OpenNebula's Approach to Edge Computing

Version 2.2 – April 2021

Abstract

Latency reduction and bandwidth improvement are becoming critical factors for forthcoming technologies. New applications and use cases such as vehicle control, augmented and virtual, IoT, online gaming and new services for the manufacturing, healthcare and energy sectors are emerging at a fast pace. They all have a common requirement: a high-performance real-time response. Edge Computing helps companies and providers to transfer computational power and resources closer to their end-users by increasing the number of endpoints and locating them near the consumers—be they users or devices.

In the context of its **ONEedge** project,[1] OpenNebula is bringing new innovative edge features to enable companies to build their own private, lightweight and nimble Edge Computing environments based on highly-dispersed edge nodes in close proximity to users, machines, and sources of data. Not only are companies able to easily create their own edge deployments, and to manage them with the utmost simplicity, but they are also able to create these environments on-demand to meet peaks of demand, implementing fault-tolerant strategies or latency requirements without having to provide or own those underlying resources at all.

This white paper describes the principles that the **ONEedge** project is following and the main architectural elements of the solution that OpenNebula's new Edge Computing platform provides. This document also includes a number of use cases and practical resources that demonstrate how Edge Computing applies to different industries and the benefits it delivers, and the solution OpenNebula offers to companies interested in increasing their competitiveness by implementing their own edge deployments.

Contents

Glossary

CORD	Central Office Re-architected as a DC
DC	Datacenter
IoT	Internet of Things
MEC	Multi-Access Edge Computing
NFV	Network Function Virtualization
OTT	Over-the-Top
SDN	Software-Defined Networking
VDI	Virtual Desktop Infrastructure
VM	Virtual Machine

[1] https://ONEedge.io

1. The Need for Low Latency

So many companies now—both long-established and newly emerging, from multiple industries like gaming, Internet of Things (IoT), social networking and telecommunications—are focusing their business strategies on being able to provide innovative services and capabilities with absolute immediacy for their customers. Consequently, **latency** has become the key factor for quality of experience. Even a delay of milliseconds is perceptible if an application requires instantaneous interaction.

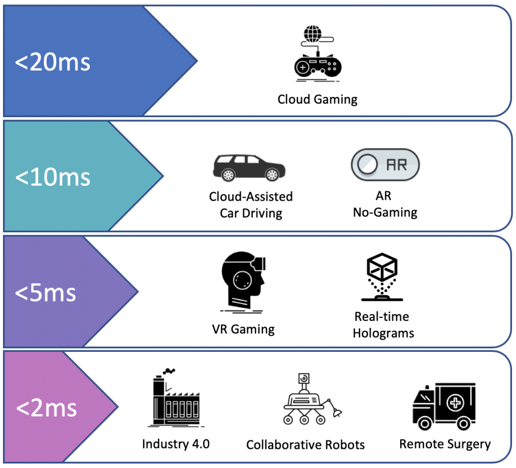

Figure 1. Estimated latency requirements in milliseconds for various types of applications.

"Edge Computing" is the progressively maturing paradigm of shifting away from centralized processing towards carrying out the processing as close in proximity as possible to mobile devices, sensors, and end-users. Yet, despite the urgent need to define a foundational platform that can deliver these edge services, and the dramatic growth in public interest and industry investment, at present, the companies that are developing next-generation solutions simply do not have a viable platform on which their low latency applications can become mainstream.

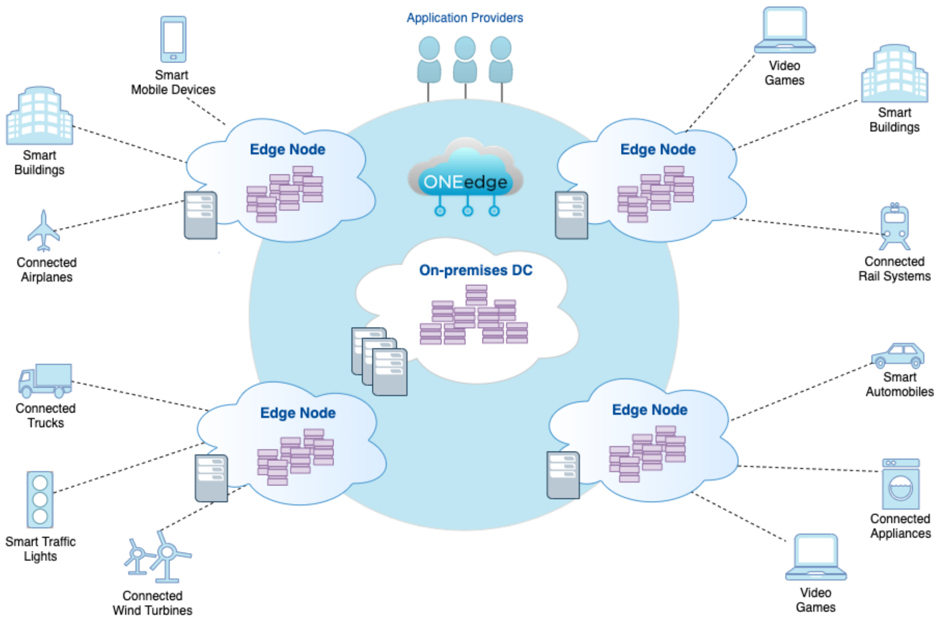

Figure 2. Architecture for disaggregated edge management.

2. 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, LXD system containers for container clouds, and Firecracker microVMs for serverless deployments.

This white paper describes the new features for Edge Computing that are being developed by OpenNebula as part of its ONEedge initiative. 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**.[2] If you are interested in designing and deploying an OpenNebula cloud on top of VMware vCenter, please refer to our VMWare Cloud Reference Architecture.[3] If you are interested in an OpenNebula cloud fully based on open source platforms and technologies, please refer to our Open Cloud Reference Architecture.[4]

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, 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.[5]

3. Edge Cloud Use Cases

This section provides a number of use cases that demonstrate how Edge Computing applies to different industries and the benefits it delivers. These examples are based on relatively simple deployments that show how OpenNebula's new edge features offer an effective solution for different scenarios. These use cases should be able to be reproduced quite simply, as we have included a screencast video for each case that walks the reader through the actual set up, as well as an accompanied guide to help with the step-by-step implementation. If you would like any assistance in deploying any of these edge use cases at scale, don't hesitate to reach out to our OpenNebula Services Team for consultation and support.

[2] https://support.opennebula.pro/hc/en-us/articles/360050302811-Edge-Cloud-Architecture-White-Paper
[3] https://support.opennebula.pro/hc/en-us/articles/206652953-VMware-Cloud-Reference-Architecture-White-Paper
[4] https://support.opennebula.pro/hc/en-us/articles/204210319-Open-Cloud-Reference-Architecture-White-Paper
[5] https://opennebula.io/discover/

Multiplayer Gaming

The massive growth of online multiplayer gaming across the globe has been driving the ever-growing need for **low latency**. Higher latency hurts the players' gaming experience and harms the reputation of gaming companies providing the service. Needless to say, having to deal with unpredictable lags when engaged, for instance, in a competitive multiplayer game can be especially frustrating. Hence, the need for low latency in online gaming requires moving beyond a "centralized cloud approach" and adopting a "distributed edge model" instead.

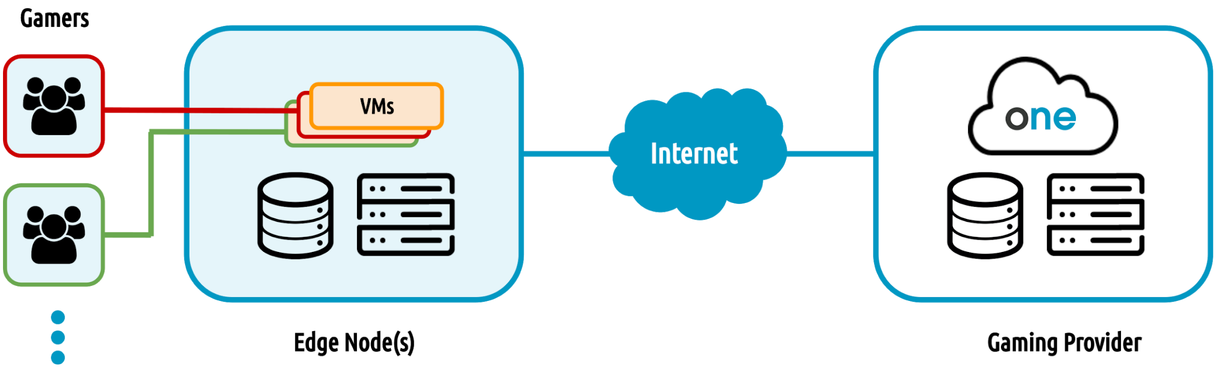

Figure 3. Online Video Gaming Architecture.

Today, the most innovative gaming and entertainment companies—and their customers—are starting to experience first-hand the many benefits derived from setting up a distributed infrastructure. This new model provides players with a much better gaming experience, especially in fast-paced environments such as First Person Shooter (FPS) and Multiplayer Online Battle Arena (MOBA) video games. As an example, we have shown how OpenNebula can easily deploy multiple Kubernetes clusters at the edge for multiplayer gaming, thus reducing latency and providing an overall better user experience. We have used Agones, a platform developed by Google and Ubisoft, to deploy at the edge a dedicated game server for Xonotic, the well-known open source FPS multiplayer game.

OpenNebula and Kubernetes - Deploy Multiplayer Gaming Servers on the Edge

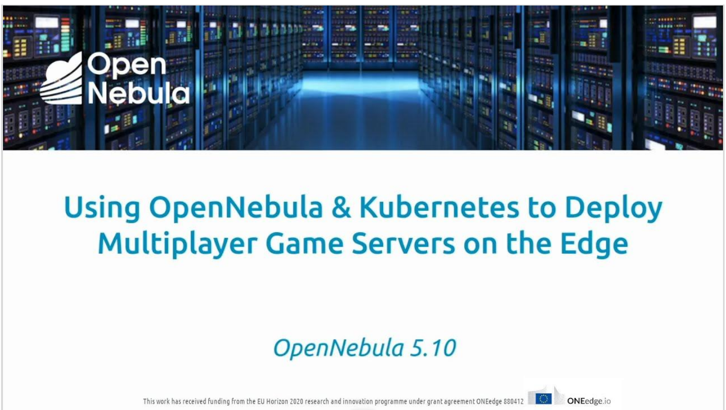

Screencast Video:	youtube.com/watch?v=2boo_vPFxO8
Guide:	support.opennebula.pro/hc/en-us/articles/360041948632

Internet of Things

The ubiquitous collection and sharing of data, and the "digitalization" of almost everything around us is transforming the way production systems are designed. The comprehensive, real-time data collection and analysis we are seeing results in the magnification of the amount of compute processing capabilities, along with broad distribution of these "things" that are transformed into active processors of information. And organizations need the flexibility and speed to deploy and manage an infrastructure that supports this highly-distributed, innovative architecture.

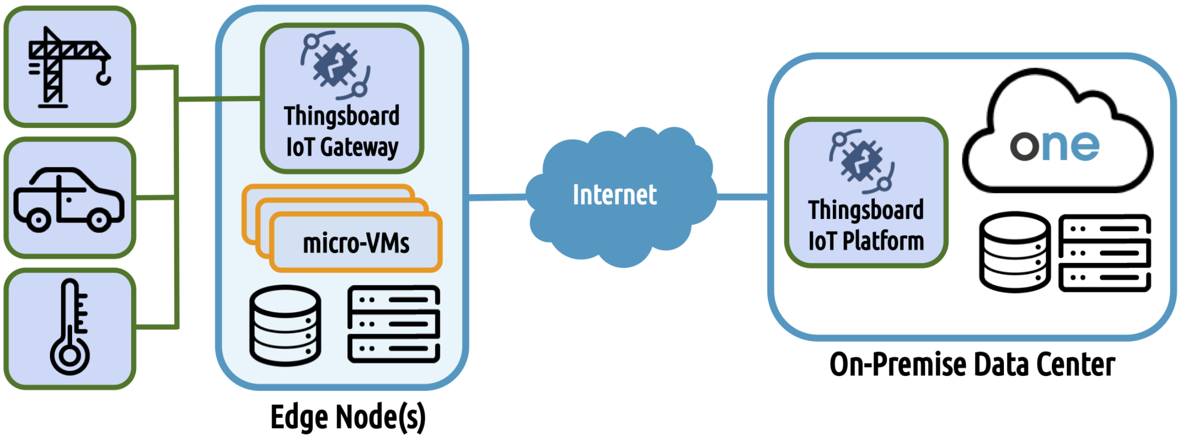

Figure 4. IoT architecture.

OpenNebula and its seamless integration with Firecracker make it a clear-cut and agile platform to be able to deploy containerized applications on cloud resources at the edge, utilizing bare-metal infrastructure providers. As a way to illustrate how to use OpenNebula to bring an IoT application to the edge, we show how to deploy a ThingsBoard IoT framework in a central location and then a ThingsBoard IoT Gateway with MQTT brokers on the edge.

OpenNebula as a Platform for IoT Applications

Screencast Video:	youtube.com/watch?v=JbNzwXz0xHc
Guide:	support.opennebula.pro/hc/en-us/articles/360045122532

Video Streaming

While the over-the-top (OTT) media market is revolutionizing the entertainment industry, with live events being broadcast directly over the internet and HTTP streaming formats gaining wide acceptance, the huge technical limitations and challenges that accompany this new broadcasting method cannot be ignored. Live events are being broadcast directly over the internet, allowing for immediate and widespread participation. The technology infrastructure that a broadcasting service company chooses influences its degree of success in the increasingly competitive OTT video market.

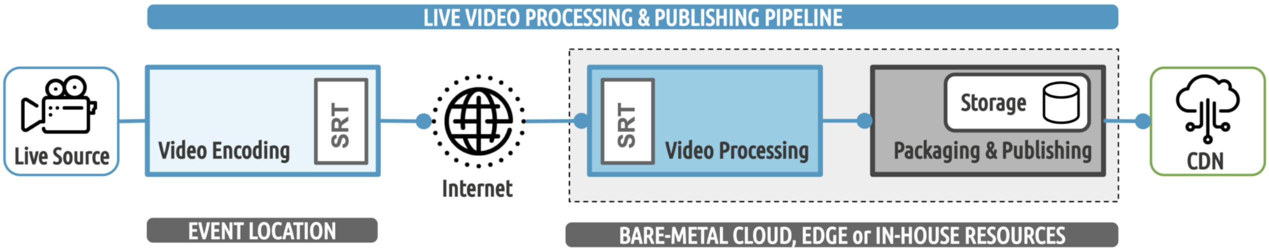

Figure 5. Video Streaming Architecture.

OpenNebula allows companies to build on-demand, elastic private clouds using resources from pay-per-use bare-metal cloud providers to run high-definition video encoding and live video streaming services. This innovative approach offers the agility, scalability and simplicity of the public cloud, with the greater levels of flexibility, performance, and security of a single-tenant, dedicated environment. We show below how to build an elastic media service to run live and on-demand video workflows based on packaging and encoding workflows.

OpenNebula Elastic Private Media Cloud

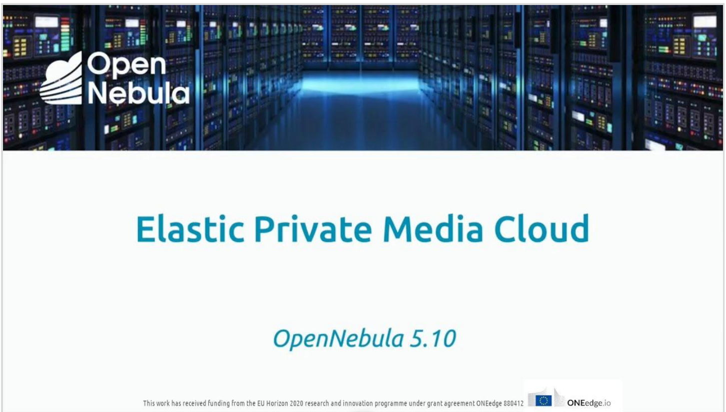

Screencast Video:	youtube.com/watch?v=70K-DJvhSiE
Guide:	opennebula.io/an-elastic-private-cloud-opennebulas-solution-for-media-services

Virtual Desktop Infrastructure

Virtual desktops are seeing broad adoption across all industries, and that adoption is predicted to increase significantly over the coming years with the consolidation of WFH models. The use cases for desktop virtualization are innumerable, and so just about every industry potentially stands to benefit from the flexibility provided by VDI deployments. The Edge Computing features provided by OpenNebula allow for the easy setup of a fully open source, pay-as-you-go Virtual Desktop Infrastructure (VDI) at the edge using on-demand bare-metal resources. A decentralized Edge Cloud model is especially suitable for companies that need to extend the VDI paradigm in order to be compliant with local data protection regulations (like the European GDPR) or whose employees require low-latency remote access to corporate applications.

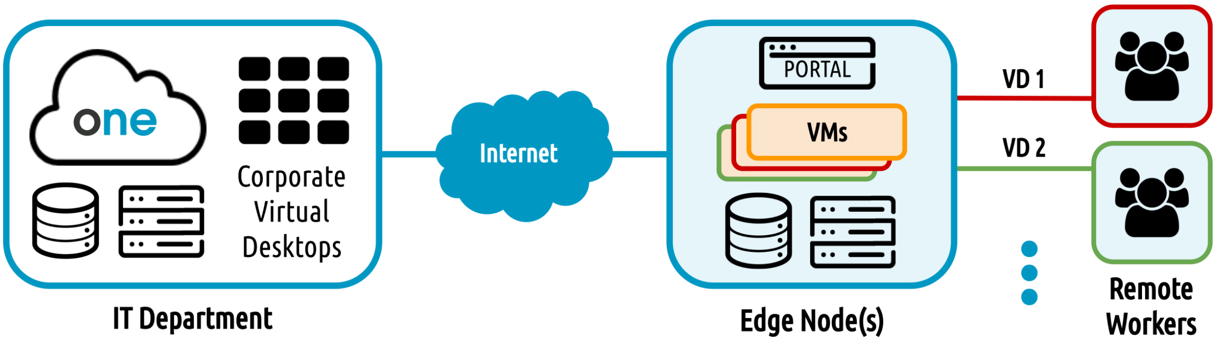

Figure 6. Virtual Desktop Infrastructure Architecture.

The resource elasticity and automatic provisioning capabilities at the edge provided by OpenNebula dramatically reduce the infrastructure burden of desktop management. Below, we show how OpenNebula can be used to set up a stable, easy-to-deploy open source VDI solution at the edge, as an alternative to complicated VDI legacy stacks or expensive cloud solutions. The use of on-demand bare-metal resources at the edge offers a powerful mechanism for companies to make sure that they observe local data protection regulations while, at the same time, improving the end-user experience and reducing costs associated with cloud infrastructures.

Virtual Desktop Infrastructure (VDI) at the Edge with OpenNebula	

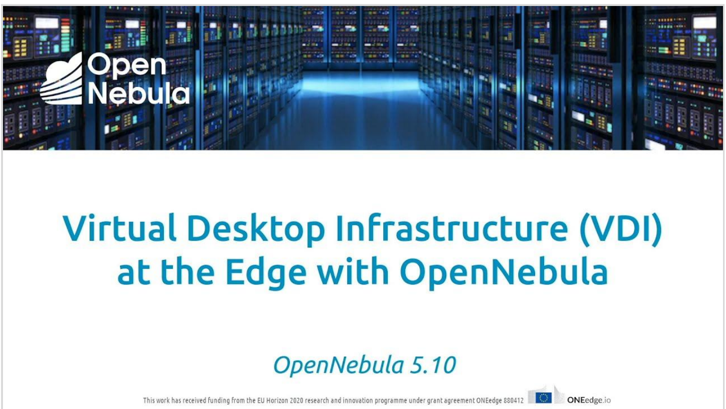

Screencast Video:	youtube.com/watch?v=oR8bjx-nEqc
Guide:	opennebula.io/opennebula-for-vdi-at-the-edge

Telecom Multi-Access Edge Computing

Telco companies are starting to develop their own solutions to address the deficiencies that current centralized cloud services suffer. These include security, network cost and, above all, latency. Multi-Access Edge Computing (MEC) is a telecommunications network architecture that involves the placement of cloud and IT resources in data centers spread across a telecom's operator network. A Central Office Re-architected as a Datacenter (CORD)[6] scenario typically integrates network function virtualization (NFV) and software-defined networking (SDN) in order to reduce costs, bring cloud agility to the Telco Central Office, and refine control to the network.

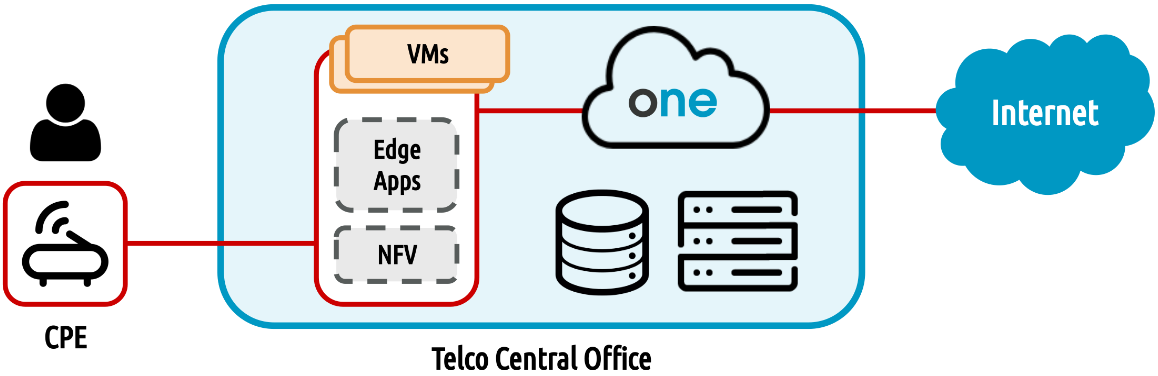

Figure 7. Multi-Access Edge Computing Architecture.

CORD's reference architecture is based on commodity hardware and a virtualization management platform such as OpenNebula to create and control the virtualized network functions. With MEC, telcos can use OpenNebula to make a new set of compute, storage and network capabilities available to customers at the edge of communications networks. In this way, services and applications are brought significantly closer to end customers, improving the user experience and enabling new applications and offers.

Telecom Multi-Access Edge Computing with OpenNebula

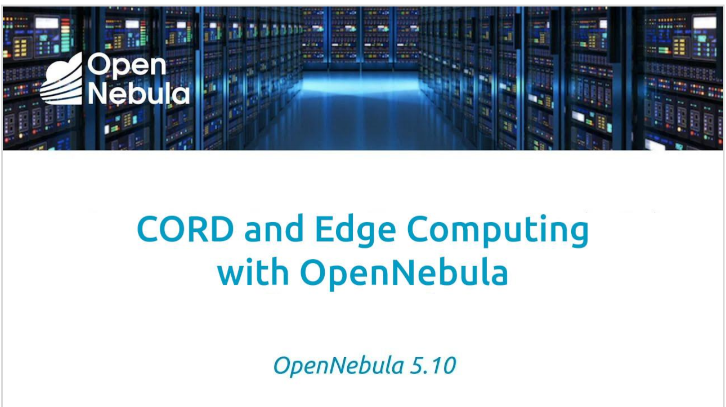

Screencast Video:	youtube.com/watch?v=Hx9EqB6RGOk
Guide:	opennebula.io/whitepaper-telefonica-edge-computing

[6] https://www.opennetworking.org/cord/

4. Edge Cloud Architectures

There are two different approaches for building Edge Computing environments with OpenNebula: the Centralized Control Plane and Distributed Control Plane models.

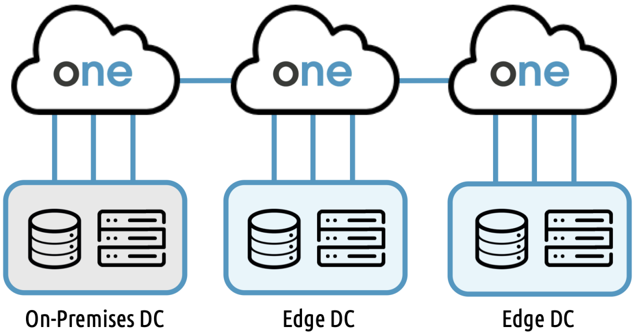

Figure 8. OpenNebula Distributed Control Plane Model.

The **Distributed Control Plane** model uses Datacenter Federation functionality offered by OpenNebula. Each OpenNebula instance of the federation is called a "zone", one of them configured as primary, on a centralized location, and the others as secondaries, on edge locations. This is a tightly coupled integration, with all the instances sharing the same user accounts, groups and permission configurations. Federation allows end-users to consume resources allocated by the federation administrators regardless of their geographic location. The integration is seamless, meaning that a user logged into the Sunstone GUI of a zone will not have to log out and enter the address of another zone. Sunstone allows users to change the active zone at any time, and it will automatically redirect the requests to the right OpenNebula instance.

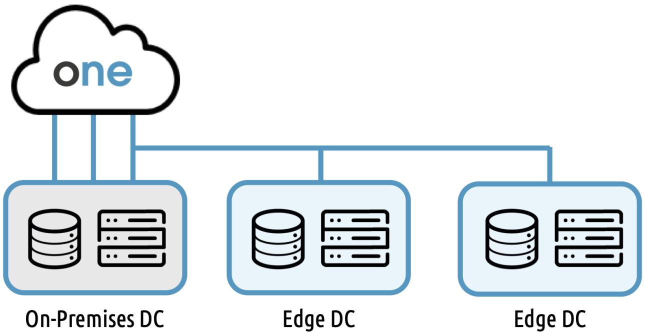

Figure 9. OpenNebula Centralized Control Plane Model.

The **Centralized Control Plane** model is built as a traditional single data center environment which is geographically distributed with WAN connections between the controller and compute nodes. This approach faces very similar challenges to operating large-scale distributed data centers. Due to the constraints of this model, the edge nodes rely heavily on the centralized data center to carry the burden of management and orchestration of the edge compute, storage and networking services because they run all the controller functions.

The Distributed model allows edge sites to be more fully functional on their own, being in principle more resilient to network connectivity issues as well as being able to minimize disruption caused by latency between edge sites. However, they have a much greater deployment and management complexity, and there is an orchestrational overhead to synchronize between these data centers and manage them both individually and as part of a larger, connected environment simultaneously. Our experience is that the

Distributed Control model is mostly used by big companies that own large/medium edge data centers, or by those interested in creating a single world-scale public cloud along the edge.

In this white paper we focus on the Centralized Control model where companies can extend private cloud orchestration capabilities to resources at the edge and manage their geo-distributed edge cloud on-demand. It provides a centralized view of the infrastructure as a whole, which has its advantages from an operational perspective. In this case, a cloud management platform manages both the physical resources of the on-premises cloud data center, if any, along with the resources of multiple, highly-dispersed edge nodes, providing users with a uniform access to this disaggregated resource pool. The new OpenNebula monitoring system allows applications to continue running on edge resources even in case of network outages if the use case requires the workload to be highly available. The OpenNebula Centralized Control model can manage small to medium Edge Computing infrastructures comprising several tens to hundreds of geographically distributed edge resources.

5. The OpenNebula Approach

OpenNebula follows these principles:

- Reuse proven **cloud management capabilities** and **open source** software components.
- Enable **true multi-cloud interoperability** by providing a uniform view of the underlying resources from different cloud providers. An application can be deployed anywhere on the edge infrastructure without performing any additional configuration or setup.
- Leverage the existing ecosystem of hyperscale and edge clouds.
- Where other solutions require an expensive and laborious process to set up and configure an edge node, as if it were another data center, ONEedge is completely **software-defined, programmatic and delivered 'as a service'**.
- Even better, the physical resources that comprise these edge cloud solutions do not need to be owned or provisioned by the edge cloud owner/administrator. ONEedge takes an **opportunistic approach to its edge node provisioning model** by taking advantage of already-available resources offered by third-party bare-metal providers.
- Simplicity of use is met by providing **access to existing marketplaces with pre-built edge applications** such as PaaS environments, serverless platforms, container orchestration, etc.

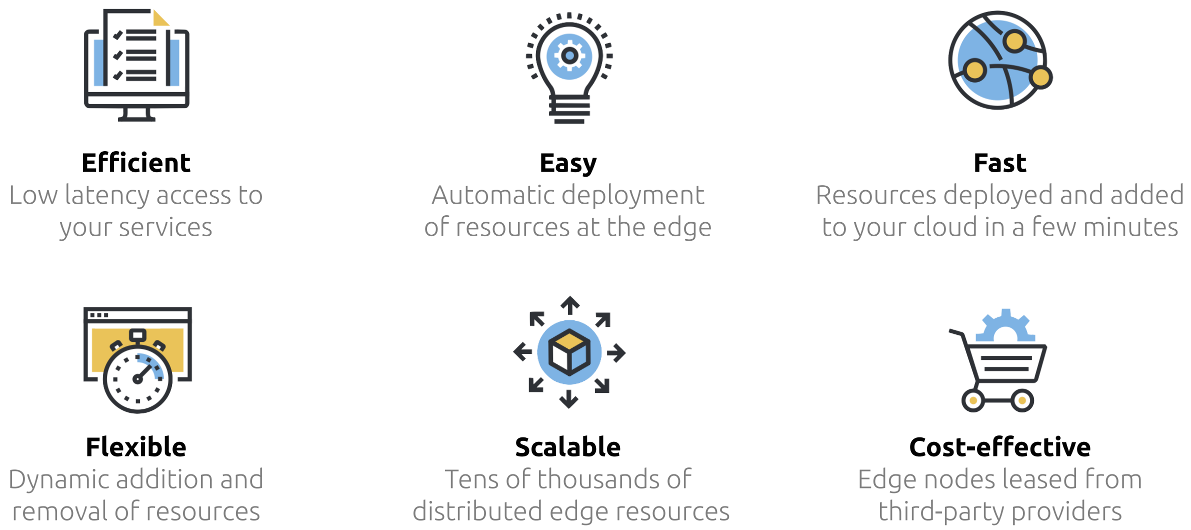

Figure 10. Benefits of the OpenNebula approach to Edge Computing.

6. Technology Considerations

As can be seen from the above use cases, there are functionalities that become even more crucial in edge environments. The focus is largely on reducing latency and mitigating bandwidth limitations. Further similarity between the different use cases, regardless of the industry they come from, is in the following features:

- **Single-Root Input/Output Virtualization (SR-IOV)**: This technology allows VMs and containers to share PCI hardware resources.
- **Data Plane Development Kit (DPDK)**: DPDK allows higher network packet throughput by using offloading and scheduling techniques.
- **Non-Uniform Memory Access (NUMA)**: This method increases throughput by allocating dedicated memory blocks to an instance.
- **CPU Pinning**: For further optimization, processor affinity enables the binding and unbinding of an instance to a central processing unit (CPU) or a range of CPUs, so that the instance will execute only on the designated CPU close to the allocated memory.
- **Hugepages**: This is a mechanism that allows the Linux kernel to utilize the multiple page size capabilities of modern hardware architectures.
- **Graphics Processing Unit (GPU)**: GPUs have a high number of cores which may be utilized by a wide variety of parallel-processing intensive workloads in the edge use cases, such as IoT and gaming.
- **Serverless Hypervisors**: MicroVMs offers a complete isolated environment with a reduced virtual hardware and feature set. This hypervisor is of special relevance for containers, and serverless and function-as-a-service scenarios.
- **Secure and Scalable Distributed Monitoring**: High distributed edge systems require a secure, reliable and scalable communication in the control plane at both the transport and application layers. They also require a mechanism to easily subscribe third party apps to alarm and warning states.
- **3-Tier Replica Storage with Cache Datastore**: This is a storage model that extends the existing OpenNebula datastore model so images can be cached to local storage areas in the edge nodes.
- **Container Workflows**: Containerized application workflows that include multiple microVMs with deployment dependencies and auto-scaling rules are easily defined, shared, instantiated, and executed with high scalability and low response time.
- **Application Marketplaces**: Integration with existing marketplaces, like Docker Hub, for backward compatibility with existing VM & container appliances.
- **Automatic Provisioning**: Automatic deployment of Edge Clusters in the remote cloud and edge data centers.
- **Kubernetes at the Edge**: Operational and security features for managing multiple Kubernetes clusters across cloud and edge locations.

7. Automatic Provision of Edge Clusters

The physical resources that comprise OpenNebula edge cloud solutions do not need to be owned or provisioned by the edge cloud owner/administrator. OpenNebula takes an opportunistic approach to its edge node provisioning model by taking advantage of already-available resources offered by third-party bare-metal providers. New Edge Clusters can be **leased on-demand from existing bare-metal cloud providers** and automatically installed and configured with the software stack needed to enroll them in the private edge cloud instance.

OpenNebula brings the provisioning tools and methods 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 edge-like cloud environments. Companies are able to grow their private cloud with resources at edge data center locations to meet 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.

The OneProvision[7] tool allows the deployment of a fully operational OpenNebula cluster in a remote provider and the management of its full life-cycle, starting with their provision and maintenance, until the unprovision. Each edge location (the "**provision**") is defined as a group of physical hosts allocated on 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. Except for the physical hosts, each provision comes with dedicated virtual networks and datastores. Every single edge location is an independent and complete computing environment.

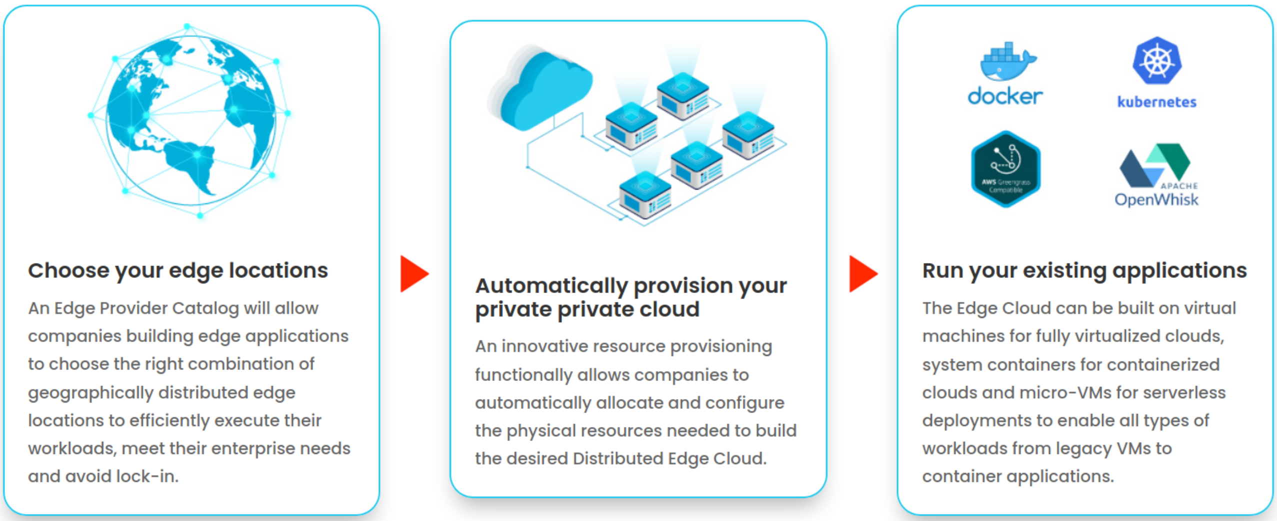

8. True Hybrid and Edge Cloud

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 for their edge applications in terms of cost, capacity, latency, bandwidth, etc. Although the automatic provision feature in our latest version (OpenNebula 6.0 "Mutara") only currently supports **AWS** and **Equinix Metal**, OpenNebula is developing drivers for most widely-used cloud and edge providers. Moreover, we are building the tools and processes to allow any other provider to develop the necessary drivers and to join this catalog.

The Edge Computing field is growing at an accelerated pace. Hence, new edge providers are publishing their offerings continuously. We have identified the following types of edge providers that can be included in the Resource Provider Catalog, based on the latency they are able to deliver to end-users.

[7] http://docs.opennebula.io/stable/advanced_components/ddc/index.html

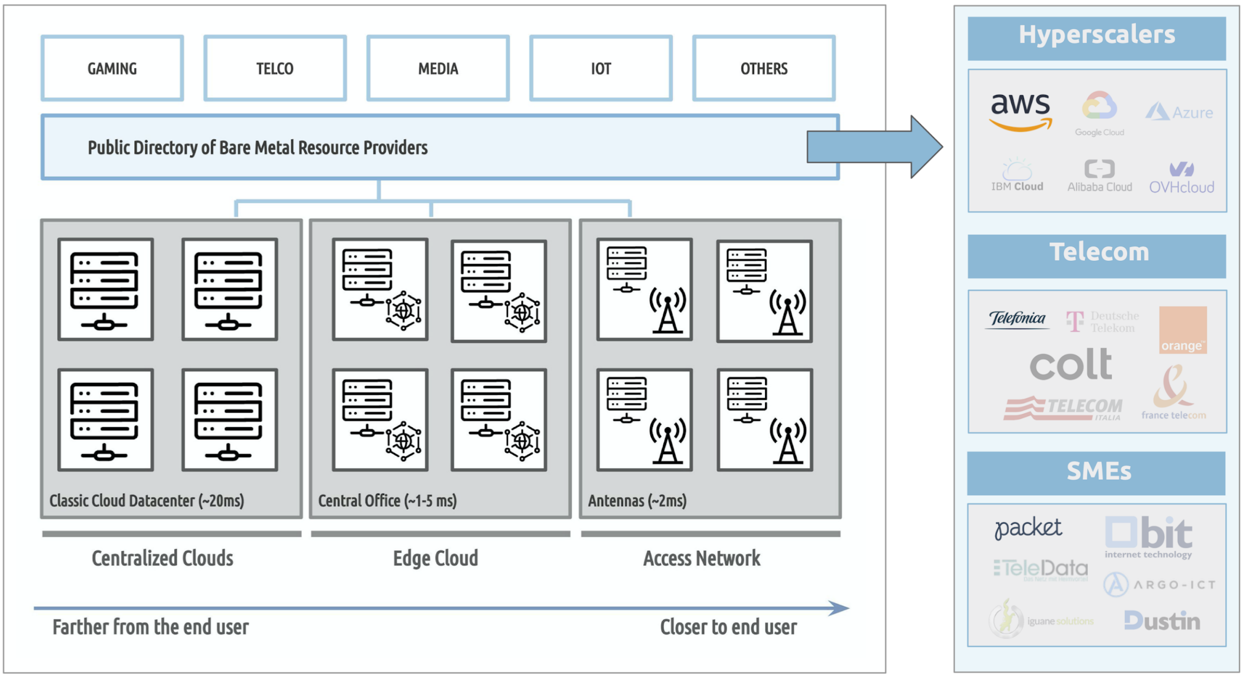

Figure 11. Different types of Edge Resource Providers by proximity to their end-users.

Centralized Cloud Providers

This offering includes public cloud providers with a bare-metal offering. As opposed to their classic cloud provider offering (i.e. the ability to manage Virtual Machines), the bare-metal offering opens the possibility to define different edge architectures. For OpenNebula, this means being able to offer a common edge architecture across different public cloud providers. They can typically offer a latency for edge services above 20 ms.

Examples of centralized public cloud providers with bare-metal offerings include classic public cloud providers like AWS, Google Cloud Engine, IBM Cloud; as well as smaller infrastructure providers like Equinix, Linode, Scaleaway, OVHcloud, and many others.

Edge Cloud Providers

This type encompasses any company with real estate within a population, housing servers that can be consumed through an API. The greater proximity to the end-user (compared with a centralized cloud) enables a lower latency for the edge application, in the order of 1-5 ms.

Examples of edge cloud providers are traditional Telecommunication companies through the CORD model (Central Office Re-architected as a Datacenter) like, for instance, Telefónica, AT&T, Deutsche Telekom, and many others. Also, this category includes infrastructure providers with presence in different population centers, such as, for instance, well-established companies like Equinix as well as members of the Kinetic Edge Alliance; and new entries in the cloud provider arena, including Walmart.

Access Network Providers

This category aggregates providers with specialized antennas able to house computing resources. This renders the closest computation to the end-user, and can deliver latencies in the order of 2ms.

An example is AWS Wavelength, which provides specific infrastructure deployments that embed AWS compute and storage services right within the data centers at the edge of the growing 5G network of communications service providers (CSP) like Verizon. Some members of the Kinetic Edge Alliance, such as Vapor IO, also fit into this category, as well as specialized companies like Smarter City Technology.

9. 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.

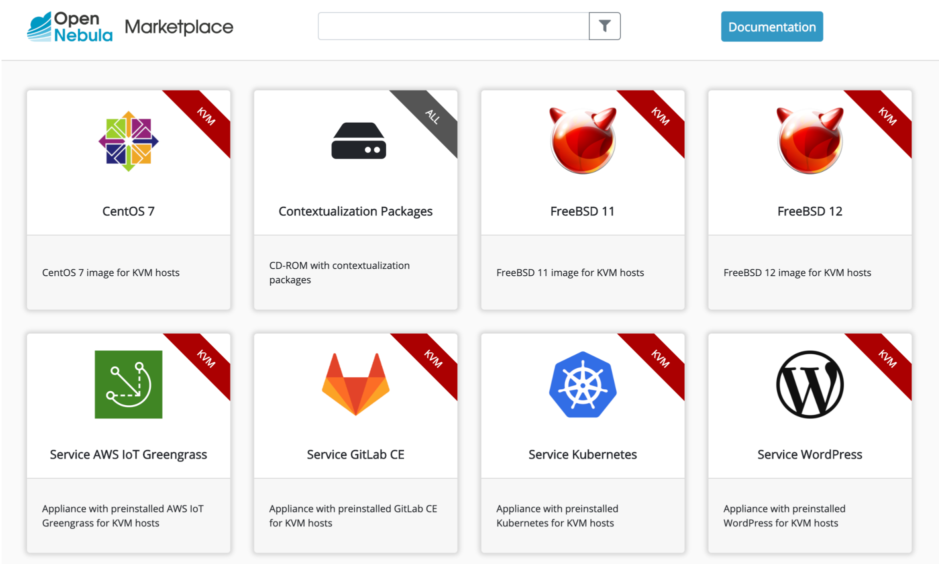

Figure 12. Some of the virtual appliances available in the OpenNebula Public Marketplace.

OpenNebula users can easily download, contextualize, and add virtual appliances from other public marketplaces, including **Linux Containers**[8] and **TurnKey Linux**.[9]

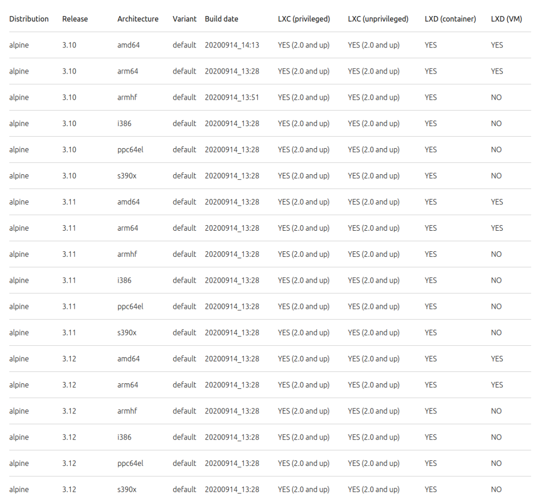

Distribution	Release	Architecture	Variant	Build date	LXC (privileged)	LXC (unprivileged)	LXD (container)	LXD (VM)
alpine	3.10	amd64	default	20200914_14:13	YES (2.0 and up)	YES (2.0 and up)	YES	YES
alpine	3.10	arm64	default	20200914_13:28	YES (2.0 and up)	YES (2.0 and up)	YES	YES
alpine	3.10	armhf	default	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

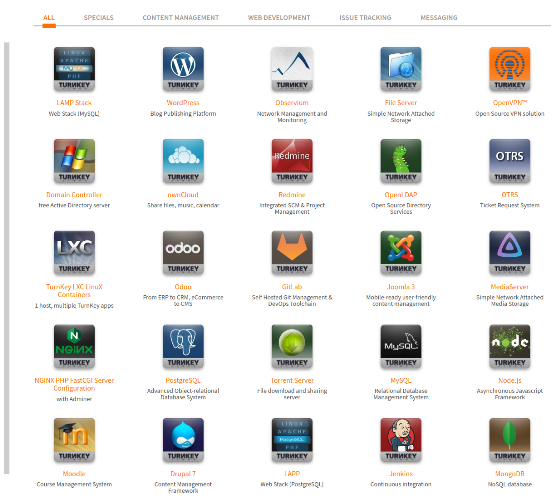

Figure 13. Some of the images provided by Linux Containers (left) and TurnKey Linux (right).

[8] https://images.linuxcontainers.org
[9] 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**,[10] 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.

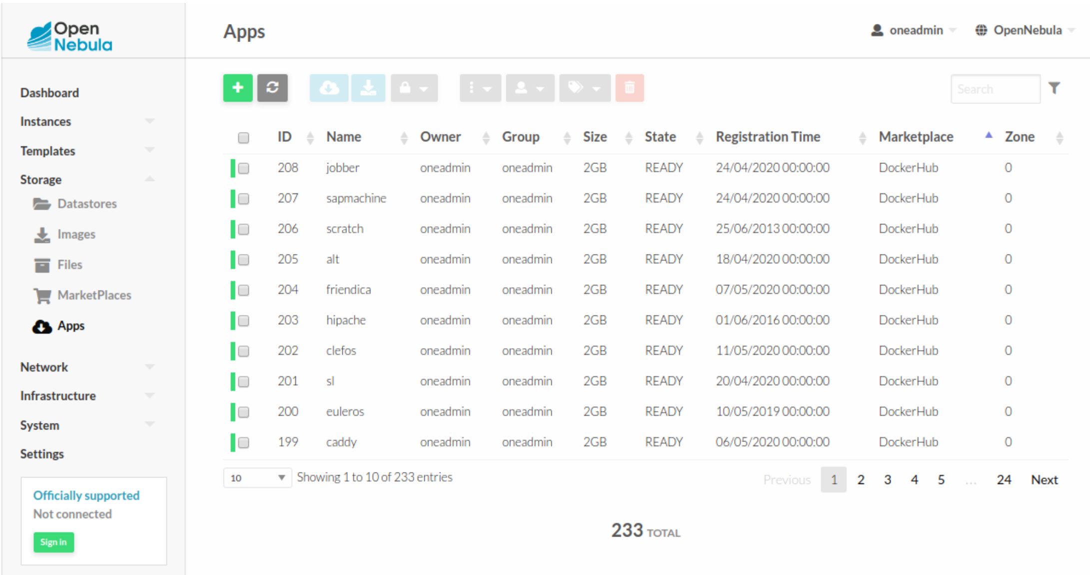

Figure 14. OpenNebula's native integration with the Docker Hub marketplace.

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 integrate 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 implies the deployment of a Kubernetes controller as a **managed service**, offering the possibility to add more worker nodes to the Kubernetes clusters on demand.

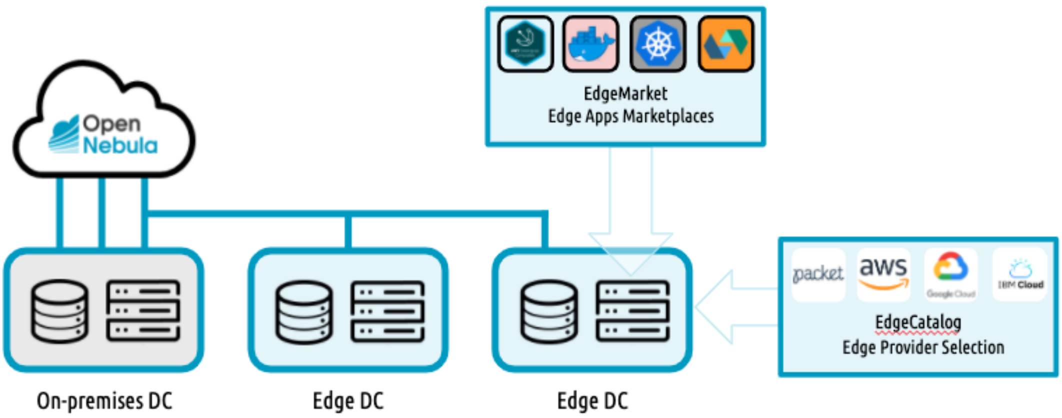

Figure 15. OpenNebula's general edge application deployment model applicable to Kubernetes.

[10] https://hub.docker.com

10. Ready for a Test Drive?

You can evaluate OpenNebula and build a cloud in just a few minutes by using **miniONE**,[11] our deployment tool for quick installation of an OpenNebula Front-end inside a Virtual Machine or a physical host, and then adding a remote Edge Cluster based on KVM, LXC or Firecracker.

https://opennebula.io/mutara/

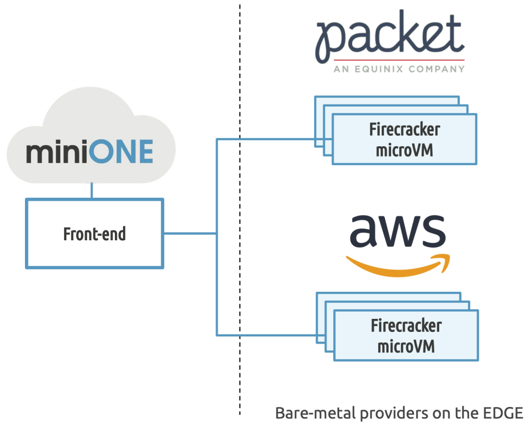

11. Conclusions

This document has been created within the context of the **ONEedge** project, OpenNebula's initiative for making Edge Computing available for companies that want to build their own private environments based on highly-dispersed edge nodes in close proximity to users, machines, and sources of data. This white paper has described the principles that the ONEedge project is following and the main architectural elements of the solution that OpenNebula provides. This paper shows just the beginning of a process that leads to the development of a truly open source Edge Computing platform, available to everyone.

This document includes a number of use cases and practical resources that demonstrate how Edge Computing applies to different industries and the benefits it delivers, and the solution OpenNebula offers to companies interested in increasing their competitiveness by implementing their own edge deployments. These use cases should be able to be reproduced quite simply, but if you require any assistance in deploying any of these edge scenarios at scale, don't hesitate to reach out to our OpenNebula Services Team for consultation and support.

[11] 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

LINUX FOUNDATION SILVER MEMBER · CLOUD NATIVE COMPUTING FOUNDATION SILVER MEMBER · LF EDGE GENERAL MEMBER · GAIA-X data-infrastructure.eu

Rev2.2_20210415