



Nutanix and Netris Announce Strategic Collaboration to Deliver a Multitenant Infrastructure Stack from Compute to Network for Neoclouds and NVIDIA Cloud Partner Environments

October 21, 2026

News Summary

- Nutanix and Netris have announced a strategic collaboration to deliver a jointly engineered multitenant infrastructure stack designed specifically for neoclouds and enterprises — combining Nutanix Cloud Platform and its proven capabilities with Netris Network Automation, Abstraction, and Multitenancy (NAAM).
- The joint solution will enable neoclouds to launch secure, NVIDIA GPU-accelerated IaaS offerings faster, create new revenue opportunities and meet NVIDIA's technical requirements for delivering multitenant services to neocloud customers.
- Enterprises will be able to run business-critical AI workloads on multitenant GPU infrastructure with hard isolation (enforced on networking hardware) between tenants.
- The combined solution will deliver automation and abstraction across compute and networking, enabling infrastructure to be provisioned and consumed through a cloud-like experience.

Why It Matters

An offering that integrates Nutanix technology with Netris technology will enable neoclouds and enterprises to bring secure, multitenant GPU services to market faster, expand their service offering, and accelerate time to revenue. By simplifying how AI infrastructure is deployed and consumed across bare-metal and virtualized environments, neoclouds will be able to scale more efficiently and improve margins as their business grows.

BERLIN, Oct. 21, 2026 (GLOBE NEWSWIRE) -- [NVIDIA GTC](#) - Nutanix (NASDAQ: NTNX), a hybrid cloud leader and AI platform company, and Netris, the leading provider of Network Automation, Abstraction, and Multitenancy (NAAM) for AI infrastructure, today announced a strategic collaboration to deliver a jointly engineered infrastructure stack designed specifically for neoclouds and enterprises. By combining Nutanix Cloud Platform and Nutanix's expertise in cloud infrastructure and compute orchestration with the Netris NAAM platform, the joint solution will provide neoclouds and enterprises with a scalable, secure, and multitenant solution across the infrastructure.

As the demand for AI infrastructure accelerates, neoclouds and enterprises are leading the deployment of large GPU clusters to support a broad range of AI workloads. Delivering this infrastructure as a service requires both soft isolation and hard isolation (enforced on networking hardware) that extends across every layer of the infrastructure, from the compute layers and data processing units to the physical switches.

An Automated Multitenant Infrastructure Stack for AI Clouds

Bringing together Nutanix's cloud infrastructure capabilities with Netris network automation and multitenancy, the joint solution will provide neoclouds and enterprises with a secure, multitenant AI infrastructure stack across bare-metal and virtualized environments, enabling providers to instantly onboard customers, maintain strict isolation between tenants from compute through the physical network, and deliver AI services through a cloud-like self-service experience.

Key highlights of the joint solution for neoclouds and enterprises will include:

- **Automated Network Fabric Operations:** Simplifies AI infrastructure deployment by automatically configuring the physical network as tenant resources are provisioned, reducing manual effort, accelerating time to value, and improving operational consistency.
- **Unified Multitenant Networking:** Delivers a consistent networking experience across physical and virtual environments by translating Nutanix tenant requirements into the appropriate underlying network services for bare-metal and virtualized GPU nodes.
- **Self-Service Cloud Networking:** Provides cloud-grade networking constructs — VPCs, elastic IPs, NAT, and L4 load balancers — through a web console, REST API, and Terraform, enabling infrastructure to be consumed through a self-service, cloud-like experience.
- **High-Performance AI Factory Networking:** Supports the connectivity for demanding AI workloads, enabling efficient edge, East-West, North-South, scale-up, and host-based networking traffic flows.
- **End-to-End Tenant Isolation:** Extends secure hard and soft tenant isolation across the entire network — from virtual workloads through the physical switching layer — supporting secure, reliable GPU-as-a-Service offerings on shared infrastructure.

Executive Commentary

"Neoclouds and [NVIDIA Cloud Partners](#) are at the forefront of the AI transformation, but building the cloud operating stack required to monetize GPU clusters at scale is remarkably complex," said **Ketan Shah, Vice President, Product Management, Nutanix**. "By pairing Nutanix's Agentic AI solution with the Netris NAAM platform, we will provide AI cloud operators with an out-of-the-box solution to accelerate their journey from metals to margins. We chose to partner with Netris, a leader in network automation, abstraction, and multitenancy, to give operators true self-service networking across both bare-metal and virtualized environments. This will enable our partners to deliver secure, multitenant IaaS environments quickly while maintaining the simplicity and enterprise reliability that Nutanix is known for."

"AI infrastructure is scaling incredibly fast, and the network has to scale with it," said **Alex Saroyan, CEO and Co-Founder of Netris**. "An AI cloud

needs two layers working as one — Nutanix's compute orchestration and virtualization, backed by deep operational leadership and experience, and Netris network automation, abstraction, and multitenancy. What makes this collaboration matter is that those layers stay continuously in sync, giving neoclouds and enterprises a single automated stack for secure, multitenant GPU services."

Analyst Commentary

"AI infrastructure providers are evolving beyond GPU capacity to deliver trusted inference environments, workload orchestration, agent lifecycle management, and managed AI services. The differentiator will be a cloud-like operating model that combines performance with control, isolation, and security at every layer of the stack."

- **Steve McDowell, Principal Analyst, NAND Research**

Availability

The joint solution is planned to be available in tech preview in early 2027.

About Nutanix

Nutanix is a hybrid cloud leader and AI platform company, offering organizations a unified infrastructure software platform to safely run applications, data, and AI anywhere. Trusted by customers worldwide, Nutanix empowers more than 50% of the Global 2000 to innovate faster with AI, while modernizing infrastructure, simplifying operations, and controlling costs. Learn more at www.nutanix.com or follow us on social media.

About Netris

Netris is the leading provider of network automation and multitenancy for AI infrastructure, purpose-built for neoclouds, sovereign AI operators, and AI factories. The Netris NAAM (Network Automation, Abstraction, and Multitenancy) platform delivers network orchestration for AI clusters, hard multitenancy enforced on networking hardware, and cloud-provider-grade automation across Ethernet, InfiniBand, DPUs, and virtual and edge networking — from Day Zero design and simulation through live operations. Netris enables operators to launch GPU clouds in weeks, provision tenants instantly with isolation configured automatically, and maximize GPU utilization by dynamically reallocating capacity across tenants. Netris is the most widely deployed network automation platform for AI clusters, with 70+ live deployments. Founded in 2018 and backed by Andreessen Horowitz, Netris is headquartered in Santa Clara, California. Learn more at netris.io.

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

This content may contain express and implied forward-looking statements, including but not limited to statements regarding our plans and expectations relating to new product features and technology under development, the capabilities of such product features and technology, and our plans to release product features and technology. Such statements are not historical facts and are instead based on our current expectations, estimates and beliefs, including statements about The accuracy of such statements involves risks and uncertainties and depends upon future events, including those that may be beyond our control, and actual results may differ materially and adversely from those anticipated or implied by such statements, including, among others: failure to develop, or unexpected difficulties, delays or disruptions in developing, releasing or distributing, new products, services, product features or technology in a timely or cost-effective basis. Any forward-looking statements included speak only as of the date hereof and, except as required by law, we assume no obligation to update or otherwise revise any such forward-looking statements to reflect subsequent events or circumstances. Certain products and features or functionalities described herein remain in varying stages of development and will be offered on a when-and-if-available basis. The development, release, and timing of any such products, features or functionalities are subject to change. Nutanix will not have any liability for any failure to deliver or delay in the delivery of any such products, features or functionalities. Any future product or product feature information is intended to outline general product directions, and is not a commitment, promise, or legal obligation for Nutanix to deliver any functionality. This information should not be used when making a purchasing decision.

Contacts Nutanix Ken Lotich pr@nutanix.com Netris Jason Klein pr@netris.io