

Lenovo Workstation Solutions

Lenovo AI Developer

NVIDIA AI Workbench Solution Guide

Lenovo



Challenges with AI Development

Commercial off the shelf software has its place in the enterprise, used extensively by millions of end customers from all over the world for standardized industry workflows tasks.

Artificial intelligence on the other hand is different, with many AI workflows truly bespoke and unique for the customers, industries and geographies that they are being developed and deployed in.

Although commercial off the shelf AI software exists, the innovation in development, use and deployment of artificial intelligence is happening within the world of open- source software, libraries, models and tools—a world where AI developers and practitioners have the freedom to choose the very best tools and techniques available to create their company's new AI, in the fastest and most efficient manner possible.

This vast array of open-source software, libraries, models and tools typically comes from a variety of different sources, with an even greater list of software, libraries and hardware setup dependencies. This results in major headaches for both data scientists and their IT execution teams when attempting to track, trace, secure, deploy and scale these new AI applications and workflows across their organizations.

As teams start to collaborate on larger, ever-evolving business transformation projects that are built on the foundations of open-source artificial intelligence, this challenge is only set to get harder.

Pain points of today's typical AI development workflows



Slow & Inefficient to Deploy and Scale



Poor Collaboration Tools



Painful Eco-System Dependencies



Security / Data Privacy Concerns



Cost & Complexity to Scale



The innovation in development & deployment of artificial intelligence is happening within the world of open-source software, libraries, models and tools.

Lenovo AI Developer

Lenovo AI Developer is a new suite of full-stack AI development solutions—expertly designed blueprints, that empower data scientists and AI developers to build, scale, and secure their AI workflows with confidence.

Built for the complexity of modern AI, Lenovo AI Developer simplifies the development journey by integrating best-in-class open-source tools, AI libraries, and frameworks with Lenovo's industry-leading workstation hardware.

Why Lenovo AI Developer



End-to-End AI Development Support

Solutions for every stage of the AI lifecycle—from data prep, machine learning, and training to deep learning, fine-tuning, and inference.



Certified & Future-Ready

NVIDIA and Linux®-certified workstations, built with cutting-edge CPUs and GPUs, ensuring compatibility and performance unlike some outdated cloud options.




Improved Efficiency

Deliver better performance, lower total cost of ownership, and faster time to market with right-sized solutions for each stage of your workflow.



Enhanced Security

Keep intellectual property protected with on-premises, sandboxed environments.



Smarter AI Development

To solve today's AI development workflow challenges, Lenovo AI Developer pairs our state-of-the-art workstations with NVIDIA AI Workbench—offering right-sized, high-performance hardware for each stage of your AI journey.

AI Workbench is a free AI development platform that takes a hybrid approach to data science, machine learning, and generative AI workloads.

NVIDIA AI Workbench helps with:



Automated Setup: Click through installation and configuration sets up your GPU accelerated ThinkStation or ThinkPad workstation in minutes.



User Friendliness: Command-line or graphical interfaces handle the details without getting in the way.



Faster Development: Start quickly in JupyterLab or Visual Studio Code with streamlined Git versioning and containerization.



Customization: Configure interfaces like Gradio and Streamlit apps to run locally while being able to modify them on the fly.



Project Collaboration: Connect to GitHub, GitLab, or self-hosted GitLab to collaborate with team members.



One-click Scaleup: Move projects from ThinkPad P Series mobile workstations with NVIDIA RTX PRO™ GPUs, and scale to more powerful, multi-GPU configured ThinkStation P Series desktop workstations, and on to the data center.

Smarter AI Deployment

AI Workbench



Setup / Configuration:

- Automated for GPU-enabled environments



Automation of workflow tasks:

- Reproduce, collaborate, scale-up/down
- Seamless portability across environments
- Supports versioning, credentials, location changes



Management for projects:

- Installation and management of GenAI resources
- Creation and tracking of project resources
- Intuitive UI plus powerful CLI

Certified Platforms

Lenovo
ThinkPad



Lenovo
ThinkStation



Lenovo
ThinkSystem
Cloud

Language Models

 **Meta**
Llama 3

 **MISTRAL**
AI_

Falcon 180B

 **Qwen**

Developer Environments:

 **docker**

 **jupyter**

AI Workbench

Frameworks:

 **TensorFlow**

 **scikit learn**

 **python™**

 **PyCharm**

 **pandas**

RAPIDS

Code / Repositories:

 **GitHub**

 **GitLab**

 **Hugging Face**

 **NVIDIA. NGC**


...and many more

Optimizing AI Workloads

Artificial intelligence and generative AI workloads can be extremely hardware intensive, with rigorous performance requirements placed on CPU, GPU, memory, and storage components. AI developers and data scientists must continuously prepare evolving datasets from a growing variety of data silos—wrangling, cleansing, and validating the data, selecting AI models, training them, and fine-tuning for maximum impact.

Data scientists, AI developers, and machine learning engineers are at the heart of many companies' digital transformation strategies, all powered by artificial intelligence and generative AI; however, efficient use of available employees, project budgets and computing resources are crucial to deliver a sustainable return on investment.

Data Science & Generative AI Development Workstreams



The **NVIDIA RTX PRO 6000 Blackwell Workstation Edition** is the most powerful desktop GPU ever created, redefining performance and capability for professionals. Dominate the most advanced models and tackle the most demanding creative workflows with unmatched AI processing power. Built on the groundbreaking NVIDIA Blackwell architecture and equipped with 96GB of ultra-fast GDDR7 memory, the RTX PRO 6000 Workstation Edition delivers unparalleled speed, precision, and efficiency, empowering you to achieve the impossible.



Up to
4000
TOPS^{1,2}

1. Peak rates based on GPU Boost Clock.
2. Theoretical FP4 TOPS using sparsity.

Revolutionizing AI Development: The AI Workbench Impact

The implementation of AI Workbench marks a transformative shift from the traditional bottlenecks in AI development to a streamlined and efficient process.

Before AI Workbench, developers often faced slow deployment times, poor collaboration, system breakdowns, security issues, and difficulties in scaling. With AI Workbench, these challenges are addressed by providing an environment that is easy to set up, fosters one-click collaboration, ensures reliability, upholds security and privacy, and seamlessly scales from laptops to data centers.

It is a comprehensive solution that redefines the efficiency and potential of AI development workflows.

Efficient AI Projects: End-to-End

With ever increasing costs associated with cloud computing, it is imperative to evaluate not just where, but when and how you compute your AI workloads. Lenovo AI Developer solutions powered by ThinkStation & ThinkPad P Series workstations with NVIDIA RTX PRO GPUs, paired with NVIDIA AI Workbench software, delivers an easy out of the box blueprint that can enable:



A simple and efficient start to your AI journey



Easy management of your cost of ownership



Better scaling of AI projects



Secure sandbox AI environments



Faster delivery of positive business outcomes

Choosing the Right Hardware for AI Workloads

Data scientists are empowered to fail fast and fail cheap utilizing the power and performance of local computing resources. This allows them to pivot quickly and course correct without overly relying on cost prohibitive GPU resources in the cloud.



Lenovo ThinkStation and ThinkPad P Series workstations deliver the world's fastest and most powerful AI workflow experiences for AI developers, data scientists, and AI practitioners. Purpose-built and engineered to not only meet but exceed the rigorous performance demands of today's most advanced AI workloads, these systems empower professionals to operate at their best—transforming data into actionable business insights and accelerating AI projects from concept to PoC, and into production with exceptional speed and efficiency.

ThinkPad
P1

 **NVIDIA**
Certified

Lenovo
ThinkShield

 **Linux™**



Lenovo Workstations for NVIDIA AI Workbench

Co-designed with Aston Martin, our flexible workstations—available direct or as a Service—deliver scalable performance for data prep, model training and generative AI. Fully configurable for software, dataset size, workflow or budget to meet diverse computational and budgetary demands.


x86 Platform + NVIDIA RTX PRO GPUs						
						
	ThinkPad P1	ThinkPad P16	ThinkStation P5	ThinkStation P7	ThinkStation P8	ThinkStation PX
CPU » Up to	» Intel® Core™ Ultra (Series 2) H-Series processors » 16 cores, 5.4GHz	» Intel Core Ultra (Series 2) HX-Series processors » 24 cores, 5.5GHz	» Intel Xeon® W-2500 processor » 26 cores, 4.8GHz	» Intel Xeon W-3500 processor » 60 cores, 4.8GHz	» AMD Ryzen Threadripper Pro 7000 WX-Series » 96 cores, 5.3GHz	» Dual 5th Gen Intel Xeon Scalable processors » 128 cores, 4.1GHz
Memory » Up to	» 64GB LPDDR5x 7467MT/s	» 192GB DDR5 5600MT/s	» 512GB DDR5 4800MT/s	» 1TB DDR5 4800MT/s	» 1TB DDR5 4800MT/s	» 2TB DDR5 5600MT/s
GPU » Up to	» NVIDIA RTX PRO 2000 Blackwell Generation Laptop GPU (8GB)	» NVIDIA RTX PRO 5000 Blackwell Generation Laptop GPU (24GB)	» 2x NVIDIA RTX PRO 6000 Blackwell Max-Q Workstation Edition GPU (192GB)	» 3x NVIDIA RTX PRO 6000 Blackwell Max-Q Workstation Edition GPU (288GB)	» 3x NVIDIA RTX PRO 6000 Blackwell Max-Q Workstation Edition GPU (288GB)	» 4x NVIDIA RTX PRO 6000 Blackwell Max-Q Workstation Edition GPU (384GB)
Storage » Up to	» 8TB M.2 PCIe Gen5 NVMe Performance SSD	» 12TB M.2 PCIe Gen5 NVMe Performance SSD	» 48TB with 6 drives	» 32TB with 7 drives	» 52TB with 7 drives	» 60TB with 7 drives

Lenovo ThinkStation PGX

Introducing the first **NVIDIA GB10 Grace Blackwell Superchip** accelerated Lenovo workstation. Purpose-built in a **compact form factor** to complement your AI workflow, including classic statistics, machine learning, computer vision, generative AI/LLM, robotics, and agentic design. Featuring **128GB** of unified system memory, **20 CPU cores**, up to **4TB** of local storage, and **NVIDIA ConnectX®-7 networking** for clustering two devices together.

Offering a controlled, sandboxed environment for local prototyping, fine-tuning, and inferencing, the PGX integrates seamlessly into existing workstation setups, unlocking even greater compute capacity for advanced AI model development.


Preconfigured with the **NVIDIA DGX™ OS** and **NVIDIA AI software stack**, plus popular tools like PyTorch and Jupyter notebooks, PGX delivers a familiar, efficient development experience. An ideal companion device for developers, researchers, and students who need who need scalable AI capabilities without the complexity or cost of cloud infrastructure.



Pair your **Lenovo workstation** with the **ThinkStation PGX** to unlock even greater compute capacity for advanced AI model development


Lenovo
ThinkStation

Discover more about Lenovo AI Developer and NVIDIA



lenovo.com/workstations

lenovo.com/AI



[NVIDIA AI Workbench](https://nvidia.com/ai-workbench)

[NVIDIA Generative AI](https://nvidia.com/generative-ai)



2025 ©Lenovo. All rights reserved.

Lenovo is not responsible for photographic or typographic errors. Lenovo makes no representation or warranty regarding third-party products or services.

LENOVO, ThinkStation, and ThinkPad are trademarks of Lenovo. NVIDIA, NVIDIA RTX PRO, NVIDIA DGX, and NVIDIA ConnectX are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Intel, Intel vPro, Core, Xeon, and Thunderbolt are trademarks of Intel Corporation or its subsidiaries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Microsoft and Windows are registered trademarks of Microsoft Corporation. Other company, product, and service names may be trademarks or service marks of others.

Lenovo