



True Hybrid and Multi-Cloud with OpenNebula

Version 1.2 – July 2024

Abstract

To support digital transformation initiatives, IT departments need the right blend of cloud environments—on-premises, public, and edge—tailored for a variety of existing and emerging use cases. There are many reasons why organizations make use of cloud services from multiple providers. Avoiding vendor lock-in, increased agility and service availability, more efficient costs, and the promise of obtaining the best possible solution from each provider are all too great to ignore. However, organizations face many challenges in the quest for complete freedom in their application workloads across different clouds, including, among others: the heterogeneity of interfaces and formats, with significant numbers of applications requiring considerable refactoring to migrate to the cloud; security concerns; cost control; lack of resources and expertise; the need for cultural changes in IT teams; and management complexity.

Multi-cloud is undoubtedly here to stay. Our goal is to make its more sophisticated incarnations easier to deploy and manage. Since version 6.0, OpenNebula offers a powerful **multi-cloud architecture** composed of Edge Clusters that can run **any application**—Virtual Machines workloads and containerized applications from Kubernetes—on **any resource**—bare-metal or virtualized—**anywhere**—on-premises, in the cloud, or at the edge. This innovative model enables true hybrid and multi-cloud computing by combining public and private cloud operations with workload portability and unified management of IT infrastructure and applications.

This white paper describes the principles at the base of our multi-cloud architecture, and lays the groundwork for how, with OpenNebula, you can exploit the many benefits of a multi-cloud model.

Contents

1. The Need for Multi-Cloud
2. Challenges of Multi-Cloud
3. Benefits of a True Multi-Cloud
4. What is OpenNebula?
5. Criteria for True Multi-Cloud
6. An Easy Button for Multi-Clouds
7. A Growing Catalog of Cloud Providers
8. Run Any Application, Anywhere
9. Ready for a Test Drive?
10. Conclusions

Glossary

CapEx	Capital Expenditure
DC	Data Center
HA	High Availability
OpEx	Operating Expenditure
VM	Virtual Machine

1. The Need for Multi-Cloud

Today there is no doubt that for organizations, the hybrid cloud model, and the use of multiple cloud providers, are the norm. In the new era of cloud, cost optimization and interoperability have become crucial factors. Multi-cloud strategies enable companies to achieve independence from providers, while addressing the risk of concentration and expanding service availability. As organizations continue their digital transformation journeys, modernizing business processes and optimizing IT infrastructures, it's clear that cloud computing services, along with automation and orchestration, are becoming increasingly important.

According to Gartner, cloud concentration and its ensuing vendor dependency constitute one of the top five risks for many organizations.¹ In recent years organizations have overwhelmingly embraced multiple cloud strategies, as indicated among others by Flexera's 2024 State of the Cloud Report.² In this survey of 753 cloud decision-makers and users, 89% of respondents affirm having opted for a multi-cloud approach, comprising 73% hybrid cloud, 14% multiple public cloud, and 2% multiple private cloud.

Moreover, the multi-cloud approach should include on-premises resources. In 2020, Forrester's IBM-commissioned report *The Key To Enterprise Hybrid Multicloud Strategy*, a survey of 350 global enterprise IT decision-makers found that in the following two years, 47% of data-intensive workloads and over 50% of mission-critical workloads would continue to run either on-premises or on an internal private cloud.³ Three years later Gartner issued a similar figure, forecasting that 50% of critical applications will run outside of centralized public clouds through 2027.⁴

2. Challenges of Multi-Cloud

Despite the obvious benefits of deploying any workload across different public cloud and edge providers, there remain a number of challenges that IT organizations must overcome to take full advantage of a truly multi-cloud approach. The greatest impediments to adopting Multi-Cloud Computing are:

- **Governance:** policies to control costs, minimize security risks, improve efficiency, and accelerate deployment.
- **Interoperability:** workloads located in different providers can communicate with each other to create an aggregated service.
- **Portability:** deployment and migration of workloads between different cloud and edge providers.
- **Orchestration:** arrangement and coordination of automated tasks for service deployments.

The evolution of the modern cloud has led to the creation of **highly complex systems, often based on proprietary technologies**. Traditionally, several major vendors (e.g. Red Hat, Nutanix, VMware) have offered their own orchestration solutions that expand private clouds with resources from cloud providers. These solutions support a limited number of cloud providers, do not offer automatic enrollment of new nodes, and are complex and expensive to deploy and maintain.

On the other hand, the **hyperscalers offer pre-configured appliances** (e.g. AWS Outposts, Google Anthos) that promise to bring the power of the public cloud to the private cloud. While they offer their customers the simplicity of using the same interfaces both on the public cloud and on their private DC, these proprietary solutions undermine digital sovereignty and can be very expensive in the long run. It is important to note that none of these approaches offers a comprehensive solution for workload portability,

¹ <https://www.gartner.com/en/newsroom/press-releases/2023-10-30-gartner-says-cloud-concentration-now-a-significant-emerging-risk-for-many-organizations>

² <https://info.flexera.com/CM-REPORT-State-of-the-Cloud>

³ https://www.blueconnections.com.au/wp-content/uploads/2020/12/IBM_Forrester_Report_The_Key_To_Enterprise_Hybrid_Multicloud_Strategy.pdf

⁴ <https://www.gartner.com/en/newsroom/press-releases/2023-10-30-gartner-says-50-percent-of-critical-enterprise-applications-will-reside-outside-of-centralized-public-cloud-locations-through-2027>

since they do not provide an overlay on top of the different providers, and hence the workload needs to adapt to the particularities of each provider.

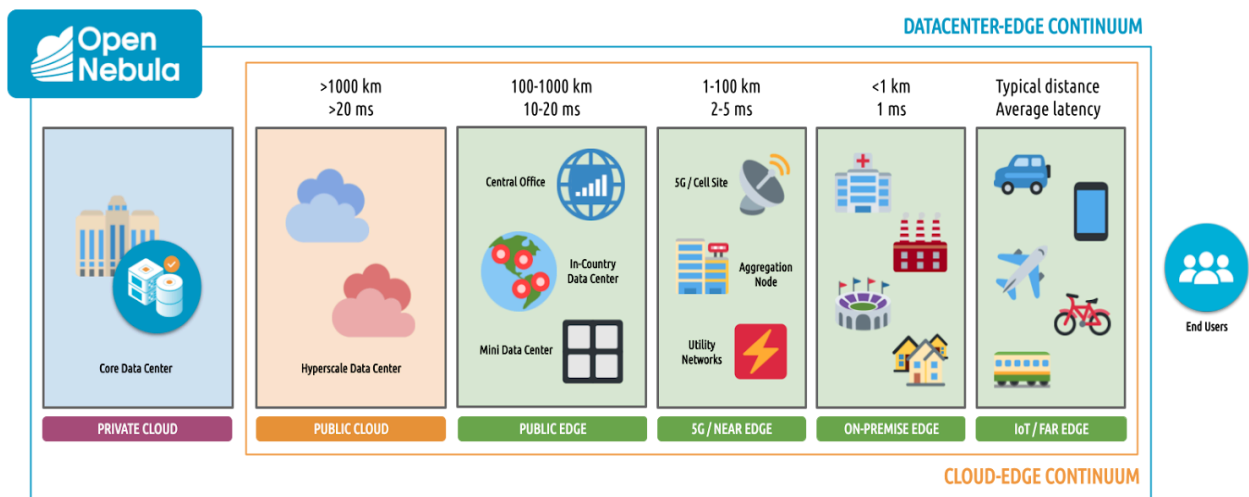


Figure 1. Edge computing brings new opportunities for organizations using a multi-cloud architecture.

Multi-cloud is familiar to many of our users who have implemented their private cloud infrastructure with OpenNebula and now effectively operate their workloads in multi-cloud environments, benefiting from the frictionless experience provided by our innovative multi-cloud architecture. OpenNebula delivers the industry's first cloud platform that offers the flexibility, simplicity, and cost efficiency needed to manage applications and infrastructure in private, public, and edge clouds, operated as a single cloud. This simplicity is critical if we consider that talent shortage is also becoming a major challenge. Based on our experience, lack of skills is holding businesses back from innovating across multiple cloud platforms.

3. Benefits of a True Multi-Cloud

OpenNebula gives you the benefits of agility, scalability and simplicity of the pay-per-use public cloud, plus the greater levels of flexibility, performance, and security of a dedicated private environment based on a geo-distributed multi-cloud infrastructure.



Single Management Interface

Combine VMs and Kubernetes containerized applications across your infrastructure.



Easy Automated Deployment

Minimize the operational overhead of your multi-cloud environment.



Reduced Vendor Lock-in

Enjoy workload portability between any cloud provider and your on-premises resources.

**Enhanced Security**

Manage dedicated, isolated resources with improved security, privacy, and control.

**Expanded Service Availability**

Deploy resources across geographic zones and cloud providers to meet your latency or HA requirements.

**Optimized Cloud Costs**

Scale up your infrastructure on-demand with a pay-as-you-go model and flexible pricing.

Minimizing Total Cost of Ownership with OpenNebula

The increasing demand for multi-cloud solutions is driven mostly by financial reasons. Many of our users try to optimize their infrastructure costs to ensure that they only use as many resources as needed, at optimal cost. Using a multi-cloud architecture—based on combining public cloud resources with a cost-effective private cloud—is essential to achieve infrastructure cost optimization. According to several studies, multi-cloud platforms can save companies 40-60% on IT expenses over time. OpenNebula can help your organization achieve this goal by offering:

- **Flexible Cost Management.** Multi-cloud offers flexible cost options. Depending on the application, companies can choose the CapEx model of the on-premises resources or the OpEx model of the cloud providers.
- **No Vendor Lock-In.** Because OpenNebula is 100% software-based, organizations can use the hardware platforms, virtualization technologies, and cloud services of their choice. This gives them total control over the cost and performance of their IT infrastructure.
- **Lower Operational Costs.** Using public and private clouds reduces the management burden on IT. OpenNebula automates the deployment of a multi-cloud architecture and offers a uniform control and management layer. Your staff does not need to learn the specifics of multiple providers.
- **Easy Repatriation of Workloads.** Using public clouds through OpenNebula allows you to very easily run your workloads across cloud providers and repatriate all or part of your workloads back to your on-premises infrastructure when cloud costs start to catch up. This is a modest investment early on that considerably reduces the work needed to repatriate workloads in the future.
- **Incrementally Repatriate Workload.** OpenNebula allows repatriation of workloads incrementally and in a hybrid fashion. Our experience is that for most companies, optimal cost is achieved by retaining a percentage of workload in the cloud.

4. What is OpenNebula?

OpenNebula⁵ is a simple, but powerful, open source solution to build and manage Enterprise Clouds. It combines existing virtualization with advanced features for multi-tenancy, automatic provision and elasticity to offer on-demand applications and services.

OpenNebula provides a single, feature-rich and flexible platform with **unified management of IT infrastructure and applications that avoids vendor lock-in and reduces complexity, resource consumption and operational costs.** OpenNebula manages:

- **Any Application:** Combine containerized applications from Kubernetes with Virtual Machine workloads in a common shared environment to offer the best of both worlds: mature virtualization technology and orchestration of application containers.

⁵ <https://support.opennebula.pro/hc/en-us/articles/360036935791-OpenNebula-Overview-Datasheet>

- **Any Infrastructure:** Open cloud architecture to orchestrate compute, storage, and networking driven by software.
- **Any Cloud:** Unlock the power of a true hybrid, edge and multi-cloud platform by combining your private cloud with infrastructure resources from third-party virtual and bare-metal cloud providers such as AWS and Equinix Metal, and manage all cloud operations under a single control panel and interoperable layer.
- **Any Time:** Add and remove new clusters automatically in order to meet peaks in demand, or to implement fault tolerance strategies or latency requirements.

OpenNebula provides the necessary tools for running containerized applications from Kubernetes while ensuring enterprise requirements for your DevOps practices. It helps organizations to easily embrace Multi-cloud, Hybrid, and Edge Computing, allowing them to grow their Enterprise Cloud on-demand with infrastructure resources from third-party public cloud and bare-metal providers such as AWS and Equinix Metal. This white paper describes our approach for true hybrid and multi-cloud computing. If you want to find out more about how to implement an edge cloud infrastructure based on OpenNebula, please refer to our **Edge Cloud Architecture**.⁶ If you are interested in an OpenNebula cloud fully based on open source platforms and technologies, please see our **Open Cloud Reference Architecture**.⁷

The development of OpenNebula follows a bottom-up approach driven by the real needs of sysadmins, DevOps, and corporate users. OpenNebula is an **open source product** with a healthy and active community, commercially supported by OpenNebula Systems through its **OpenNebula Subscription**. Updated versions of OpenNebula are released regularly, and delivered as a single package with a smooth migration path. More information on the benefits of running an OpenNebula cloud is available on the Key Features page.⁸

5. Criteria for True Multi-Cloud

The truth is that managing and supporting multi-cloud is not an easy task. In an ideal world, application workloads—whatever their heritage—should be able to move seamlessly between, or be shared among, cloud service providers, and to be deployed wherever the optimal combination of performance, functionality, cost, security, compliance, availability, resilience, and so on, is to be found—while avoiding the dreaded “vendor lock-in.” The following sections describe the key criteria that a solution must meet to guarantee the successful adoption of multi-cloud:

Manage Workload across Every Cloud from a Single Pane of Glass

The most crucial criteria to manage multi-cloud complexity is interoperability (defined as the ability to manage your workload across every cloud from a single pane of glass), and portability (defined as the execution of your workloads with the same images and templates on any infrastructure and their mobility across clouds and on-premises infrastructure).

OpenNebula implements a multi-cloud architecture based on a single Front-end that controls one or several interconnected Edge Clusters that can run in multiple geographically-distributed DC locations, and on public cloud or edge infrastructure providers. This approach offers a complete end-to-end solution that leverages proven open source technologies to provide uniform configuration management, network automation, performance optimization, and capacity and cost management. An OpenNebula multi-cloud model enables:

- **Interoperability**, by providing a uniform view of the underlying resources of all your computing infrastructure—on-premises, in public clouds, or at the edge—with a single console.

⁶ <https://support.opennebula.pro/hc/en-us/articles/360050302811-Edge-Cloud-Architecture-White-Paper>

⁷ <https://support.opennebula.pro/hc/en-us/articles/204210319-Open-Cloud-Reference-Architecture-White-Paper>

⁸ <https://opennebula.io/discover/>

- **Portability**, by using Edge Clusters based on a common hyper-converged architecture designed to be deployed on any resource—from bare-metal to virtualized—enabling workload portability and mobility across clouds with absolutely no code changes needed.

Automate Resource Allocation and Align the Right Workloads with the Right Cloud

Deciding where to place an application is a complex decision based on infrastructure costs, data fees, performance, uptime, and latency.

In OpenNebula, Edge Clusters can be **added and removed dynamically** to meet peaks in demand, or to implement fault tolerance strategies or latency requirements. Once the new cluster is part of your cloud, OpenNebula offers advanced capacity planning features to fine-tune workload allocation, including:

- Application-level **scheduling and placement** constraints based on arbitrary Edge Cluster metrics or properties.
- Scheduling for multi-component applications that often require rules for specific **affinity** (place in the same node or cluster) or **anti-affinity** (place in a different node or cluster) to improve their fault-tolerance or performance.
- Lastly, in a multi-tenant environment admins need a **quota** system to limit resource usage in order to implement the capacity plan.

Submit Any Workload on Any Cloud

Due to its complexity and dynamic nature, the multi-cloud environment poses challenges in assessing the suitability of on-premises apps for migrating to the cloud.

OpenNebula **combines Virtual Machine workloads and Kubernetes containerized applications on a single platform**, and integrates with existing Virtual Machine and container images hubs and marketplaces. An application can be deployed anywhere on the multi-cloud infrastructure without performing any additional configuration or setup. While new workloads based on containers, microservices, and functions (serverless) should generally be stateless and ephemeral, almost all business applications require data persistence in some form. This is why our architecture provides support for both non-persistent and persistent VMs and Kubernetes containerized applications.

Combine On-Premises and Public Cloud Resources

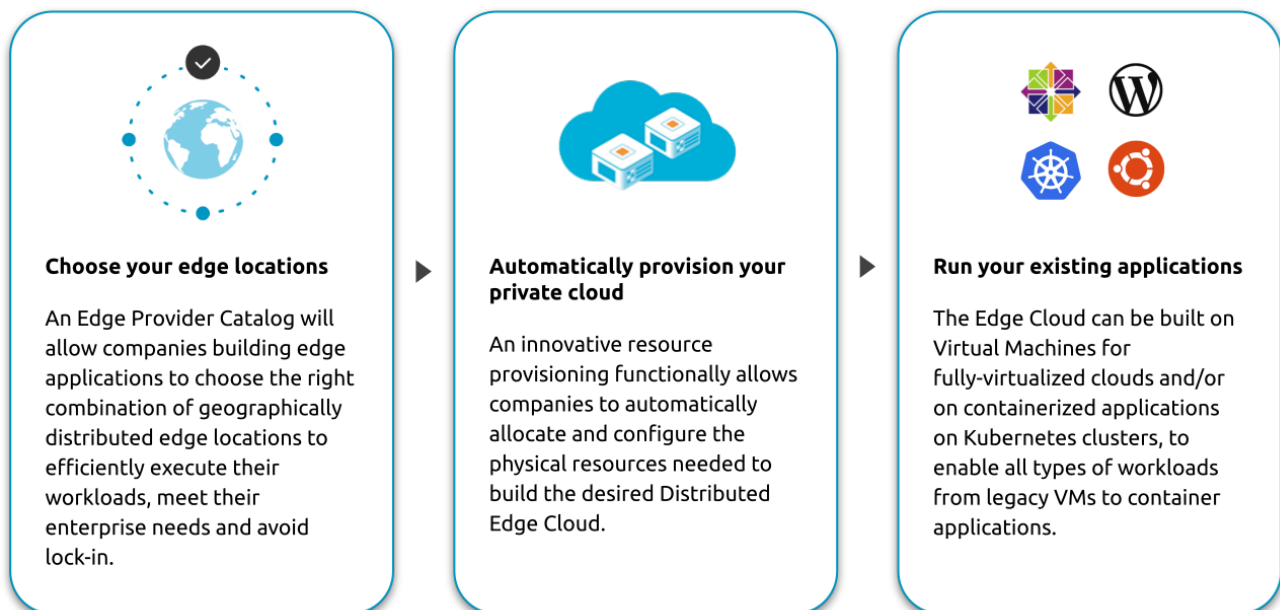
Although multi-cloud will become the norm, there's still a place for the on-premises DC, at least in the near term, either as part of a hybrid cloud strategy or to host legacy applications that, for whatever reason, are not suitable for migration to the cloud. Some of the main reasons to keep using on-premises resources to host workloads include cost, control, security, and performance.

OpenNebula allows you to **leverage your on-premises infrastructure by deploying Edge Clusters locally**. These on-prem Clusters can be hosted and automatically configured in your core DC in the same way as if they were deployed on a public cloud or at the edge, offering a uniform control and management layer for your Enterprise Cloud and isolating your corporate users from the underlying heterogeneity.

6. An Easy Button For Multi-Clouds

OpenNebula brings the provisioning tools needed to dynamically grow a private cloud infrastructure with resources running on remote bare-metal as well as virtual cloud and edge providers. This disaggregated cloud approach allows a seamless transition from centralized private clouds to distributed multi-cloud environments. Companies are able to grow their on-prem private cloud with resources at cloud or edge locations in order to meet their latency and bandwidth requirements. This approach involves a single management layer where organizations can continue using the existing OpenNebula images and templates, keep complete control over the infrastructure, and avoid vendor lock-in.

OpenNebula allows you to deploy a fully operational Edge Cluster in a remote provider, and to manage its full life cycle from provisioning and maintenance to unprovisioning. Each cloud or edge location (the “provision”) is defined as a group of physical hosts allocated from the remote bare-metal or virtual provider. They are fully configured with the user-selected hypervisor and enabled in the cloud stack, available for end-users.



7. A Growing Catalog of Cloud Providers

Edge Clusters can be deployed on on-premises infrastructures as well as on bare-metal and virtualized cloud environments to enable powerful hybrid and edge cloud computing, and support all major clouds. Infrastructure teams can choose their preferred hardware platform and cloud provider, and deliver an exceptional OpenNebula experience. Similarly, IT teams can seamlessly manage applications across cloud and edge providers, and leverage the growing ecosystem of hyperscale and edge clouds.

We are building a **Resource Provider Catalog** that will maintain a list of edge resource providers certified to work with OpenNebula. This catalog will allow users to easily select which providers, locations, and instances are better suited to their edge applications in terms of cost, capacity, latency, bandwidth, etc. The automatic provision feature, available since OpenNebula 6.0, currently supports **AWS** and **Equinix Metal**. OpenNebula will continue to expand these options by incorporating new drivers for additional public cloud and edge providers. Moreover, we are building the tools and processes to allow any cloud service provider to develop the necessary drivers by themselves and join this public catalog.

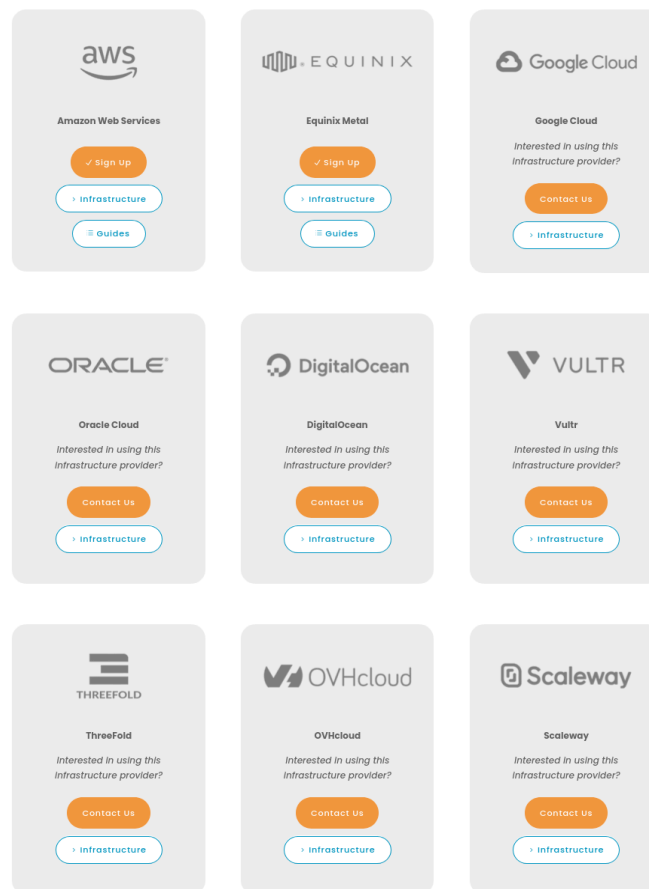


Figure 2. OpenNebula's Provider Catalog offers a growing number of multi-cloud connectors.

8. Run Any Application, Anywhere

A critical aspect of an OpenNebula cloud is its ability to support modern applications, combine Kubernetes containerized applications and Virtual Machines on a single platform, and integrate with existing Virtual Machine and container image hubs and marketplaces. OpenNebula maintains its own Marketplace, an indexed repository of virtual appliances tailored for varying infrastructure needs. Appliances are curated, configured and tested by OpenNebula. Users can easily import appliance images into their datastores, to instantiate and use at any time. The Marketplace currently features over 70 appliances, ranging from virtual hardware to pre-configured VMs to a full-fledged multi-master Kubernetes cluster (OpenNebula's OneKE appliance), which users can deploy in minutes through the Sunstone GUI.⁹

OpenNebula's Marketplace subsystem also offers users the possibility of accessing third-party public or private marketplaces, or creating their own marketplace. Additionally, OpenNebula is currently developing a new Community Marketplace to enable users to contribute their own appliances. The Community Marketplace will host applications ready for use in any OpenNebula deployment, requiring minimal configuration.

To ensure the quality and relevance of new appliances, submissions will undergo an exhaustive Contribution Process designed specifically for the Community Marketplace. The process includes evaluating appliances' documentation, their compliance with OpenNebula standards, and the capacity of the contributor to ensure its long-term maintenance. The end result will be that Community Marketplace

⁹ <https://support.opennebula.pro/hc/en-us/articles/6554989538717-Enterprise-Kubernetes-Made-Simple-White-Paper>

Appliances will be curated and tested under the same rigorous requirements as those in the OpenNebula Marketplace, and will be certified to work against the stable and latest Long-Term Support (LTS) versions of OpenNebula.

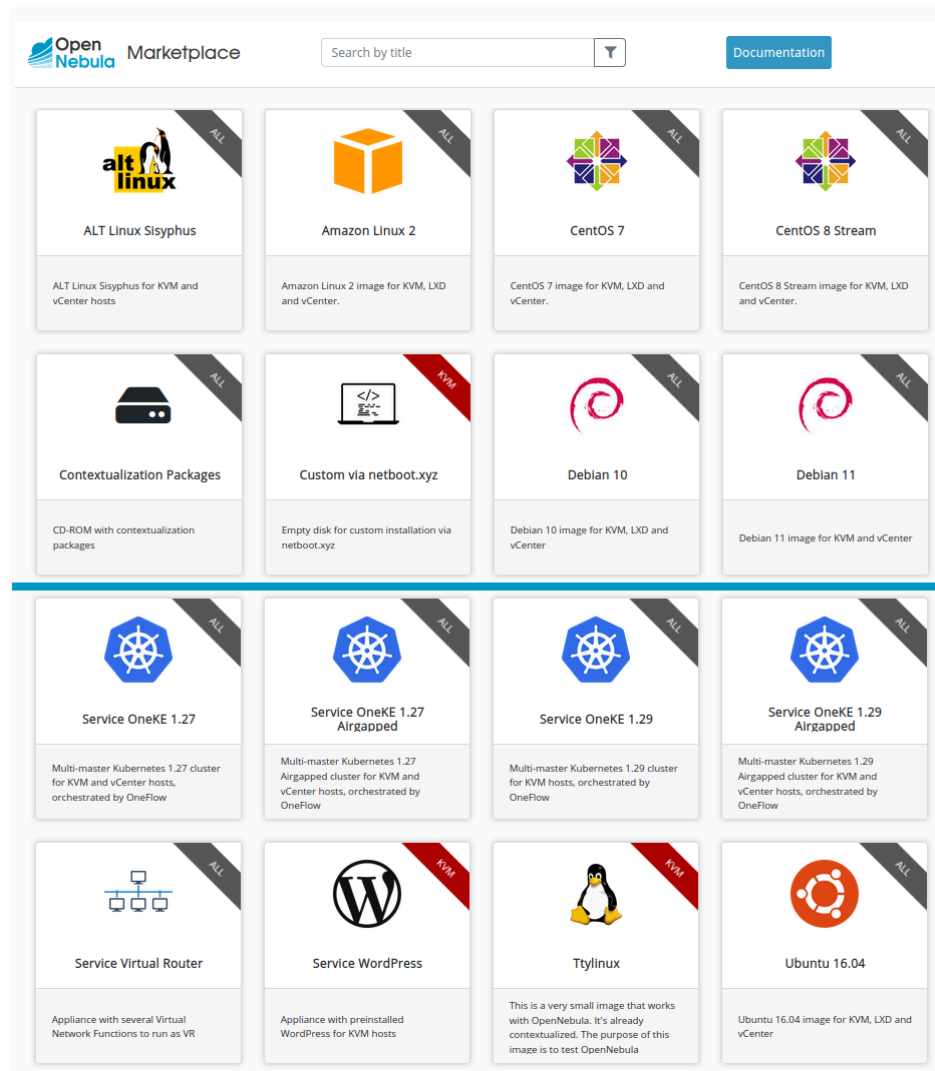


Figure 3. Some of the virtual appliances available in the OpenNebula Public Marketplace.

In relation to application container technologies, OpenNebula offers complete container orchestration services based on Kubernetes, packaged in a single, easy-to-deploy virtual appliance: OneKE (mentioned above), available on the OpenNebula Marketplace. OneKE is an enterprise-grade Kubernetes distribution certified by the Cloud Native Computing Foundation (CNCF). It provides a Kubernetes cluster ready for production environments, with the option of integrating several CNI plugins and load balancers, Longhorn distributed storage, and HAProxy/Traefik for exporting HTTP/HTTPS applications.

Thanks to the combination of OneFlow and the Sunstone UI, users can provision managed Kubernetes clusters on demand with just one click. The cluster can be configured as a multi-master cluster for High Availability, and can be easily deployed on-premises, at the edge or in the cloud. Every node in the cluster is managed by OpenNebula as a regular VM. OpenNebula does not manage containers or pods in the cluster—the cluster exposes the Kubernetes API, enabling access through the `kubectl` command-line tool or the Kubernetes dashboard to create pods, deployments and services. Implemented as a OneFlow service,

OneKE provides elasticity features to scale up/down the Kubernetes Cluster for HA, or for specialized workloads by adding more worker nodes to the cluster on demand.

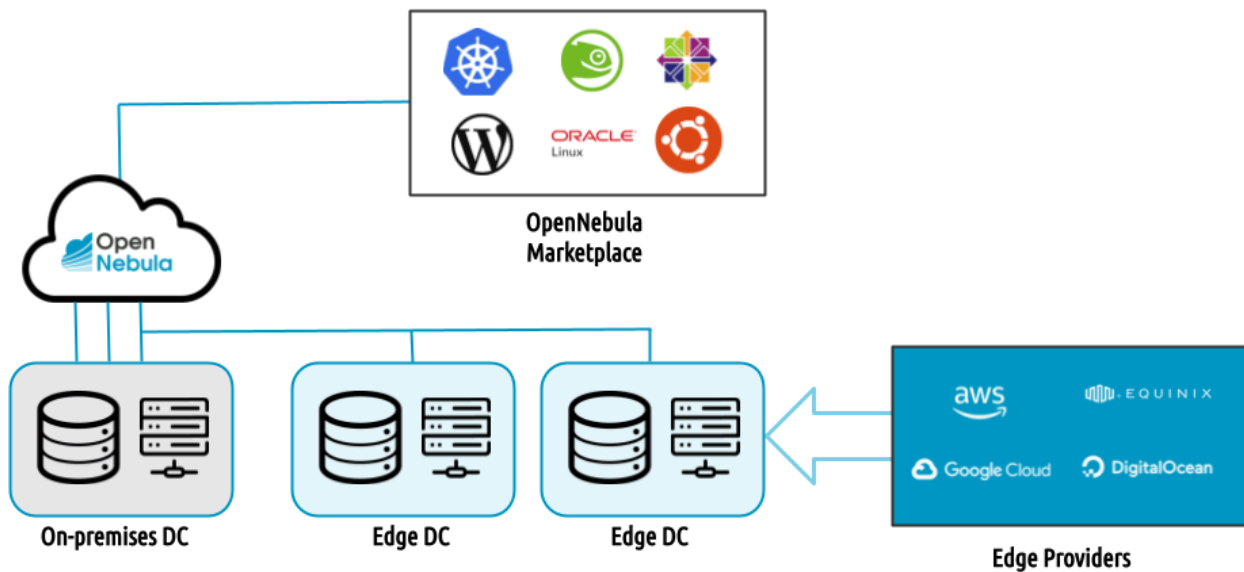


Figure 4. OpenNebula's general edge application deployment model applicable to Kubernetes.


9. Ready for a Test Drive?

You can evaluate OpenNebula and build a multi-cloud environment in just a few minutes by using **miniONE**,¹⁰ our deployment tool for quickly installing an OpenNebula Front-end inside a Virtual Machine or a physical host, which you can then use to easily add remote Edge Clusters based on KVM, Virtual Machines and Kubernetes Clusters, on multiple cloud providers.



**Preventing Vendor Lock-in
with an OpenNebula Multi-Cloud**

<https://opennebula.io/multi-cloud/>



¹⁰ <https://minione.opennebula.io>

10. Conclusions

Multi-cloud helps enterprises avoid the pitfalls of single-vendor reliance. Spreading workloads across multiple cloud providers gives organizations the flexibility to use (or stop using) a specific cloud service whenever they want. This document outlines how OpenNebula's choice for simplicity and open source technologies brings a consistent foundation to any cloud deployment, whether public, private, hybrid, or edge. The strategy is based on a powerful **multi-cloud architecture** composed of Edge Clusters. They are built on-demand using storage and networking technologies that already exist in the Linux operating system. They can run any **application**—Virtual Machine workloads or containerized applications using Kubernetes—on any **resource**—bare-metal or virtualized—**anywhere**—on-premises, in the cloud, or at the edge.

OpenNebula enables true hybrid and multi-cloud computing by combining public and private cloud operations with workload portability and unified management of IT infrastructure and applications. Now, you can enjoy a single vendor experience: we offer Enterprise support for the complete software stack through our OpenNebula **Software Subscription**, and managed cloud services through a new OpenNebula **Managed Subscription**, so your team can forget about infrastructure and focus on business workloads. [Contact us](#)—we look forward to helping you at any stage of your cloud computing journey.

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 and to our 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



This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement [ONEedge 880412](#)



[ONEedge](#) is an OpenNebula project developing innovative features to bring private cloud computing to the Edge ([ONEedge.io](#))

Sign up for updates at OpenNebula.io/getupdated

© OpenNebula Systems 2024. This document is not a contractual agreement between any person, company, vendor, or interested party, and OpenNebula Systems. This document is provided for informational purposes only and the information contained herein is subject to change without notice. OpenNebula is a trademark in the European Union and in the United States. All other trademarks are property of their respective owners. All other company and product names and logos may be the subject of intellectual property rights reserved by third parties.



Rev1.2_20240702