

OpenNebula Systems Joins VAST Cosmos to Deliver Optimized Infrastructure for AI Factories and Gigafactories

Collaboration combines cloud infrastructure orchestration with a high-performance data platform to support sovereign, large-scale AI environments.

MADRID, SPAIN, August 10, 2026 /EINPresswire.com/ -- OpenNebula Systems, the leading provider of sovereign cloud and AI Factory management solutions, today announced it has joined the [VAST Cosmos](#) Community, a global community of developers, builders and experts in innovative AI solutions, as a Technology Partner. The strategic collaboration with [VAST Data](#), the AI Operating System company, will provide an integrated infrastructure solution for AI Factories, Gigafactories, and next-generation Neoclouds.



This integration brings together the sovereign AI Factory management platform OpenNebula and the high-performance [VAST AI Operating System](#), enabling organizations to deploy, operate, and scale AI infrastructures for model training, fine-tuning, inference, and scientific computing across GPU-intensive environments.

As enterprises, research organizations, cloud providers, and national AI initiatives invest in increasingly large AI infrastructures, they require more than compute capacity alone. They need a unified operating model for managing heterogeneous accelerated hardware, virtual machines, Kubernetes clusters, bare metal, networking, storage, and AI services. These multi-tenant environments also require a powerful data platform capable of delivering consistent performance, resilience, isolation, and scale across different workload deployment models, while integrating seamlessly with the cloud platform's tenant model.

This joint architectural approach is designed to support several integration patterns depending on the operational requirements of each AI Factory. For performance-sensitive AI workloads, customers can provide direct, high-throughput access between compute resources and the data



AI Factories need an integrated platform, not just GPU orchestration. Together, OpenNebula and VAST Data deliver secure, multi-tenant AI infrastructure with cloud-like services.”

*Dr. Alberto P. Martí, Director
of Strategic Alliances at
OpenNebula Systems*

platform. For shared environments, storage can be integrated through the virtualization layer, allowing OpenNebula to manage tenant-aware VM lifecycle, isolation, and infrastructure policies while VAST provides the underlying data services. This integration can also expose shared file services directly to workloads, supporting cloud-like file consumption models for virtual machines, containers, and AI applications.

This flexibility allows the same infrastructure platform to support dedicated high-performance workloads, secure multi-tenant services, and shared file-based access without forcing customers into a single storage consumption model. Organizations can therefore align the integration

approach with workload sensitivity, tenant isolation, performance requirements, and multi-tenant operational requirements.

By combining OpenNebula’s cloud infrastructure orchestration and lifecycle management capabilities with the VAST AI Operating System, this joint solution provides a consistent operational layer across GPU infrastructures and AI data services. OpenNebula manages the provisioning and lifecycle of compute resources, while VAST delivers the scalable data foundation required by large-scale training, inference, retrieval-augmented generation, checkpointing, model repositories, and data-intensive scientific workloads.

The combined solution is designed for modern AI platforms based on NVIDIA accelerated computing systems, including NVIDIA Grace Blackwell environments. It supports deployment models spanning on-premises infrastructure, sovereign cloud, hybrid environments, commercial GPU clouds, research infrastructures, and geographically distributed AI Factory deployments.

“AI Factories and Gigafactories require much more than GPU orchestration,” said Dr. Alberto P. Martí, Director of Strategic Alliances at OpenNebula Systems. “They need an integrated operating model for compute, networking, virtualization, Kubernetes, bare metal, data services, and AI workloads with a consistent operational multi-tenant model from compute to data. Our collaboration with VAST Data gives customers the flexibility to combine high-performance data access with secure multi-tenancy infrastructure and cloud-like service delivery across large-scale AI infrastructures.”

The joint offering is aimed at organizations developing national AI infrastructure, enterprise AI platforms, scientific computing environments, AI cloud services, and HPC centers evolving into sovereign AI Factories. It is also designed for Neocloud providers that must securely serve multiple customers while maintaining performance, efficiency, and resource isolation.

“Our customers are building some of the world’s most advanced AI environments,” said John Mao, Vice President of Global Business Development and Strategic Alliances at VAST Data. “By combining the VAST AI Operating System with OpenNebula’s AI Factory management capabilities, we can help organizations simplify infrastructure operations while supporting the different data access patterns required by large-scale and multi-tenant AI environments.”

The collaboration will also support the definition of a joint reference architecture and specific deployment blueprints bringing together accelerated computing, AI-native data services, virtualization, Kubernetes, workload orchestration, and infrastructure automation. These resources will help customers move from isolated GPU clusters toward AI Factory operating models in which compute and data resources are managed as a shared, policy-driven, multi-tenant infrastructure optimized for AI services.

OpenNebula Systems and VAST Data are already collaborating on opportunities across Europe and international markets, with a special focus on AI Factories, AI Gigafactories, HPC deployments and next-generation Neoclouds.

About OpenNebula Systems

OpenNebula Systems is the developer of OpenNebula, an open, vendor-neutral enterprise cloud and virtualization platform designed for sovereign clouds and AI factories. The company provides SLA-based support and professional services to enterprises, service providers, and public-sector organizations worldwide.

OpenNebula is increasingly adopted as a VMware replacement platform, enabling organizations to modernize and consolidate their infrastructure while preserving existing hardware investments and operational practices. At the same time, the platform has evolved to support AI Factories and AI Gigafactories, acting as a unified control plane for GPU-accelerated infrastructure, Kubernetes platforms, and large-scale AI workloads.

With a global presence and offices in Europe and the United States, OpenNebula Systems is a trusted partner for large-scale, mission-critical deployments, supporting more than 5,000 cloud deployments worldwide and infrastructures that scale to thousands of hosts and tens of thousands of GPUs.

Anastasiia Rachkova
OpenNebula Systems
community-manager@opennebula.io

Visit us on social media:

[LinkedIn](#)

[Bluesky](#)

[Facebook](#)

[YouTube](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/931331489>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.