Edge Cloud Reference Architecture

Version 2.4 – September 2024

Abstract

To support digital transformation initiatives, IT departments need to take advantage of the right blend of cloud environments—on-premises, public, and edge—tailored for a variety of existing and emerging use cases, while avoiding vendor lock-in and enabling cost optimization. They also need to combine Virtual Machines workloads with containerized Kubernetes applications in a shared environment, in order to get the most out of both—mature virtualization technologies plus secure orchestration of application containers.

This document presents a powerful distributed Edge Cloud Architecture for OpenNebula, composed of Edge Clusters that can run **any application**—Virtual Machine workloads and containerized applications from Kubernetes—on **any resource**—bare-metal or virtualized—**anywhere**—on-premises and on a cloud provider. Our Edge Cloud Architecture 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.We have defined this architecture to be much simpler than traditional cloud computing architectures, which are usually composed of complex, proprietary general-purpose software systems for storage and networking.

This architecture has been created from the collective information and experiences of hundreds of users and client engagements over the last ten years. It builds on storage and networking technologies that already exist in the Linux operating system and on modern storage hardware that is available from existing cloud and edge providers, leading to a greatly simplified design. Our Edge Cloud Architecture implements enterprise-grade cloud features for performance, availability, and scalability with a very simple design that avoids vendor lock-in and reduces complexity, resource consumption, and operational costs.

Contents

Glossary

DFS	Distributed File System
HA	High Availability
I/O	Input/Output
IaC	Infrastructure as Code
QCOW	QEMU Copy On Write
VM	Virtual Machine

1. 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 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.
- **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. OpenNebula combines support for KVM virtual machines and Kubernetes clusters with federation, multi-tenancy, and automatic provisioning of cloud and edge infrastructure resources for building private, public, and hybrid clouds. If you want to find out more about OpenNebula's innovative approach towards hybrid cloud and multi-cloud, see our white paper True Hybrid and Multi-Cloud with OpenNebula.[1] If you are interested in building an OpenNebula cloud using your choice of open source platforms and technologies, please refer to our Open Cloud Reference Architecture.[2][3]

The development of OpenNebula follows a bottom-up approach driven by the real need of sysadmins, DevOps, and corporate users. OpenNebula is an **open source product** with a healthy and active community, and is 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.[4]

2. Design Principles for Distributed Edge Cloud

The pressures faced by IT Departments and Operations Teams are increasing at a fast pace. Teams, and in fact organizations as a whole, are expected to:

- **Innovate** to bring new applications and services to market faster.
- **Be agile** to meet the increasing needs of developers for new development tools and frameworks, such as containers or infrastructure as code (IaC), while continuing to offer infrastructure on demand to application administrators and ensuring enterprise requirements for DevOps practices.
- **Adapt to stay relevant** in a new reality where hybrid cloud and multi-cloud are quickly growing, and edge computing is the next IT transformation.

[1] https://opennebula.io/white-papers/true-hybrid-and-multi-cloud-with-opennebula-white-paper/
[2] https://support.opennebula.pro/hc/en-us/articles/204210319-Open-Cloud-Reference-Architecture-White-Paper
[3] https://support.opennebula.pro/hc/en-us/articles/206652953-VMware-Cloud-Reference-Architecture-White-Paper
[4] https://opennebula.io/discover/

Based on our experience of working with hundreds of users and client engagements, we have defined an Edge Cloud Architecture that builds a single distributed cloud platform to run any workload—from virtualized to containerized—across multiple clusters that can run anywhere—on-premises or at the edge—and on any resource—from bare-metal to virtual—with unparalleled availability, performance, and simplicity.

- The architecture defines a **complete end-to-end solution based on proven open source** storage and networking technologies that already exist in the Linux operating system, which allows to avoid vendor lock-in while minimizing IT complexity, resource consumption, and operational costs; maximizing performance, availability, and reliability; and simplifying the automation of its deployment and management.
- The architecture implements **a distributed approach** with a single cloud Front-end that controls one or several interconnected Edge Clusters that can run in multiple geographically-distributed data center locations, and cloud and edge resource providers.
- **Clusters can be added and removed dynamically** in order to meet peaks in demand, or to implement fault tolerance strategies or latency requirements.
- The **Edge Clusters are based on a common reference architecture** that has been designed to be deployed on any resource, from bare metal to virtual, enabling **workload portability**.
- The **Edge Cluster's internal design follows a hyper-converged approach** that uses the optimal infrastructure configurations in each environment to ensure isolation and performance, to easily scale-in and scale-out nodes to match compute and storage needs, and to migrate workload.
- The edge architecture **enables true hybrid and multi-cloud computing,** and leverages the growing ecosystem of cloud, edge, and access network cloud providers.
- The **Edge Clusters are capable of operating autonomously** in case of no network connectivity, to allow applications running in local data centers to continue their operation without any management service.
- The system supports modern applications, **combines application containers and Virtual Machines on a single platform**, and integrates with existing Virtual Machine and container images hubs and marketplaces. 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 the Edge Cloud Architecture provides support for both non-persistent and persistent VMs and Kubernetes containerized applications.

Lastly, it produces a **single vendor experience** because the complete cloud stack is fully supported and optionally managed by OpenNebula Systems. This means a simplified experience for procurement, consulting, and support. You can accelerate and streamline roadmap development, migrations, and upgrade paths with a single vendor.

3. High-Level Reference Architecture

Modern cloud environments have evolved to highly complex backends that limit their reliability and hinder their operability. They are composed of proprietary, costly general-purpose software systems for storage and networking that are unnecessarily complex because they were designed to solve too many problems at once. The OpenNebula Edge Cloud Architecture departs from this trend by using lightweight systems and software components with modest hardware requirements that are easy to maintain and operate—while also providing maximum application performance—and by implementing an innovative design specifically conceived to manage virtualized applications in distributed multi-site environments.

OpenNebula's Edge Cloud Architecture is articulated around the concept of Edge Cluster. An Edge Cluster is a hyperconverged functional set of managed objects that include storage, network, and host resources. An Edge Cluster provides all the resources needed to run virtualized or containerized applications.

OpenNebula's management services, including scheduling, monitoring and life-cycle management, run in the cloud Front-end and orchestrate the Edge Clusters. The Front-end also provides access to the administration tools, user interfaces, and API. Although the requirements may vary depending on the number and size of the clusters and API load, the Front-end node only requires 8 GB of main memory and 4 cores. The basic building blocks of the architecture, including the Edge Clusters, are depicted in Figure 1.

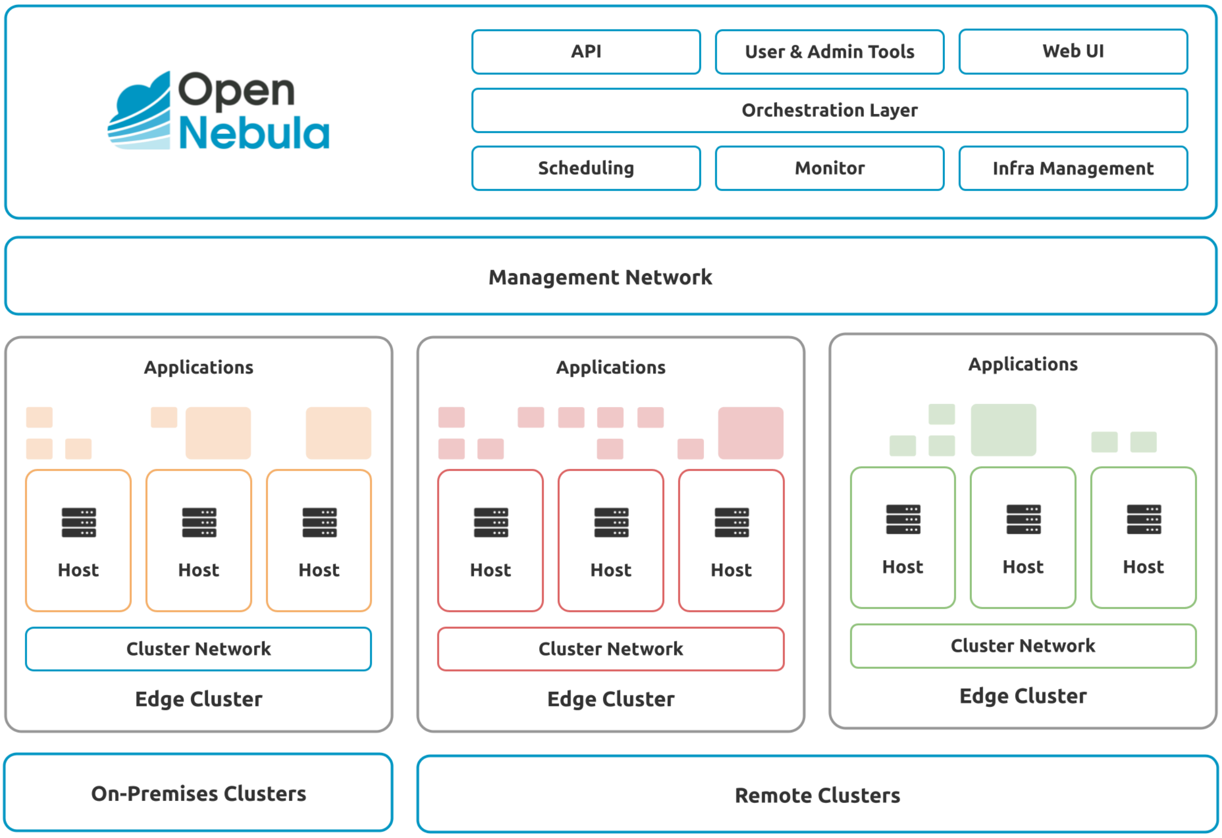

Figure 1. Main components of our Edge Cloud Architecture.

The Edge Clusters are built on virtualization hosts and an interconnection network. The virtualization hosts are responsible for providing applications with execution resources (e.g. CPU, memory, or network access) through a suitable form of virtualization, such as KVM Virtual Machines or containerized Kubernetes applications.. The actual virtualization technology depends on how the cluster hosts are provisioned (see Section 4) and the profile of the workload. Additionally, hosts supply the cluster with the storage space needed to run the applications.

Hosts in a cluster share the same configuration for installed software components, the OpenNebula administration user, and accessible storage. The Front-end nodes must connect to the cluster hosts to receive status and monitoring updates, as well as to initiate management operations. Cluster nodes are interconnected through one or more private networks, which are normally used for storage transfers as well as for application communication across virtualization hosts. Finally, access to a public link is required in order to provide applications with internet connectivity.

The Edge Cloud Architecture is able to provide a lightweight and easy-to-use storage platform for medium-sized clusters consisting of tens of nodes that can run on-premises and on-cloud, and on physical or virtualized resources. Overall, OpenNebula's Edge Cloud Architecture is able to manage hundreds of these clusters, as they operate autonomously in terms of networking and storage, and handle thousands of virtualized hosts and tens of thousands of virtualized applications.[5]

4. Edge Cluster Deployment Models

The main advantage of the OpenNebula Edge Cloud Architecture is the ability to deploy clusters anywhere. This provides application mobility and true multi-cloud computing in the most literal sense. Clusters can be deployed on on-premises infrastructure, on public bare-metal providers, and on virtualized cloud environments to enable powerful hybrid and multi-cloud computing. Infrastructure teams have the flexibility to choose their preferred hardware platform and cloud provider, and deliver an exceptional OpenNebula experience.

In particular clusters can be provisioned in three different forms:

- **On-premises** assumes full management capabilities of the resources, and has no restrictions in terms of networking. Usually, this type of provision uses in-house data center resources.
- **Metal Remote** presents some restrictions on IP addressing and connectivity of resources. A metal provision usually requires interacting with a cloud/edge provider API and uses bare-metal instances.
- **Virtual Remote** presents limitations on the capabilities of the cluster hosts as well as network connectivity. A virtual remote cluster deployment is usually based on Virtual Machine instances from a cloud or edge provider.

Note that some cluster deployments or cloud providers may impose restrictions on the functionality available in the cluster (e.g. supported hypervisors or host connectivity). As shown in Figure 1, a single Front-end can manage clusters from multiple locations at the same time. While the cluster architecture has been tested on the major cloud providers, the automatic provision feature currently only supports Equinix Metal and AWS. OpenNebula is developing drivers for other widely-used cloud and edge providers.

5. Storage Architecture

Our Edge Cloud Architecture is based on our own implementation of a local storage driver, developed for the efficient management of disk images in highly-distributed cloud environments, and fully supported by OpenNebula Systems. The local storage driver was designed to meet the following requirements:

- **Access to marketplaces** that act as global image repositories, such as the OpenNebula Marketplace or private HTTP repositories.
- **Minimize image transferring and maximize application I/O performance**. Clusters can use a cloud/edge deployment model over public internet links. Moreover, on-premises provisions can scale to a high number of hosts. Storage should not be a bottleneck for any of these situations.
- **Simple deployment.** Reduce the solution’s complexity and footprint, by using lightweight technologies that already exist in the Linux operating system to accommodate any deployment model—on both physical and virtual resources—and by increasing the reliability of the backend.

The Edge Cloud Architecture combines a 3-tier global architecture for image distribution with an enhanced Filesystem Datastore that implements replica caching, snapshotting, and backups within each Edge Cluster.. The 3-tier storage architecture (see Figure 2) consists of:

[5] https://docs.opennebula.io/6.8/installation_and_configuration/large-scale_deployment/scalability.html

- **Tier 1 - Marketplace**. This tier consists of the remote servers and storage implementing the global application image repositories, including the OpenNebula Public Marketplace and third-party HTTP public or private marketplaces. Additionally, it will include the OpenNebula Community Marketplace, currently in development, which will offer OpenNebula Appliances developed by third-party contributors and certified by OpenNebula.
- **Tier 2 - Image Datastores**. This tier consists of the zone image datastores, and provides the primary image storage location for the OpenNebula zone. This storage area is provided for one or more dedicated nodes. Its contents are cached and replicated on demand to each cluster. It supports deploying to Edge Clusters on any location, as well as scaling on-premises infrastructure.
- **Tier 3 - Cluster Replicas.** Application images are cached within a cluster in dedicated replica hosts, in order to minimize image transfers. Replica hosts use a specialized distribution system to make images available to all cluster hosts, and support snapshotting for recovering from host failures.

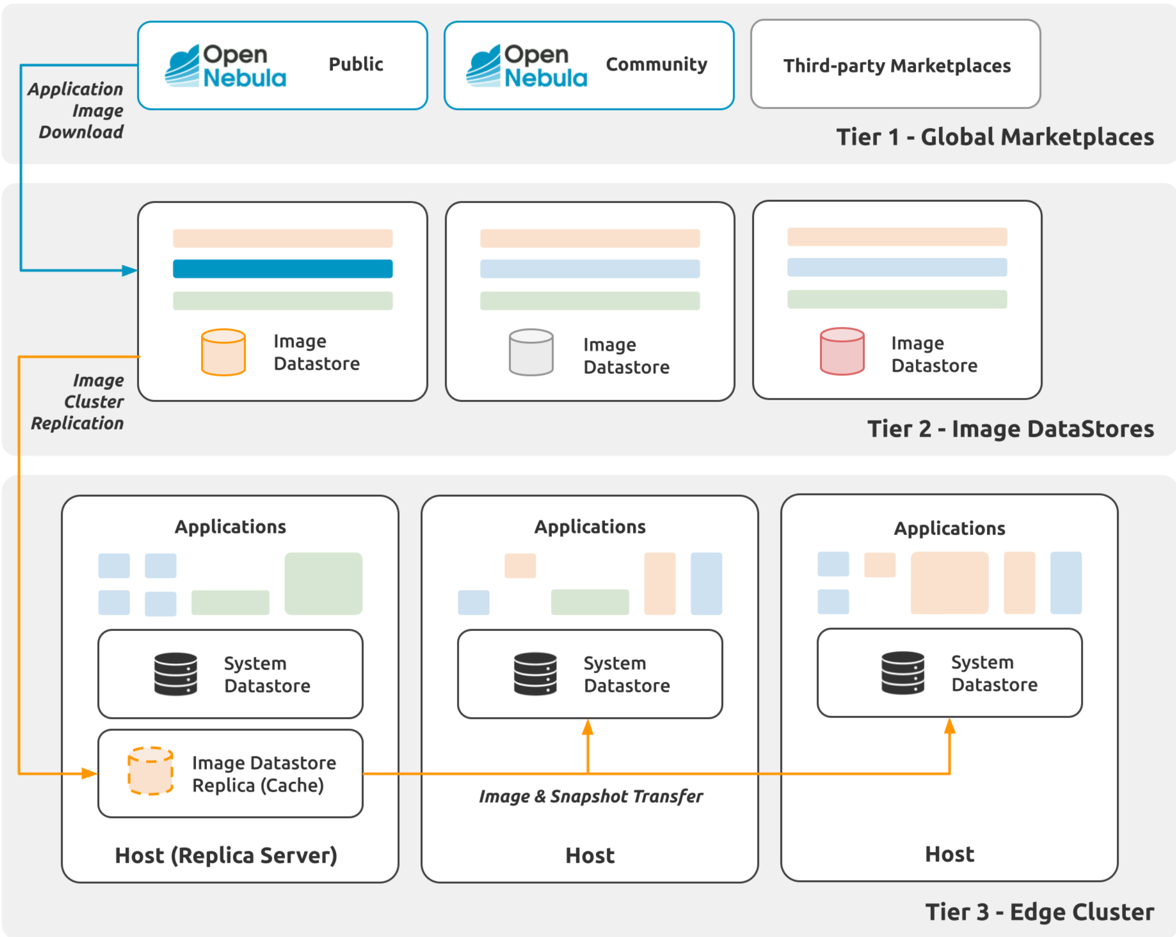

Figure 2. The 3-tier storage system of the Edge Cloud Architecture.

Within each tier-3 cluster, disk images are transferred between the Image and System datastores by using an enhanced SSH transfer mode that includes replica caching and snapshotting, which greatly improves its scalability, performance, and reliability. Images are cached in each cluster to ensure they are available and in close proximity to hypervisors, thus reducing bandwidth requirements for the tier-2 Image datastore servers and considerably reducing deployment time. This is especially important in highly-distributed edge deployments, where copying images from the tier-2 Front-end to the tier-3 cluster hypervisors could be very slow.

The edge storage architecture has been implemented with lightweight technology components that translate into modest hardware requirements, such as SATA SSDs and 10-gigabit networks. Moreover, its deployment follows a hyperconverged approach which does not require dedicated servers to implement a distributed storage system. This reduces the complexity of the solution, enabling the use of the local storage area of the cloud cluster hosts.

Application Deployment Times and I/O Performance

To reduce file transfer and instantiation times, and to transfer with minimal overhead, application images are stored in the qcow2 file format. Using qcow2 files to back up disk images also facilitates backup solutions, reduces image transfer times, and allows for efficiently implementing advanced features such as snapshotting. Expected deployment times will depend on disk sizes and the interconnection links between tiers.

To maximize I/O performance for applications, applications run from the direct-attached storage on the hosts. The result is that I/O performance is close to that of the native host, and is only impacted by the virtualization layer. Tests on a t1.small metal instance on Equinix and a VM on the same server showed a difference in performance of less than 5%.

Application Snapshot Performance

To improve the availability of the Edge Cluster blocks, live migration is supported within the same cluster. Application snapshots are also kept within the Edge Cluster (tier 3) to enable fast recovery from the last application checkpoint.

A recovery operation can potentially impact two different areas:

- I/O noise generated by the snapshot operation could reduce the I/O of neighbor applications. In our case this cost is negligible as it is based on the QEMU Redirect-on-Write feature.
- Network bandwidth to move the recovery snapshots to the cluster replica server. In this case we use a delta-transfer algorithm to reduce the information transferred to the server. However, this time will increase as the contents of the disk diverges from the original.

Another important aspect to consider is VM recovery time. Recovery times for VMs with or without snapshots are similar, since the VM base images are already located on the cluster replica (tier 3). Hence the only additional overhead results from transferring the disk snapshot, which is already available on the Edge Cluster as well.

Alternative Storage Backends

OpenNebula provides support for most popular enterprise-class SDS (Software Defined Storage) backends, such as Ceph or Gluster, designed to offer high scalability and availability. Although these platforms are rich in features and offer high performance, their operation requires extensive experience, higher financial investment, dedicated hardware (whether or not they are hyperconverged), and a significant amount of resources. In general, their cost, complexity, and resource requirements prevent their practical use in Edge Clusters.

6. Networking Architecture

Clusters use four types of networks:

- A **Storage network** dedicated to the distribution system of application images.
- A **Private network** to implement application interconnection networks.
- A **Management network** to interconnect the hosts to the Front-end services.
- A **Public network** to interconnect applications to the Internet. Usually requires a predefined set of public IPs available in the Edge Cluster.

Note that the Management traffic may be routed through the public network if the clusters are deployed remotely.

The characteristics of each network strongly depend on the deployment model used for a cluster. The on-premises network model provides full capabilities, and is based on standard Linux bridging. It supports VLAN tagging (802.1Q) for isolation of private networks over shared links and devices, and VXLAN technology designed for dealing with large cloud deployments. VXLAN solves the 4096 VLAN limitation problem by encapsulating Ethernet frames within UDP packets.

Cloud or Edge Clusters, on the other hand, may introduce certain limitations to the network topologies that are available for each network type, depending on the provider. In general, the OpenNebula network stack uses specific provider drivers to register elastic public IPs and private networks to applications.

In both cases—whether on-premises or cloud/edge—applications benefit from the features implemented by OpenNebula's network kernel, including security groups, automatic network scheduling, and user self-provisioning models.

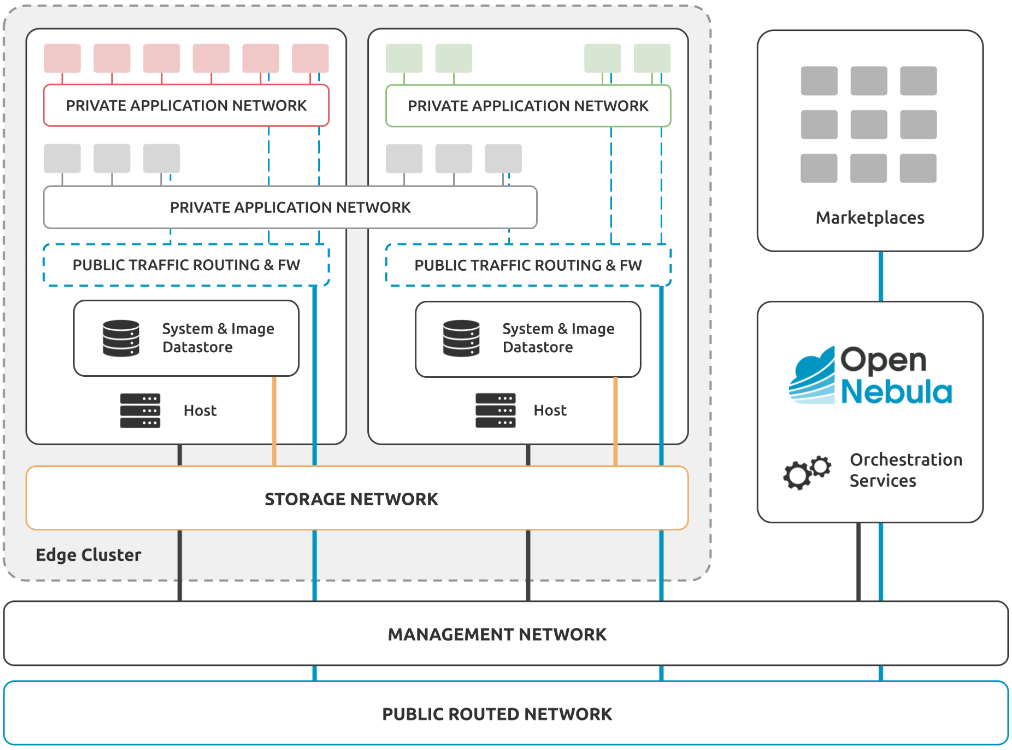

Figure 4. Overview of the networks used in the Edge Cloud Architecture.

7. High Availability

Availability is a critical aspect of cloud architecture design, specially for application data recovery and integrity. The Edge Cloud Architecture provides multiple mechanisms and built-in capabilities to achieve high availability and guarantee that your data and services are always available:

- **Front-end**: OpenNebula uses a distributed consensus protocol, based on Raft, to provide fault tolerance and state consistency across OpenNebula Front-end services. To support one node failure, a minimum of three Front-ends must be deployed. When running the Front-end on cloud resources, the Front-end HA clusters can be deployed across countries or even continents.

- **Multi-cluster**: The Edge Cloud Architecture helps you to automatically deploy geo-distributed cloud and edge environments that enable multi-cloud availability patterns by allowing applications to span across multiple clusters. Thanks to this approach, you can make your applications more highly available, scalable, and fault tolerant than traditional single data center infrastructures.

- **Cluster**: OpenNebula implements autonomous operation of remote clusters, in case clusters running on cloud and edge locations lose their connection with the central centers. Applications running in a cluster can continue their operation without any management service.

- **Host**: The 3-Tier Replica Storage architecture in the Edge Clusters implements an availability system based on periodic snapshots. These are used to recover from VM and host failures by soft-fencing the node to prevent split-brain conditions, then automatically restarting the application in another node. Application snapshots are kept within the cluster (tier 3) to enable fast recovery from the last application checkpoint. Within each node, availability may be improved by replicating hardware (e.g. RAID, NICs) and network paths. Availability within applications may be improved by implementing application-level HA across multiple clusters, when data and application state integrity is required.

- **Application:** The networking architecture of the Edge Cluster, combined with OpenNebula's capabilities for orchestration, allow developers to use application-specific replication mechanisms where data and application state integrity is required. Moreover, maintenance work can be performed while the system is in operation, and host failures can be mitigated thanks to support for live migration of applications within cluster hosts.

8. Automatic Provision of Edge Clusters

OpenNebula brings the provisioning tools and methods needed to dynamically grow on demand a private cloud infrastructure with resources running on remote cloud and edge providers. This disaggregated cloud approach allows for a seamless transition from centralized private clouds to distributed edge-like cloud environments. Companies can grow their private cloud with resources at cloud and edge data center locations to meet peaks in demand or the latency and bandwidth needs of their workload. 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's OneProvision tool 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 for the end-users.

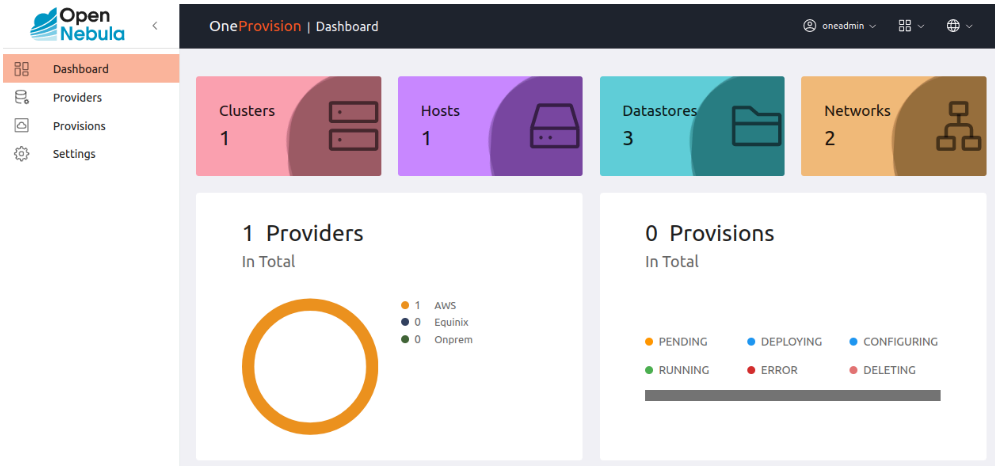

Figure 5. Visual interface for the OneProvision tool.

9. 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.[6]

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 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.

[6] https://support.opennebula.pro/hc/en-us/articles/6554989538717-Enterprise-Kubernetes-Made-Simple-White-Paper

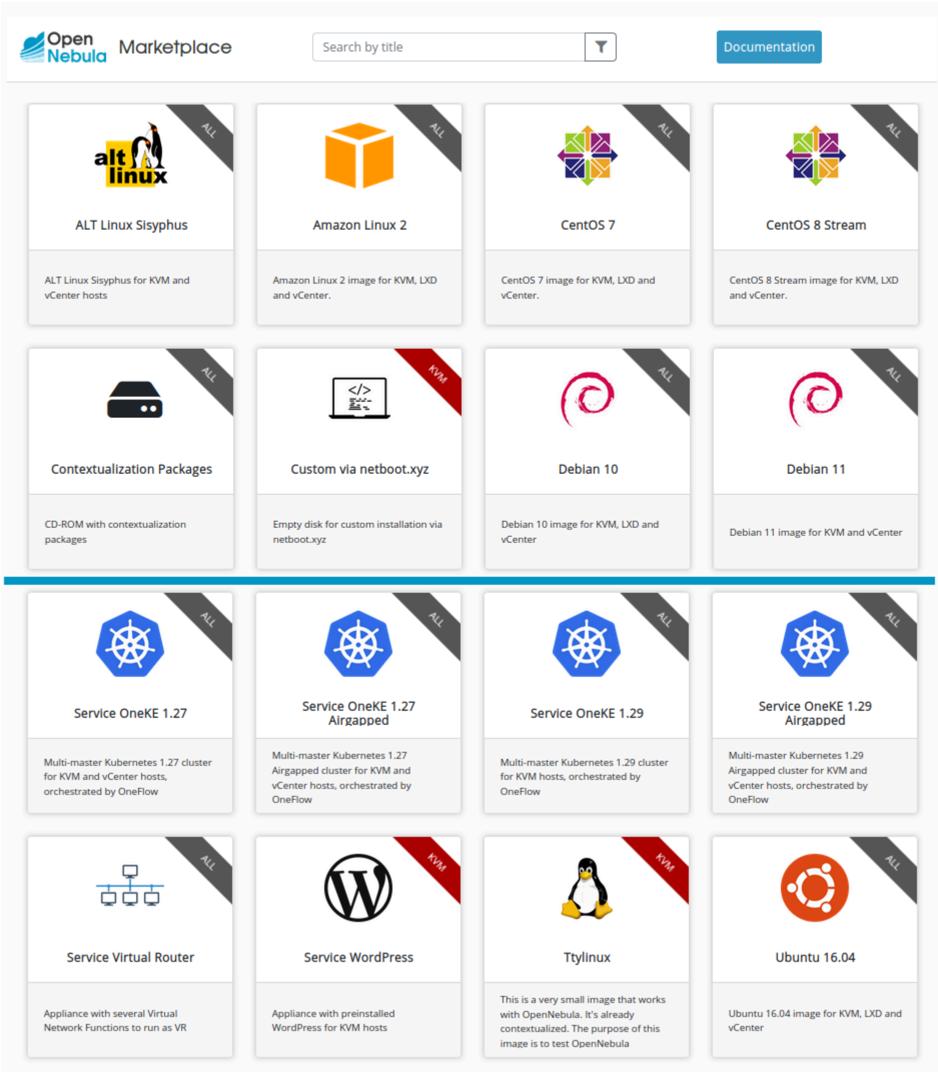

Figure 6. Some of the virtual appliances available in the OpenNebula Public Marketplace.

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.

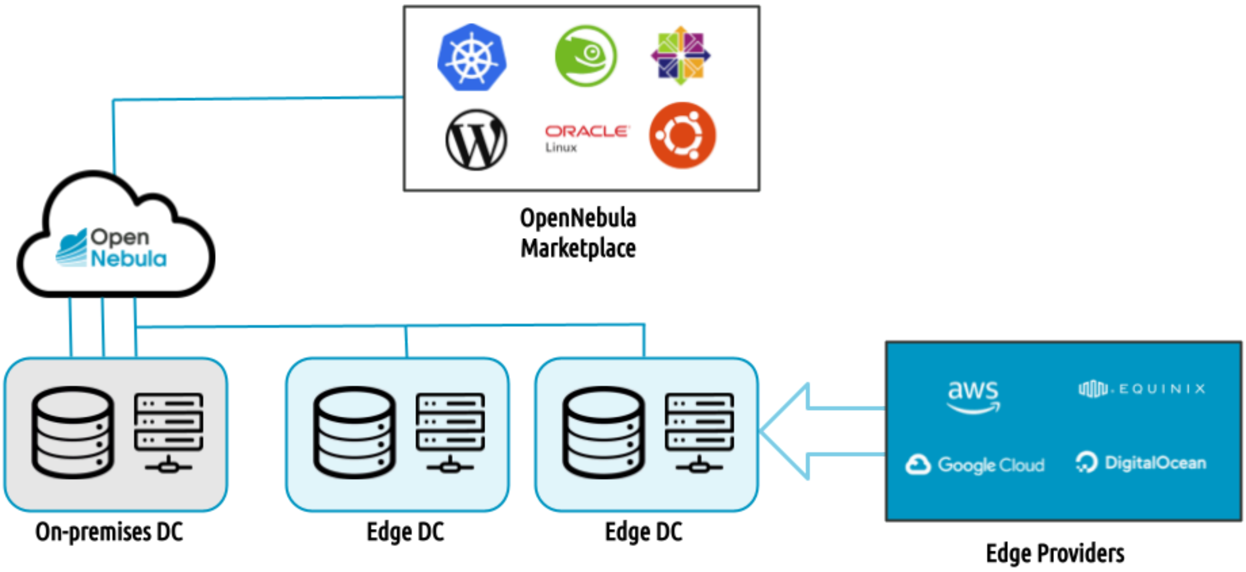

Figure 7. OpenNebula's general edge application deployment model applicable to Kubernetes.

10. Ready for a Test Drive?

You can evaluate OpenNebula and build a cloud in just a few minutes by using **miniONE**,[7] 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.

miniONE

11. Conclusions

It is clear that the evolution of the modern cloud has led to the creation of highly complex systems, often based on proprietary technologies. This document outlines OpenNebula's choice for simplicity and open source, and presents a powerful Edge Cloud Architecture composed of Edge Clusters. They are built on demand using storage and networking technologies that already exist in the Linux operating system, and can run **any workload**—both virtual machines and application containers—**on any resource**—bare-metal or virtual—**anywhere**—on-premises, on the cloud, or at the edge.

Our Edge Cloud Architecture 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.

[7] 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

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

LINUX FOUNDATION SILVER MEMBER

CLOUD NATIVE COMPUTING FOUNDATION SILVER MEMBER

LF EDGE GENERAL MEMBER

GAIA-X data-infrastructure.eu

Rev2.4_20240910