White Paper
OpenStack Enhanced Platform Awareness
Enterprise Business

OpenStack* Enhanced Platform Awareness

Enabling Virtual Machines to Automatically Take Advantage of Advanced Hardware Capabilities

Abstract

Enhanced Platform Awareness (EPA) contributions from Intel and others to the OpenStack* cloud operating environment enable fine-grained matching of workload requirements to platform capabilities, prior to launching a virtual machine. For example, EPA can automatically launch a cryptographic workload on a platform with a hardware-based crypto-accelerator.

For workloads requiring particular CPU and/or I/O capabilities, EPA helps OpenStack assign virtual machines (VMs) to run on the optimal platforms. EPA can benefit VM performance and operation, such as for Software Defined Networking (SDN) and Network Function Virtualization (NFV). EPA also enables Cloud Service Providers (CSPs) to offer premium, revenue-generating services based on specific hardware features. This white paper describes OpenStack EPA and how it can deliver business value to Telcos and CSPs.

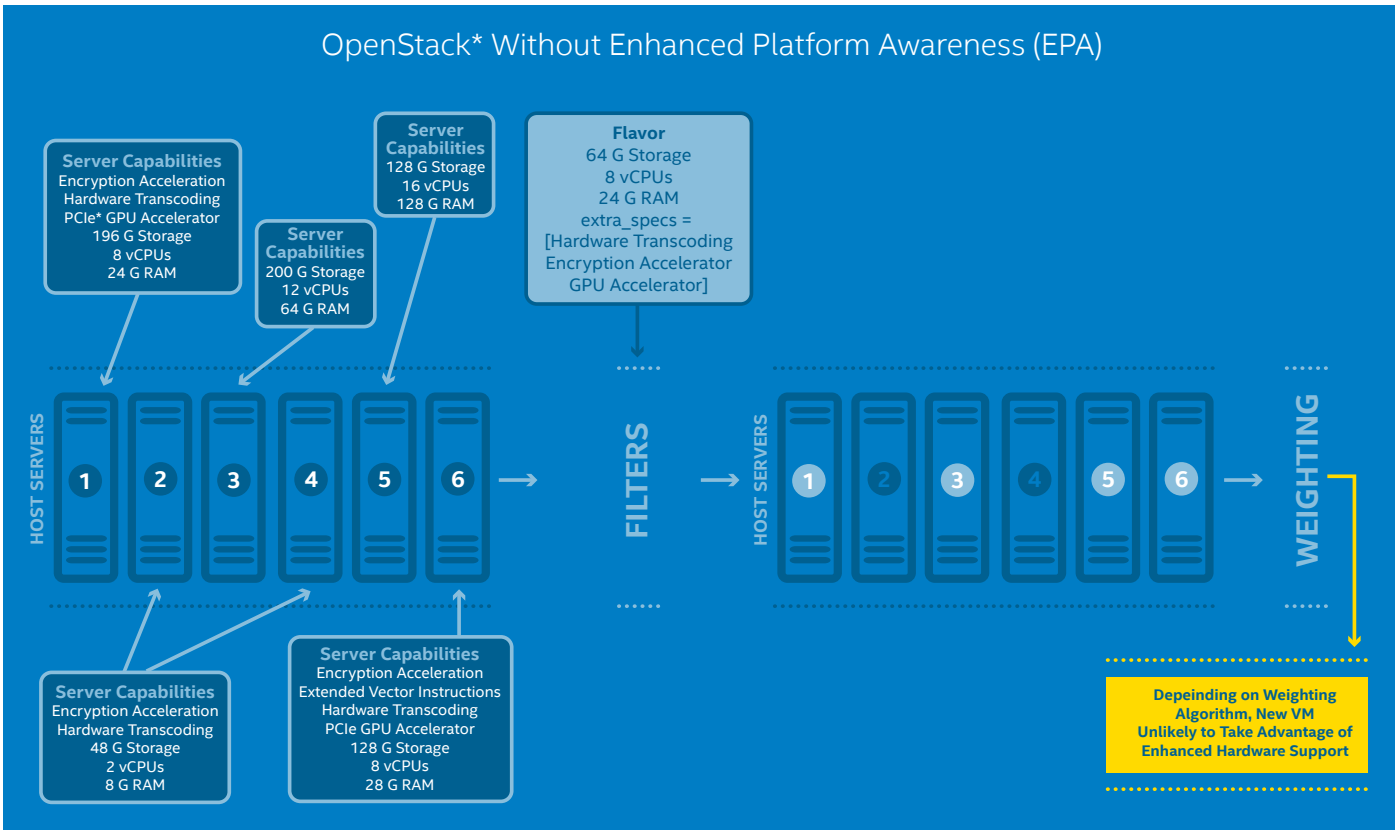

Figure 1. Before Enhanced Platform Awareness a new VM might not be launched on enhanced hardware.

Table of Contents

OpenStack* Enhanced Platform Awareness

Matching Virtual Machines to Hardware Capabilities

OpenStack,* the open source cloud operating environment, automates VM provisioning by fulfilling a user's request to install a VM onto server hardware. The resources allocated to the VM are governed by "flavors" that specify basic operating parameters, such as required virtual CPUs, desired memory, and needed storage space. The filter scheduler in the OpenStack Nova* module then matches the flavor to an available server with the required characteristics.

With today's advanced technologies embedded in processors and chipsets, integrated on server boards, and installed in PCIe* slots, hardware platforms offer many more capabilities now than in the recent past. OpenStack, however, could not take advantage of these enhancements. It was unaware of more than the basic set of server features and functions. OpenStack could not proactively load an application onto specialized hardware to accelerate performance (Figure 1), such as assigning an IPsec VPN workload to a server with built-in cryptographic acceleration.

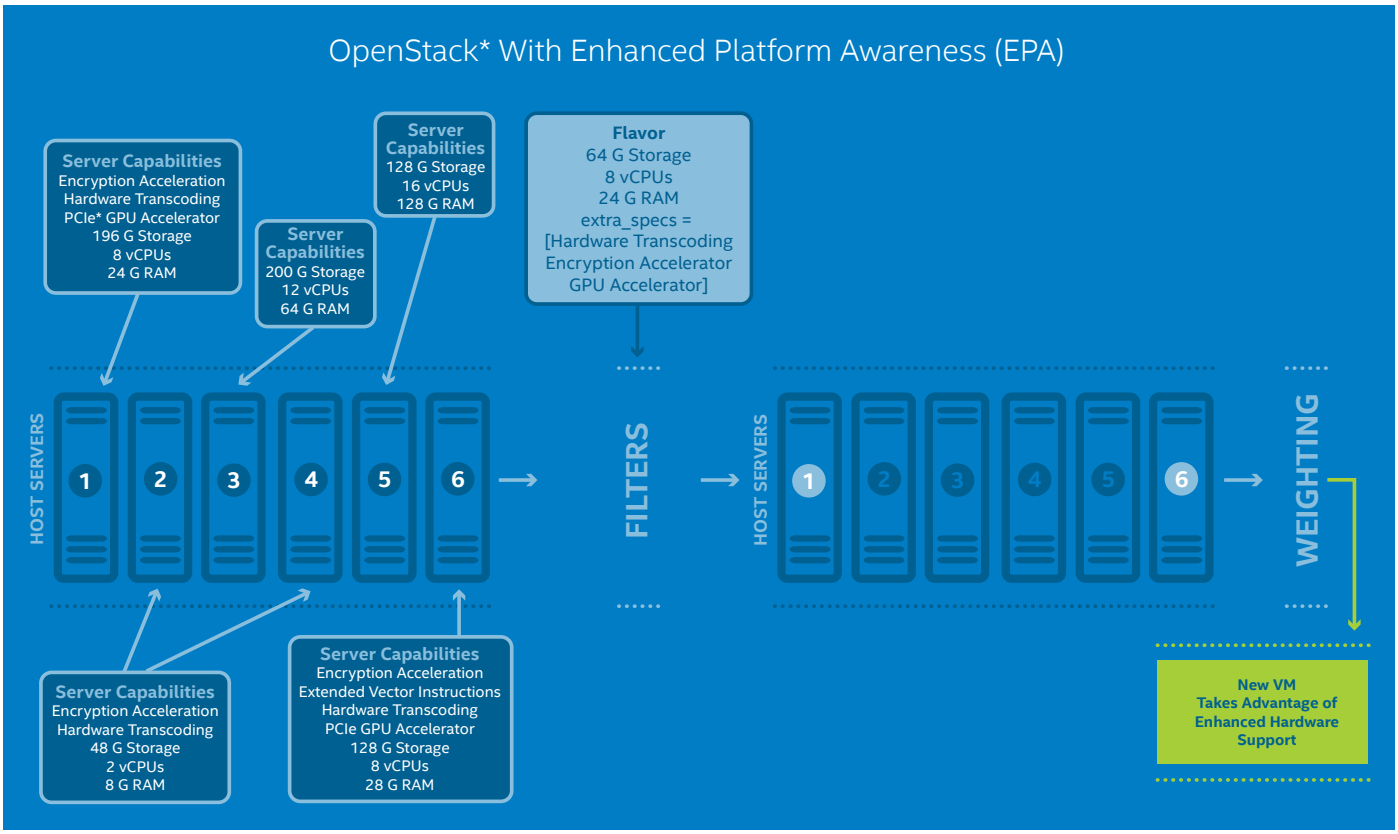

Figure 2. With Enhanced Platform Awareness, the Nova Scheduler matches available hardware to an Instance Type (Flavor).

Today, Enhanced Platform Awareness (EPA) contributions to OpenStack Icehouse* version enable the Nova filter scheduler to match a flavor having specific hardware requirements to a server that can meet them (Figure 2). EPA adds the following enhancements to Nova:

1. Detect platform capabilities through the discovery, tracking, and reporting of enhanced features in the CPU and PCIe slots (new PCIe contributions are currently being reviewed for future inclusion in Nova).
2. Filter and match available platforms with the specific capabilities to an instance type requesting the desired features.
3. Schedule and install the instance onto the selected platform with the enabled features.

Nova assumes the image provider has selected and configured a VM image that can take advantage of the functionality requested in the flavor. For example, if the flavor specifies Intel® Advanced Vector Extensions (Intel® AVX),[1] Nova assumes the image is optimized to take advantage of the expanded Streaming SIMD Extensions (SSE) in the Intel® processor instruction set. If the image does not use Intel AVX, Nova still may schedule the image to install on an Intel AVX-enabled platform, but the VM would not take advantage of those added capabilities.

Potential Benefits

Telcos and Network Operators

EPA additions will help to scale up performance of virtualized networking functions, adding momentum to the growing movement in the telecommunications industry to transform the network using Software Defined Networking (SDN) and Network Functions Virtualization (NFV). SDN and NFV will allow Telcos and Cloud Service Providers to build powerful, flexible network functionality using standards-based, commercial off-the-shelf servers rather than proprietary networking appliances. The benefits to these organizations include:

- Lower operator capital expenses (CapEx) due to the use of commercial off-the-shelf servers.
- Reduced dependence on specific equipment vendors and greater choice of solutions.
- Greater customer agility through automated deployment of virtual machines providing the desired functionality, instead of installing and configuring application-specific hardware.
- Virtual appliance software enables equipment providers to expand their offerings.

The successful adoption of NFV and SDN, however, depends on the ability of a virtualized cloud infrastructure to deliver sufficient performance comparable to today's proprietary network appliances. Advanced processor and networking technologies embedded in silicon and in PCIe devices offer performance that can help increase workload efficiency when the VM has access to them. EPA enables that access.

Cloud Service Providers (CSPs)

EPA additions will enable CSPs to market the use of enhanced technologies, which likely have been in their servers all along. By offering premium services, CSPs can generate new revenues. Some examples of the enhanced services might include:

- Speed-up secure e-commerce transactions to increase sales throughput with hardware-based cryptographic acceleration.
- Return results faster in a customer's rendering engine with floating-point instruction acceleration using a processor's instruction set extensions.
- Improve VM efficiency and increase the throughput of a customer's workload by directly assigning it to Single Root I/O Virtualization (SR-IOV) PCIe acceleration devices.
- Enhance the upload and streaming experiences for users of a video service by taking advantage of a server with embedded transcoding hardware acceleration in the processor.
- Reduce time-to-solution for parallel computing operations using PCIe-based accelerators (currently being reviewed for future inclusion in OpenStack).

EPA Usage Examples

Many cloud data centers today are built around heterogeneous platforms offering a variety of technologies—some with only basic computing functions and others with accelerated performance and higher efficiency. EPA enables system administrators and data center architects to take advantage of enhanced platform capabilities, as follows:

- Increase efficiency by delivering more capabilities with existing resources.
- Improve effectiveness offering new and valuable services with underutilized resources; that is, take advantage of existing capabilities to generate new revenue streams.

With EPA, when the customer requests a new VM with specific hardware support, Nova will filter the servers with the requested features and deploy the new VM onto one of those platforms. The operation is seamless and automatic.

Compute with EPA

CSPs using OpenStack to deploy SSL-enabled web sites can use EPA to launch those sites on VMs with encryption acceleration, if the server hardware has the appropriate capability. CSPs can offer that service at a premium.

Many of today's encryption software products, such as OpenSSL,* are optimized to use Intel® Advanced Encryption Standard – New Instructions (Intel® AES – NI). Intel® Xeon® processors with Intel® Data Protection Technology[2] include hardware acceleration to speed-up encryption and decryption with Intel AES – NI, making those transactions many times faster. The benefit to the end-user is less time waiting to complete a transaction. For the owner of the site, it can mean higher return on investment, because more transactions can be processed with the same resources in the same time, earning higher revenues. It can also improve completed sales because faster encryption processing helps result in fewer abandoned shopping carts. Finally, using Intel AES – NI can improve security, because instruction-based advanced encryption standard versus a coded table-based implementation is less subject to side-channel attacks.

There are many public benchmarks illustrating the acceleration potentials of Intel AES – NI. To gain these benefits, CSPs need to:

1. Specify Intel AES – NI-enabled processors for their fleet of servers.
2. Add these servers to their infrastructure.
3. Ensure the Nova filter scheduler uses the ComputeCapabilitiesFilter.
4. Ensure the workload uses Intel AES – NI (for example, current versions of OpenSSL).
5. Specify AES – NI in the flavor.

Accelerated Encryption/Compression I/O with EPA

CSPs using OpenStack to provision VMs for the highest performance encryption/decryption and compression/decompression can accelerate those tasks on Intel® Xeon® processor E5 and E7 family-based servers with Intel® QuickAssist technology. The embedded encryption/decryption accelerators are available on select Intel® Atom™ processors, as well.

Intel QuickAssist is a set of hardware acceleration modules offered in specific Intel processors, chipsets such as Intel® Communications Chipset 89xx series, and add-in cards to speed-up encryption/decryption and compression/decompression. Intel QuickAssist uses industry-standard application programming interfaces (APIs) for greater software flexibility and longevity.

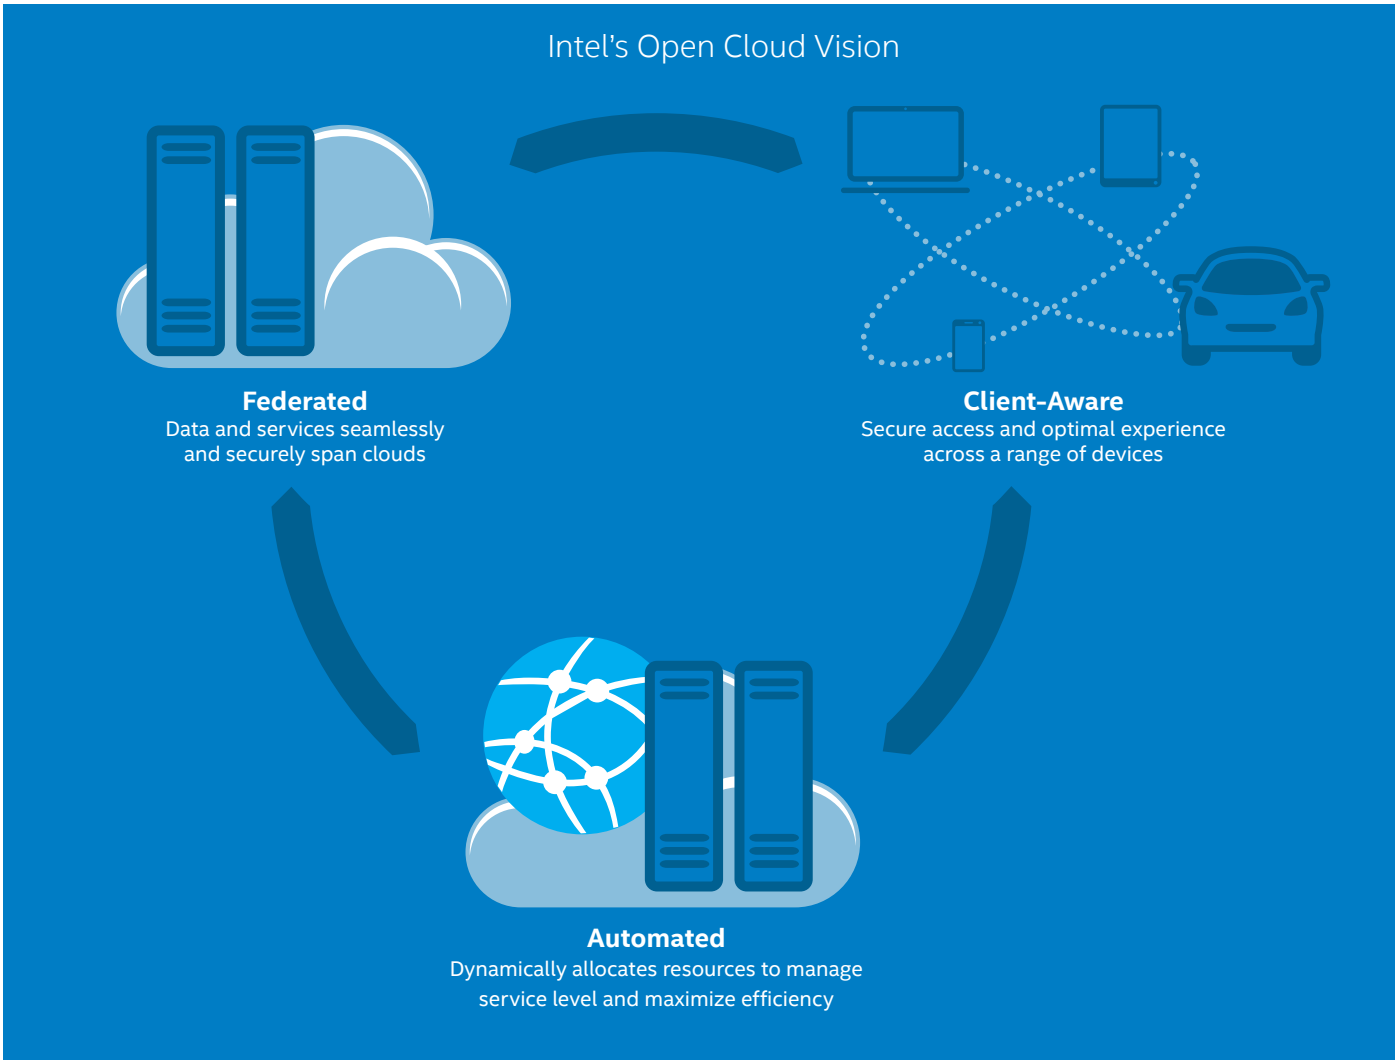

EPA enables Intel's open cloud vision by ensuring the right workloads run on the right hardware with the right compute, storage, communications, and security capabilities.

Intel® Technologies for Enhanced Platform Awareness

Embedded Processor/Chipset Technologies

- Intel® Advanced Encryption Standard—New Instructions (Intel® AES – NI) in Intel® Data Protection technology—accelerates encryption and decryption when Intel AES – NI instructions are included in the software code. The technology is provided in 4th generation Intel® Core™ processors and Intel® Xeon® processors.
- Intel® Advanced Vector Extensions—accelerates floating-point operations with an expanded instruction set for Intel® Streaming SIMD Extensions. This technology is provided in 4th generation Intel Core processors and Intel Xeon processors.
- Intel® Quick Sync Video technology—accelerates transcoding of certain video codecs. This technology is provided in 4th generation Intel Core processors.
- Intel® QuickAssist technology—accelerates encryption/decryption and compression/decompression of I/O operations. This technology is provided in select Intel processors and chipsets. Intel Quick Assist technology may also be found embedded in certain accelerator PCIe* boards.

PCIe-based Accelerators

- Intel® Xeon Phi™ coprocessor—accelerates codes using a large number of cores on the PCIe*-based card.

To ensure OpenStack launches VMs on hardware with Intel QuickAssist technology, CSPs should:

1. Specify Intel Xeon processors (for both encrypt/decrypt and compress/decompress acceleration) and/or Intel Atom processors (for encrypt/decrypt acceleration) with Intel QuickAssist technology, or use accelerator cards with Intel QuickAssist technology on them.
2. Configure Nova and the instance type to use Intel QuickAssist technology.
3. Ensure the Nova filter scheduler uses the *ComputeCapabilitiesFilter* method.
4. Ensure the chosen VM image uses encryption/decryption and/or compression/decompression.
5. Specify Intel QuickAssist in the flavor.

Implementation

With the EPA contributions to OpenStack, implementation is straightforward:

1. Build the OpenStack compute node with Nova.
2. Ensure the Nova filter scheduler uses the *ComputeCapabilitiesFilter* and the pci_passthrough_filter.
3. Add or modify a flavor to include the appropriate key pairs that call out the desired hardware features.

As noted above, the Nova scheduler assumes the selected instance supports the desired functionality requiring the hardware features.

Filter Scheduler

All compute nodes periodically publish their status, resources available, and hardware capabilities to the Nova database. The Nova scheduler uses that data to make decisions when a request comes in. There are many filtering strategies for the scheduler to support.

The *ComputeCapabilitiesFilter* filters hosts meeting the *"extra_specs"* specified in the flavor and passes the list of hosts to the weighting function.

Flavor (Instance Type)

Flavors define a number of parameters in the instance type, resulting in the user having a choice of what type of VM to run. The parameters include an *"extra_specs"* listing of key-value pairs that identify specific hardware features desired to support functions within the VM, such as I/O accelerators.

Summary

With EPA in OpenStack, a new level of control and configuration is available for Telcos and CSPs. The benefits extend from the Telco and cloud operator to their customers, with improved web site throughput, faster network performance, and new revenue streams.

EPA implementation merely requires simple additions to VM flavors to enable launch of specific VMs with enhanced hardware capabilities. While EPA currently addresses enhancements only in the processor, the OpenStack community is considering adding awareness for PCIe-based devices, such as GPUs and coprocessors (see https://wiki.openstack.org/wiki/Enhanced-platform-awareness-pcie).

For more information about EPA, visit the OpenStack web site.

Appendix A – Additional Resources

OpenStack

- IDF presentation on OpenStack
  - Optimizing Workloads in OpenStack Public Cloud Environments
  - Configuration & Deployment Guide for OpenStack Swift Object Storage
- Intel® Developer Zone OpenStack page

SDN & NFV

- ONP Switch Reference Design product brief
- ONP Server Reference Design audio-enabled IDF presentation
- Intel DPDK-optimized open vSwitch
- Wind River* Open Virtualization Profile
- Growth of SDN, NFV, and how the Intel® Open Network Platform speeds development
- Open, Simplified Networking based on SDN and NFV
- Open Networks Provide Needed Flexibility in IT Market

1 Intel® Advanced Vector Extensions (Intel® AVX)* are designed to achieve higher throughput to certain integer and floating point operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you should consult your system manufacturer for more information.

2 No computer system can provide absolute security. Requires an enabled Intel® processor and software optimized for use of the technology. Consult your system manufacturer and/or software vendor for more information.

Printed in USA 0514/DBF/HBD/PDF ♻ Please Recycle 330584-001US