OpenNebula Enhanced Platform Awareness

Version 1.0 – May 2022

Abstract

Network functions, like wireless nodes, Broadband Network Gateways (BNG), routers, firewalls, etc., have traditionally been delivered in dedicated hardware appliances. Virtual Network Functions (VNF) are now emerging to replace that hardware-centric approach through software appliances that provide more agility, flexibility, simplicity, efficiency, and scalability than legacy infrastructure, while also reducing costs and allowing greater innovation.

The implementation of the Enhanced Platform Awareness (EPA) features in OpenNebula enables the fine-grained matching of processor capabilities to Virtual Machine (VM) and Kubernetes workloads prior to launching the applications. The implementation of EPA allows OpenNebula to improve the VM packet forwarding performance (throughput, latency, jitter) by exposing low-level CPU and NIC acceleration components to the VNF.

This white paper presents the current capabilities of OpenNebula to address the requirements for NFV deployments.
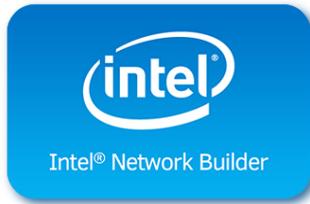
Contents

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 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 this 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**: Automate operations, and deploy and manage Kubernetes clusters, containers, and VMs on a single shared environment.
- **Any Infrastructure**: Choose from VMware ESXi deployments to software-centric HCI architectures based on KVM or LXC.
- **Any Cloud**: Unlock the power of a true hybrid and multi-cloud platform by combining edge, public, and private cloud operations.
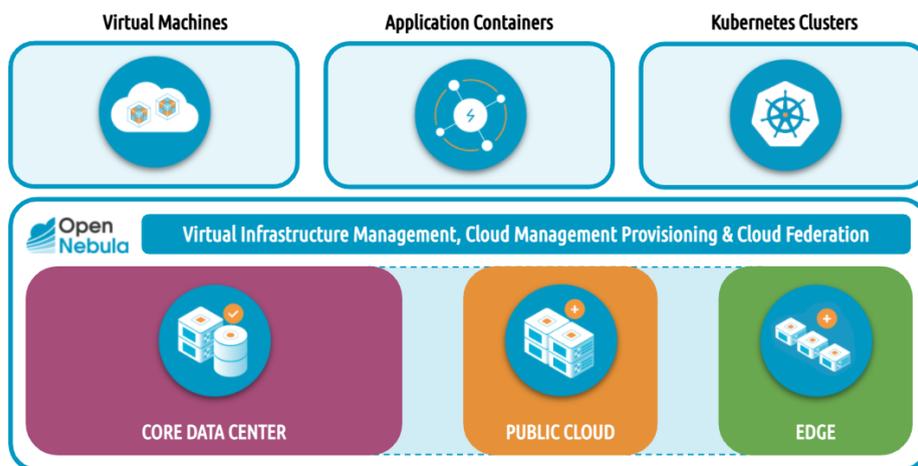

Figure 1. Any application and infrastructure on any cloud.

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, and LXC system containers for container clouds.

This white paper describes the OpenNebula support for EPA features. 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 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 consult our **VMWare Cloud Reference Architecture.**[3] If you are interested in an OpenNebula cloud fully based on open source platforms and technologies, please see our **Open Cloud Reference Architecture.**[4]

The development of OpenNebula follows a bottom-up approach driven by the real needs of sysadmins, DevOps, and corporate users. OpenNebula is an **open source product** with a healthy and active community, commercially supported by OpenNebula Systems through its **OpenNebula Subscription**. Releases are produced on a regular basis and delivered as a single package with a smooth migration path. More information on the benefits of running an OpenNebula cloud can be found on the Key Features page.[5]

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 gets 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, it also allows admins to select which devices can be hotplugged. Network devices are also integrated with the Network stack. This allows guests to identify passthrough devices.

3. Hugepages support
Description: Use memory pages larger than the standard size
OpenNebula Compliance: Full
More Info: Using Hugepages
Rationale: Scheduling includes hugepage requirements and availability in its placement algorithm.

[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/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/

4. NUMA awareness

Description: Ensures 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 the pinning of VMs to physical processors

OpenNebula Compliance: Full

More Info: CPU & NUMA Pinning

Rationale: OpenNebula supports several pinning policies, it 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 above

4. Conclusions

The EPA features described in this document have been 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 taken into account, like 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.

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

Sign up for updates at OpenNebula.io/getupdated

Rev1.0_20220520