

Lenovo
ThinkStation

Lenovo AI Workstation Solutions with 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 companies 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 and collaborate on larger and 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

Smarter AI Development

To solve this challenge, Lenovo is collaborating with NVIDIA® to bring NVIDIA® AI Workbench to Lenovo Workstations.

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™ GPUs, and scale to more powerful Multi-GPU configured ThinkStation P Series Workstations, and on to the datacenter.

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



Falcon 180B



Developer Environments:



AI Workbench

Frameworks:



RAPIDS

Code / Repositories:



...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 & Storage components. AI Developers and Data Scientists are having to prepare ever evolving datasets, from a growing variety of data silos, de-wrangling, cleansing and validating this data, selecting AI models, model training, and then fine tuning these for maximum impact.

“Over ~70% of a data scientist's time can be spent on data preparation tasks, so it pays to keep them productive”

Anaconda State of Industry - Data Science Report, 2023


Data Scientists, AI Developers and Machine Learning Engineers are at the heart of many company's 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



Edge Inference & Re-Training

Built on the NVIDIA CUDA® compute platform and leveraging CUDA-X libraries to speed up data processing, NVIDIA® has created a Workbench Project to leverage the NVIDIA RAPIDS™ cuDF library, which can accelerate the workflow of nearly 10 million data scientists using pandas software by up to 110x on an NVIDIA RTX™ 6000 Ada GPU without requiring any code changes. This Workbench Project, along with others, can be found on [NVIDIA's GitHub](#).



110x Accelerated workflow

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 Solutions powered by ThinkStation & ThinkPad P Series Workstations with NVIDIA RTX™ GPUs, paired with NVIDIA® AI Workbench software delivers an easy out of the box solution 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 instead of over relying on cost prohibitive GPU resources in the cloud.



Lenovo ThinkStation & ThinkPad P Series 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 just meet but exceed the rigorous performance requirements of today's most advanced AI workloads, enabling AI professionals to operate at their best, turning data into business insight and enabling AI projects to move quickly from concept, to PoC and on into production in the fastest possible time.

“Lenovo Workstations are built on the three fundamental pillars of Power, Performance & Reliability - Backed by years of legendary Think enterprise IT heritage.”



Lenovo Workstations for NVIDIA® AI Workbench

Lenovo Workstations are co-designed with Aston Martin and are the epitome of form and function. They can be configured to suit a wide variety of industry software applications, dataset sizes, customer workflows or project budgets. Available direct or through our flexible Workstation as a Service procurement model.

Recommended Lenovo Workstations for Data Science and Generative AI workloads with NVIDIA® AI Workbench

Featuring maximum achievable configurations. Workstations with lower specifications are also proficient, offering scalable performance to meet diverse computational and budgetary demands

Lenovo



The World's Most Powerful AI Developer Workstation

	ThinkPad P1 Gen7	ThinkPad P16 Gen2	ThinkStation P5	ThinkStation P7	ThinkStation P8	ThinkStation PX
CPU	» Intel® Core™ Ultra 9 processors with Intel vPro® » 16 cores, 5.1GHz	» 14 th Gen Intel® Core™ i7 HX-series processors » 20 cores, 5.5GHz	» Intel® Xeon® W-2400 processor » 24 cores, 4.8GHz	» Intel® Xeon® W-3400 processor » 56 cores, 4.8GHz	» AMD Ryzen™ Threadripper™ Pro 7000 WX-Series » 96 cores, 5.3GHz	» Dual 4 th Gen Intel® Xeon® Scalable processors » 120 cores, 4.1GHz
Memory	» 64GB LPDDR5x 7467MT/s	» 192GB DDR5 5600MHz	» 512GB DDR5 4800MHz	» 2TB DDR5 4800MHz	» 2TB DDR5 4800MHz	» 4TB DDR5 4800MHz
GPU	» NVIDIA RTX™ 3000 Ada Generation mobile GPU	» NVIDIA RTX™ 5000 Ada Generation mobile GPU	» 2x NVIDIA RTX™ 6000 Ada Generation GPU	» 3x NVIDIA RTX 6000 Ada Generation GPU	» 3x NVIDIA RTX™ 6000 Ada Generation GPU	» 4x NVIDIA RTX™ 6000 Ada Generation GPU
Storage	» 8TB M.2 PCIe Gen4 NVMe Performance SSD	» 8TB M.2 PCIe Gen4 NVMe Performance SSD	» 48TB with 6 drives	» 52TB with 7 drives	» 52TB with 7 drives	» 60TB with 7 drives

Desktop & Rackmount Optimized


Remote Workstation Ready

NVIDIA® AI Workbench & AI Enterprise Certified/Qualified

✓ Data Prep/Wrangling ✓ ~3B+ LLM Training ✓ AI Model Fine Tuning	✓ Data Prep/Wrangling ✓ ~7B+ LLM Training ✓ AI Model Fine Tuning	✓ Data Prep/Wrangling ✓ ~10B+ LLM Training ✓ AI Model Fine Tuning	✓ Data Prep/Wrangling ✓ ~13B+ LLM Training ✓ AI Model Fine Tuning	✓ Data Prep/Wrangling ✓ ~20B+ LLM Training ✓ AI Model Fine Tuning ✓ Multi-User Appliance	✓ Data Prep/Wrangling ✓ ~70B+ LLM Training ✓ AI Model Fine Tuning ✓ Multi-User Appliance
--	--	---	---	---	---


Lenovo
ThinkStation

Discover more from Lenovo and NVIDIA®



lenovo.com/workstations

lenovo.com/AI



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

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



2024 ©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® is a trademark and/or registered trademark 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