OpenNebula vs OpenStack: Cost, Operational Complexity, and Modern Cloud Readiness

Version 2.0 – May 2026

Scope

This report evaluates Opennebula as an alternative to OpenStack, with a focus on enterprise and service provider environments. While OpenStak has traditionally been used to build large-scaled Infrastructure-as-a-Service (IaaS) platforms, evolving infrastructure requirements-such as support for AI and GPU workloads, hybrid and multi-site deployments, edge computing, and data sovereignty-are introducing increased complexity and operational challenges in OpenStack-based architectures.

The purpose of this report is to assess OpenNebula not only from a cost perspective, but also in terms of operational simplicity, flexibility, and alignment with modern cloud use cases. The analysis highlights key differences in architecture, deployment, lifecycle management, and scalability, as well as the ability to support emerging workloads such as AI/ML platforms and GPU accelerated services.

This document is intended for organizations currently operating OpenStack environments, evaluating alternatives to VMware-based infrastructure, or designing sovereign and future-ready cloud platforms. It provides a clear understanding of how OpenNebula simplifies cloud management and enables a more efficient and adaptable infrastructure model.

Contents

Glossary

AD	Active Directory
ACL	Access Control List
COW	Copy on Write
DB	Database
DC	Data Center
HA	High Availability
NFS	Network File System
NIC	Network Interface Card
VDC	Virtual Data Center
VM	Virtual Machine
KVM	Kernel-based Virtual Machine

1. What is OpenStack?

OpenStack is an open-source cloud computing platform designed to manage compute, storage, and networking resources through APIs and a centralized dashboard. It provides Infrastructure-as-a-Service (IaaS) capabilities for building private and public clouds.

Over time, OpenStack has evolved into a broad ecosystem composed of multiple independent services responsible for areas such as compute orchestration, networking, storage, identity, and image management. While this modular architecture offers flexibility, it also introduces significant integration and operational complexity.

In practice, most organizations do not deploy OpenStack directly from upstream. Instead, they rely on enterprise distributions from vendors such as Red Hat, which package and support selected components. These distributions often include additional dependencies that can reduce interoperability and introduce elements of vendor lock-in.

Operating OpenStack environments typically requires specialized expertise due to the number of interacting services and the complexity of lifecycle management. Upgrades, maintenance, and troubleshooting often demand dedicated resources and careful coordination across components.

Although OpenStack remains a powerful platform for traditional IaaS use cases, newer requirements—such as GPU-accelerated workloads, AI/ML platforms, edge computing, and hybrid deployments—can increase operational complexity and frequently require additional layers or external tools.

As a result, OpenStack environments often evolve into multi-layered architectures where infrastructure operations depend on the coordination of multiple technologies rather than a single integrated system.

2. What is OpenNebula?

OpenNebula is an open-source cloud and virtualization management platform designed to simplify the deployment and operation of private, hybrid, and edge cloud infrastructures. It provides a unified control layer to manage compute, storage, and networking resources across distributed environments.

Unlike multi-component cloud frameworks, OpenNebula delivers a streamlined architecture that integrates infrastructure management, orchestration, and automation within a single platform. This reduces operational overhead and minimizes the need for external dependencies. The platform supports multiple workload models, including virtual machines, containers, and Kubernetes environments, allowing organizations to select the most appropriate execution model without adding management complexity.

OpenNebula is designed to operate across on-premises data centers, edge locations, and distributed multi-site environments. Built-in federation and hybrid cloud capabilities enable geographically dispersed infrastructure to be managed as a single logical cloud, supporting use cases such as disaster recovery and workload mobility.

It also provides strong support for modern infrastructure requirements, including GPU-accelerated and AI-driven workloads. Organizations can build multi-tenant AI environments with resource isolation and lifecycle management, supporting use cases such as GPU-as-a-Service, inference platforms, and large-scale training workloads. The platform integrates with modern AI ecosystems and supports flexible execution models, allowing workloads to run on GPU-enabled virtual machines or within Kubernetes clusters while maintaining consistent governance and operational control.

OpenNebula is built with a focus on openness, sovereignty, and vendor neutrality. It enables full control over infrastructure, data, and lifecycle management, supporting on-premises and air-gapped deployments as well as strict regulatory and data residency requirements.

By combining simplicity, flexibility, and support for modern workloads, OpenNebula enables organizations to build cloud infrastructures that are easier to operate, more efficient, and better aligned with evolving technological and regulatory demands.

3. Subscription Structure and Hidden Operational Costs

The cost of a cloud platform extends beyond software licensing. It includes infrastructure management, lifecycle operations, and the human resources required to run and maintain the environment. For this reason, total cost of ownership (TCO) is the most relevant metric when evaluating cloud platforms.

OpenNebula follows a single, unified subscription model based on the number of managed infrastructure nodes and the selected level of enterprise support. This approach eliminates the need for multiple layered subscriptions and provides a predictable cost structure. The subscription includes enterprise software, long-term support releases, security updates, and expert support under defined SLA agreements.

In contrast, OpenStack-based platforms—particularly enterprise distributions—typically involve a multi-layered subscription model. Organizations often need separate subscriptions for the cloud platform, operating system, container platforms, and additional services such as networking, storage, and lifecycle management.

Modern deployments increasingly combine OpenStack with Kubernetes platforms such as OpenShift, adding further layers to the architecture. Each component introduces its own subscription model and lifecycle, often tied to hardware characteristics such as CPU sockets or node capacity. As infrastructure scales, additional subscriptions must be added across multiple layers. This layered approach increases both cost variability and operational complexity. Managing multiple tightly coupled components requires specialized expertise, larger teams, and more effort for deployment, upgrades, and troubleshooting.

OpenNebula, by contrast, reduces operational complexity through its unified architecture. Organizations can operate production environments with smaller teams, simplified lifecycle management, and reduced dependency on external tools. This leads to faster deployment, improved efficiency, and lower operational overhead.

As a result, the primary advantage is not only in the subscription model itself, but in the overall reduction of operational effort. The combination of architectural simplicity and streamlined operations enables a significantly lower total cost of ownership compared to multi-layered cloud platforms.

4. Simple Case Subscription Pricing Review

Both OpenNebula and enterprise OpenStack distributions require a commercial subscription to provide enterprise-grade support. While Red Hat OpenStack pricing is not publicly standardized, available reseller listings and market references indicate that subscriptions are typically structured on a per-socket-pair basis, with annual costs generally in the range of $5,000 to $7,500 per two sockets, depending on support level and configuration.

In addition, deployments require a separate Red Hat Enterprise Linux (RHEL) subscription for the underlying operating system. RHEL pricing varies depending on the deployment model and support tier, but commonly observed enterprise subscriptions for physical servers typically range from approximately $700 to $1,500 per socket pair annually, increasing to $2,500–$4,500 per socket pair for datacenter or virtualization-oriented subscriptions with broader entitlements.

Both OpenStack Platform and RHEL follow a socket-pair-based, stackable subscription model, meaning that additional subscriptions must be purchased as infrastructure capacity increases (e.g., more CPU sockets or nodes). This results in a layered pricing structure, where costs scale with both infrastructure size and hardware configuration. As capacity grows, organizations must stack multiple subscriptions across the cloud platform and operating system layers, increasing overall cost variability and operational complexity.

OpenNebula follows a per-host subscription model, independent of the number of CPUs within each host. The annual subscription is priced at $875 per host for a Standard SLA and $1,375 per host for a Premium SLA.
This subscription also includes support for the host operating system (Ubuntu or CentOS), eliminating the need for separate OS subscriptions.

In addition to the compute hosts, OpenNebula requires a cloud front-end component. A single front-end deployment is priced at $8,750 for a Standard SLA and $13,750 for a Premium SLA per instance In contrast, the cost of deploying the control plane in OpenStack environments depends on the number of infrastructure nodes required. Typical OpenStack architectures require a minimum of multiple nodes—commonly around six hosts—to support controller and core services, which further contributes to overall infrastructure and subscription requirements.

The per-host pricing model provides greater transparency and cost control compared to per-socket pricing. With per-socket licensing, costs increase as hardware density grows, requiring additional subscriptions when scaling compute capacity. In contrast, per-host pricing allows organizations to adopt more powerful hardware without increasing licensing costs, enabling more efficient infrastructure scaling.

Category	Red Hat OpenStack	OpenNebula
Front-end	$39,000	$13,750
10 x 2-CPU hosts	$65,000	$11,000
TOTAL	**$104,000**	**$24,750**

This scenario demonstrates savings of more than **$80,000 per year** when using OpenNebula compared to Red Hat OpenStack. As infrastructure scales—either in number of nodes or CPU density—the cost difference increases significantly due to the per-socket pricing model used in OpenStack environments.

5. Consulting Services

OpenNebula is designed to enable rapid deployment of production-ready cloud environments with minimal architectural overhead. Its unified platform approach significantly reduces the complexity typically associated with cloud infrastructure design and implementation.

To accelerate adoption and ensure optimal architecture, OpenNebula provides professional services covering cloud design, deployment, and optimization. These services deliver a validated, production-ready implementation based on reference architectures, including pre-configured environments, sample workloads, and comprehensive documentation such as architecture design reports, implementation guides, and verification procedures.

Due to the streamlined architecture of OpenNebula, consulting engagements are typically short and focused. Even for enterprise environments, deployment timelines are significantly reduced, enabling faster time-to-production and lower initial project costs. For smaller environments, engagements typically start from a few days, with entry-level projects starting at approximately $15,000–$17,000.

In contrast, OpenStack environments require extensive planning, integration, and validation due to their multi-component architecture. Deployment typically follows structured phases such as discovery, design, and implementation, often extending over several weeks or months. This results in higher consulting costs, commonly estimated at approximately $10,300–$11,500 per week, as well as longer time-to-value.

As a result, organizations adopting OpenNebula benefit from faster deployment cycles, reduced project risk, and significantly lower professional services costs compared to OpenStack-based solutions.

This results in faster time-to-value and significantly lower project risk compared to complex multi-component cloud deployments.

6. Key Differentiators

Architecture Model

OpenNebula provides a unified architecture, while OpenStack relies on a multi-layered stack combining infrastructure services, Kubernetes orchestration, and external components. This fundamental difference impacts deployment speed, operational complexity, and long-term scalability.

Subscription Costs

OpenNebula delivers a significantly more efficient subscription model compared to OpenStack-based platforms. Its per-host pricing ensures predictable and transparent costs, independent of hardware configuration. In contrast, OpenStack enterprise distributions typically rely on layered and hardware-dependent licensing models, which increase costs as infrastructure scales. This results in substantially lower subscription costs for OpenNebula, often exceeding 70% savings in comparable environments.

Operational and Maintenance Costs

Operational complexity is one of the primary drivers of long-term cost. OpenStack environments require specialized expertise across multiple services, including compute, networking, storage, and orchestration layers. This typically leads to the need for dedicated engineering teams and higher ongoing operational expenditure.

OpenNebula's unified architecture significantly reduces operational overhead. Deployment, configuration, and day-to-day management are streamlined, enabling organizations to operate large-scale environments with smaller teams. This results in improved operational efficiency, reduced staffing requirements, and lower total cost of ownership over time.

Software Upgrades and Lifecycle Management

OpenNebula provides a controlled and predictable upgrade process with minimal disruption to running environments. Its integrated architecture allows upgrades to be performed in a consistent and automated manner.

OpenStack upgrades, by contrast, are complex and risk-prone due to the interdependencies between multiple services. Upgrade procedures often require extensive planning, validation, and coordination, and in some cases may lead to partial redeployments. This increases operational risk and introduces additional maintenance effort over the lifecycle of the platform.

Workload Support (Virtualization, Containers, and AI)

OpenNebula natively supports virtual machines, containerized workloads, and Kubernetes environments within a unified platform. This eliminates the need for additional control layers and enables consistent management across different workload types.

The platform also provides a strong foundation for GPU-accelerated and AI-driven workloads, supporting modern use cases such as AI/ML platforms and high-performance computing environments without introducing additional architectural complexity.

OpenStack environments typically rely on external platforms to support containerized and AI workloads, increasing system complexity, resource consumption, and operational overhead.

Hybrid, Multi-Site, and Edge Deployments

OpenNebula provides built-in capabilities for hybrid cloud, multi-site federation, and edge computing. Infrastructure across data centers, remote locations, and edge sites can be managed as a single logical cloud, with integrated tools for provisioning and workload mobility.

OpenStack does not provide these capabilities natively and typically requires additional components or external platforms to achieve similar functionality. This increases both deployment complexity and operational effort.

7. Summary

This report has presented a detailed comparison between OpenNebula and OpenStack, with a focus on architecture, cost structure, and operational impact.

Selecting the appropriate cloud management platform is a strategic decision with long-term implications. Architectural complexity, operational overhead, and licensing models directly influence time-to-market, scalability, and total cost of ownership. Platforms that require extensive integration, specialized expertise, and multi-layered dependencies can introduce significant operational risk and limit long-term flexibility.

The analysis demonstrates that OpenNebula provides a more efficient and streamlined approach to building and operating cloud infrastructure. Its unified architecture reduces complexity across deployment, lifecycle management, and day-to-day operations, enabling organizations to achieve faster implementation timelines and lower operational costs.

In addition to cost efficiency, OpenNebula is designed to support modern infrastructure requirements. It enables consistent management of virtual machines, containerized workloads, and Kubernetes environments within a single platform, while also providing native capabilities for hybrid, multi-site, and edge deployments. The platform further supports emerging use cases such as GPU-accelerated and AI-driven workloads without introducing additional architectural layers.

By combining simplicity, flexibility, and support for modern workloads, OpenNebula enables organizations to build cloud environments that are easier to operate, more cost-effective, and better aligned with evolving technological and regulatory requirements.

Organizations evaluating alternatives to OpenStack or seeking to modernize their infrastructure can benefit from adopting a platform that reduces operational complexity while maintaining full control over their cloud environment.

In this context, OpenNebula represents a future-ready platform that simplifies cloud operations while enabling organizations to address emerging infrastructure demands with confidence.

8. Ready for a Test Drive?

OpenNebula can be evaluated and deployed quickly using a set of purpose-built tools designed to minimize setup complexity and accelerate time-to-value.

For quick testing and initial validation, **miniONE** provides a lightweight deployment method that installs a fully functional OpenNebula environment on a single virtual machine or physical host within minutes. This environment can then be extended by attaching additional resources or remote clusters, enabling users to explore core platform capabilities with minimal effort.

In addition, the **OpenNebula ISO Community Edition** delivers a ready-to-use appliance that includes a preconfigured operating system and OpenNebula stack. This option is particularly suited for rapid evaluations, proof-of-concept environments, and standardized deployments.

For production-grade deployments, **OneDeploy** offers an automated and scalable installation framework that simplifies the rollout of OpenNebula across multi-node environments. It provides a consistent and repeatable deployment process aligned with recommended architectures, reducing setup time and operational complexity.

Together, these tools enable organizations to move from evaluation to production efficiently, reducing deployment time and allowing teams to quickly validate OpenNebula in real-world scenarios.

LET US HELP YOU DESIGN, BUILD, AND OPERATE YOUR CLOUD

CONSULTING & ENGINEERING

Our experts will help you design, integrate, build, and operate an OpenNebula cloud infrastructure

OPENNEBULA SUBSCRIPTION

Get access to our Enterprise Edition, support and exclusive services for Corporate Users

CLOUD DEPLOYMENT

Focus on your business and let us take care of setting up your OpenNebula cloud infrastructure

Sign up for updates at OpenNebula.io/getupdated