

Lenovo



Navigating Enterprise IT Modernization & Virtualization

A practical guide for business leaders

Explore Enterprise IT modernization and virtualization challenges along with strategies and insights from Lenovo and AMD.

Choose where to begin, go deeper into the issues that matter most, and move forward with confidence.





Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Incremental change is no longer enough

Enterprise IT has never mattered more. Today, infrastructure determines how quickly enterprises innovate, how efficiently they can scale, and how resilient they remain.

The question isn't whether to modernize, but how - while balancing performance, efficiency, security, and long-term flexibility.

Get practical guidance from Lenovo and AMD to help you prioritize what matters most and take the next step forward.

Which challenge comes first?

Every organization faces these pressures differently. Begin with the challenge that has the greatest impact on your business today.



Scaling the Enterprise IT foundation



Optimizing cost and energy efficiency



Stronger security and resilience



Future-ready AI foundations



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD



Performance at enterprise scale

Enterprise workloads are growing in scale, complexity, and business criticality. Core systems—from virtualization to analytics—must deliver higher performance and density, without compromising stability or predictability.

Traditional architectures can limit progress. As core counts rise, VMs, and memory demands increase, many organizations see higher cost and operational complexity—without meaningful performance gains.

Modernization enables innovation and scale without adding complexity.



Pressure points to consider

Why can't legacy systems keep up?

As analytics, digital services, and data-intensive workloads expand, performance expectations rise. When compute density and memory bandwidth cannot scale efficiently, core systems become innovation bottlenecks.

Why are enterprises assessing new virtualization platforms?

Rising VM counts, changing licensing models, and increasing performance needs are prompting organizations to re-evaluate their environments. Virtualization is foundational, but the economics and architecture are under renewed scrutiny.

Why are global enterprises spending more on IT than ever before?

As workload demands increase, worldwide IT spending is set to exceed \$6 trillion for the first time.¹ Much of that is driven by innovation – new services, AI initiatives – but if infrastructure investments cannot scale efficiently, costs can rise faster than capability.

How do today's infrastructure decisions shape tomorrow's flexibility?

Platform choices determine how easily organizations can adopt hybrid models, integrate emerging technologies, or prepare for AI-driven workloads. Limited flexibility today can create long-term constraints that are costly to reverse.

Which workloads are you focused on?

Drill down into the area that reflects your current priorities.



Core enterprise workloads: Databases, analytics, and other mission-critical systems.



Virtualized environments: VM growth, consolidation pressure, and shifting virtualization economics.



Hybrid and distributed IT: Coordinating data center, cloud, and edge without losing control or flexibility.

¹ Gartner. 2025, October 22. Gartner Forecasts Worldwide IT Spending to Grow 9.8% in 2026, Exceeding \$6 Trillion For the First Time. <https://www.gartner.com/en/newsroom/press-releases/2025-10-22-gartner-forecasts-worldwide-it-spending-to-grow-9-point-8-percent-in-2026-exceeding-6-trillion-dollars-for-the-first-time>.



Identifying modernization needs

Structured data, analytics platforms, ERP systems, and other mission-critical workloads sit at the heart of enterprise operations.

Your enterprise workloads	Priorities to consider
Databases: Larger datasets and more transactions.	Higher memory bandwidth and faster data processing
Analytics and BI: Real-time reporting and heavier queries.	Greater throughput and consistent performance
ERP and business systems: More users and more integrated services.	Stable performance during peak demand
Mixed enterprise workloads: More applications running on shared platforms.	Strong performance even as workloads grow





Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo


AMD

Modern Enterprise IT architecture helps:

- ✓ Increase performance per server rather than increasing footprint
- ✓ Improve memory bandwidth to support data-intensive applications
- ✓ Scale core density to enable server consolidation
- ✓ Reduce infrastructure sprawl while increasing throughput

How to move forward

-  See how Lenovo and AMD deliver the performance modern workloads need
-  Explore another challenge area
-  Discuss your modernization needs





Scalable foundations

Cost and energy efficiency

Security and resilience


Future-ready AI

Lenovo


AMD

How Lenovo and AMD support core workload performance

Lenovo ThinkSystem servers powered by AMD EPYC™ processors are designed to support high-density, data-intensive enterprise workloads with:



Up to 150% more cores than prior generations
enabling greater workload consolidation.²



DDR5 memory that delivers up to 2.6x higher memory streaming bandwidth
compared to DDR4.³



190+ world record benchmark results
across a broad range of enterprise workloads.⁴

² In the ThinkSystem SR665 V3 server, which supports up to 2x160-core processors (320 cores max), compared to the previous generation system (SR665) which supported up to 2x64-core processors (128 cores max). See <https://lenovopress.lenovo.com/lp1608-thinksystem-sr665-v3-server>.

³ Based on Lenovo internal testing comparing the memory bandwidth of Lenovo ThinkSystem SR665 V3 servers with AMD Turin (12 channels of DDR5 @ 6400MHz) to Lenovo ThinkSystem SR665 servers with AMD Milan (8 channels of DDR4 @ 3200MHz).

⁴ As of November 14, 2025. See <https://www.amd.com/en/products/processors/server/epyc/epyc-world-records.html>.



A smarter virtualization strategy

Modernizing virtualization is not just about meeting higher capacity needs. Enterprises need to improve density, flexibility, and cost predictability while maintaining performance, avoid vendor lock-in and gain efficiency and control.

70% of organizations are migrating VMs to a new hypervisor.

Key drivers are cost, complexity, and fears about vendor lock-in.⁵

“A mass exodus”

Businesses facing recent price rises are looking for alternatives to traditional hypervisors.⁶

Why virtualization?

Organizations are re-evaluating virtualization for several reasons:

- ✓ Rising licensing and subscription costs
- ✓ Growing VM counts that increase infrastructure density pressure
- ✓ Greater scrutiny on infrastructure ROI
- ✓ The need to support hybrid and distributed environments
- ✓ Concerns about vendor concentration and long-term flexibility
- ✓ Pressure to simplify management and reduce operational overhead

⁵ The New Stack (2025, July 16). 'State of Virtualization' Report Reflects Shifting Strategies. <https://thenewstack.io/state-of-virtualization-report-reflects-shifting-strategies/>.

⁶ Forbes (2025, February 6). Life After VMware: Which Alternative Is Right For You? <https://www.forbes.com/sites/stevemcdowell/2025/02/06/life-after-vmware-which-alternative-is-right-for-you/>.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Virtualization challenges?

Lenovo and AMD can help.

Escalating licensing costs

Lenovo infrastructure powered by AMD EPYC™ processors can enable higher VM density and stronger consolidation, helping reduce the number of physical servers required.

- **Up to 100% more VMs** compared to prior generations.⁷
- **Up to 61% reduced infrastructure costs** with hyperconverged Lenovo ThinkAgile HX solutions.⁸

Higher consolidation can help offset licensing exposure while improving operational efficiency.


Performance bottlenecks

With high core density and the large data bandwidth of DDR5 memory, modern Lenovo and AMD platforms deliver **up to 88% better performance for virtualization applications**, compared to previous generations.⁹

Vendor lock-in and platform inflexibility

Enterprises choose the virtualization platform that fits their strategy including hypervisors from Microsoft Azure, Nutanix, Red Hat, and VMware – in Lenovo Enterprise IT solutions powered by AMD EPYC™ processors.

Continue your evaluation



Explore another challenge area



Discuss your virtualization challenges

⁷ Internal testing comparing V1 servers to V3. 100% more VMs means faster simultaneous performance.

⁸ An ESG economic validation found that the Lenovo ThinkAgile HX series with AMD could provide a five-year Total Cost of Ownership (TCO) that is up to 61% lower than continuing with a traditional three-tier infrastructure. See: <https://lenovopress.lenovo.com/lp1990-hybrid-cloud-on-the-edge-optimizing-for-ai-with-nutanix-and-thinkagile-hx#:~:text=According%20to%20a%202024%20study,Authors.>

⁹ VMmark 3.1.1 results, as of 11/14/25. Two Lenovo ThinkSystem SR665 V3 servers, each with two AMD EPYC 9654 processors, scored 40.66 @ 42 tiles. See <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2023-06-13-Lenovo-ThinkSystem-SR665V3.pdf> for further details. Two previous generation Lenovo ThinkSystem SR665 servers, each with two AMD EPYC 7763 processors, scored 21.58 @ 24 tiles. See <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2021-05-04-Lenovo-ThinkSystem-SR665.pdf> for further details. To find out more about VMmark, visit <https://www.vmware.com/products/vmmark.html>. VMware® and VMmark® are trademarks or registered trademarks of VMware, Inc. VMware VMmark is a product of VMware, Inc.



Scaling across hybrid and edge IT

Hybrid and distributed models allow workloads to span the data center, cloud, and edge, while AI introduces new infrastructure demands.

The challenge is not only where to run workloads – it's also how to scale them efficiently, maintain control, and simplify operations across environments seamlessly.

Rising public cloud costs

are leading enterprises to reconsider where they deploy AI workloads.¹⁰

Edge and hybrid deployments bring AI to the data, instead of transporting data to public cloud.

Hybrid cloud is a strong foundation for AI, enabling up to:

- 74% faster business change execution
- 38% higher development team productivity
- 60% faster new application development lifecycles¹¹

Designing for flexibility and scale

Organizations are re-evaluating virtualization for several reasons:

- ✓ **Aligning cost models** with workload behavior
- ✓ **Simplifying infrastructure management** across environments
- ✓ **Building AI-ready foundations** that integrate private and public resources
- ✓ **Preserving flexibility across platforms** and deployment models

¹⁰ Deloitte. 2025, December 10. The AI infrastructure reckoning: Optimizing compute strategy in the age of inference economics. <https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/ai-infrastructure-compute-strategy.html>.

¹¹ Red Hat. 2025, December 8. Key considerations for 2026 planning: Insights from IDC. <https://www.redhat.com/en/blog/key-considerations-2026-planning-insights-idc>.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Explore three practical strategies

Aligning cost model with consumption

The challenge: Public cloud offers flexibility, but unpredictable costs can erode financial control.

The approach: Adopt consumption-based infrastructure models that deliver cloud-like economics with on-premises control.

Enterprises can shift to cloud economics while maintaining the control of on-premises infrastructure, with **Lenovo TruScale Infrastructure as a Service**, solutions are backed by Lenovo's global scale and support, and enable customers to scale infrastructure flexibly.

Simplifying operations across distributed IT

The challenge: Hybrid and distributed environments introduce operational fragmentation. Separate tools, inconsistent visibility, and manual lifecycle management increase overhead.

The approach: Consolidate management through unified, cloud-based platforms that provide visibility and lifecycle control across edge, data center, and cloud.

Lenovo XClarity One management software delivers cloud-based infrastructure management. It simplifies deployment, monitoring, and lifecycle management across environments.

Building seamless hybrid AI foundations

The challenge: AI workloads introduce new demands across inference, training, and data pipelines. Running them efficiently requires integration across private and public environments.

The approach: Deploy pre-integrated, full-stack infrastructure that supports AI use cases across deployment models.

Lenovo Hybrid AI solutions integrate AI seamlessly across private and public models, leveraging an open AI ecosystem, full-stack integration, and the high performance of AMD Instinct™ GPU options in Lenovo ThinkSystem servers.

Take your next step



Explore another challenge area



Learn more about Lenovo and AMD solutions



Return to the start



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Optimizing cost and energy efficiency

Global IT spending is accelerating quickly, driven chiefly by investment in AI and data center systems as well as rising memory prices (CIO, 2026).¹²

But the challenge isn't simply to spend less. When infrastructure growth outpaces improvements in performance-per-watt, consolidation, or utilization, cost becomes a constraint on innovation. Enterprises want every investment to deliver measurable returns — efficiently.

¹² CIO (2026, April 23). Gartner(r) ups IT spending growth to 13.5% in revised forecast. <https://www.cio.com/article/4162831/gartner-has-raised-its-forecast-for-global-it-spending-growth-in-2026-by-nearly-three-percentage-points.html>



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Diagnosing IT inefficiencies

Review the symptoms below to see which ones reflect your current environment and learn how modernization can help.

**IT footprint is
expanding – faster than
performance gains**

Shift the focus from adding servers to increasing performance per node.

**Energy and
cooling demands
are rising**

Power-efficient processors and modern cooling strategies, such as liquid cooling, can materially reduce energy intensity.

**Compute and
memory are
underutilized**

Higher core density and smarter workload placement can unlock existing value.

**Infrastructure is designed
for peak demand – not
typical demand**

Flexible consumption models can align cost more closely with real usage.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Where can efficiency improve most?

Choose the priority that aligns with your efficiency goals, then return to explore further opportunities.



Consolidating infrastructure: Reducing server footprint through higher core density and improved workload consolidation.



Reducing energy intensity: Addressing energy consumption levels and costs.



Aligning cost with consumption: Finding ways to scale and finance IT more flexibly.






Increasing efficiency with consolidation

In many enterprises, infrastructure has expanded incrementally over time. New workloads are added, refresh cycles replace aging hardware one system at a time, and virtualization grows steadily. The result is often a larger server footprint without a proportional increase in usable performance.

By increasing core density, raising VM density, and improving overall utilization, organizations can run more workloads on fewer physical systems, improving cost control, simplifying operations, and creating headroom for future growth without increasing IT footprint.

What would consolidation change in your business?

Consolidation is designed to deliver multiple, measurable improvements in cost, capacity, and operational efficiency. Decisions about core density, virtualization strategy, and refresh cycles shape not only infrastructure design, but how effectively IT supports the wider business.





Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Reduce data center footprint

Running more workloads on fewer servers reduces rack count and overall floor space.

This helps to lower colocation and real estate costs, simplifies capacity planning, and reduces the need to expand facilities.

Lower energy consumption

Running fewer physical systems can mean smaller power and cooling needs.

Lenovo Neptune® cooling options on ThinkSystem servers powered by AMD EPYC™ processors help deliver improved energy efficiency.

Simplify daily operations

Fewer devices to patch, monitor, secure, and support.

This reduces administrative effort and lowers ongoing operational overhead.

Power to run bigger workloads


Higher core counts and larger memory capacity can run bigger databases, analytics platforms, and business systems on a single server.

Readiness for modern workloads like AI

A smaller server footprint means fewer devices to patch, monitor, secure, and support, helping to reduce administrative effort and lower ongoing operational overhead.

Next: How much consolidation could you achieve?

Click below to learn more about how consolidation works, or explore a different topic.



How consolidation works – and what's possible



Explore a different efficiency topic



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

How consolidation works – and what's achievable

Modern platforms designed to allow more workloads to run per node.

How does consolidation help?



Run more workloads per server

Modern processors with greater core counts allow higher VM and workload density per node, increasing useful compute without expanding footprint.



Support heavier workloads without adding nodes

Stronger per-core performance enables demanding applications to run efficiently at higher utilization levels, reducing the need for additional physical systems.



Sustain data-intensive applications at scale

Expanded memory capacity and high-bandwidth DDR5 memory help maintain performance for analytics, databases, and AI-driven workloads as consolidation ratios increase.



Maintain performance at higher VM density

Systems engineered for dense virtualization environments help sustain stability as consolidation increases.



What level of consolidation could you achieve?

Lenovo solutions powered by AMD EPYC™ processors can deliver up to 3:1 server consolidation¹³ when modernizing older-generation systems. In practical terms, this means replacing three older servers with one modern system.

The Lenovo ThinkSystem SR665 V3 servers tested, each with two 4th Gen AMD EPYC™ 9654 processors, delivered through:

- ✓ High core counts in AMD EPYC™ 9654 processors
- ✓ Strong per-core performance
- ✓ Large capacity and bandwidth of DDR5 memory (up to 6 TB per server)

Explore more ways to improve efficiency




Reducing energy use and costs



Aligning costs to consumption



Explore a different efficiency topic



¹³ VMMark 3.1.x results, as of 11/14/25. Two Lenovo ThinkSystem SR665 V3 servers, each with two AMD EPYC 9654 processors, scored 40.66 @ 42 tiles. See <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2023-06-13-Lenovo-ThinkSystem-SR665V3.pdf> for further details. Two HPE servers, each with two AMD EPYC 7702 processors, scored 12.78 @ 14 tiles. See <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2019-08-07-HPE-ProLiant-DL385Gen10.pdf> for further details. To find out more about VMmark, visit <https://www.vmware.com/products/vmmark.html>. VMware® and VMmark® are trademarks or registered trademarks of VMware, Inc. VMware VMmark is a product of VMware, Inc. Actual consolidation results will vary based on many factors.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Reducing energy intensity in modern Enterprise IT

Energy consumption is becoming a defining constraint in enterprise infrastructure strategy. As AI adoption accelerates and data volumes grow, compute demand is rising faster than traditional power and cooling models were designed to support.

See how IT energy consumption is changing

How much energy does AI need?

AI data centers use a lot of energy, and globally this amounted to **90 terawatt-hours in 2026**. That is 1/7 of all data center energy – and a **10x increase** since 2022.¹⁴


How much of a data center's energy is used for cooling?

Air-based cooling uses much more data center energy than many people realize. A Deloitte report found that air-based cooling alone uses **up to 40%** of a typical data center's electricity consumption – almost half of the energy costs enterprises pay.¹⁵

How fast is data center energy consumption growing?

Data centers are using more energy than before. According to an International Energy Agency (IEA) forecast, electricity usage across all data centers **doubled** from 2024 to 2026.¹⁶

Explore considerations that can impact your efficiency strategy.



The value of liquid cooling



Why performance per watt matters

¹⁴ Deloitte. 2024, November 19. As generative AI asks for more power, data centers seek more reliable, cleaner energy solutions. <https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2025/genai-power-consumption-creates-need-for-more-sustainable-data-centers.html>. "Worldwide, AI data centers' annual power consumption is expected to reach 90 terawatt-hours by 2026 (or roughly one-seventh of the predicted 681 TWh of all data centers globally), roughly a tenfold increase from 2022 levels."

¹⁵ Deloitte. 2024, November 19. As generative AI asks for more power, data centers seek more reliable, cleaner energy solutions. <https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2025/genai-power-consumption-creates-need-for-more-sustainable-data-centers.html>. "Two broad areas drive most of the electricity consumption in a data center: computing power and server resources like server systems (roughly 40% data center power consumption) and cooling systems (consume 38% to 40% power)"

¹⁶ Data Centre Dynamics (2024). Global data center electricity use to double by 2026 - IEA report. <https://www.datacenterdynamics.com/en/news/global-data-center-electricity-use-to-double-by-2026-report/>.



Why liquid cooling is valuable

How much additional compute can your current facility support before power or cooling becomes the limiting factor?

Liquid cooling changes the equation. Instead of pushing chilled air through an entire room, coolant is circulated directly across cold plates attached to CPUs, GPUs, and other high-heat components.

Water transfers heat far more efficiently than air, it can capture and remove thermal energy at source, supports higher rack densities, and operates with warmer inlet temperatures. Resulting in greater cooling efficiency, improved thermal stability, and more usable compute per rack.

Choose your focus



Discuss your energy strategy with Lenovo specialists



Explore a different efficiency topic

Lenovo Neptune brings liquid cooling to the enterprise

Lenovo Neptune Direct Water Cooling can be deployed in enterprise solutions, delivering measurable efficiency improvements in standardized implementations.



Improve energy efficiency by up to 40% compared to air-cooled systems¹⁷



Capture up to 100% of system heat¹⁸



Maximize compute density to support demanding workloads for technical computing, analytics, simulation, life sciences, energy, and engineering

¹⁷ See <https://lenovopress.lenovo.com/lp1612-lenovo-thinksystem-sd665-v3-server>.

¹⁸ Depending on specific environment. See <https://lenovopress.lenovo.com/lp1612-lenovo-thinksystem-sd665-v3-server>.



Performance per watt – what difference does it make?

When processors deliver more compute for each watt consumed, the impact can compound across IT with:

- ✓ **More workloads per server**
- ✓ **Lower total energy consumption**
- ✓ **Reduced cooling demand**
- ✓ **Better management of operating costs**

In environments scaling virtualization, analytics, or AI, performance per watt becomes a financial lever — not just a technical metric.

Exceptional efficiency with AMD EPYC™ processors

AMD EPYC™ processors deliver outstanding performance per watt, at the heart of some of the industry's most energy-efficient x86 servers.¹⁹

Define your next step

-  Discuss your energy strategy with Lenovo specialists
-  Explore a different efficiency topic

¹⁹Lenovo ThinkSystem servers with AMD processors have eight of the top ten SPECpower_ssj®2008 results | [Rank - overall ssj_ops/watt] #3 - 42,312; #4 - 42,182; #5 - 41,526; #6 - 41,427; #7 - 40,785; #8 - 40,528; #9 - 40,278; #10 - 40,013 results as of 11/14/2025. See <http://www.spec.org> for additional information. SPEC® and SPECpower_ssj are registered trademarks of the Standard Performance Evaluation Corporation.



Does your cost structure match real usage?

Capacity is provisioned for end-of-quarter spikes, seasonal surges, and projected growth. But peak conditions are temporary. As workloads become more dynamic, the gap between provisioned capacity and actual usage can become more visible.

IT modernization presents an opportunity to ensure that spending reflects real demand, rather than worst-case scenarios.

Move forward, efficiently



Discuss your efficiency goals with Lenovo specialists



Explore a different efficiency topic



Return to the start

Move from fixed capacity to flexible economics

Lenovo TruScale Infrastructure as a Service enables enterprises and SMB customers to freely select their financing option, be it predictable- or consumption-based financing.

With TruScale, organizations can:



Align cost with real demand rather than provisioned peak capacity



Scale infrastructure up or down as workloads change




Retain on-premises control and governance with cloud-like economics




Reduce lifecycle and operational complexity through modular services

This approach brings financial flexibility to Enterprise IT — allowing infrastructure strategy and business demand to move in step.



83%

of security leaders
witnessed a surge in
cyberattacks in 2025
(KPMG, 2025)²⁰



93%

of leaders are concerned
about the financial and
organizational impacts
of outages²¹

Modernizing security and resilience strategies

Security and uptime are not new priorities for Enterprise IT leaders. What has changed is the scale of impact.

As organizations consolidate workloads, centralize data, and depend on continuous digital services, businesses can be affected by a single outage.

Modernization creates an opportunity to address this challenge.

²¹ CockroachDB (2025). The State of Resilience 2025. <https://apexassembly.com/wp-content/uploads/2025/06/State-of-Resilience-2025-Report-FINAL.pdf>. "Leaders are worried: 93% of leaders are concerned about the financial and organizational impacts of outages."

²² Lenovo ranks 8th in the Gartner® Supply Chain Top 25 for 2025. See <https://news.lenovo.com/pressroom/press-releases/lenovo-ranks-8th-in-the-gartner-supply-chain-top-25-for-2025/>.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

What's at risk in your environment?

Reducing server downtime

Downtime risks can increase as servers age. Lenovo is #1 in reliability²² with predictive monitoring to identify issues early and reduce the likelihood of service interruption.

Protecting against ransomware

Attacks increasingly target firmware and system layers below the OS. Hardware-rooted security and automated protection mechanisms help contain threats before they spread.

Keeping data safe from threats

As more workloads share infrastructure, having robust security capabilities becomes critical. Hardware-enforced encryption and virtualization security features help protect data and maintain separation between workloads.

Making server management more secure

The management layer of Enterprise IT controls the entire infrastructure. Secure access controls, encryption, and hardened management platforms reduce exposure at this critical control point.

Increasing trust in IT hardware

Infrastructure integrity begins before deployment. Secure manufacturing processes and hardware-based validation help ensure systems arrive trusted and uncompromised.

Go deeper into security and resilience



See how modern infrastructure embeds security




Explore infrastructure reliability

²² ITIC 2025 Global Server Hardware, Server OS Reliability Report. See <https://lenovopress.lenovo.com/lp1117-itic-reliability-study>

[Scalable foundations](#)[Cost and energy efficiency](#)[Security and resilience](#)[Future-ready AI](#)[Lenovo](#)[AMD](#)

Where should enterprise security strategy begin?

Modern Enterprise IT embeds security features across multiple layers — from silicon and firmware through virtualization and system management. This layered approach is designed to reduce the impact of a single vulnerability, strengthens workload isolation, and helping contain threats before they propagate.





Building a multi-layer security strategy

Security is strongest when it is embedded across the entire infrastructure stack.

Explore each infrastructure layer to see how modern Enterprise IT builds in security features, with examples from Lenovo and AMD.

Manufacturing and supply chain

Infrastructure integrity must begin before systems are deployed.
Lenovo hardware is built across 30+ secure global manufacturing sites, within a secure supply chain ranked 8th in the Gartner® Supply Chain Top 25 for 2025.²³

Enterprise servers

Lenovo and AMD solutions offer an immutable hardware and software root of trust, designed to prevent unauthorized or malicious code from gaining control of your systems.
AMD Infinity Guard™, built into AMD EPYC™ processors, offers advanced capabilities required to help defend against internal and external threats and keep your data safe.²⁴

Virtualized workloads

Strong workload isolation is required to protect sensitive data in virtualized environments.
AMD Infinity Guard™ includes Secure Encrypted Virtualization (SEV), a multi-layered, silicon-based security architecture that enables encrypted virtual machines and stronger workload isolation.²⁴

Infrastructure management

Management systems must be protected to keep the entire IT infrastructure secure.
Lenovo XClarity One software delivers comprehensive security features for user account management, encryption, security standards compliance, and communication for all managed data storage solutions.

Data and applications

Ransomware attacks target both data integrity and application availability.
Lenovo and AMD solutions can offer Automated Ransomware Protection, leveraging machine learning models and tamper-proof snapshots to detect and mitigate ransomware attacks and protect against unauthorized data deletions.

²³ Lenovo ranks 8th in the Gartner® Supply Chain Top 25 for 2025. See <https://news.lenovo.com/pressroom/press-releases/lenovo-ranks-8th-in-the-gartner-supply-chain-top-25-for-2025/>.

²⁴ AMD Infinity Guard features vary by AMD EPYC™ processor generations and/or series. Infinity Guard security features must be enabled by server OEMs and/or Cloud Service Providers to operate. Check with your OEM or provider to confirm support of these features. Learn more about Infinity Guard at <https://www.amd.com/en/technologies/infinity-guard.GD-183A>



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Modernize your security foundation

Security architecture decisions made today determine your exposure tomorrow. Engage Lenovo and AMD specialists to assess your current environment and define a multi-layer strategy aligned to your risk profile and compliance requirements.



Discuss your security strategy with Lenovo



Explore reliability and uptime



Explore a different Enterprise IT security topic





Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Why reliability matters more than ever

Downtime is often measured in minutes. Its impact can be felt for much longer.

When systems fail, the cost is not limited to lost transactions. It ripples across the organization: disrupting operations, diverting IT teams into reactive firefighting, delaying projects, and eroding user confidence.

The impacts of downtime can include:

- ✓ **Lost revenue and service interruption** - Customer transactions halt, digital services become unavailable, and revenue-generating systems go offline.
- ✓ **Operational disruption** - Upstream and downstream processes are affected, creating bottlenecks beyond the original failure point.
- ✓ **Expensive emergency remediation** - IT teams shift from strategic initiatives to urgent troubleshooting and recovery efforts.
- ✓ **Regulatory or compliance exposure** - Service outages can trigger reporting requirements, penalties, or contractual breaches.
- ✓ **Reputational damage** - Reliability issues weaken confidence among customers, partners, and stakeholders.

When modernizing IT, reliability matters. The platform you choose affects how often systems stay up – and how well the business keeps running.





Maximizing server reliability

Lenovo ThinkSystem is consistently rated the most reliable x86 server line in global customer surveys. According to the ITIC Global Server Hardware Reliability Report, ThinkSystem is ranked:

- ✓ **#1 for reliability** among x86 servers for 12 consecutive years²⁵
- ✓ **#1 for x86 server security design** for five consecutive years²⁵
- ✓ **Lowest downtime** from security attacks²⁵

This reliability is reinforced through predictive maintenance capabilities and Lenovo Premier Support. By identifying potential issues early, organizations can often resolve problems before they lead to service disruption.

For enterprises modernizing infrastructure, this translates into fewer unplanned outages, lower remediation costs, and greater operational continuity.

Build infrastructure you can rely on

Engage Lenovo and AMD specialists to assess your current infrastructure and define a reliability strategy that reduces downtime risk and supports continuous operations.



Discuss your IT resilience challenges with Lenovo



Explore Enterprise IT security strategies



Return to the start

²⁵ Lenovo ThinkSystem servers are consistently ranked #1 for reliability and the lowest downtime among x86 server lines in the ITIC Global Server Hardware, Server OS Reliability Report. See <https://lenovopress.lenovo.com/lp1117-iticreliability-study>.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Building the right foundation for AI

AI has become a central business priority. Enterprises are already using AI in production to automate processes, improve decisions, and unlock new value from data.

Implementing AI requires infrastructure designed for new compute demands, data pipelines, and deployment models.

At the same time, organizations want to move forward carefully – matching investment to real needs and ROI at every stage with the right infrastructure for evolving demands.

The #1 global business priority in 2026 is to enhance, innovate, or reinvent business with AI, overtaking productivity and profit growth.²⁶

96% of organizations plan to increase AI investment in the next 12 months, with spending growth averaging +13%.²⁶



²⁶ IDC CIO Playbook 2026 Survey, commissioned by Lenovo, n=3,120.

[Scalable foundations](#)[Cost and energy efficiency](#)[Security and resilience](#)[Future-ready AI](#)[Lenovo](#)[AMD](#)

How should your infrastructure evolve for AI?

Adopting AI at your own pace

Many organizations begin with AI inference rather than large-scale model training. Inference-optimized platforms allow teams to experiment and deploy AI workloads, using pre-trained or fine-tuned models, without making large upfront investments.

Using AI at the edge and across hybrid cloud

When AI runs close to where data is created and used, instead of transporting data to and from an AI processing center, enterprises can improve performance and efficiency. Modern hybrid and edge AI platforms enable consistent infrastructure and operations across data centers, cloud environments, and edge locations.

Scaling AI as workloads grow

Platforms designed with scalability in mind allow organizations to increase compute capacity and performance without rearchitecting environments or disrupting existing operations.

[Explore inference-ready infrastructure for early-stage AI](#)[Explore scalable, high performance AI platforms](#)[Explore hybrid and edge AI infrastructure](#)



Building early value from AI

Many organizations want to move forward with AI, but they are also cautious about committing to large investments before clear business value is proven.

A common assumption is that adopting AI requires training custom models and deploying specialized infrastructure from the outset. In practice, much of the value from AI comes from running inference – applying existing models to enterprise data to generate insights, predictions, and automated actions.

What AI use cases could you deploy today?

Early AI successes often focus on targeted business problems where models can be applied to operational data. Many of these use cases rely on inference: using trained models to generate predictions, recommendations, or automated decisions.

Real-time decision support

AI models can analyze events as they occur, helping organizations detect fraud, flag anomalies, or guide operational decisions in real time.

Personalized experiences

Inference models can tailor recommendations, content, and interactions for each user, improving engagement and conversion.

Process automation

AI can automate repetitive tasks such as document classification, claims processing, and customer service workflows.

Predictive operations

AI models can analyze operational data to anticipate equipment failures, forecast demand, or identify emerging risks.

Quality inspection and defect detection

Computer vision models can analyze images or video streams to identify defects and ensure quality at scale.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Starting with low-risk, inference-optimized AI

Lenovo infrastructure powered by AMD EPYC™ processors enables organizations to run AI inference efficiently on CPUs. This approach is designed to allow enterprises to begin deploying AI workloads without immediately investing in large GPU environments, while still delivering real business value.

For example, the Lenovo ThinkSystem SR665 V3 server powered by AMD EPYC™ processors provides:

- ✓ **CPU-based AI inference**
- ✓ **Seamless scalability**, allowing AI workloads to expand from CPU-based inference to larger GPU-accelerated systems as requirements grow
- ✓ **DDR5 memory bandwidth** that helps handle data-intensive AI and analytics tasks

Take the next step



Discuss your AI infrastructure strategy with Lenovo and AMD specialists



Explore hybrid and edge AI



Prepare for large-scale AI workloads





Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Deploying AI across hybrid and edge environments

As AI adoption grows, many organizations find that workloads do not run in a single location. Data may be generated at the edge, processed in the data center, and integrated with cloud services. Supporting AI across these environments requires infrastructure that can operate consistently across distributed locations.



80%

**of enterprises will
deploy distributed edge
AI infrastructure by 2027**

to improve the latency
and responsiveness
of AI applications²⁷



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Strategic directions for AI deployment

Deploying AI closer to the data

Some AI use cases require decisions to be made where data is created – in stores, factories, telecommunications networks, or remote operational environments.

Lenovo ThinkEdge SE455i V3 servers powered by AMD EPYC™ processors bring AI inference directly to these locations. The platform enables organizations to run real-time AI workloads close to operational data while maintaining consistency with enterprise infrastructure. Enterprises can:



Enable real-time decisions where data is generated, by running AI inference directly at the edge



Extend AI beyond the data center, supporting retail, industrial, and telecom deployments



Deploy powerful AI in space-constrained environments, with high core density and DDR5 memory



Maintain operational consistency across environments, using platforms aligned with data center infrastructure

Deploying an end-to-end hybrid AI platform

Many organizations need a platform that supports AI across its entire lifecycle – from experimentation and development through deployment and scaling.

Lenovo Hybrid AI provides a portfolio of full-stack solutions designed to integrate AI across data center, cloud, and edge environments. Built on open architectures and optimized Lenovo infrastructure, it enables organizations to deploy scalable AI solutions across multiple environments while maintaining operational consistency. Enterprises can:



Support the full AI lifecycle, from experimentation and development to enterprise-scale deployment



Run AI across data center, cloud, and edge environments, choosing the right location for each workload



Scale AI initiatives confidently, with infrastructure designed for growing analytics and AI workloads



Maintain flexibility through an open AI ecosystem, integrating preferred tools, frameworks, and software stacks

Where next?



Starting small with inference-focused AI infrastructure



Explore infrastructure for large-scale AI workloads



Explore a different AI readiness topic



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

Scaling AI infrastructure for advanced workloads

AI infrastructure needs can increase rapidly.

Larger datasets, more complex models, and broader deployment across the enterprise all place greater demands on compute, memory, and system architecture. Meeting these demands requires a platform designed to scale.

Modern AI infrastructure must support high-performance compute, integrate specialized GPUs where needed, and expand capacity without forcing disruptive changes to the underlying environment.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Lenovo

AMD

The priorities that matter for large-scale AI infrastructure

AI initiatives place different demands on infrastructure.

Scaling more easily

AI programs grow fastest when they can expand on a consistent platform rather than introducing new environments at each stage.

Lenovo ThinkSystem servers powered by AMD EPYC™ processors provide a unified infrastructure foundation that supports AI development, inference, and large-scale deployments across the enterprise.

Accelerating advanced workloads

Many advanced AI workloads require specialized acceleration to process large models and complex data efficiently.

Lenovo ThinkSystem servers support a range of GPU options including AMD Instinct™ GPUs, enabling high-performance processing for generative AI, large-scale analytics, and scientific computing workloads.

Running large AI models

Large models and high-volume inference workloads require dense compute environments that maximize performance.

Platforms such as the Lenovo ThinkSystem SR675i V3 supports up to eight PCIe Gen5 GPUs in a 3U design, delivering powerful AI processing capacity for demanding workloads.

Scaling infrastructure flexibly

AI infrastructure requirements can expand quickly as projects move from experimentation to production.

Lenovo TruScale Everything as a Service (XaaS) enables organizations to scale infrastructure capacity while shifting from large upfront capital investments to a flexible pay-for-use model.

Maintaining flexibility with an open AI ecosystem

AI teams rely on a wide range of frameworks, tools, and deployment environments.

Lenovo infrastructure powered by AMD technology supports an open ecosystem across operating systems, hypervisors, and AI software stacks, allowing organizations to deploy AI workloads without platform lock-in.



Scalable foundations

Cost and energy efficiency

Security and resilience

Future-ready AI

Continue your AI infrastructure journey



Discuss your large-scale AI infrastructure strategy with Lenovo and AMD specialists



Start small with inference-focused AI infrastructure



Explore hybrid and edge AI deployments



Return to the start

Lenovo

AMD

