BRING AI PERFORMANCE CLOSER TO THE WORK — WITH THE FLEXIBILITY TO SCALE

Hover or click on this icon throughout the infographic to learn more.



For suitable workloads, local AI on Lenovo Workstations with NVIDIA gives IT another option within Lenovo Hybrid AI Advantage — with dedicated, predictable capacity close to users, applications, or data.

ACCELERATE AI DEVELOPMENT

Give developers predictable access to the performance they need to build, test, and refine.

IMPROVE WORKFLOW PRODUCTIVITY

Keep both local inference and demanding professional applications responsive with a configuration sized for the complete workflow.

SUPPORT DATA CONTROL AND GOVERNANCE

Give IT greater control over where processing occurs and how data moves, while meeting approved data access, security, and AI governance requirements.

MAKE MORE INTENTIONAL INFRASTRUCTURE INVESTMENTS

Match dedicated and shared capacity to workload requirements, utilization, economics, and scale.

THE RIGHT FIT STARTS WITH THE WORKLOAD

40% of enterprise applications are forecast to include task-specific AI agents by the end of 2026.[1]

FIVE FACTORS TO HELP PLACE AN AI WORKLOAD

Start with how the workload needs to operate — then evaluate which environment fits best.

1	2	3	4	5
Access and sharing	Performance and responsiveness	Data and proximity	Operating requirements	Economics

No single factor determines placement. Evaluate the workload across all five to determine whether workstation compute or another Hybrid AI environment is the stronger fit.

54% of AI developers surveyed (n=77) say memory limitations slow iteration.[2]

SEE WHERE LOCAL AI ON WORKSTATIONS FITS WITHIN LENOVO HYBRID AI ADVANTAGE WITH NVIDIA

Workstation compute is one purposeful part of Hybrid AI. Depending on workload requirements, AI can also run across edge, data center, and cloud environments.

Place each workload in the environment that best matches its performance, operating, and economic requirements.

LENOVO WORKSTATIONS WITH NVIDIA

Give IT greater control over where processing occurs and how data moves, while meeting approved data access, security, and AI governance requirements.

EDGE

Processing close to equipment, operational systems, or distributed data

DATA CENTER

Shared capacity for centralized workloads, services, and operations

CLOUD

Flexible capacity and cloud services when the workload or operating model calls for them

62% of surveyed organizations prefer hybrid as their primary AI deployment model.[3]

ADAPT AS REQUIREMENTS EVOLVE

As AI needs change, the environment can change with them. These are parallel choices — not sequential stages.

STAY

Stay on workstation compute when dedicated performance, predictable access, and the operating model continue to meet requirements.

EXPAND

Increase workstation capability when requirements grow but workstation compute still fits the work.

CONNECT

Connect to edge, data center, or cloud when shared access, availability, orchestration, governance, or scale call for a different operating model.

WHY LENOVO WORKSTATIONS WITH NVIDIA

Lenovo Workstations with NVIDIA give IT options for two growing AI needs: developers building, testing, and refining AI and agentic workflows, and professionals using AI alongside demanding applications. Across the portfolio, IT can match each workload to an appropriate workstation configuration and supported NVIDIA software stack while maintaining security, manageability, and flexibility within a broader Hybrid AI environment.



LENOVO THINKSTATION® AND THINKPAD® P SERIES WORKSTATIONS WITH NVIDIA RTX PRO™

A portfolio of desktop and mobile workstations gives IT options to match performance, form factor, mobility, and capability to different local AI, inference, and professional workflows.

FULL-SYSTEM WORKSTATION ENGINEERING

CPU and GPU capability, memory, storage, thermals, expansion, and software compatibility work together to support demanding AI and professional workflows.

THINKSHIELD SECURITY + IT MANAGEABILITY

ThinkShield helps protect Lenovo Workstations, while IT can deploy, manage, and support workstation-based AI within established practices.

NVIDIA ACCELERATED COMPUTING + AI SOFTWARE

Support development, inference, retrieval, and agentic AI with NVIDIA RTX PRO accelerated computing and software including NVIDIA NIM™ microservices, Nemotron™ open models, NeMo™ Retriever, and TensorRT™ for RTX — with options to connect workloads to broader Lenovo infrastructure as requirements evolve.

FIND THE RIGHT FIT FOR YOUR AI WORKLOADS

Map your AI workloads to the compute environments that best support productivity, control, economics, and scale — and identify where Lenovo Workstations fit within Lenovo Hybrid AI Advantage with NVIDIA.

Assess a local AI or inference workflow with a Lenovo Workstation specialist

Landing page TBD

Sources
1 Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," August 2025
2 Lenovo, "Workstation AI Developer Influencer Research," September 2026. Lenovo AI Survey quantitative analysis, n=77 AI developers
3 IDC, "IDC Playbook 2026: The Race for Enterprise AI," January 2026. IDC CIO Playbook 2026 Survey, commissioned by Lenovo, n=1,120 IT and business decision-makers globally

© Lenovo 2026. All rights reserved. V1.00 October 2026

Smarter technology for all

BRING AI PERFORMANCE CLOSER TO THE WORK — WITH THE FLEXIBILITY TO SCALE

NVIDIA

Hover or click on this icon throughout the infographic to learn more.



For suitable workloads, local AI on Lenovo Workstations with NVIDIA gives IT another option within Lenovo Hybrid AI Advantage — with dedicated, predictable capacity close to users, applications, or data.

ACCELERATE AI DEVELOPMENT

Give developers predictable access to the performance they need to build, test, and refine.

IMPROVE WORKFLOW PRODUCTIVITY

Keep both local inference and demanding professional applications responsive with a configuration sized for the complete workflow.

SUPPORT DATA CONTROL AND GOVERNANCE

Give IT greater control over where processing occurs and how data moves, while meeting approved data access, security, and AI governance requirements.

MAKE MORE INTENTIONAL INFRASTRUCTURE INVESTMENTS

Match dedicated and shared capacity to workload requirements, utilization, economics, and scale.

THE RIGHT FIT STARTS WITH THE WORKLOAD

40% of enterprise applications are forecast to include task-specific AI agents by the end of 2026.[1]

FIVE FACTORS TO HELP PLACE AN AI WORKLOAD

Start with how the workload needs to operate — then evaluate which environment fits best.

1. Access and sharing
2. Performance and responsiveness
3. Data and proximity
4. Operating requirements
5. Economics

No single factor determines placement. Evaluate the workload across all five to determine whether workstation compute or another Hybrid AI environment is the stronger fit.

ACCESS AND SHARING

Workstation compute may fit better when:

- A user or small team benefits from dedicated, predictable access
- Waiting for shared capacity slows development or workflow responsiveness
- The workload is tied closely to one role, team, or application workflow

Another environment may fit better when:

- Many users need access to the same compute environment
- Capacity needs to be pooled across teams
- Shared services, centralized scheduling, or broad access are priorities

PERFORMANCE AND RESPONSIVENESS

Workstation compute may fit better when:

- The workload benefits from dedicated local resources
- Interactive development, inference, or AI-assisted workflows need responsive performance
- Model size, memory, runtime, and software requirements fit the workstation configuration

Another environment may fit better when:

- The workload exceeds practical workstation capacity
- Larger models, higher concurrency, or longer runtimes require broader infrastructure
- Elastic or production-scale resources are needed

DATA AND PROXIMITY

Workstation compute may fit better when:

- Processing close to the user, application, or approved data improves the workflow
- Data movement creates friction, delay, or added complexity
- Local control over where work happens supports business or policy needs

Another environment may fit better when:

- Processing needs to occur near equipment or distributed data sources
- Data, applications, or services are already centralized
- Governance, residency, or shared-access requirements call for processing in another environment

OPERATING REQUIREMENTS

Workstation compute may fit better when:

- Workstation-level security, governance, monitoring, deployment, management, and support controls meet the workload's needs
- IT can manage approved models, software updates, data access, logging, and support within established practices
- Workstation availability meets the workload's requirements

Another environment may fit better when:

- The workload requires higher availability, centralized orchestration, production SLAs, or broader operational controls
- Centralized governance, monitoring, or shared services are required by the operating model
- Service management needs extend beyond workstation-based operations

ECONOMICS

Workstation compute may fit better when:

- Dedicated capacity is used often enough to justify the investment
- Tokenomics favor dedicated capacity for recurring inference or agentic workloads
- Faster iteration, predictable access, or less waiting improves productive user time

Another environment may fit better when:

- Token or compute demand is bursty or highly variable
- Shared or usage-based capacity provides a better utilization and cost profile
- Another environment better balances cost, usage, operations, and scale

Example:

A designer using local AI alongside project files and professional applications may favor workstation compute when **access and sharing** needs are limited to one user or a small team and **performance and responsiveness** requirements fit the configuration. **Data and proximity** may also favor keeping files and applications close to the user, while **operating requirements** can be met within established workstation practices. If **economics** support regularly used dedicated capacity, workstation compute may be the stronger fit. If the AI workload grows to require broad shared access, higher concurrency, centralized services, or production-scale availability, broader Hybrid AI infrastructure may become the better fit.

FIND THE RIGHT FIT FOR YOUR AI WORKLOADS

Map your AI workloads to the compute environments that best support productivity, control, economics, and scale — and identify where Lenovo Workstations fit within Lenovo Hybrid AI Advantage with NVIDIA.

Assess a local AI or inference workflow with a Lenovo Workstation specialist

Landing page TBD

Sources

1 Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," August 2025

2 Lenovo, "Workstation AI Developer Influencer Research," September 2026. Lenovo AI Survey quantitative analysis, n=77 AI developers

3 IDC, "IDC Playbook 2026: The Race for Enterprise AI," January 2026. IDC CIO Playbook 2026 Survey, commissioned by Lenovo, n=1,026 IT and business decision-makers globally

© Lenovo 2026. All rights reserved. V1.00 October 2026

Smarter technology for all

NVIDIA

BRING AI PERFORMANCE CLOSER TO THE WORK — WITH THE FLEXIBILITY TO SCALE

Hover or click on this icon throughout the infographic to learn more.



For suitable workloads, local AI on Lenovo Workstations with NVIDIA gives IT another option within Lenovo Hybrid AI Advantage — with dedicated, predictable capacity close to users, applications, or data.

ACCELERATE AI DEVELOPMENT

Give developers predictable access to the performance they need to build, test, and refine.

IMPROVE WORKFLOW PRODUCTIVITY

Keep both local inference and demanding professional applications responsive with a configuration sized for the complete workflow.

SUPPORT DATA CONTROL AND GOVERNANCE

Give IT greater control over where processing occurs and how data moves, while meeting approved data access, security, and AI governance requirements.

MAKE MORE INTENTIONAL INFRASTRUCTURE INVESTMENTS

Match dedicated and shared capacity to workload requirements, utilization, economics, and scale.

THE RIGHT FIT STARTS WITH THE WORKLOAD

40% of enterprise applications are forecast to include task-specific AI agents by the end of 2026.[1]

FIVE FACTORS TO HELP PLACE AN AI WORKLOAD

Start with how the workload needs to operate — then evaluate which environment fits best.

1. Access and sharing
2. Performance and responsiveness
3. Data and proximity
4. Operating requirements
5. Economics

No single factor determines placement. Evaluate the workload across all five to determine whether workstation compute or another Hybrid AI environment is the stronger fit.

54% of AI developers surveyed (n=77) say memory limitations slow iteration.[2]

SEE WHERE LOCAL AI ON WORKSTATIONS FITS WITHIN LENOVO HYBRID AI ADVANTAGE WITH NVIDIA

Workstation compute is one purposeful part of Hybrid AI. Depending on workload requirements, AI can also run across edge, data center, and cloud environments.

Place each workload in the environment that best matches its performance, operating, and economic requirements.

LENOVO WORKSTATIONS WITH NVIDIA

Give IT greater control over where processing occurs and how data moves, while meeting approved data access, security, and AI governance requirements.

EDGE

Processing close to equipment, operational systems, or distributed data

DATA CENTER

Shared capacity for centralized workloads, services, and operations

CLOUD

Flexible capacity and cloud services when the workload or operating model calls for them

62% of surveyed organizations prefer hybrid as their primary AI deployment model.[3]

ADAPT AS REQUIREMENTS EVOLVE

As AI needs change, the environment can change with them. These are parallel choices — not sequential stages.

STAY

Stay on workstation compute when dedicated performance, predictable access, and the operating model continue to meet requirements.

EXPAND

Increase workstation capability when requirements grow but workstation compute still fits the work.

CONNECT

Connect to edge, data center, or cloud when shared access, availability, orchestration, governance, or scale call for a different operating model.

STAY

- The workload still fits the workstation's capacity and operating model
- Dedicated access continues to support productivity and control
- Security, governance, and economics still support the choice

EXPAND

- Move to a higher-capability workstation
- Add memory, storage, or GPU capability where supported by the system
- Support growing model, data, or concurrency demands while keeping a workstation-based approach

CONNECT

- More users need access to the environment
- Availability, orchestration, or governance requirements increase
- Another operating model becomes a better fit for workload economics or scale

NVIDIA ACCELERATED COMPUTING + AI SOFTWARE

Support development, inference, retrieval, and agentic AI with NVIDIA RTX PRO accelerated computing and software including NVIDIA NIM™ microservices, Nemotron™ open models, NeMo™ Retriever, and TensorRT™ for RTX — with options to connect workloads to broader Lenovo infrastructure as requirements evolve.

FULL-SYSTEM WORKSTATION ENGINEERING

CPU and GPU capability, memory, storage, thermals, expansion, and software compatibility work together to support demanding AI and professional workflows.

FIND THE RIGHT FIT FOR YOUR AI WORKLOADS

Map your AI workloads to the compute environments that best support productivity, control, economics, and scale — and identify where Lenovo Workstations fit within Lenovo Hybrid AI Advantage with NVIDIA.

Assess a local AI or inference workflow with a Lenovo Workstation specialist

Landing page TBD

Sources

1. Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," August 2025.
2. Lenovo, "Workstation AI Developer Influencer Research," September 2026. Lenovo AI Survey quantitative analysis, n=77 AI developers.
3. IDC, "IDC Playbook 2026: The Race for Enterprise AI," January 2026. IDC CIO Playbook 2026 Survey, commissioned by Lenovo, n=1,120 IT and business decision-makers globally.

© Lenovo 2026. All rights reserved. V1.00 October 2026

Smarter technology for all

BRING AI PERFORMANCE CLOSER TO THE WORK — WITH THE FLEXIBILITY TO SCALE

NVIDIA

Hover or click on this icon throughout the infographic to learn more.



For suitable workloads, local AI on Lenovo Workstations with NVIDIA gives IT another option within Lenovo Hybrid AI Advantage — with dedicated, predictable capacity close to users, applications, or data.

ACCELERATE AI DEVELOPMENT

Give developers predictable access to the performance they need to build, test, and refine.

IMPROVE WORKFLOW PRODUCTIVITY

Keep both local inference and demanding professional applications responsive with a configuration sized for the complete workflow.

SUPPORT DATA CONTROL AND GOVERNANCE

Give IT greater control over where processing occurs and how data moves, while meeting approved data access, security, and AI governance requirements.

MAKE MORE INTENTIONAL INFRASTRUCTURE INVESTMENTS

Match dedicated and shared capacity to workload requirements, utilization, economics, and scale.

THE RIGHT FIT STARTS WITH THE WORKLOAD

40% of enterprise applications are forecast to include task-specific AI agents by the end of 2026.[1]

FIVE FACTORS TO HELP PLACE AN AI WORKLOAD

Start with how the workload needs to operate — then evaluate which environment fits best.

1	2	3	4	5
Access and sharing	Performance and responsiveness	Data and proximity	Operating requirements	Economics

No single factor determines placement. Evaluate the workload across all five to determine whether workstation compute or another Hybrid AI environment is the stronger fit.

54% of AI developers surveyed (n=77) say memory limitations slow iteration.[2]

SEE WHERE LOCAL AI ON WORKSTATIONS FITS WITHIN LENOVO HYBRID AI ADVANTAGE WITH NVIDIA

Workstation compute is one purposeful part of Hybrid AI. Depending on workload requirements, AI can also run across edge, data center, and cloud environments.

Place each workload in the environment that best matches its performance, operating, and economic requirements.

LENOVO WORKSTATIONS WITH NVIDIA

Give IT greater control over where processing occurs and how data moves, while meeting approved data access, security, and AI governance requirements.

EDGE

Processing close to equipment, operational systems, or distributed data

DATA CENTER

Shared capacity for centralized workloads, services, and operations

CLOUD

Flexible capacity and cloud services when the workload or operating model calls for them

62% of surveyed organizations prefer hybrid as their primary AI deployment model.[3]

ADAPT AS REQUIREMENTS EVOLVE

As AI needs change, the environment can change with them. These are parallel choices — not sequential stages.

STAY

Stay on workstation compute when dedicated performance, predictable access, and the operating model continue to meet requirements.

EXPAND

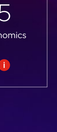
Increase workstation capability when requirements grow but workstation compute still fits the work.

CONNECT

Connect to edge, data center, or cloud when shared access, availability, orchestration, governance, or scale call for a different operating model.

WHY WORKSTATIONS WITH NVIDIA

Lenovo Workstations with NVIDIA RTX PRO™ options for AI development, data science, model building, fine-tuning, and inferencing for local workflows — including demanding professional applications. IT can match the right workstation to the job, with NVIDIA software, manageability, and flexibility within a broader Hybrid AI environment.



NVIDIA ACCELERATED COMPUTING + AI SOFTWARE

Relevant NVIDIA software for supported workstation AI workflows can include:

- NVIDIA NIM™ microservices: Help deploy optimized inference capabilities through consistent APIs
- NVIDIA Nemotron™ open models: Support reasoning and agentic AI use cases
- NVIDIA NeMo™ Retriever: Connect AI to approved enterprise content for retrieval and RAG workflows
- NVIDIA TensorRT™ for RTX: Optimize AI inference deployment on NVIDIA RTX-powered systems

Software availability, support, and licensing requirements vary by system, operating system, model, and workload.

LENOVO THINKSTATION® AND THINKPAD® P SERIES WORKSTATIONS WITH NVIDIA RTX PRO™

A portfolio of desktop and mobile workstations gives IT options to match performance, form factor, mobility, and capability to different local AI, inference, and professional workflows.

FULL-SYSTEM WORKSTATION ENGINEERING

CPU and GPU capability, memory, storage, thermals, expansion, and software compatibility work together to support demanding AI and professional workflows.

THINKSHIELD SECURITY + IT MANAGEABILITY

ThinkShield helps protect Lenovo Workstations, while IT can deploy, manage, and support workstation-based AI within established practices.

NVIDIA ACCELERATED COMPUTING + AI SOFTWARE

Support development, inference, retrieval, and agentic AI with NVIDIA RTX PRO accelerated computing and software including NVIDIA NIM™ microservices, Nemotron™ open models, NeMo™ Retriever, and TensorRT™ for RTX — with options to connect workloads to broader Lenovo infrastructure as requirements evolve.

FIND YOUR

Map your product view against parameters, data scales … and find out where they're

Assess a

Landing

LENOVO THINKSTATION® AND THINKPAD® P SERIES WORKSTATIONS WITH NVIDIA RTX PRO™

From desk-side performance to mobile productivity, the portfolio gives IT options to support different AI use cases, user roles, and work environments. Select configurations with NVIDIA RTX PRO™ Blackwell GPUs support demanding AI and professional workflows. ThinkStation® PGX adds a purpose-built AI development option powered by the NVIDIA GB10 Grace Blackwell Superchip.

FULL-SYSTEM WORKSTATION ENGINEERING

Sustained AI performance depends on the complete workstation — including compute, memory, storage, thermals, expansion, and software compatibility.

THINKSHIELD SECURITY + IT MANAGEABILITY

ThinkShield helps protect Lenovo Workstations across the lifecycle. Depending on the configuration and services selected, IT can apply protections such as encryption, identity and access controls, endpoint monitoring, firmware integrity, threat detection and response, and device policies.

Sources
1. Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," August 26, 2025.
2. Lenovo, "2026 Hybrid AI Advantage Survey," September 2025, n=77.
3. IDC, "2026 Hybrid AI Survey," January 2026.

© Lenovo 2026