OpenNebula Enhanced Platform Awareness

Version 1.4 – June 2024

Abstract

Traditionally, network functions such as wireless nodes, Broadband Network Gateways (BNGs), routers or firewalls were delivered in dedicated hardware appliances. Today, Virtual Network Functions (VNFs) are replacing that hardware-centric approach, using software appliances that provide enhanced agility, flexibility, simplicity and scalability when compared to legacy infrastructure, at the same time reducing costs and allowing for greater innovation.

OpenNebula's implementation of Enhanced Platform Awareness (EPA) enables fine-grained matching of processor capabilities to Virtual Machine (VM) and Kubernetes workloads prior to launching the applications. Implementing EPA allows OpenNebula to improve VM packet forwarding performance (throughput, latency, jitter) by exposing low-level CPU and NIC acceleration components to the VNF. Besides network performance, OpenNebula provides the GPU support, bare-metal automation and multi-cluster features needed to address edge and Telco use cases such as distributed User Plane Function (UPF) in 5G Core, CDN or O-RAN.

This white paper presents the current capabilities of OpenNebula to address the requirements for Networks Functions Virtualization (NFV) deployments.

Contents

Glossary

BNG	Broadband Network Gateways
DC	Datacenter
EPA	Enhanced Platform Awareness
HCI	Hyper-converged Infrastructure
NFV	Network Functions Virtualization
VM	Virtual Machine
VNF	Virtual Network Function

1. Introduction

OpenNebula's Enhanced Platform Awareness (EPA) implementation enables fine-grained matching of workload requirements to platform capabilities. EPA features provide OpenNebula with an improved understanding of the underlying platform hardware (HW), which allows it to accurately assign the workload to the best HW resource.

This white paper explains the main OpenNebula EPA features. Each feature is covered individually with a brief description, the level of compliance (applicable to **OpenNebula 6.2+**), a link to the configuration details to enable the feature, and a short discussion of its benefits.

2. What is OpenNebula?

OpenNebula is a **powerful, but easy-to-use, open source solution to build and manage Enterprise Clouds and Edge environments.** 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.

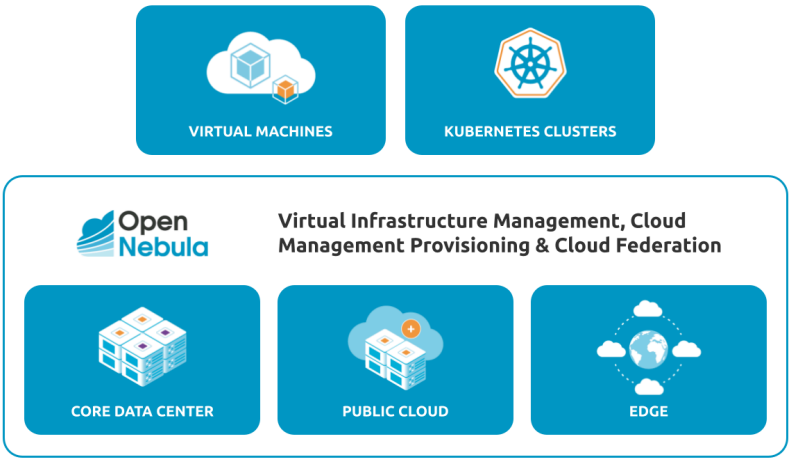

Figure 1. Any application and infrastructure on any cloud.

OpenNebula provides the necessary tools for running containerized applications from Kubernetes, while ensuring enterprise requirements for your DevOps practices. It helps organizations to easily embrace Multi-cloud, Hybrid, and Edge Computing, allowing them to grow their Enterprise Cloud on-demand with infrastructure resources from third-party public cloud and bare-metal providers such as AWS and Equinix Metal.

This white paper describes support for EPA features in OpenNebula. If you are interested in our approach for true hybrid and multi-cloud computing, please refer to **True Hybrid and Multi-cloud with OpenNebula.**[1] If you are interested in an OpenNebula cloud fully based on open source platforms and technologies, please see our **Open Cloud Reference Architecture**.[2]

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**. New versions are released 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.[3]

3. Features Breakdown and Analysis

1. Host CPU feature request
Description: Expose host CPU features to managed guests
OpenNebula Compliance: Full
More Info: CPU_MODEL attribute
Rationale: OpenNebula also retrieves information about which models and CPU features are available at the hypervisors.

2. SR-IOV & PCI passthrough
Description: Provide direct access to a physical or virtual PCI device
OpenNebula Compliance: Full
More Info: PCI Passthrough
Rationale: OpenNebula tracks and allocates devices to guests, and also allows admins to select which devices can be hotplugged. Network devices are also integrated with the Network stack, allowing guests to identify passthrough devices.

3. Hugepages support
Description: Use memory pages larger than the standard size
OpenNebula Compliance: Full
More Info: Using Hugepages

[1] https://support.opennebula.pro/hc/en-us/articles/4403620426129-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://opennebula.io/discover/

Rationale: Scheduling includes hugepage requirements and availability in its placement algorithm.

4. NUMA awareness
Description: Ensures that virtual CPUs' (vCPUs') executing processes and the memory used by these processes are on the same NUMA node
OpenNebula Compliance: Full
More Info: NUMA Topology
Rationale: OpenNebula offers a great deal of flexibility to define virtual NUMA topologies and map them to the physical configuration of the host, including several pinning policies and support for asymmetric configurations.

5. IO-based NUMA scheduling
Description: Creates an affinity that associates a VM with the same NUMA nodes as the PCI device passed into the VM
OpenNebula Compliance: Full
More Info: CPU & NUMA Pinning - PCI Passthrough
Rationale: The scheduling process is slightly modified when a pinned VM includes PCI passthrough devices. In this case, the NUMA nodes where the PCI devices are attached are prioritized to pin the VM vCPUs and memory to speed up I/O operations. No additional configuration is needed.

6. CPU pinning
Description: Supports pinning of VMs to physical processors
OpenNebula Compliance: Full
More Info: CPU & NUMA Pinning
Rationale: OpenNebula supports several pinning policies, and also includes a NUMA scheduler in order to match virtual topologies and hypervisor configurations as closely as possible.

7. CPU threading policies
Description: Provides control over how guests can use the host hyper-thread siblings
OpenNebula Compliance: Full
More Info: CPU & NUMA Pinning
Rationale: OpenNebula can work in CORE (whole core to virtual CPU), THREAD (each vCPU is assigned to a hyper-thread), SHARED (hyper-threads are shared by vCPUs of the guest) and NONE (no pinning).

8. OVS-DPDK
Description: An industry standard virtual switch accelerated by DPDK

OpenNebula Compliance: Full
More Info: OpenNebula and DPDK
Rationale: In this mode OpenNebula will create and configure OVS bridges and ports. This mode is recommended to be used with NUMA+Hugepages described above

4. Other Features for NFV and Telco Cloud

9. GPU hardware support
Description: As customers across industries start to leverage GPU-intensive workloads, such as AI, to transform their business, telcos are uniquely positioned to expand into the role of edge AI service providers.
OpenNebula Compliance: Full
More Info: PCI Passthrough - NVIDIA vGPU proprietary drivers
Rationale: To support GPU-based workloads on your instances, OpenNebula manages virtual GPU (vGPU) resources according to the available physical GPU devices and the hypervisor type.

10. Bare Metal automation provisioning
Description: Automatic deployment and configuration of cloud clusters, on hardware clusters on-demand from public cloud and edge providers or from on-prem edge locations.
OpenNebula Compliance: Full
More Info: Automatic Cluster Deployment
Rationale: Out of the box, OpenNebula's Edge Cloud Architecture provides a comprehensive bare-metal provisioning mechanism by fully automating the lifecycle of OpenNebula clusters, using on-prem hardware and bare-metal offerings from various cloud/edge infrastructure providers.

11. Multi-clustering deployment
Description: Management of multiple highly-distributed clusters, for instance to run large numbers of Kubernetes clusters isolated in different 5G sites that are susceptible to connectivity issues.
OpenNebula Compliance: Full
More Info: OpenNebula Kubernetes Engine (OneKE)
Rationale:Through the OneProvision engine, OpenNebula supports distributed automation and management of multiple clusters simultaneously across different bare-metal cloud/edge locations, as well as automated deployment of front-end instances using GitOps techniques (e.g. through our Edge-as-a-Service platform Edgify). Managing multi-cluster Kubernetes is supported through the OpenNebula Kubernetes Engine (OneKE). Coupled with the OneProvision engine, this makes it possible to manage several Kubernetes clusters in different locations, from a single OpenNebula instance.

5. Conclusions

The EPA features described in this document were created from the collective information and experiences during our collaboration with several network operators. This document recommends software configurations for a smooth OpenNebula installation that supports VNF workloads. However, in many cases there are other aspects to be considered, such as infrastructure platforms and pre-existing services in the datacenter, as well as specific provisioning processes within the company. In these scenarios, OpenNebula can be easily adapted to fit your datacenter and corporate policies. Contact us—we look forward to helping you at any stage of your cloud computing journey.

Rev1.4_20240601