

Nutanix and HPE: A Sustainable, Consumption-Based Hybrid Cloud

A turnkey hybrid multicloud platform powered by AMD EPYC™, delivered on HPE ProLiant and GreenLake.

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 HPE answer with an open, high-density hybrid multicloud foundation powered by AMD. The Nutanix Cloud Platform comes factory pre-installed on HPE ProLiant servers running AMD EPYC™ processors, available as a service through HPE GreenLake for consumption-based economics. Up to 96 cores per socket and high performance-per-watt consolidate aging server estates, reclaim trapped power and floor space, and lower total cost of ownership. The result is a sustainable, turnkey platform that runs any application anywhere and prepares the enterprise for secure 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
- Shift to flexible, consumption-based economics with HPE GreenLake, paying for capacity as you use it.
- Run any application anywhere on factory pre-installed HPE ProLiant, with one-click, non-disruptive lifecycle management.
- Protect workloads with HPE Silicon Root of Trust and silicon-level AMD SEV-SNP confidential computing.

Modernize Sustainably with Flexible HPE GreenLake Economics

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 HPE 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.

The Nutanix Cloud Platform ships factory pre-installed on HPE ProLiant servers, so teams deploy quickly and manage the full hardware and software lifecycle

through one-click, non-disruptive updates. Underneath, AMD EPYC™ processors deliver high core density and strong performance-per-watt, replacing multiple racks of aging hardware with a compact, efficient cluster that lowers Power Usage Effectiveness and supports corporate sustainability goals.

HPE GreenLake adds flexible, consumption-based economics, so organizations pay for capacity as they use it rather than over-provisioning. The reclaimed capital, power, and floor space, combined with pay-as-you-grow flexibility, can then fund the next phase of enterprise AI.

Accelerate Transformation with Resilient, Cloud-Native IT

Beyond consolidation, Nutanix and HPE deliver an integrated platform that simplifies IT operations and drives business performance. HPE industry-leading hardware combines with Nutanix cloud infrastructure software to reduce complexity, mitigate risk, and accelerate time-to-market, so teams spend less time managing infrastructure and more time innovating.

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.

HPE GreenLake and Nutanix Prism streamline operations across the estate, giving teams the visibility to make faster, data-driven decisions. With greater scalability and optimized resources, enterprises hold a competitive edge, navigate market shifts with confidence, and drive sustainable, long-term growth, all on infrastructure powered by AMD EPYC™.

Operationalize Secure, Open AI Anywhere

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 HPE 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 GPUs, then scale

to AMD Instinct™ accelerators as demand grows. Because the platform runs any application anywhere, those AI services move across data center, edge, and cloud without re-architecting.

Data sovereignty holds throughout: the Nutanix AI Gateway governs multi-agent workflows, HPE Silicon Root of Trust anchors platform integrity, and 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 Sustainable Foundation for Modernization and AI

Nutanix, HPE, and AMD give enterprises one sustainable 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 HPE team to scope a deployment.

[Nutanix.com/HPE](https://nutanix.com/HPE)

[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).