

Nutanix and Lenovo: A High-Density, AI-Ready Backbone

A turnkey hybrid multicloud platform powered by AMD EPYC™, delivered on Lenovo ThinkAgile HX.

IT leaders face strict datacenter power limits, rising core-based software licensing costs, and growing pressure to operationalize agentic AI on sensitive corporate data. Nutanix and Lenovo answer with an open, high-density hybrid multicloud foundation powered by AMD. The Nutanix Cloud Platform runs on 5th Gen AMD EPYC™ processors across Lenovo ThinkAgile HX V3 integrated systems, factory pre-validated with a single point of contact for purchase, deployment, and support. Up to 160 cores per system and PCIe Gen 5 I/O consolidate aging server estates, reclaim trapped power and floor space, and lower total cost of ownership. The result is a turnkey, GPU-ready backbone for virtualization modernization and secure enterprise AI.

Key Benefits

- **Consolidate up to 8:1**, from 1,000 x86 servers to 127 2P AMD EPYC™ 9965 nodes,¹ cutting TCO up to 79%.²
- Reclaim trapped power and floor space with up to **66% better energy efficiency per watt**.³
- Deploy faster with factory pre-validated ThinkAgile HX integrated systems and a single Lenovo support contact.
- Support GPU-rich, liquid-cooled configurations for GenAI and agentic AI with high thermal efficiency.
- Protect sensitive AI workloads with silicon-level confidential computing using AMD Infinity Guard and SEV-SNP.

Modernize Virtualization on a Turnkey HX Foundation

The shift away from legacy virtualization, accelerated by recent licensing changes, has enterprises seeking an exit from the virtualization tax and the rigidity of locked-in hypervisors. Nutanix and Lenovo provide a clear, low-risk path. Nutanix Move migrates critical workloads from legacy platforms with minimal downtime and no application refactoring, preserving license portability and long-term cost predictability.

Lenovo ThinkAgile HX V3 systems arrive factory integrated and pre-validated with Nutanix software, giving teams a single Lenovo contact for purchase, deployment, and support. They support 5th Gen EPYC™ 9005 (Turin) processors, with up to 160 cores

per system. The 1U HX645 V3 targets rack-dense virtualization, VDI, and ROBO sites, while the 2U HX665 V3 handles heavy enterprise applications and high-performance databases such as Oracle 21c and Microsoft SQL Server.

All-flash, hybrid, and storage-dense configurations span workloads from VDI to big data analytics, with the HX665 V3 Storage variant adding expanded drive capacity. Because higher performance-per-core shrinks the socket and core footprint, enterprise software licensing for SQL, Oracle, and SAP drops as organizations consolidate aging racks. The reclaimed capital, power, and floor space can then fund the next phase of enterprise AI.

Stay Agile with Resilient, Cloud-Native IT

Beyond consolidation, Nutanix and Lenovo deliver a flexible platform built to streamline modernization and reduce complexity. Lenovo industry-leading hardware and lifecycle services combine with the Nutanix Cloud Platform to create a resilient environment that simplifies IT management, strengthens data security, and accelerates time-to-market.

The platform is built for cloud-native applications that are inherently scalable, resilient, and easy to manage. Teams deploy and operate containerized workloads on Nutanix Kubernetes Platform alongside virtual

machines, with a consistent operating model from core to edge. Native data services, including snapshots, replication, and self-healing from drive and node failures, help ensure business continuity without bolt-on tooling.

Because the same platform scales from a single ROBO node to large hybrid cloud environments, organizations optimize resources and grow without re-architecting or retraining teams. The result is an agile, cost-effective, future-ready foundation that keeps enterprises competitive and ready to adopt secure AI, all on infrastructure powered by AMD EPYC™.

Power Secure, Open AI on a GPU-Ready Backbone

Moving from simple chatbots to agentic AI means autonomous systems that reason across multiple steps, retrieve sensitive data, and execute workflows on their own. That demands constant access to protected memory, and proprietary AI stacks tend to inflate cost, create governance gaps, and lock customers in.

Nutanix and Lenovo take a different approach. The 12-channel memory bandwidth of AMD EPYC™ processors runs fast local inference and data preprocessing on the CPU, letting teams begin agentic work before committing to dedicated accelerators,

then scale to AMD Instinct™ GPUs as demand grows. The GPU-rich ThinkAgile HX650a V4, cooled by Lenovo Neptune liquid cooling, sustains high GPU density with strong thermal efficiency for GenAI and agentic workloads.

Data sovereignty holds throughout: the Nutanix AI Gateway governs multi-agent workflows while AMD SEV-SNP encrypts and isolates individual VMs and AI threads at the silicon level. Standardizing on the open AMD ROCm™ ecosystem and Nutanix Kubernetes Platform keeps the environment free of proprietary lock-in.

One Turnkey Foundation for Modernization and AI

Nutanix, Lenovo, and AMD give enterprises one turnkey foundation to modernize virtualization, consolidate infrastructure, and operationalize secure AI, all without lock-in. Reclaim your power, your budget, and your AI runway. Talk with your Nutanix or Lenovo team to scope a deployment.

[Nutanix.com/Lenovo](https://nutanix.com/lenovo)

[Nutanix.com/AMD](https://nutanix.com/amd)

¹AMD EYPC™ Claim: See 9xx5TCO-002a.

²AMD EYPC™ Claim: See 9xx5TCO-018.

³AMD EYPC™ Claim: See 9xx5-134.

©2026 Nutanix, Inc. All rights reserved. Nutanix, the Nutanix logo, and all product and service names mentioned herein are registered trademarks or trademarks of Nutanix, Inc. in the United States and other countries. All other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s).