPC022+P3CMI48	NVIDIA vPC Subscription License 4 Years, 1 CCU
7S02003GWW	S835	711-VPC022+P3CMI60	NVIDIA vPC Subscription License 5 Years, 1 CCU
7S020047WW	SDHC	711-VPC021+P3CMI00	NVIDIA vPC Perpetual License, 1 CCU
7S020009WW	B1MV	711-VPC021+P3CMI00	NVIDIA vPC SUMS 5Yr ONLY, 1 CCU (required for perpetual license)
NVIDIA RTX vWS			
7S02000GWW	B1N2	711-DWS022+P3CMI12	NVIDIA RTX vWS Subsc Lic 1Yr 1 CCU
7S02000HWW	B1N3	711-DWS022+P3CMI36	NVIDIA RTX vWS Subscription License 3 Years, 1 CCU
7S02000XWW	S6YJ	711-DWS022+P3CMI48	NVIDIA RTX vWS Subscription License 4 Years, 1 CCU
7S02000YWW	S6YK	711-DWS022+P3CMI60	NVIDIA RTX vWS Subscription License 5 Years, 1 CCU
7S02000MWW	B1N7	711-DWS022+P3EDI12	NVIDIA RTX vWS EDU Subscription License 1 Year, 1 CCU
7S02000NWW	B1N8	711-DWS022+P3EDI36	NVIDIA RTX vWS EDU Subscription License 3 Years, 1 CCU
7S02003BWW	S830	711-DWS022+P3EDI48	NVIDIA RTX vWS EDU Subscription License 4 Years, 1 CCU
7S02003CWW	S831	711-DWS022+P3EDI60	NVIDIA RTX vWS EDU Subscription License 5 Years, 1 CCU
7S020048WW	SDHD	711-DWS021+P3CMI00	NVIDIA RTX vWS Perpetual License, 1 CCU

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
7S02000FWW	B1N1	711-DWS021+P3CMI00	NVIDIA RTX vWS SUMS ONLY 5Yr, 1 CCU (required for perpetual license)
7S020049WW	SDHE	711-DWS021+P3EDI00	NVIDIA RTX vWS EDU Perpetual License, 1 CCU
7S02000LWW	B1N6	711-DWS021+P3EDI00	NVIDIA RTX vWS EDU SUMS ONLY 5Y, 1CCU (required for perpetual license)
NVIDIA RTX vWS Support & Services			
7S020015WW	S6YS	712-DWSA24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS Production SUMS, 1CCU, 1 Year
7S02005CWW	SDZB	712-DWSA24+P3CMI60	24X7 Support Services for NVIDIA RTX vWS Production SUMS 1CCU 5 Years
7S020016WW	S6YT	712-DWSA24+P3EDI12	24X7 Support Services for NVIDIA RTX vWS Production SUMS, 1CCU, EDU, 1 Year
7S02005DWW	SDZC	712-DWSA24+P3EDI60	24X7 Support Services for NVIDIA RTX vWS Production SUMS 1CCU EDU 5 Years
7S02005EWW	SDZD	712-DWSB24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS SUMS 4 CCU 1 Year
7S020017WW	S6YU	712-DWSD24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS Subscription License, 1CCU, 1 Year

Recommended Lenovo vGPU platforms

The following Lenovo ThinkSystem servers are recommended platforms for running NVIDIA vGPU software:

- **ThinkSystem SR645 V3** – 1U Dual-AMD processor server ideal for dense GPU deployment.
- **ThinkSystem SR665 V3** – 2U Dual-AMD processor server with support for up to 4 double-wide GPUs.
- **ThinkSystem SR650 V4** – Best in class mainstream server with latest Intel Xeon 6 processors and NVIDIA GPUs

vGPU Deployment scenarios

The following are example deployment scenarios for vGPU applications:

- **AI-Enabled VDI Labs:**

  Virtualize GPUs to support multiple data scientists working on training/inference tasks without requiring dedicated GPU nodes.

- **Global Design Collaboration:**

  Enable real-time, GPU-accelerated design reviews and simulations from anywhere using Omniverse-ready virtual desktops.

- **Secure Remote Access for Regulated Industries:**

  Healthcare, financial services, and government sectors benefit from centralized data control and encrypted access to GPU-accelerated workstations.

vGPU Business value

The use of vGPU applications on ThinkSystem servers has the following business value:

- Lower TCO – Reduces the need for dedicated physical GPU workstations.
- Higher GPU Utilization – Shares underutilized GPU capacity across users.
- Workforce Flexibility – Supports hybrid and remote work without compromising performance.
- Simplified IT Operations – Centralized management of GPU resources and user profiles.
- Secure AI Enablement – Delivers AI capabilities within secure, virtual environments.

vGPU Deployment considerations

Before deployment of vGPU, organizations should carefully evaluate their infrastructure readiness through below requirements:

- **Hypervisor Support:** Compatible with VMware vSphere, Citrix Hypervisor, and Red Hat Virtualization.
- **Management Tools:** Integrated with NVIDIA vGPU Manager and Lenovo XClarity.
- **Network Requirements:** Recommend low-latency connectivity (10–25GbE) for VDI performance.
- **Storage Consideration:** Use NVMe SSDs or hybrid storage for fast user data access.

NVIDIA AI Enterprise

NVIDIA AI Enterprise (NVAIE) is a production-grade, cloud-native software platform purpose-built to accelerate the development, training and deployment of AI, machine learning (ML), and generative AI models in enterprise environments.

The best companion for the Lenovo Hybrid AI 285 platform and AI nodes from the ThinkSystem family, NVAIE simplifies AI infrastructure by delivering over 50 frameworks, pre-trained models, and enterprise-grade support —enabling organizations to confidently scale from pilot to production, whether on-prem or hybrid.

The key capabilities of NVAIE offerings include the following:

- **Full-Stack AI Platform** – Includes frameworks like TensorFlow, PyTorch, RAPIDS, XGBoost, Triton Inference Server, and more.
- **Generative AI Enablement** – Provides pre-built pipelines for Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), and fine-tuning.
- **Enterprise-Grade Support** – Backed by NVIDIA with 24/7 support, performance optimization, and validated infrastructure.
- **Security & Governance** – Supports secure, multi-tenant environments with integrated access control and policy enforcement.
- **Hybrid-Cloud Ready** – Deployable across VMware, Red Hat OpenShift, Kubernetes, or bare metal, with full container support.

Topics in this section:

NVAIE Deployment Architecture

Refer to **Lenovo Hybrid AI 285 Platform Guide** to get deeper details on the deployment architecture. NVAIE supports flexible, enterprise-ready architectures. Options include:

1. **VMware vSphere with NVIDIA GPU Operator**

   For enterprises with virtualized infrastructure and on-prem security requirements.

2. **Bare Metal Kubernetes (K8s)**

   For performance-focused AI workloads with containerized development workflows.

3. **Hybrid Cloud with OpenShift / Tanzu**

   For multi-location or hybrid deployment of AI models and pipelines.

Supported NVIDIA GPUs

- **H100 / H200 / B200** (HGX or PCIe) – LLMs, training, and fine-tuning
- **L40S** – Balanced performance for training and inference
- **RTX PRO 6000 Blackwell Server Edition** – Multi-model-class GenAI and model development
- **L4** – Energy-efficient GPU for inference and analytics

Why NVIDIA AI Enterprise on Lenovo?

Deploying AI at scale in the enterprise is no longer a vision. it's a competitive necessity. But building and operationalizing an AI platform across distributed teams, complex IT environments, and security-sensitive applications remains one of the most significant challenges for organizations today.

NVIDIA AI Enterprise (NVAIE) is engineered to bridge this gap, delivering a fully optimized, production-ready AI software suite built for real-world enterprise needs. When combined with Lenovo's AI-ready infrastructure, such as the ThinkSystem SR675 V3 SR680a, SR685 and our state of the art the SR780a Neptune liquid cooling system, the result is a solution that accelerates AI adoption while reducing risk and complexity.

Unlike open-source stacks that require heavy integration, configuration, and maintenance, NVAIE provides an end-to-end, SLA-supported software platform—including pretrained models, GPU-optimized frameworks, security features, and seamless MLOps integration.

This makes it the go-to choice for IT leaders seeking scalable AI governance, and for data scientists and developers who want to spend more time building models and less time troubleshooting environments.

Customer Success Spotlight: AISHA Transforms Medical Imaging with Lenovo + NVIDIA AI Enterprise

One powerful example of NVAIE in action is AISHA, a healthcare organization that dramatically enhanced patient care by deploying AI to analyze medical imaging. AISHA trained an AI model to analyze MRI scans using a Lenovo ThinkSystem SR675 V3 server powered by NVIDIA H100 NVL GPUs and the NVIDIA AI Enterprise software suite. The result: faster, more accurate diagnostics, delivered at scale.

? Watch the success story on YouTube:
AISHA | Empowering Healthcare with AI on Lenovo + NVIDIA

This deployment exemplifies how healthcare institutions can benefit from secure, scalable, production-grade AI without sacrificing compliance or operational control—especially when powered by Lenovo's GPU-accelerated servers and the NVAIE software stack.

Lenovo part numbers for NVAIE

Note: A 5-year NVAIE subscription is included with H100, H200, and B200 PCIe double-wide GPUs on Lenovo systems. Customer can redeem the license thru this link: https://www.nvidia.com/en-us/data-center/activate-license/

All NVIDIA AI Enterprise subscriptions include NVIDIA Business Standard Support and can be purchased as either a perpetual license, as an annual or multi-year subscription, and on an hourly consumption basis via cloud marketplaces. NVIDIA AI Enterprise with perpetual licenses must be purchased in conjunction with five-year support services. A one-year support service is also available for renewals.

The following YouTube video playlists provide additional information:

- NVAIE Registration w/Hardware purchases
- NVAIE Registration for Standalone purchases

The following table lists the ordering part numbers and feature codes.

Table 4. NVIDIA AI Enterprise Software (NVAIE)

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
AI Enterprise Perpetual License			
7S02001BWW	S6YY	731-AI7004+P3CMI60	NVIDIA AI Enterprise Perpetual License and Support per GPU Socket, 5 Years
7S02001EWW	S6Z1	731-AI7004+P3EDI60	NVIDIA AI Enterprise Perpetual License and Support per GPU Socket, EDU, 5 Years
AI Enterprise Subscription License			
7S02001FWW	S6Z2	731-AI7003+P3CMI12	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 1 Year
7S02001GWW	S6Z3	731-AI7003+P3CMI36	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 3 Years
7S02001HWW	S6Z4	731-AI7003+P3CMI60	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 5 Years
7S02001JWW	S6Z5	731-AI7003+P3EDI12	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 1 Year
7S02001KWW	S6Z6	731-AI7003+P3EDI36	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 3 Years
7S02001LWW	S6Z7	731-AI7003+P3EDI60	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 5 Years
AI Enterprise Support Services			
7S02001MWW	S6Z8	731-AI7007+P3CMI12	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, 1 Year
7S02001NWW	S6Z9	731-AI7007+P3CMI36	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, 3 Years
7S02001PWW	S6ZA	731-AI7007+P3CMI60	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, 5 Years
7S02001QWW	S6ZB	731-AI7007+P3EDI12	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, EDU, 1 Year
7S02001RWW	S6ZC	731-AI7007+P3EDI36	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, EDU, 3 Years
7S02001SWW	S6ZD	731-AI7007+P3EDI60	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, EDU, 5 Years

Find more information in the NVIDIA AI Enterprise Sizing Guide.

NVIDIA Run:ai

As AI adoption accelerates, many organizations encounter a new bottleneck: underutilized or fragmented GPU resources across teams and workloads. NVIDIA Run:ai resolves this by introducing an intelligent orchestration and scheduling platform purpose-built for AI infrastructure. It enables multi-tenant GPU sharing, workload prioritization, and centralized resource management, dramatically improving ROI on AI infrastructure.

When deployed on Lenovo ThinkSystem servers, Run:ai transforms traditional GPU clusters into fully virtualized AI infrastructure, capable of allocating GPUs dynamically based on policies, project needs, or user demand: on-prem, in hybrid environments, or across Kubernetes clusters.

The key capabilities of Run:ai include the following:

- **Dynamic GPU Scheduling** – Allocates GPU resources to users and projects in real time—full, fractional, or virtualized.
- **Fair Share Quotas** – Ensures teams get guaranteed GPU access without overprovisioning or idle capacity.
- **Multi-Tenancy** – Supports multiple departments or teams with isolated workloads, policies, and quotas.
- **Kubernetes-Native Integration** – Runs seamlessly on K8s, OpenShift, or any CNCF-compliant container environment.
- **Visibility & Dashboards** – Offers real-time metrics, cost tracking, and usage visualization for IT and leadership.
- **AI Workload Prioritization** – Automatically prioritizes high-value, production workloads over experimentation.

Topics in this section:

How Run:ai works on Lenovo infrastructure

Implementing Run:ai on Lenovo infrastructure has been simplified through our partnership. Following steps below would enable compute enhancements through Run:ai

- Users submit jobs via familiar tools (Kubeflow, JupyterHub, CLI).
- Run:ai schedules those jobs to available GPUs based on defined policies.
- Admins use the Run:ai dashboard to monitor usage, create policies, and allocate quotas.
- Works across ThinkSystem SR685a V3, SR675 V3, and SR650a V4 servers, enabling shared, multi-user AI platforms.

Recommended Lenovo Run:ai platforms

The following Lenovo ThinkSystem servers are recommended platforms for running NVIDIA Run:ai software:

- **ThinkSystem SR685a V3 and SR680a V3** – High-performance GPU server, ideal for large multi-tenant clusters
- **ThinkSystem SR675 V3** – AI training, inference, and hybrid orchestration environments
- **ThinkSystem SR650a V4** – Balanced compute and GPU density for shared team clusters

Supported GPUs

- **H200 / B200** (HGX or PCIe) – Large-scale training and inferencing
- **L40S** – Versatile performance for model training, RAG, and GenAI
- **RTX PRO 6000 Blackwell Server Edition** – Workstation-grade GPU for experimentation and shared inference

Key use cases for Run:ai

In the table below we have mapped out the personas that can benefit from use cases of Run:ai improved GPU orchestration and the associated business impact.

Table 5. Key use cases

Persona	Scenario	Business Impact
Data Scientists	Need access to GPUs for model training without delays or bottlenecks	Instant, on-demand GPU access via self-service
MLOps Teams	Require efficient resource pooling for multiple pipelines	Simplified orchestration, pipeline scaling
IT Managers	Must monitor and control infrastructure costs and usage	Dashboards, policy-based controls, chargeback models
Executives / Leaders	Need visibility into AI investment and ROI	Analytics on usage, productivity, and business impact

Lenovo part numbers for Run:ai

The following table lists the licensing options.

Table 6. NVIDIA Run:ai

Part number	Feature 7S02CTO1WW	NVIDIA part number	Description
Software subscription			
7S02004UWW	SDYT	744-RA7001+P3CMI12	NVIDIA Run:ai Subscription per GPU 1 Year
7S02004XWW	SDYW	744-RA7001+P3CMI36	NVIDIA Run:ai Subscription per GPU 3 Years
7S020050WW	SDYZ	744-RA7001+P3CMI60	NVIDIA Run:ai Subscription per GPU 5 Years
7S02004VWW	SDYU	744-RA7001+P3EDI12	NVIDIA Run:ai Subscription per GPU EDU 1 Year
7S02004YWW	SDYX	744-RA7001+P3EDI36	NVIDIA Run:ai Subscription per GPU EDU 3 Years
7S020051WW	SDZ0	744-RA7001+P3EDI60	NVIDIA Run:ai Subscription per GPU EDU 5 Years
7S02004WWW	SDYV	744-RA7001+P3INI12	NVIDIA Run:ai Subscription per GPU INC 1 Year
7S02004ZWW	SDYY	744-RA7001+P3INI36	NVIDIA Run:ai Subscription per GPU INC 3 Years
7S020052WW	SDZ1	744-RA7001+P3INI60	NVIDIA Run:ai Subscription per GPU INC 5 Years
Support Services subscription			
7S020053WW	SDZ2	744-RA7002+P3CMI12	24x7 Support Services for NVIDIA Run:ai Subscription per GPU 1 Year
7S020056WW	SDZ5	744-RA7002+P3CMI36	24x7 Support Services for NVIDIA Run:ai Subscription per GPU 3 Years
7S020059WW	SDZ8	744-RA7002+P3CMI60	24x7 Support Services for NVIDIA Run:ai Subscription per GPU 5 Years
7S020054WW	SDZ3	744-RA7002+P3EDI12	24x7 Support Services for NVIDIA Run:ai Subscription per GPU EDU 1 Year
7S02005AWW	SDZ9	744-RA7002+P3EDI60	24x7 Support Services for NVIDIA Run:ai Subscription per GPU EDU 5 Years
7S020057WW	SDZ6	744-RA7002+P3EDI36	24x7 Support Services for NVIDIA Run:ai Subscription per GPU EDU 3 Years
7S020055WW	SDZ4	744-RA7002+P3INI12	24x7 Support Services for NVIDIA Run:ai Subscription per GPU INC 1 Year
7S020058WW	SDZ7	744-RA7002+P3INI36	24x7 Support Services for NVIDIA Run:ai Subscription per GPU INC 3 Years
7S02005BWW	SDZA	744-RA7002+P3INI60	24x7 Support Services for NVIDIA Run:ai Subscription per GPU INC 5 Years

Run:ai + NVAIE integration: Unified E2E AI Platform at Scale

When Run:ai and NVIDIA AI Enterprise (NVAIE) are combined, they create a comprehensive AI platform that spans the entire AI model development and training lifecycle. This platform extends from experimentation to full-scale production, while ensuring maximum GPU efficiency, governance, and enterprise-grade support. Certified and validated by NVIDIA and Lenovo, the joint solution empowers organizations to achieve faster time-to-value and higher ROI from their GPU investments.

Together, Run:ai and NVAIE offer the following key capabilities:

- Standardized AI solutions and frameworks (via NVAIE)
- Smart Dynamic GPU scheduling and resource allocation policy control (via Run:ai)
- Multi-user, multi-project orchestration with real-time monitoring and prioritization
- Maximized resource utilization for AI model training within the shortest time
- Future proof solution for all upcoming NVIDIA integration tools

This integration is certified and validated on Lenovo ThinkSystem infrastructure, allowing IT and AI leaders to confidently scale operations across teams and geographies.

Real-World Integration Use Case: AI Center of Excellence

Scenario: A global manufacturer is building an internal AI Center of Excellence (CoE) to support cross-functional teams developing AI models for predictive maintenance, generative design, and visual inspection.

Challenge: Multiple teams using different tools (PyTorch, RAPIDS, Hugging Face) want access to limited GPUs.

- The infrastructure team needs visibility and control to avoid capacity issues.
- Business leadership requires clear ROI tracking for AI investments.

Solution: The organization deployed:

- **NVIDIA AI Enterprise** to standardize the software environment with GPU-accelerated libraries, pretrained models, and AI workflows.
- **Run:ai** to manage GPU sharing, enforce team quotas, and prioritize workloads based on business criticality.

Results:

- **AI Time-to-Production reduced by 50%** – Standardized pipelines + instant GPU access
- **80% GPU utilization** – Up from 40% before Run:ai orchestration
- **Secure, multi-tenant access** – Role-based controls + NVAIE container hardening
- **Transparent ROI tracking** – Dashboards for execs and IT with usage/cost reporting

NVIDIA Omniverse Enterprise

NVIDIA Omniverse Enterprise is a real-time collaboration and simulation platform that enables engineers, designers, and AI developers to create physically accurate digital twins, simulate complex systems, and streamline product development workflows—all in a shared 3D environment.

When deployed on Lenovo ThinkStation workstations and ThinkSystem servers with NVIDIA RTX GPUs, Omniverse accelerates workflows for industries ranging from manufacturing and AEC (architecture, engineering & construction) to automotive, energy, and telecom.

It enables multi-disciplinary teams to collaborate simultaneously across design tools and geographies, while AI-enhanced simulation drives faster innovation cycles, better decision-making, and reduced time to market.

The key capabilities of NVIDIA Omniverse Enterprise include the following:

- **USD-Based Collaboration** – Uses Pixar's Universal Scene Description (USD) to enable real-time, multi-tool collaboration across design pipelines.
- **Real-Time Physically Accurate Simulation** – Supports ray tracing, physics, materials, and environment rendering for true-to-reality simulation.
- **AI-Enabled Digital Twins** – Integrates with AI workflows to create intelligent, interactive environments for robotics, inspection, predictive maintenance.
- **Connectors for Industry Tools** – Native plugins for Autodesk, Revit, Rhino, SolidWorks, PTC Creo, Unreal Engine, Blender, and more.
- **Multi-User Collaboration** – Distributed teams can co-design and co-simulate from anywhere in the world.
- **Enterprise Deployment Support** – Includes scalability, security, licensing, and remote access via VDI and NVIDIA vGPU.

Topics in this section:

Certified Lenovo Systems for Omniverse Enterprise

To deliver the demanding performance required for real-time 3D collaboration, simulation, and digital twin development, NVIDIA Omniverse Enterprise relies on a powerful backend infrastructure. Lenovo offers a portfolio of certified mobile, desktop, and server-class systems purpose-built to run Omniverse workloads at scale.

The following table is a summary of NVIDIA-certified Lenovo systems optimized for Omniverse Enterprise, supporting a wide range of deployment scenarios—from workstation-based development to GPU-accelerated data center deployments.

Table 7. Certified Lenovo Systems

Form Factor	System	CPU	System Memory	Boot Drive	Data Drive	Networking	GPU
Mobile Workstation	ThinkStation P16 MTM: 21FA002EUS	Intel Core i9-13850HX	64GB DDR5	1TB M.2 NVMe SSD	Not Required	WiFi 6E	1× RTX 5000 Ada Mobile
Desktop Workstation	ThinkStation PX MTMs: 30EUxxCTOWW /30EVxxCTOWW	2× Xeon Silver 4416+	256GB ECC DDR5	1TB M.2 NVMe SSD	2× 2TB M.2 NVMe SSDs	10GbE + NVIDIA CX6 Dx Active*	1× RTX 6000 Ada (max 4x)
Nucleus Server	ThinkSystem SR655 V3	1× Xeon 3.6GHz / 16 cores	96GB ECC DDR5	512GB M.2 NVMe SSD	2× 1TB M.2 NVMe SSDs	2× CX7 (2×200GB)	Not applicable
Omniverse OVX Node – 4 GPU	ThinkSystem SR675 V3 CTO: 7D9ROVX3WW	2× AMD Genoa 32C	384GB DDR5 ECC	1TB M.2 NVMe SSD	2× 4TB NVMe SSD	2× CX7 (2×200GB) + BF-3 DPU*	4× NVIDIA L40S
Omniverse OVX Node – 8 GPU	ThinkSystem SR675 V3 CTO: 7D9ROVX3WW	2× AMD Genoa 64C	768GB DDR5 ECC	1TB M.2 NVMe SSD	4× 4TB NVMe SSD	4× CX7 (2×200GB) + BF-3 DPU*	8× NVIDIA L40S

* NVIDIA ConnectX-6 and ConnectX-7 network interfaces and BlueField-3 (BF-3) DPUs are recommended for enhanced throughput and data center-grade orchestration.

Omniverse Enterprise Deployment scenarios

The following are example deployment scenarios for Omniverse Enterprise:

- **Single-user simulation & design** – ThinkStation PX or P16 mobile workstation
- **Small team collaboration** – ThinkSystem SR655 V3 (Nucleus Server) with centralized data access
- **Large-scale digital twin deployment** – OVX L40S (4-8 GPU) on SR675 V3 with NVIDIA Omniverse Nucleus + Enterprise Suite

Use Case: Digital Twin in Automotive

Scenario: An automotive design firm uses Omniverse Enterprise to build a full digital twin of its latest vehicle prototype. Design, aerodynamics, AI/ADAS systems, and manufacturing engineers all collaborate in real-time —from different locations.

Deployment:

- **ThinkStation P8** for simulation and local modeling
- **ThinkSystem SR675 V3** for rendering and AI-based environmental simulation
- **NVIDIA RTX PRO 6000 Blackwell Server Edition GPUs** deliver ultra-fast ray-traced feedback

Outcome:

- 40% reduction in time to prototype
- Real-time performance testing using AI-generated simulation data
- Seamless integration of CAD, Unreal Engine, and physics engines

Lenovo part numbers for Omniverse Enterprise

The following table lists the licensing options.

Table 8. NVIDIA Omniverse Software (OVE)

Part number	Feature 7S02CTO1WW	NVIDIA part number	Description
7S02003ZWW	SCX0	721-OV7006+P3CMI12	NVIDIA Omniverse Enterprise Subscription per GPU, 1 Year
7S020042WW	SCX3	721-OV7006+P3CMI36	NVIDIA Omniverse Enterprise Subscription per GPU, 3 Years
7S020044WW	SD5T	721-OV7006+P3CMI60	NVIDIA Omniverse Enterprise Subscription per GPU, 5 Year
7S020041WW	SCX2	721-OV7006+P3INI12	NVIDIA Omniverse Enterprise Subscription per GPU, INC, 1 Year
7S020040WW	SCX1	721-OV7006+P3EDI12	NVIDIA Omniverse Enterprise Subscription per GPU, EDU, 1 Year
7S020043WW	SCX4	721-OV7006+P3EDI36	NVIDIA Omniverse Enterprise Subscription per GPU, EDU, 3 Years
7S020045WW	SD5U	721-OV7006+P3EDI60	NVIDIA Omniverse Enterprise Subscription per GPU EDU, 5 Year

NVIDIA AI Workbench

NVIDIA AI Workbench is a GPU-accelerated, collaborative, development environment that streamlines AI workflows from experimentation to deployment. Certified on Lenovo infrastructure, it unifies container, environment, and GPU management with a simple interface for consistent, reproducible results. Developers can prototype locally, then scale projects seamlessly to data center or cloud environments in just a few clicks. With integrated access to popular repositories - AI Workbench accelerates collaboration, version control, and workload portability across teams and platforms.

Deployment

NVIDIA AI Workbench is built to enable simplified use of GPUs and easy deployment.

- Installation: Quick setup on Windows, Ubuntu Linux, and macOS.
- Dependencies: Automates setup of microservices, NVIDIA Enterprise Toolkit, Git, and other required components.
- Container Management: Uses microservices to ensure consistent AI environments across systems.

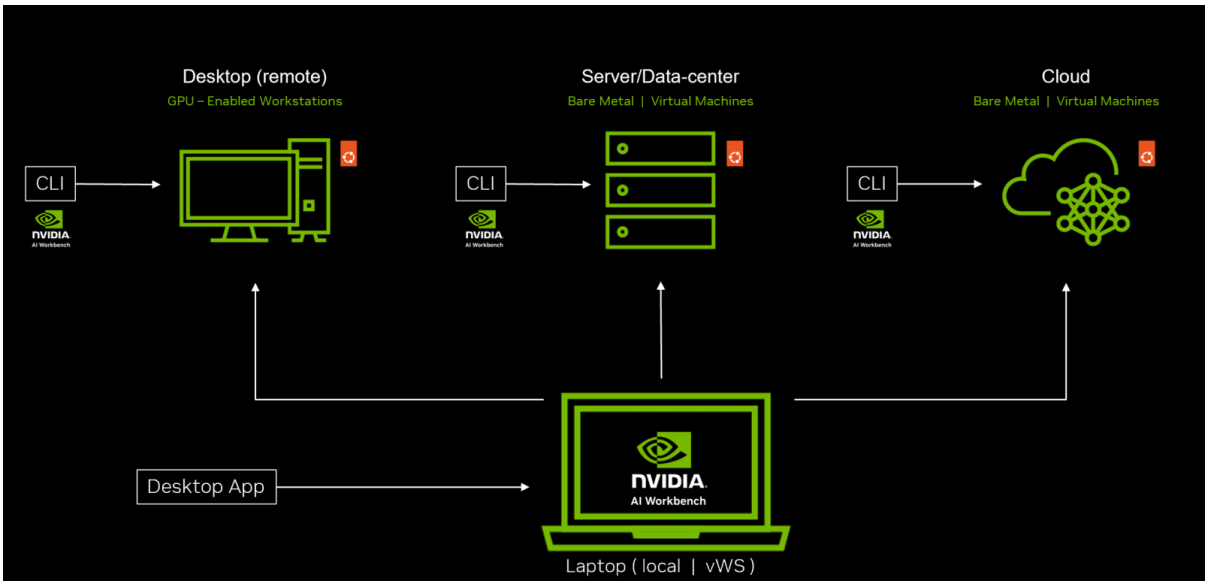

Figure 1. NVIDIA AI Workbench

Key features

The main differentiating features of NVIDIA AI Workbench include the following:

- Unified Development Experience: Manage projects, environments, and GPU resources from a single interface.
- Hybrid Computing: Seamlessly move workloads between local, data center, and cloud environments.
- Collaboration and Version Control: Integrated Git with real-time environment versioning. Share environments and codebases across teams and geographies.
- GPU Optimization: Dynamically allocates GPU resources for efficient multi-user and multi-project performance.
- Rapid Prototyping: Quickly create and iterate on AI models locally.

Frequently Asked Questions (FAQ)

FAQ topics:

General Software FAQs

Q1: What is included in the NVIDIA AI Enterprise (NVAIE) suite?
A: NVAIE includes over 50+ GPU-accelerated frameworks and tools such as TensorFlow, PyTorch, RAPIDS, Triton Inference Server, NeMo (LLMs), Riva (ASR/NLP), and TAO Toolkit for transfer learning. It comes with enterprise support and deployment documentation.

Q2: Can I deploy NVIDIA AI Enterprise in a fully on-prem environment?
A: Yes. NVAIE is designed for on-prem, hybrid, or air-gapped deployments using VMware vSphere, bare metal, or Kubernetes (including Red Hat OpenShift).

Q3: What Lenovo systems include NVAIE licensing?
A: NVAIE is bundled with a 5-year license on qualifying Lenovo servers using the PCIe/NVL variants of **H100, H200** or **B200** GPUs

Q4: How is NVAIE licensed?
A: It's licensed **per GPU socket**. Options include perpetual or subscription terms (1, 3, or 5 years), with EDU and INC variants available.

Run:ai FAQs

Q5: What problem does Run:ai solve?
A: Run:ai dynamically schedules GPU workloads across users, ensuring better utilization, multi-tenancy, and real-time quota enforcement. It prevents idle GPUs and infrastructure bottlenecks, especially in AI Centers of Excellence or multi-team environments.

Q6: Is Run:ai only for Kubernetes?
A: Yes, Run:ai is Kubernetes-native but can be layered on VMware environments running Tanzu or integrated with OpenShift. It works with Lenovo platforms that support containerized or virtualized workloads.

Q7: Is Run:ai included with any Lenovo system?
A: No, Run:ai is a standalone NVIDIA software license that must be selected separately (sold per GPU/year). Lenovo does offer bundle promotions with certain AI-ready servers.

Q8: What GPUs and Lenovo platforms are recommended for Run:ai?
A: Recommended platforms include SR675 V3, SR685, and SR650 V4. GPUs include L40S, H200, B200, and RTX 6000 Ada/Blackwell SE.

vGPU / VDI FAQs

Q9: What are the differences between vApps, vPC, and vWS licenses?
A:

- **vApps** – For application streaming via RDSH; ideal for delivering virtulized business applications to many users at scale.
- **vPC** – For full VDI with multimedia; supports rich graphics and video playback for everyday office productivity, multimedia and knowledge worker work
- **vWS (RTX Virtual Workstation)** – For high-end 3D rendering, CAD, CAE (L40s, A16); ideal for workstation level performances with GPU accelerators

Q10: Can vGPU be used for AI model training or inference?
A: Yes, especially when GPUs are virtualized with vWS. This enables AI/ML workloads in secure, centralized environments and is ideal for regulated industries.

Q11: What systems support vGPU deployments?
A: Lenovo ThinkSystem SR645, SR665, and ThinkStation P8 are optimized for vGPU use with L40S, A16, or A40 GPUs.

Omniverse Enterprise FAQs

Q12: What is Omniverse used for in an enterprise?
A: Omniverse enables real-time 3D collaboration, simulation, and digital twin development across industries like manufacturing, automotive, AEC, and telecom.

Q13: Is Omniverse GPU-intensive?
A: Yes. It requires powerful GPUs such as **RTX 6000 Ada**, **L40S**, or **Blackwell SE** to support ray tracing, simulation physics, and real-time rendering.

Q14: What Lenovo systems are certified for Omniverse?
A: Lenovo ThinkStation P16, PX (workstations), SR655 V3 (Nucleus Server), and SR675 V3 (OVX nodes for 4- or 8-GPU configurations).

Sales FAQs

Q15: Can I bundle multiple NVIDIA software licenses with a Lenovo server quote?
A: Yes. You can configure quotes with combinations of NVAIE, Run:ai, vGPU, and Omniverse licenses. Ensure proper SKUs are included per GPU type and license duration.

Q16: Are education discounts available?
A: Yes. NVIDIA offers **EDU and INC SKUs** across all software products. Lenovo also supports academic pricing for servers in qualifying institutions.

Q17: Who provides support — Lenovo or NVIDIA?
A: NVIDIA provides the **software support (8x5 and 24x7)** through NVIDIA Enterprise Support. Lenovo handles hardware support and integration.

Comparison – Run:ai and NVAIE vs Alternatives

The following table provides a comparison of the options available to customers categorized based on the key capabilities. In this analysis, all alternative solutions are categorized in last column. This includes in-house solutions by leveraging customized licenses like Kubeflow and Databricks, and open source libraries.

Table 9. Comparative analysis of the available options for implementing NVAIE solutions

Category	NVAIE + Run:ai	NVAIE	Run:ai	Alternative open-source/ in-house solutions
Target Use	Comprehensive E2E enterprise AIML solution	NVIDIA native frameworks to maximize efficiency & minimize cost	Intelligent orchestration solution to minimize idle GPU	Customized in-house solution
Scalability	E2E Future proof solution built for NVIDIA native ecosystem to maximize benefits	Over 50+ modular frameworks for enterprise scalability	Built for large scale resource utilization	Requires long-term expensive compatibility & scalability analysis with enterprises
Integration	Seamless integration with End to end NVIDIA ecosystem stack	Built-in support for MLOps. Containerized, documented, ready-to-deploy	Straightforward integration for NVIDIA GPU workload scheduling	Requiring manual comprehensive integration with ecosystem
GPU Resource Efficiency	Most effective GPU utilization solution with estimated 3x token throughput and compute cost savings up to 200M tokens/month or $1M/yr	Fully optimized with CUDA, cuDNN, TensorRT, Triton, MIG	Native; supports MIG, GPU slicing, and job-level GPU assignment	Would require manual performance tuning and maintenance
Flexibility	Future proof solution customizable for clients needs	List of 50+ frameworks to pick and choose from based on required functionality	Limited to latest Run:ai solutions and functionality	Customizable to all functionalities clients want to build
Security	Using NVIDIA native AgentIQ for security	Leveraging NVIDIA frameworks for matching enterprise security and compliance	Role-based access control, project isolation	Open-source libraries would put enterprise data at risk – security
Cost Efficiency	Predictable licensing option – preventing future extra costs	License model available	Cost efficient with a big ROI by saving idle GPU cost	High cost engineering and support expense required
Support	24/7 NVIDIA support with Lenovo bundle	24/7 NVIDIA support with Lenovo bundle	24/7 NVIDIA support with Lenovo bundle	No SLA – community support etc.
Monitoring	Run:ai standalone easy to use GPU utilization monitoring dashboard	Standalone metrics that can be customized to client's KPI through each NVAIE tools	Stand alone customizable easy-to-use dashboard	Implementing the basic metrics and measurement tool for tracking and customizing data
Time to Value	Days to weeks	Days to weeks	Immediately	Months

Key Takeaways:

- **NVAIE** is ideal for organizations that want a **secure, optimized, supported AI platform** without piecing together tools.

- **Run:ai** offers **enterprise-grade orchestration** that's purpose-built for AI teams—not generic DevOps.
- Open-source stacks are flexible but require deep technical knowledge, **ongoing integration**, and **no official support**.
- Competing tools like **Paperspace**, **K8s-native schedulers**, or **Slurm** often lack AI-focused features, creating operational friction

Seller training courses

The following sales training courses are offered for employees and partners (login required). Courses are listed in date order.

1. **Partner Technical Webinar - NVIDIA Software**
   2025-07-21 | 60 minutes | Employees and Partners

   In this 60-minute replay, Carlos Huescas, Lenovo, and Sandeep Brahmarouthu and Rob Magno of NVIDIA, presented the key software offerings of NVDIA AI Enterprise (NVAIE) and Run:ai, including a demo of Run:ai.

   Tags: Artificial Intelligence (AI)

   Published: 2025-07-21
   Length: 60 minutes

   Start the training:
   Employee link: Grow@Lenovo
   Partner link: Lenovo 360 Learning Center

   Course code: JUL1825
2. **Partner Technical Webinar - AI Vertical Spotlight Pt 2**
   2025-07-08 | 60 minutes | Employees and Partners

   In this 60-minute replay, we concluded the AI Vertical Spotlight (Pt 2) with our final two speakers. Peter Orban, AI Business Development Manager, discussed Financial and Banking, while Eric Skomra, Public Sector & Spaces AI Technologist, provided insights on State, Local, Education (SLED), and Smart Spaces.

   Tags: Artificial Intelligence (AI)

   Published: 2025-07-08
   Length: 60 minutes

   Start the training:
   Employee link: Grow@Lenovo
   Partner link: Lenovo 360 Learning Center

   Course code: JUN2725

3. **AI VTT: NVIDIA Run:ai**
   2025-07-02 | 75 minutes | Employees Only

   NVIDIA Run:ai is a GPU orchestration and optimization platform designed to help organizations maximize their GPU compute resources for AI workloads. It accelerates AI development, reduces costs, and improves AI development cycles by enabling dynamic allocation and scheduling of GPU resources, as well as workload submission and sharing. Essentially, it provides a centralized interface to manage AI compute infrastructure, making it easier for AI teams to access and utilize GPUs effectively.

   Join Carlos Huescas from Lenovo, Sandeep Brahmarouthu and Robert Magno from NVIDIA as they discuss NVIDIA Run:ai. Topics include:
   •What is Run:ai and its capabilities?
   •Customer segmentation for Run:ai
   •How to order, part numbers and licensing
   •Demo of Run:ai

   Tags: Artificial Intelligence (AI), NVIDIA

   Published: 2025-07-02
   Length: 75 minutes

   Start the training:
   Employee link: Grow@Lenovo

   Course code: DVAI218
4. **Partner Technical Webinar - Enterprise AI Team Intro and Vertical Spotlight Pt1**
   2025-06-17 | 60 minutes | Employees and Partners

   In this 60-minute replay, John Encizo introduced his new Enterprise AI Team. Part 1 covered three verticals: Retail with Allen Holmes, Manufacturing with Jason Hamp, and Healthcare with Janna Templin.

   Tags: Artificial Intelligence (AI)

   Published: 2025-06-17
   Length: 60 minutes

   Start the training:
   Employee link: Grow@Lenovo
   Partner link: Lenovo 360 Learning Center

   Course code: JUN1325

5. **VTT Edge: Understanding Visual AI Agents with NVIDIA June 2025**
   2025-06-16 | 60 minutes | Employees and Partners

   Join our guest speakers from NVIDIA as they discuss what's behind the scenes of visual AI Agents for Smart Cities, Smart Spaces and Manufacturing. Explore the modular approach to building a workforce of AI Agents. Topics include:

   • Sensors which feed the AI Agents
   • How AI agents improve safety and prevent accidents in Smart Spaces
   • Demo: Modular development of AI Agents

   Tags: Artificial Intelligence (AI), Technical Sales, NVIDIA

   Published: 2025-06-16
   Length: 60 minutes

   Start the training:
   Employee link: Grow@Lenovo
   Partner link: Lenovo 360 Learning Center

   Course code: DVEDG221

6. **Lenovo Cloud Architecture VTT: Supercharge Your Enterprise AI with NVIDIA AI Enterprise on Lenovo Hybrid AI Platform**
   2025-04-17 | 75 minutes | Employees and Partners

   Join us for an in-depth webinar with Justin King, Principal Product Marketing Manager for Enterprise AI exploring the power of NVIDIA AI Enterprise, delivering Generative and Agentic AI outcomes deployed with Lenovo Hybrid AI platform environments.
   In today's data-driven landscape, AI is evolving at high speed, with new techniques delivering more accurate responses. Enterprises are seeking not just an understanding but also how they can achieve AI-driven business outcomes.
   With this, the demand for secure, scalable, and high-performing AI operations-and the skills to deliver them-is top of mind for many. Learn how NVIDIA AI Enterprise, a comprehensive software suite optimized for NVIDIA GPUs, provides the tools and frameworks, including NVIDIA NIM, NeMo, and Blueprints, to accelerate AI development and deployment while reducing risk-all within the control and security of your Lenovo customer's hybrid AI environment.

   Tags: Artificial Intelligence (AI), Cloud, Data Management, Nvidia, Technical Sales

   Published: 2025-04-17
   Length: 75 minutes

   Start the training:
   Employee link: Grow@Lenovo
   Partner link: Lenovo 360 Learning Center

   Course code: DVCLD221

7. **AI VTT: GTC Update and The Lenovo LLM Sizing Guide**
   2025-03-12 | 86 minutes | Employees Only

   Please view this session that is two parts. Part one is Robert Daigle, Director, Global AI Solutions and Hande Sahin-Bahceci, AI Solutions Marketing Leader explaining the upcoming announcements for NVIDIA GTC. Part Two is Sachin Wani, AI Data Scientist explaining the Lenovo LLM Sizing Guide with these topics:

   • Minimum GPU requirements for fine-tuning/training and inference
   • Gathering requirements for the customer's use case
   • LLMs from a technical perspective

   Tags: Artificial Intelligence (AI), Technical Sales

   Published: 2025-03-12
   Length: 86 minutes

   Start the training:
   Employee link: Grow@Lenovo

   Course code: DVAI214

8. **VTT AI: Components of the AI Stack and Where Lenovo Sits November 2024**
   2024-11-26 | 75 minutes | Employees Only

   Join Per Ljungstrom, Lenovo Principal TC EMEA, as he explores AI concepts where innovations meet simplified predefined solutions which deploy at scale. Topics for this session include:
   • Associating software with the ground level of hardware
   • Attach NVIDIA AI Enterprise, Microsoft, Tiber AI Stacks and more
   • AI at the Edge and the complete solution
   • What to consider when talking AI Stack with your customer

   Tags: Artificial Intelligence (AI), Cloud, Technical Sales, Technology solutions, ThinkEdge

   Published: 2024-11-26
   Length: 75 minutes

   Start the training:
   Employee link: Grow@Lenovo

   Course code: DVAI210

9. **VTT AI: NVIDIA OVX**
   2024-10-23 | 55 minutes | Employees and Partners

   Please join this session as Steven Puzio, Global Sales Leader of NVIDIA Omniverse speaks to us about these topics:

   • OVX use cases
   • Target customers
   • OVX reference architectures
   • Parts, pieces and technical details

   Tags: Artificial Intelligence (AI), Nvidia

   Published: 2024-10-23
   Length: 55 minutes

   Start the training:
   Employee link: Grow@Lenovo
   Partner link: Lenovo 360 Learning Center

   Course code: DVAI209

10. **Think AI Weekly: Ride the NVIDIA Wave for AI**
    2024-10-07 | 60 minutes | Employees Only

    In this session, a panel including speakers from NVIDIA, Lenovo IDG and Lenovo ISG address the topics:
    •Leveraging AI workstations to start an AI journey
    •Leading an ISG sale with NVIDIA AI Enterprise
    •NVIDIA sales tools available for Lenovo sellers
    •NVIDIA training on grow@lenovo and more

    Tags: Artificial Intelligence (AI), Nvidia

    Published: 2024-10-07
    Length: 60 minutes

    Start the training:
    Employee link: Grow@Lenovo

    Course code: DTAIW121

11. **Lenovo VTT Cloud Architecture - Unlock Gen AI with VMware Private AI Foundation with NVIDIA**
2024-07-16 | 60 minutes | Employees and Partners

In today's rapidly evolving digital landscape, businesses are hungry for the transformative power of Artificial Intelligence (AI). They see AI as the key to streamlining operations and unlocking exciting new opportunities. However, widespread adoption has been hampered by concerns surrounding privacy, the complexity of implementation, and the hefty costs associated with deploying and managing AI solutions at an enterprise level.
Join Chris Gully and Baker Hull, Solutions Architects from VMware by Broadcom, as they discuss how Lenovo, NVIDIA, and VMware By Broadcom are partnering to deliver a private, secure, scalable, and flexible AI infrastructure solution that helps enterprise customers build and deploy AI workloads within their own private cloud infrastructure, ensure the control of sensitive data and compliance with regulatory requirements, ultimately driving faster time to value and achieving their AI objectives.

Tags: Artificial Intelligence (AI), Cloud, Nvidia, ThinkAgile, VMware

Published: 2024-07-16
Length: 60 minutes

Start the training:
Employee link: Grow@Lenovo
Partner link: Lenovo 360 Learning Center

Course code: DVCLD214

12. **Guidance for Selling NVIDIA Products at Lenovo for ISG**
2024-07-01 | 25 minutes | Employees and Partners

This course gives key talking points about the Lenovo and NVIDIA partnership in the Data Center. Details are included on where to find the products that are included in the partnership and what to do if NVIDIA products are needed that are not included in the partnership. Contact information is included if help is needed in choosing which product is best for your customer. At the end of this session sellers should be able to explain the Lenovo and NVIDIA partnership, describe the products Lenovo can sell through the partnership with NVIDIA, help a customer purchase other NVIDIA product, and get assistance with choosing NVIDIA products to fit customer needs.

Tags: Artificial Intelligence (AI), Nvidia

Published: 2024-07-01
Length: 25 minutes

Start the training:
Employee link: Grow@Lenovo
Partner link: Lenovo 360 Learning Center

Course code: DNVIS102

13. **NVIDIA AI Solutions and Market Trends**
    2023-10-12 | 55 minutes | Employees Only

    The purpose of this course is to help the learner recognized AI Market and trends. Also, explain NVIDIA's Computing platform, and discuss its importance for the market.

    Course Objectives:
    Recognize AI Trends
    Explain NVIDIA Computing Platform
    Discuss Industry Verticals Marketing

    Tags: Artificial Intelligence (AI), Nvidia, Sales

    Published: 2023-10-12
    Length: 55 minutes

    Start the training:
    Employee link: Grow@Lenovo

    Course code: DAINVD101

Authors

Farah Toosi is a Software Product Manager for NVIDIA enterprise software in Lenovo's Infrastructure Solutions Group. She specializes in AI/ML software infrastructure, GPU orchestration, and enterprise AI platform integrations. She has 8+ years of experience across software products and program management on several high tech companies

Carlos Huescas is the Worldwide Product Manager for NVIDIA software at Lenovo. He specializes in High Performance Computing and AI solutions. He has more than 15 years of experience as an IT architect and in product management positions across several high-tech companies.

Related product families

Product families related to this document are the following:

- GPU adapters

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION ”AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document, LP2289, was created or updated on September 8, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.lenovo.com/LP2289
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/LP2289.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
Neptune®
ThinkStation®
ThinkSystem®
XClarity®

The following terms are trademarks of other companies:

AMD is a trademark of Advanced Micro Devices, Inc.

Intel®, Intel Core®, and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Windows Server® and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.