



# True Hybrid and Multi-Cloud with OpenNebula

Version 1.0 – July 2021

## Abstract

To support digital transformation initiatives, IT departments need the right blend of on-premise, public, and edge cloud environments to support 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 each provider's best solutions 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. OpenNebula 6.0 Mutara brings with it a powerful **multi-cloud architecture** composed of Edge Clusters that can run **any workload**—both Virtual Machines and application containers—on **any resource**—bare-metal or virtualized—**anywhere**—on-premise, on 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	Datacenter
HA	High Availability
OpEx	Operating Expenditure
VM	Virtual Machine

## 1. The Need for Multi-Cloud

There is no doubt today that the hybrid cloud model, and the use of multiple cloud providers, is the norm for enterprises. In the new era of cloud, cost optimization and interoperability have become crucial factors. Multi-cloud strategies make provider independence possible, while addressing concentration risk 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<sup>1</sup>, multi-cloud strategies will reduce vendor dependency for two-thirds of organizations through 2024. The use of multiple clouds is by far the most common pattern among enterprises, with 92% adopting this strategy in Flexera's *2021 State of the Cloud Report*<sup>2</sup> (82% hybrid cloud, 10% multiple public cloud). Moreover, the multi-cloud approach should include on-premises resources. In the 2020 edition of the Forrester/IBM report *The Key To Enterprise Hybrid Multicloud Strategy*<sup>3</sup>, a survey of 350 global enterprise IT decision makers found that more than half of mission-critical workloads and 47% of data-intensive workloads will still be run either on-premises or on an internal private cloud in two years.

The COVID-19 pandemic of 2020-21 has overshadowed all aspects of life, and IT, both strategic and operational, is no exception. Flexera's 2021 survey duly showed that organizations' cloud plans and adoption did indeed shift as a result of the pandemic, with 90% reporting that cloud usage was higher than initially planned.

## 2. Challenges of Multi-Cloud

Despite the obvious benefits of deploying any workload across different public cloud and edge providers, there are still a number of challenges that IT organizations must overcome to get the most out 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:** workload 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, have no automatic enrollment of new nodes, and are complex and expensive solutions to deploy and maintain.

On the other hand, the **hyperscalers are now starting to offer pre-configured appliances** (e.g. AWS Outpost, 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, 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.

---

<sup>1</sup> <https://www.gartner.com/smarterwithgartner/4-trends-impacting-cloud-adoption-in-2020/>

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

<sup>3</sup> [https://ibm.com/downloads/cas/2EAPMGYK?cm\\_sp=ThinkDigitalResources\\_-\\_Infrastructure\\_-\\_Hybrid%20multicloud%20strategy](https://ibm.com/downloads/cas/2EAPMGYK?cm_sp=ThinkDigitalResources_-_Infrastructure_-_Hybrid%20multicloud%20strategy)

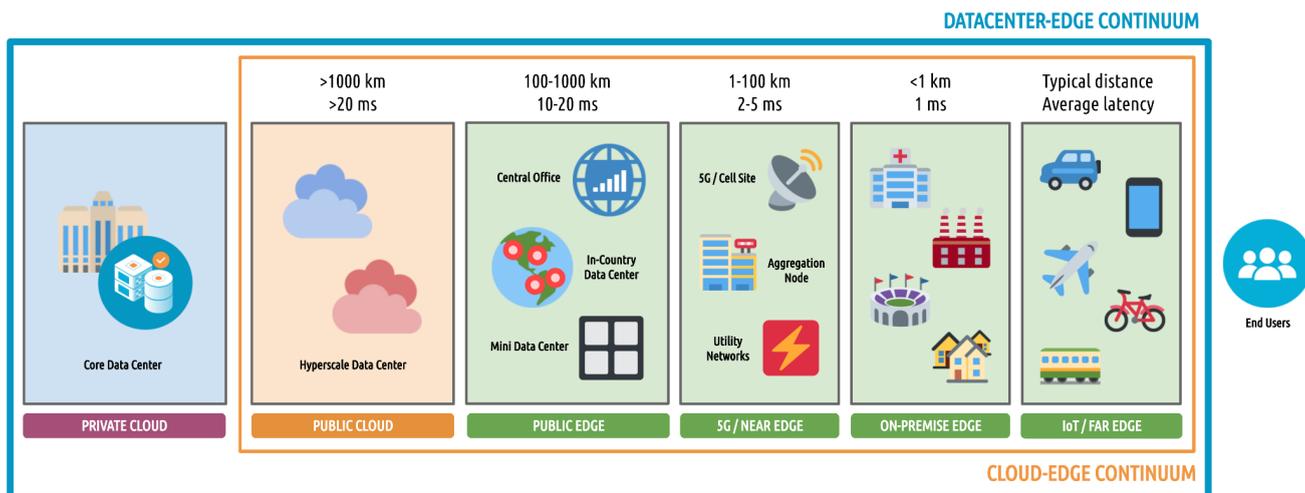


Figure 1. Edge computing brings new opportunities for organizations using a multi-cloud architecture.

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** with 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. From our experience, lack of skills is holding enterprises back from innovating across multiple cloud platforms.

### 3. Benefits of a True Multi-Cloud

With OpenNebula you can get the agility, scalability, and simplicity benefits of the pay-per-use public cloud, with 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 Virtual Machines and application containers 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

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 the 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 make sure that they only use as many resources as needed and pay less for them. 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. This is how OpenNebula can help your organization achieve that:

- **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 companies 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 the 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 incrementally and in a hybrid fashion. Our experience is that for most companies, the optimal cost is achieved by retaining a percentage of workload in the cloud.

## 4. What is OpenNebula?

OpenNebula is a **powerful, but easy-to-use, open source solution to build and manage Enterprise Clouds.** It combines virtualization and container technologies with 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 containers with Virtual Machine workloads in a common shared environment to offer the best of both worlds: mature virtualization technology and orchestration of application containers.
- **Any Infrastructure:** Unlock the power of a true hybrid and multi-cloud platform by combining edge, public, hosted, and private cloud operations.
- **Any Virtualization:** Integrate multiple types of virtualization technologies to meet your workload needs, from a fully virtualized environment to system containers and serverless deployments.

OpenNebula provides the necessary tools for running containerized applications from Kubernetes and Docker Hub, 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. OpenNebula supports a number of virtualization technologies, including VMware and KVM Virtual Machines for fully virtualized clouds, LXC system containers for container clouds, and Firecracker microVMs for serverless deployments.

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.**<sup>4</sup> If you are interested in designing and deploying an OpenNebula cloud on top of VMware vCenter, please consult our **VMWare Cloud Reference Architecture.**<sup>5</sup> If you are interested in an OpenNebula cloud fully based on open source platforms and technologies, please see our **Open Cloud Reference Architecture.**<sup>6</sup>

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**. Releases are produced on a regular basis and delivered as a single package with a smooth migration path. More information on the benefits of running an OpenNebula cloud can be found on the Key Features page.<sup>7</sup>

## 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 has to meet to make multi-cloud adoption a success:

### Manage Workload across Every Cloud from a Single Pane of Glass

The most crucial criteria to manage the 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 in order 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-premise, in public clouds, or at the edge—with a single console.
- **Portability**, by using Edge Clusters based on a common hyper-converged architecture that has been 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.

The Edge Clusters in OpenNebula can be **added and removed dynamically** in order to meet peaks in demand, or to implement fault tolerance strategies or latency requirements. Once the new cluster is part of

<sup>4</sup> <https://support.opennebula.pro/hc/en-us/articles/360050302811-Edge-Cloud-Architecture-White-Paper>

<sup>5</sup> <https://support.opennebula.pro/hc/en-us/articles/206652953-VMware-Cloud-Reference-Architecture-White-Paper>

<sup>6</sup> <https://support.opennebula.pro/hc/en-us/articles/204210319-Open-Cloud-Reference-Architecture-White-Paper>

<sup>7</sup> <https://opennebula.io/discover/>

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 specific **affinity** (place in the same node or cluster) or anti-affinity (place in a different node or cluster) rules to improve their fault-tolerance or performance.
- Finally, 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 brings challenges in assessing the suitability of on-premises apps for migrating to the cloud.

OpenNebula **combines application containers and Virtual Machines 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 containerized applications.

### Combine On-Premise and Public Cloud Resources

Although multi-cloud will become the norm, there's still a place for the on-premise 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-premise infrastructure by deploying Edge Clusters locally**. These on-prem Clusters can be hosted and automatically configured in your core DC in the same way that they would be 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.

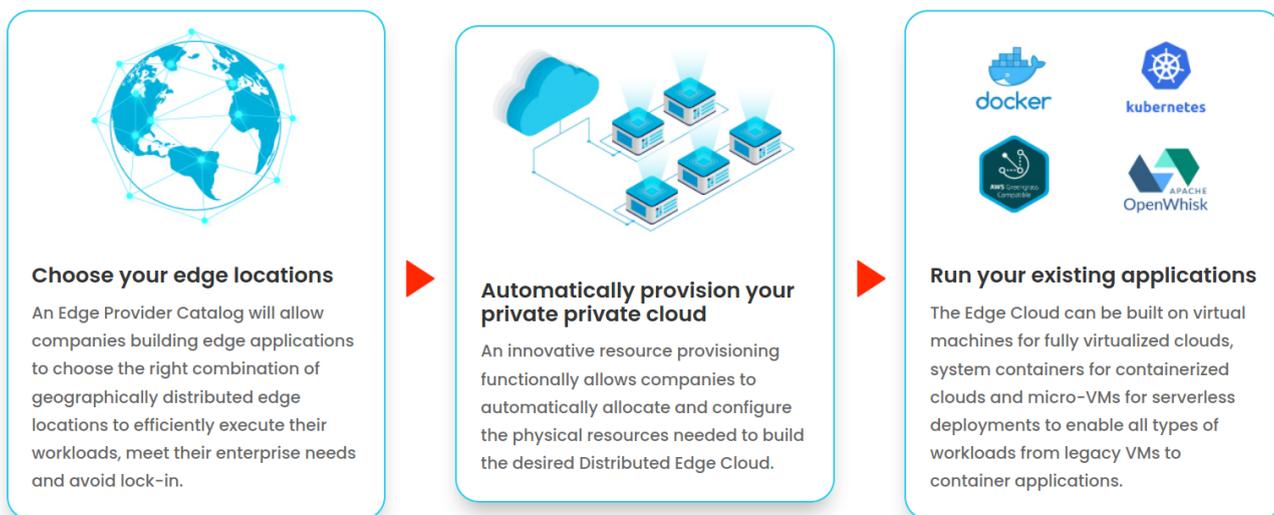
## 6. An Easy Button For Multi-Clouds

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.

The OneProvision<sup>8</sup> tool allows the deployment of a fully operational OpenNebula Cluster in a remote provider and the management of its full life-cycle, starting with its provisioning and maintenance, up to the final unprovisioning. Each remote location is defined as a group of physical hosts located in a bare-metal or virtual provider, fully configured with the selected hypervisor and enabled in the OpenNebula cloud stack for the end-users to deploy their workload on them.

---

<sup>8</sup> [http://docs.opennebula.io/stable/advanced\\_components/ddc/index.html](http://docs.opennebula.io/stable/advanced_components/ddc/index.html)



## 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 clouds 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 which are 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 in our latest version ([OpenNebula 6.0 "Mutara"](#)) currently supports **AWS**, **DigitalOcean**, **Equinix Metal**, **Google Cloud**, and **Vultr**. OpenNebula will be expanding this initial list of 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.

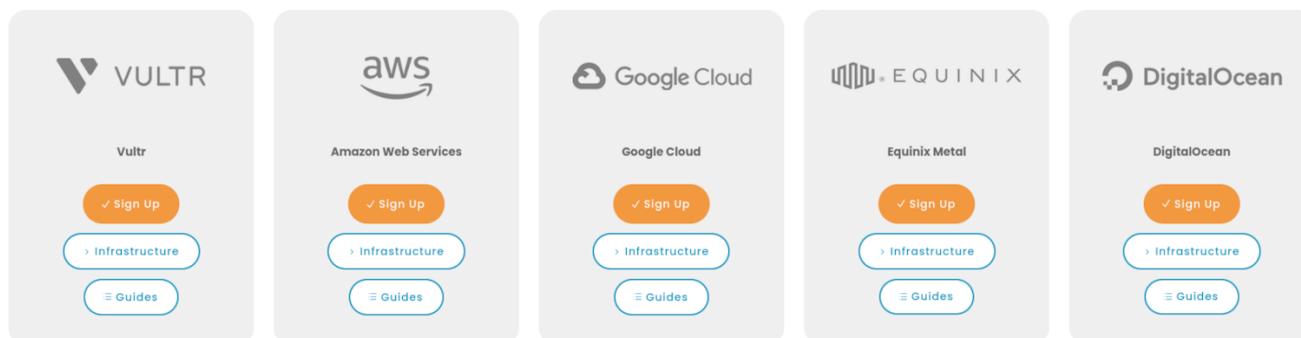


Figure 2. OpenNebula’s Provider Catalog offers a growing number of multi-cloud connectors.

## 8. Run Any Application, Anywhere

One of the critical parts of an OpenNebula cloud is its ability to support modern applications, combine application containers and Virtual Machines on a single platform, and integrate with existing Virtual Machine and container image hubs and marketplaces.

OpenNebula has inherent access to its own **Marketplace**, which enables users to import images from a public repository (containing images of common use that have been tested and certified by OpenNebula Systems) or from private repositories. These images can be added to a datastore and used at any time by existing VM templates or instances.

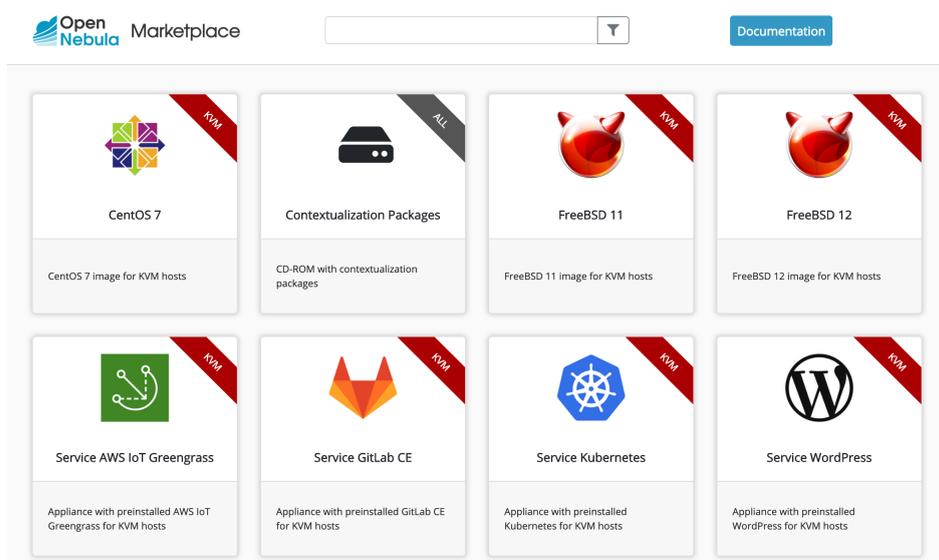


Figure 3. Some of the virtual appliances available in the OpenNebula Public Marketplace.

OpenNebula users can easily download, contextualize, and add virtual appliances from other public marketplaces, including **Linux Containers**<sup>9</sup> and **TurnKey Linux**.<sup>10</sup>

Distribution	Release	Architecture	Variant	Build date	LXC (privileged)	LXC (unprivileged)	LXD (container)	LXD (VM)
alpine	3.10	amd64	default	20200914_14:13	YES (2.0 and up)	YES (2.0 and up)	YES	YES
alpine	3.10	arm64	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	YES
alpine	3.10	armhf	default	20200914_13:51	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.10	i386	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.10	ppc64el	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.10	s390x	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.11	amd64	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	YES
alpine	3.11	arm64	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	YES
alpine	3.11	armhf	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.11	i386	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.11	ppc64el	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.11	s390x	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.12	amd64	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	YES
alpine	3.12	arm64	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	YES
alpine	3.12	armhf	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.12	i386	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.12	ppc64el	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO
alpine	3.12	s390x	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	NO

Figure 4. Some of the images provided by Linux Containers (left) and TurnKey Linux (right).

<sup>9</sup> <https://images.linuxcontainers.org>

<sup>10</sup> <https://www.turnkeylinux.org>

Application development is now increasingly relying on microservices architectures, which avoid breaking the whole stack when a particular feature is modified or added to the application. This trend in development is tightly coupled with the way these applications are deployed, usually through **application containers**. From version 5.12 “Firework” onwards, OpenNebula also comes with seamless integration with **Docker Hub**,<sup>11</sup> permitting direct execution of Docker Hub images on any hypervisor in your OpenNebula open cloud. This is a development and distribution model for applications and is especially suited for edge environments, since only a few legacy applications will ever be deployed on the edge, and new applications will most likely be developed using these modern models.

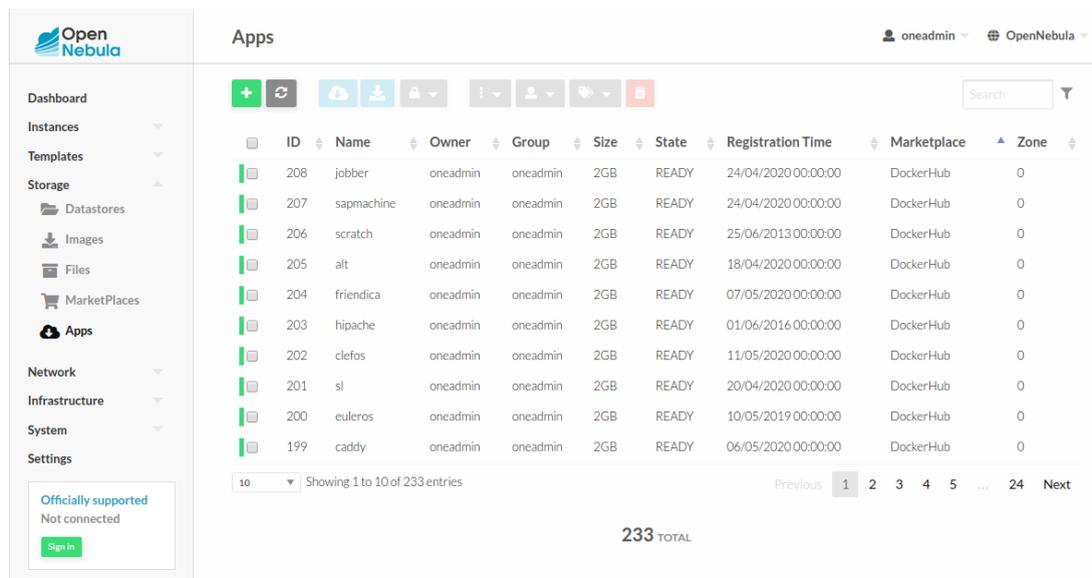


Figure 5. OpenNebula’s native integration with the Docker Hub marketplace.

OpenNebula’s native integration with Docker Hub provides an alternative, simpler method to orchestrate application containers, but users can always resort to a solution based on **Kubernetes**. OpenNebula offers two different approaches to integrating Kubernetes. The first one follows our general edge application deployment model, where a **multi-VM application** representing a full Kubernetes cluster (generally with elasticity rules) is deployed as a single entity on a particular edge deployment cloud location. The second approach involves the deployment of a Kubernetes controller as a **managed service**, offering the possibility to add more worker nodes to the Kubernetes clusters on demand.

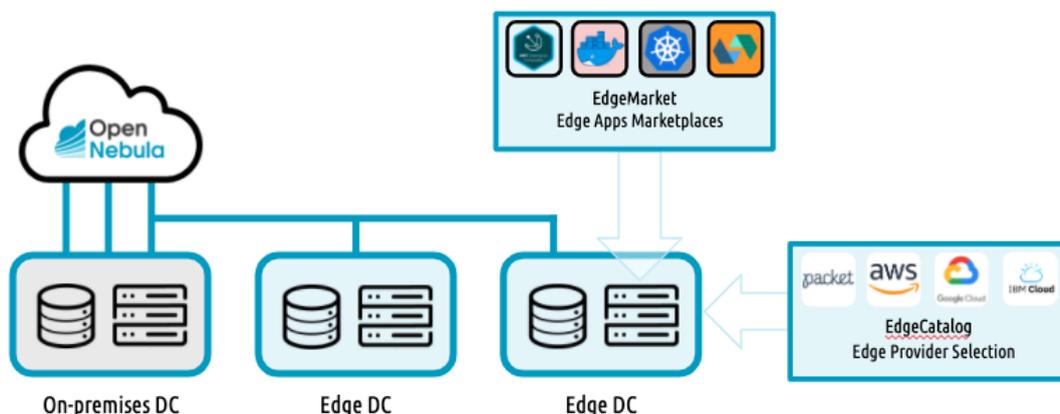


Figure 6. OpenNebula’s general edge application deployment model applicable to Kubernetes.

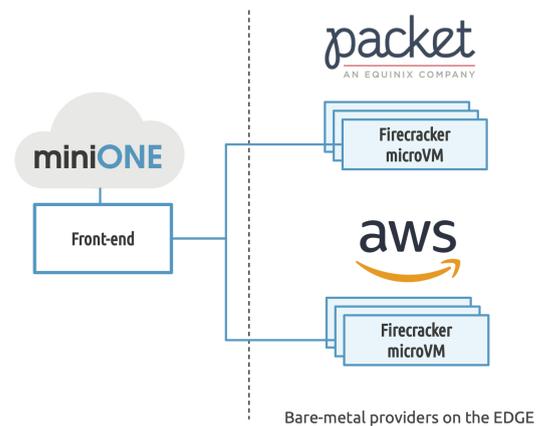
<sup>11</sup> <https://hub.docker.com>

## 9. Ready for a Test Drive?

You can evaluate OpenNebula and build a multi-cloud environment in just a few minutes by using **miniONE**,<sup>12</sup> 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, LXC, or Firecracker on multiple cloud providers.



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



## 10. Conclusions

Multi-cloud helps enterprises avoid the pitfalls of single-vendor reliance. Spreading workloads across multiple cloud providers gives enterprises 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: public, private, hybrid, and edge. The strategy is based on a powerful **multi-cloud architecture** that is 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 workload**—both Virtual Machines and application containers—**on any resource**—bare-metal or virtualized—**anywhere**—on-premise, on 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.

<sup>12</sup> <https://minione.opennebula.io>

## 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



### MANAGED SERVICES

Our team of experts can fully manage and administer your OpenNebula cloud for you



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](https://OpenNebula.io/getupdated)

© OpenNebula Systems 2021. 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.



Rev1.0\_20210706