ONEedge.io

A Software-defined Edge Computing Solution

D4.6. Deployment of Validation Cases and Demonstrations - c

Deployment of Validation Cases and Demonstrations Report

Version 1.0

3 November 2021

Abstract

This report, delivered at the end of the Third Innovation Cycle (M17-M23), summarizes the work done to generate actual demonstrations of the project's capabilities in real life situations and to deploy validation cases for different scenarios.

This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No 880412.

Deliverable Metadata

Project Title:	A Software-defined Edge Computing Solution
Project Acronym:	ONEedge
Call:	H2020-SMEInst-2018-2020-2
Grant Agreement:	880412
WP number and Title:	WP4. Demo and Operational Infrastructure
Nature:	R: Report
Dissemination Level:	PU: Public
Version:	1.0
Contractual Date of Delivery:	30/9/2021
Actual Date of Delivery:	3/11/2021
Lead Authors:	Michael Abdou, Vlastimil Holer and Alberto P. Martí
Authors:	Sergio Betanzos, Ricardo Díaz, Jim Freeman, Christian González, Alejandro Huertas, Shivang Kapoor, Jorge M. Lobo, Jan Orel and Petr Ospaly
Status:	Submitted

Document History

Version	Issue Date	Status[1]	Content and changes
1.0	3/11/2021	Submitted	First final version of the D4.6 report

[1] A deliverable can be in one of these stages: Draft, Peer-Reviewed, Submitted and Approved.

Executive Summary

A lot of the work carried out within the last several months has been focused not only on developing a sophisticated Edge Computing software solution to innovatively address the need for lower latency and provide a simple, distributed infrastructure management, but also to generate actual demonstrations of its capabilities in real life situations (Section 2), as well as to get some of these capabilities in the hands of users to put them to the test on their own (Section 1).

One of our goals is to reduce to a minimum the requirements to try out ONEedge, specifically for the resources needed to set up a Proof-of-Concept (PoC) environment. Considering the nature of the ONEedge product, which is able to deploy a provision in a variety of Cloud/Edge providers in minutes, we mostly focused on streamlining the experience for the ONEedge Frontend. We have designed a set of Try-out Guides and Tools (Section 1.1) that only requires a Cloud account (e.g. AWS or Equinix) to set up a fully functional PoC that illustrates how to deploy the highly distributed edge and run VMs, containers, and Kubernetes clusters at edge locations.

While it is easy for people interested in learning about OpenNebula to download the open source software, deploy it into their own environment, and begin testing out how to create and manage an OpenNebula cloud, we have dedicated a good amount of time and effort to accommodate the need to see it in action by an OpenNebula expert. Hence, we created and deployed an On-line Demo Request System (Section 1.2) where individuals could request and schedule a demo to be carried out by an OpenNebula Systems engineer, and provided them with a forum to ask questions and inquire about how it could meet their needs.

As this is an open source project, a great deal of our work managing the project is not simply producing code, deliverables, and informational material, but also coordinating the creation of these items and publishing them to the broader public. Since a good portion of the content material is publicly available, we decided to include easily referenced links to the pertinent material, as opposed to recopying all of the text.

Table of Contents

1. Support for Demonstrations **5**
1.1. Try-out ONEedge Edge Cloud Features in 5 Minutes 5
1.1.1 Frontend Setup 5
1.1.2 Provisioning an Edge Cluster 5
1.1.3 Running apps on the Edge Cluster 6
1.2. On-line Demo Request System 7
1.3. Guided Demonstrations to New Users 9
2. Demonstrations **10**
2.1. Deploying Kubernetes at the Edge 10
2.2. Using OpenNebula and K3s at the Edge 11
2.3. Managing Kubernetes at the Edge 12
2.4. Multi-Cloud Capabilities with Google Cloud Platform 13
2.5. On-Demand Provisioning 14
2.6. Side-Stepping Vendor Lock-in 14
2.7. Multi-Cloud Portability 15

1. Support for Demonstrations

During this cycle, we designed a set of Try-out Guides and Tools (Section 1.1), which provides corporate users with a fully functional PoC, and an On-line Demo Request System (Section 1.2), which provides corporate users with a guided demonstration by an OpenNebula expert.

1.1. Try-out ONEedge Edge Cloud Features in 5 Minutes

One of our goals is to reduce to a minimum the requirements to try out ONEedge, specifically the resources needed to set up a *Proof-of-Concept* (PoC) environment. Considering the nature of the ONEedge product that is able to deploy a provision in a variety of Cloud/Edge providers in minutes, we only need to streamline the experience for the ONEedge Frontend.

We have designed a set of *try-out* guides and tools that only requires a Cloud account (e.g. AWS or Equinix) to set up a fully functional PoC. Alternatively, you can use a local server or virtual machine with low requirements (4Gb of RAM and 20Gb of disk space) to set up a server.

1.1.1 Frontend Setup

The first step is to set up a Front-end using the minione[2] automatic installation tool. Minione can be used either in a physical or virtual server to ease the initial installation process. The whole process takes just a few minutes and requires the execution of a single line. When the installation process finishes the tool outputs the initial credentials and connection endpoints:

```
$> wget
'https://github.com/OpenNebula/minione/releases/latest/download/minione';
sudo bash minione --frontend
...
### Report
OpenNebula 6.0 was installed
Sunstone is running on:
  http://3.121.76.103/
FireEdge is running on:
  http://3.121.76.103:2616/
Use following to login:
  user: oneadmin
  password: lCmPUb5Gwk
```

Figure 1.1.1: One liner to set up a fully functional Frontend

1.1.2 Provisioning an Edge Cluster

Once the Frontend is ready you can allocate an edge cluster in one of the supported providers. In order to reduce possible options and problems, the quick start guide focuses on AWS. The steps needed to deploy an edge cluster are:

- Gather your **AWS account information**, mainly the access keys and secrets.
- Create an **AWS provider in Fireedge** for the AWS region you are interested in.
- **Provision a Virtual Edge cluster**. Bare metal instances are expensive and may prevent some users from trying the software. The quick start guide recommends the creation of a provision based on regular AWS instances. The virtualization technology used in this case is LXC.

At the end of the process an edge cluster is ready.

[2] https://github.com/OpenNebula/minione

1.1.3 Running apps on the Edge Cluster

After guiding the user in the process of installing a ONEedge Frontend and deploying the first edge cluster, users can follow three tutorials to deploy relevant workloads, namely:

- Single instance applications[3]
- Kubernetes (or K3s) clusters[4]

Single instance applications

The guide details the process of creating a WordPress server in the edge. The WordPress application is downloaded directly from the OpenNebula marketplace. This appliance is already prepared to be fully integrated with OpenNebula, so the full process just requires a couple of clicks in the Web UI.

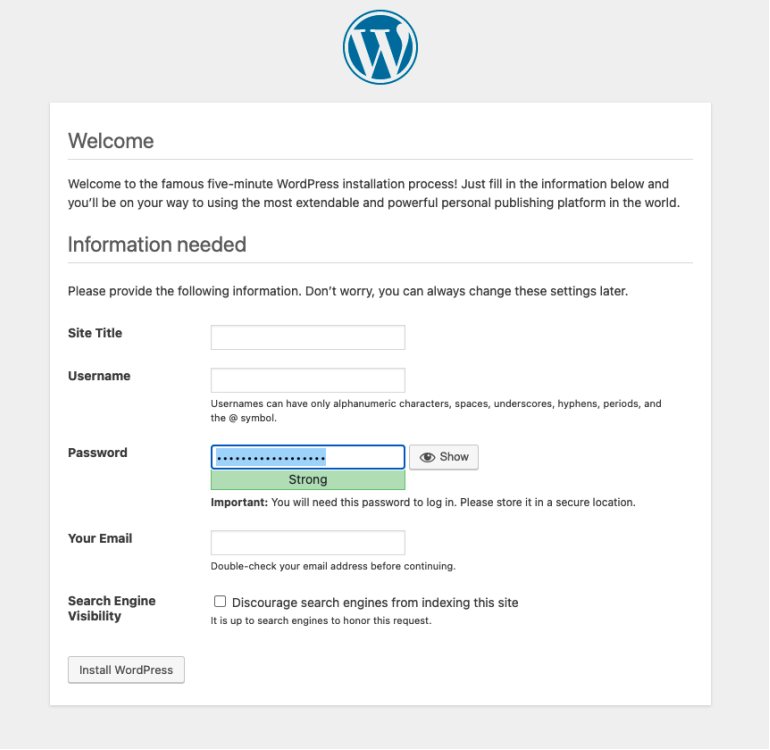

Figure 1.1.2: Welcome screen of the WordPress appliance after is created in the edge

Kubernetes Clusters

A more advanced workload is a full kubernetes cluster. The kunerketes cluster is also available in the public OpenNebula marketplace ready to use. In this case the K8s cluster is defined as a OneFlow, i.e. a composition of virtual machines.

The OneFlow template deploys a K8s Frontend and several workers, it also sets up the networking for the k8s cluster and offers multiple configuration options like the installation of a load balancer or the network Backend. The following figure shows the connection to a K8s cluster deployed at the edge:

```
[root@onekube-ip-10-0-17-190 ~]# kubectl get nodes
```

[3] https://docs.opennebula.io/stable/quick_start/usage_basics/running_virtual_machines.html
[4] https://docs.opennebula.io/stable/quick_start/usage_basics/running_kubernetes_clusters.html

```
NAME                                    STATUS   ROLES    AGE   VERSION
onekube-ip-10-0-109-134.localdomain     Ready    <none>   27m   v1.18.10
onekube-ip-10-0-17-190.localdomain      Ready    master   29m   v1.18.10
```

Figure 1.1.3: Connection to a K8s cluster deployed in ONEEdge cluster

To verify the K8s application the guide ends with the deployment of a web server (nginx) in the K8s cluster. In this case we use a NodePort network mode for the web server:

```
[root@onekube-ip-10-0-17-190 ~]# kubectl expose pod nginx --type=NodePort
--name=nginx
```

```
[root@onekube-ip-10-0-17-190 ~]# kubectl get svc
NAME         TYPE        CLUSTER-IP     EXTERNAL-IP   PORT(S)        AGE
kubernetes   ClusterIP   10.96.0.1      <none>        443/TCP        30m
nginx        NodePort    10.104.44.89   <none>        80:30317/TCP   13s
```

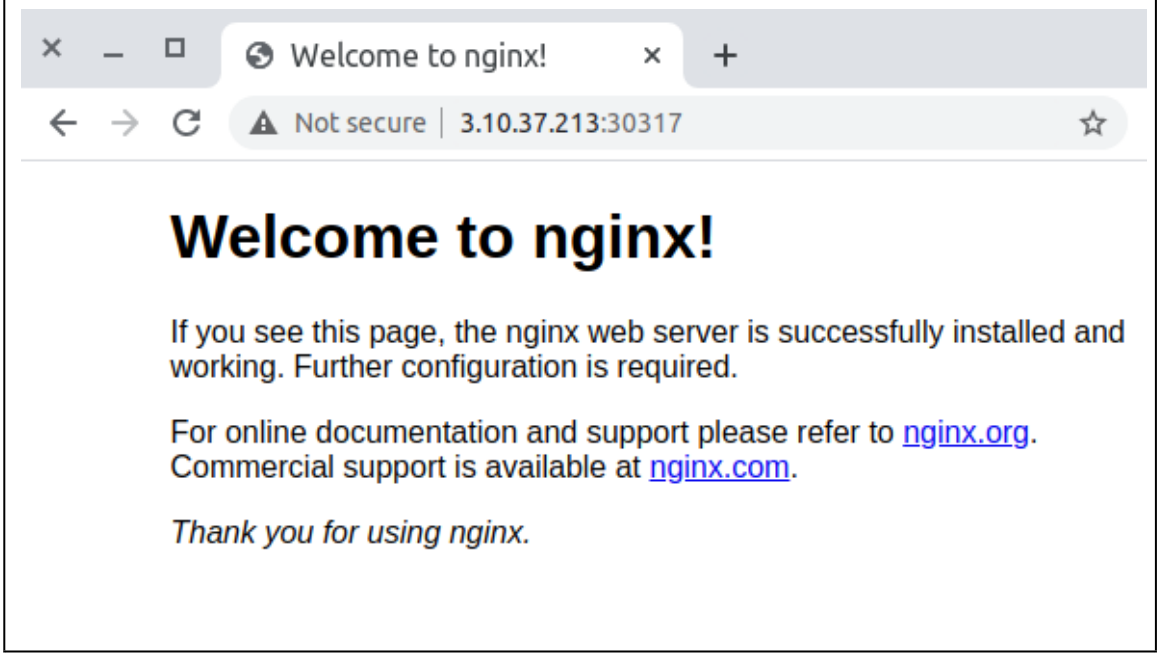

Figure 1.1.4: Deployment of an nginx web server and a sample connection to it

1.2. On-line Demo Request System

While it is easy for people interested in learning about OpenNebula to download the open source software, deploy it into their own environment, and begin testing out how to create and manage an OpenNebula cloud, we have dedicated a good amount of time and effort to accommodating the need to see it in action by an OpenNebula expert. Hence, we created and deployed an On-line Demo Request System, where individuals could request and schedule a demo to be carried out by an OpenNebula Systems engineer, and provided them with a forum to ask questions and inquire about how it could meet their needs.

Workflow

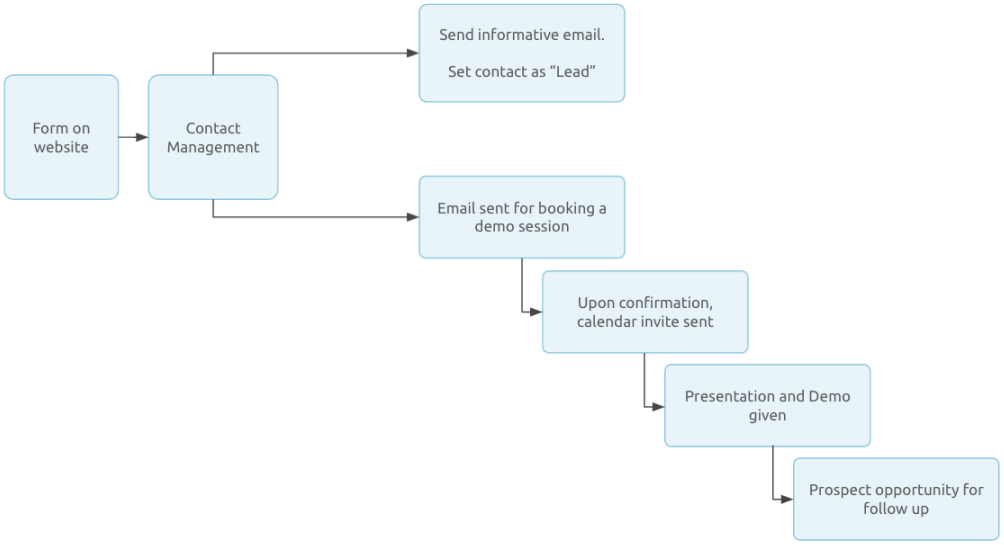

Figure 1.2.1: Demo system request workflow

We developed a simple workflow which allows users to request a demo session from a form on our website. Upon confirming their interest, we are able to capture contact information for further lead and prospect management, and subsequently schedule a presentation and demo session given by an OpenNebula Systems engineer.

Demo Script

A demo script has been prepared to provide an outline for the key features and functionalities to be reviewed during the demo. It is a tool that allows the engineering team to have a general template to ensure that key topics are reviewed and demoed, and provides a consistent structure to each of the demos provided.

Demo Slides

We have a prepared presentation to accompany the demo session for prospective users, highlighting the key features of OpenNebula developed in ONEedge and showcasing the benefits that are offered by the new upcoming "Edge as a Service" model.

1.3. Guided Demonstrations to New Users

Over the course of the five-month period in which this demo system (Section 1.2) has been in place, we have successfully given an average of **8 demos per month**, all with a focus on prospective customers interested in the edge features of the software.

These have been carried out as a supplement to the presentations and demonstrations that we have organized with a number of organizations and relevant communities, derived from our communication and dissemination efforts during these months. Some of them, like the ones for the **Dutch Cloud Infrastructure Coalition** (an initiative coordinated by TNO) and for the Eclipse Foundation's **Edge Native Working Group**, have led to stronger collaborations in related initiatives such as GAIA-X and the IPCEI in Next-Generation Cloud Infrastructure and Services.

Cloud Infra Coalition / TNO (June 2021)

Eclipse's Edge Native WG (June 2021)

Also, with the improvement in the general situation in many European countries regarding the COVID-19 pandemic, a number of physical events have started to take place across the continent. For us, the first of these events has been the **Madrid Tech Show / Cloud Expo Europe** (October 27-28, 2021), in which a total of 10 live demos were performed in our booth.

Given the level of engagement and the success of our participation in this first physical event of the year, we have decided to promote ONEedge also at the **Cloud Expo Europe Frankfurt** (December 8-9, 2021), where we will have a booth and will give a presentation as part of the multi-cloud session.

2. Demonstrations

This section describes the main actions that we have taken for demonstrating the new features developed in ONEedge. We make sure that all these actions acknowledge the support received from the European Commission, i.e. by incorporating a standard slide on ONEedge:

2.1. Deploying Kubernetes at the Edge

The Edge Computing World Conference, sponsored by Topio Networks, is a premier event for the Edge Computing space, and in March 2021 it came to Europe with a focus on highlighting edge use cases in EMEA. We took the opportunity to review and demonstrate how the new ONEedge features in OpenNebula are able to address the need to deploy and orchestrate Kubernetes clusters at the edge.

Edge Computing World EU – Deploying Kubernetes at the Edge	
Blog Post	https://opennebula.io/edge-computing-world-europe-2021/ (March 4, 2021)

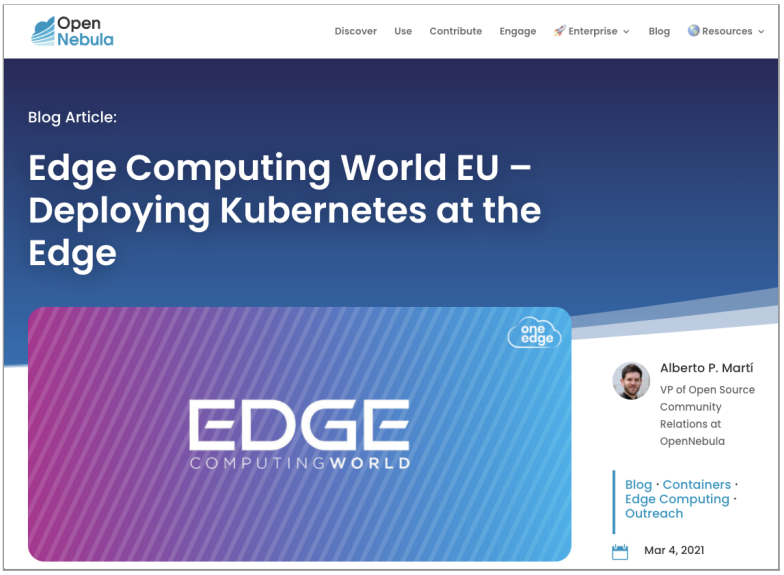

Event Details:	https://www.edgecomputingworld.com/ecwe2021/ (March 9-11, 2021)

2.2. Using OpenNebula and K3s at the Edge

Following the announcement in 2020 that OpenNebula joined the Cloud Native Computing Foundation (CNCF), we took the initiative to schedule and present an on-demand webinar to showcase how OpenNebula's ONEedge features can be used to effectively address the need for low latency and provide the tools for deploying Edge Computing solutions. This webinar focused specifically on demonstrating how the new integration with the CNCF-hosted K3s (a lightweight Kubernetes distribution for the edge) provides a simple platform for easily deploying edge applications like multiplayer games.

CNCF Webinar – K3s at the Edge for Multiplayer Gaming	
Blog Post	https://opennebula.io/cncf-webinar-kubernetes-at-the-edge/ (March 5, 2021)
Demonstration	https://youtu.be/m-3ZqHOO02I (March 11, 2021)

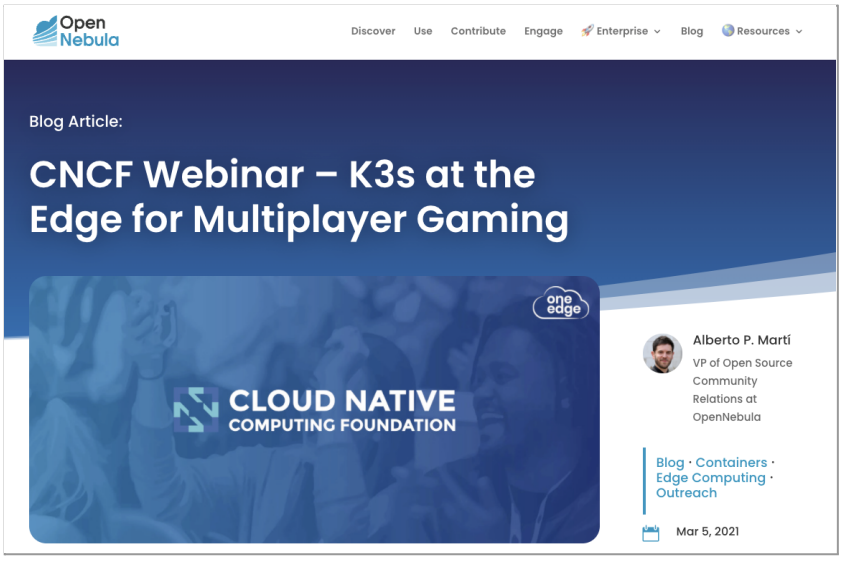

2.3. Managing Kubernetes at the Edge

SUSECON 2021 provided a great opportunity for us to highlight, in collaboration with our Austrian partner FullStackS, the exciting and innovative ONEedge features in OpenNebula to offer a solution for managing large-scale Kubernetes deployments at the edge. This was a coordinated effort in putting OpenNebula to work, highlighting its integration with SUSE's Rancher and K3s, and showing how easy it is to build a solution to manage and orchestrate resources from one's core data center all the way out to the far edge.

Large-scale Kubernetes Deployments on the Edge with OpenNebula, Rancher, and K3s	
Blog Post	https://opennebula.io/susecon-2021-managing-kubernetes-at-the-edge/ (May 31, 2021)
Demonstration	https://events.rainfocus.com/widget/suse/susecon21/catalog21/session/1613736702047001lpf1 (May 19, 2021)

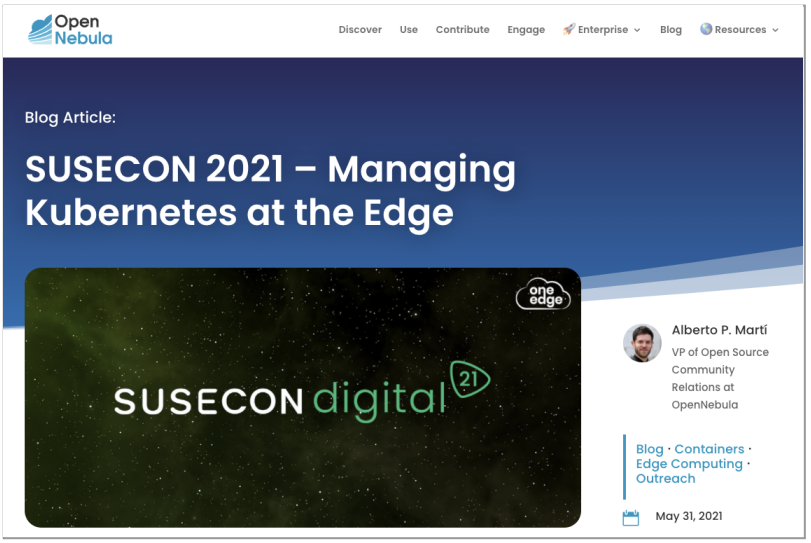

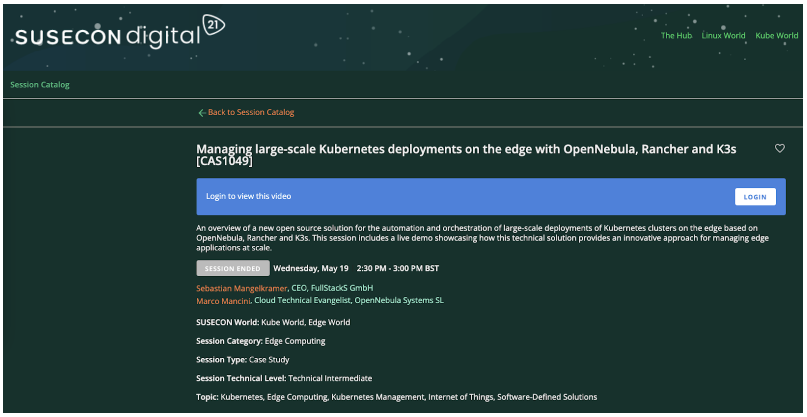

2.4. Multi-Cloud Capabilities with Google Cloud Platform

As we have been working to broaden the catalog of Cloud Providers and Partners, we took the opportunity to demonstrate the ever-growing flexibility of OpenNebula and how its newly developed integration with Google Cloud Platform makes expanding one's enterprise cloud to the edge a simple, yet powerful, capability for all users.

Expanding Your OpenNebula Multi-Cloud with Google Cloud	
Blog Post	https://opennebula.io/expanding-your-opennebula-multi-cloud-with-google-cloud/ (June 7, 2021) 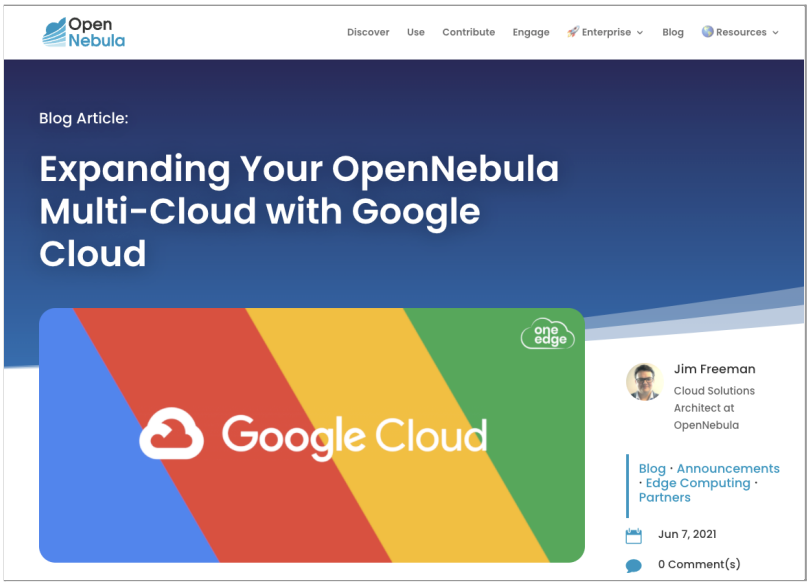
Demonstration	https://youtu.be/a0nNbPGQZ94 (May 30, 2021) 

2.5. On-Demand Provisioning

With the OpenNebula 6.0 release, one of the key enhancements included was the new FireEdge user interface, providing a new dimension to OpenNebula allowing users the flexibility to deploy and provision clusters to cloud and edge resources providers of their choice, in a completely on-demand protocol. We took the opportunity to demonstrate this seamless integration of this new FireEdge UI, and show what it is capable of providing.

Presenting FireEdge: OpenNebula's New Provisioning Tool	
Demonstration	https://youtu.be/gw2ODAcQX1k (April 7, 2021)

2.6. Side-Stepping Vendor Lock-in

With OpenNebula 6.0 and the revolutionary improvements that have been made in broadening users' reach to the Multi-Cloud, we demonstrated, in a screencast video, the ease with which one can build and deploy a fully functional and easily manageable Multi-Cloud solution with OpenNebula, and maintain the flexibility to choose from the catalog of cloud and edge resource providers with complete independence and freedom from vendor lock-in pressures.

Preventing Vendor Lock-in with an OpenNebula Multi-Cloud	
Demonstration	https://youtu.be/0Db21Qf1Oxs (July 5, 2021)

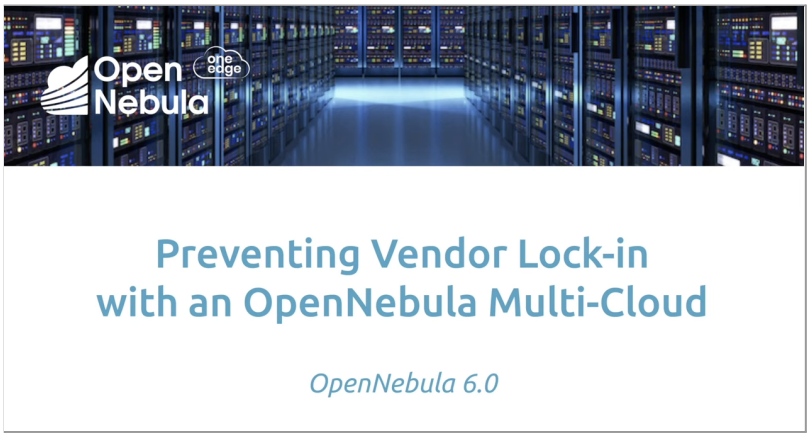

2.7. Multi-Cloud Portability

Not only does OpenNebula provide the freedom to choose between several cloud and edge providers for resourcing an enterprise cloud, but it also provides users with the complete flexibility to “change their mind” and migrate workloads from one provider to another. This feature provides a level of adaptability that other solutions simply do not provide. We took the opportunity to showcase this simple workload migration use case, and to demonstrate just how valuable this “lift and load” capability is.

OpenNebula - Multi-Cloud Portability: Migrating Workloads between Google and AWS	
Demonstration	https://youtu.be/IopA_deQK4M (September 8, 2021)