OPENSTACK* ENHANCED PLATFORM AWARENESS

Enabling Virtual Machines to Automatically Take Advantage of Advanced Hardware Capabilities

Executive Summary

Enhanced Platform Awareness (EPA) contributions from Intel and others to the OpenStack* Cloud Operating System enable fine-grained matching of server capabilities to virtual machine (VM) workload requirements prior to launching a VM. For workloads requiring particular CPU and/or I/O capabilities, EPA helps OpenStack assign VMs to run on the most appropriate platforms and gain additional benefits from features built into the system. For example, OpenStack with EPA can specifically and automatically launch a cryptographic workload onto a platform with a hardware-based crypto-accelerator to improve throughput performance.

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. For enterprise data centers, IT can use EPA to automatically assign policy-controlled, sensitive workloads to trusted platforms that provide enhanced capabilities, such as stronger security, compliance, and data protection. This allows IT to further optimize existing resources from which the organization can benefit the most.

This white paper describes OpenStack EPA and how it can deliver business value to enterprise IT, Telcos, and CSPs.

Table of Contents

OpenStack Enhanced Platform Awareness

Matching Virtual Machines to Hardware Capabilities

The OpenStack Cloud Operating System is a scalable, open-source cloud computing platform for private, public, and hybrid clouds used by enterprise IT, Telcos, and Cloud Service Providers (CSPs). OpenStack includes a collection of interoperable, open-source software modules that orchestrate large pools of compute, storage, and networking resources. All of these resources can be managed through a web-based dashboard, command line tools, and RESTful APIs, so organizations have a great deal of flexibility in how they deploy, operate, and manage their cloud. The Enhanced Platform Awareness (EPA) addition to OpenStack enables fine-grained matching of resources available in the data center to varying types of workloads to better optimize the cloud services being provided.

OpenStack fulfills a user's request to provision a virtual machine (VM) by automatically installing it 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 (Figure 1).

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. Until now, OpenStack was unaware of more than the basic set of server features and functions. OpenStack could not proactively load an application onto enhanced hardware to accelerate performance, such as assigning an IPsec VPN workload to a server with built-in cryptographic acceleration to speed cryptography tasks.

Today, with EPA, OpenStack can take specific advantage of these enhancements to benefit the functionality of the VM's services. EPA enables the Nova filter scheduler to match a flavor with specific hardware requirements to a server that can deliver them (Figure 1).

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 specific capabilities to a flavor requesting the desired features.
3. Schedule and install the VM onto the selected platform with the enabled features.

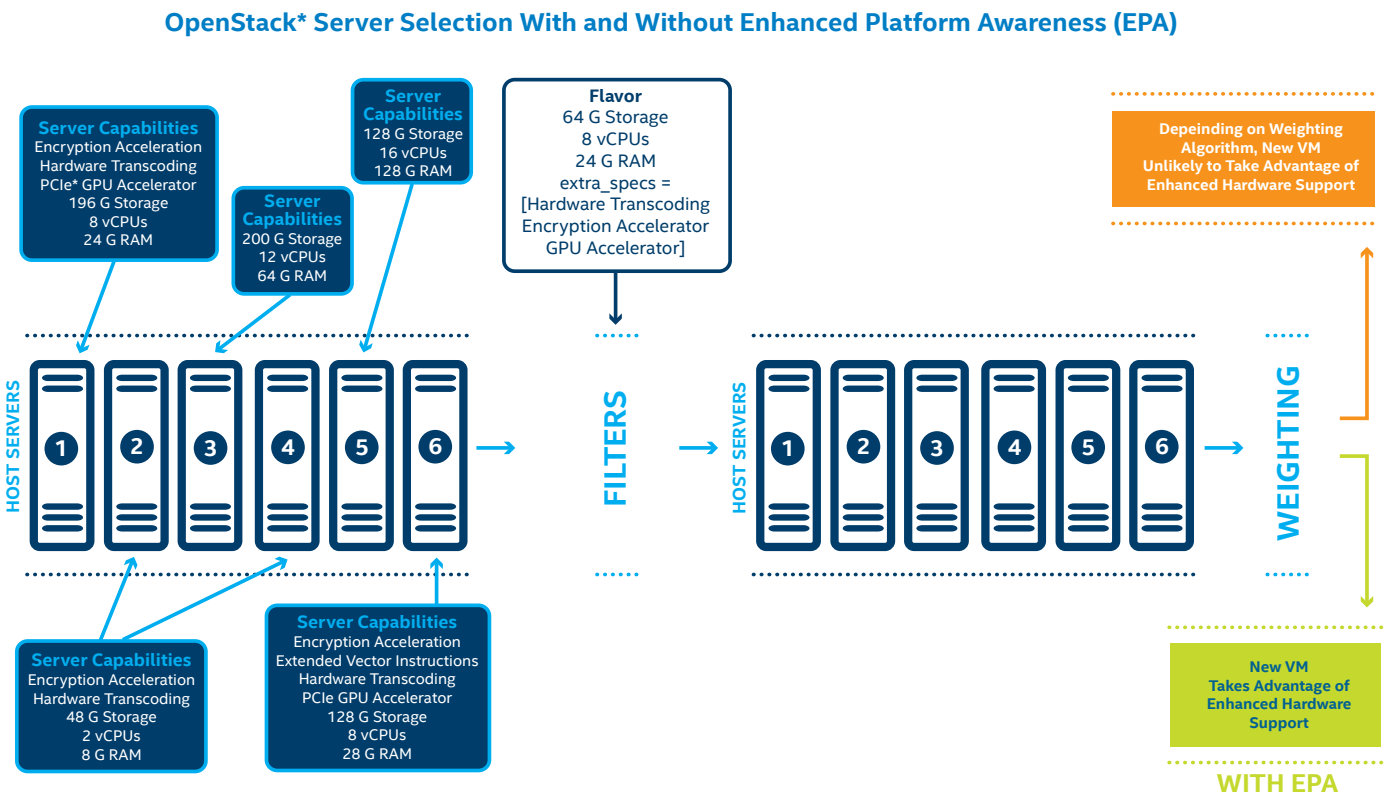

Figure 1. Enhanced Platform Awareness guides selection of servers for new VM.

Images

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 2 (Intel® AVX2),[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 AVX2, Nova still may schedule the image to install on an Intel AVX2-enabled platform, but the VM would not take advantage of those added capabilities.

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.

Nova 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* method filters those hosts meeting the *"extra_specs"* specified in the flavor and passes the list of hosts with those capabilities to the weighting function.

EPA additions to OpenStack help further enable Intel's Open Cloud Vision (Figure 2).

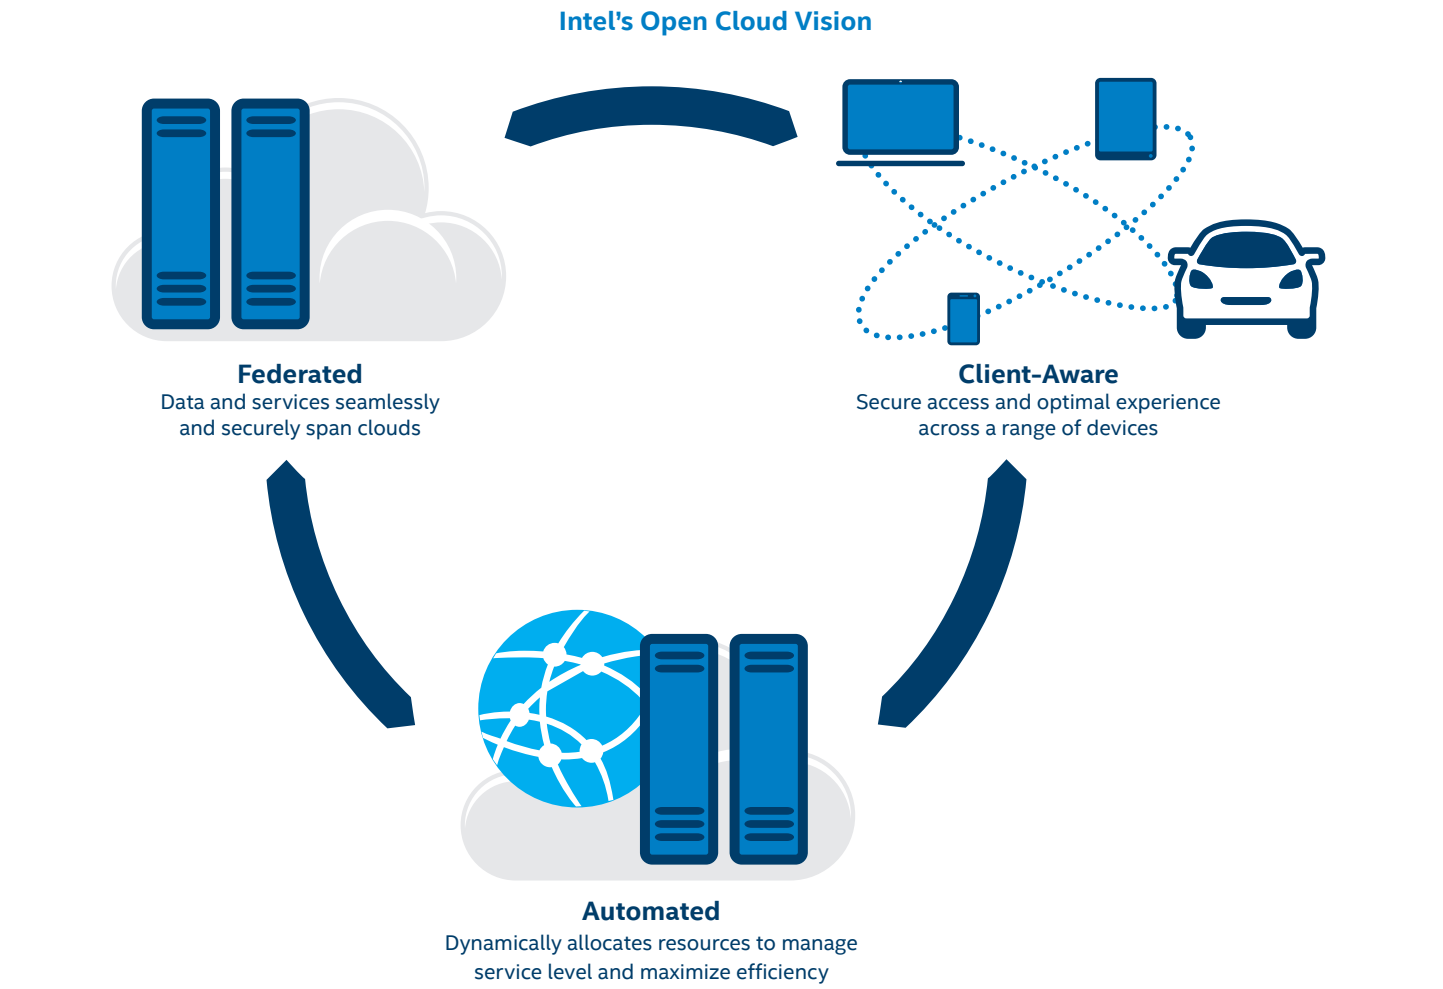

Figure 2. EPA enables Intel's open cloud vision by ensuring the right workloads run on the right hardware with the right compute, storage, communications, and security capabilities.

Potential Benefits

EPA offers considerable benefits to enterprise, Telco, and CSPs, by enhancing the functionality of VM performance, security, data protection, and compliance, determined by the flavor and enhanced capabilities in the hardware. These benefits can be offered to customers for sell-up revenue generation opportunities, as features to differentiate the business from competitors, or to better optimize workloads with unused resources.

Enterprise Data Centers

EPA enables system administrators and data center architects to automatically take advantage of enhanced platform capabilities and thus further optimize existing resources in the data center.

Enterprise IT is under increasing constraint to deliver more and differentiated services, enhance security and compliance, and simplify rollout of new services without expanding their operational budgets. OpenStack alone automates orchestration of new services, easing the burdens IT faces with provisioning. But when new services comprise policy-controlled, sensitive workloads that require specific capabilities, EPA allows IT to restrict these types of workloads to trusted platforms without impairing application performance or throughput. With OpenStack and EPA, IT gains both simplified rollouts and confidence in the type of infrastructure being assigned to support those rollouts.

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 software-based network functions and run them on standards-based, commercial off-the-shelf servers rather than proprietary networking appliances. The benefits to these organizations include:

- Lower capital expenses (CapEx) due to the use of commercial off- the-shelf servers.
- Reduced dependence on specific equipment vendors and greater choice of solutions.
- Faster time to innovative new services driven by software rather than waiting for hardware vendors to launch new products.
- Greater customer agility through automated deployment of virtual machines that provide the desired functionality, instead of installing and configuring application-specific hardware.
- Opportunities for equipment providers to expand their offerings with marketable virtual appliance instances.

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

EPA additions will enable CSPs to market richer services built on enhanced technologies, which likely have been in their servers all along, as sell-up opportunities to customers. By offering premium services with EPA, CSPs can generate new revenues and differentiate their business in the marketplace. 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 enhanced 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

EPA allows data center architects and IT managers to better optimize resources and enhance cloud services.

Enhanced Security with EPA

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

Many of today's encryption software products, such as OpenSSL,* are optimized to use Intel® Advanced Encryption Standard – New Instructions[2] (Intel® AES – NI) to accelerate encryption/decryption operations. Intel® Xeon® processors with Intel® Data Protection Technology[2] include Intel AES – NI, making SSL 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.

Enhanced Compute with EPA

Enterprises with private clouds delivering VMs to run advanced vectorized codes, such as multimedia, scientific, and financial software, can use EPA to significantly accelerate code runs and reduce times to solution.

Many modern applications utilizing SSE for SIMD processing take advantage of Intel® Advanced Vector Extensions 2 (Intel® AVX2) built into several Intel® processors, including the Intel® Xeon® processors E5 and E7 v3 families. Intel AVX2 has shown to improve performance of floating-point and SIMD operations.[3] When servers with these processor families are available in the data center, IT can use OpenStack EPA to provision VMs specifically for enhanced performance of vectorized codes.

Enhanced Encryption/Compression I/O with EPA

Telcos and 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 as add-on PCIe cards.

Intel QuickAssist technology 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 technology uses industry-standard application programming interfaces (APIs), enabling easy integration into applications and for greater software flexibility and longevity.

Enhanced Video Processing with EPA

Telcos, CSPs, and enterprise data centers that support video transcoding from one video format to another, can accelerate this service using Intel® Quick Sync Video, dedicated hardware embedded in the silicon of select Intel® Xeon® processor E3 v1, v2, and v3 families. Whether the services are part of a transcoding farm that supports video production customers, to convert consumer video formats to an internal format for social media operations, or other types of tasks, accelerating them with Intel Quick Sync Video can shorten the overall transcode time per job.

Using OpenStack EPA to provision VMs specifically on servers supporting Intel Quick Sync Video can not only service customers faster, but possibly allow more jobs to run on a particular VM. By increasing transcoding performance, operators can potentially reduce the number of servers required to support their customer base and without the need of a dedicated GPU. These kinds of economies can benefit costs of operations.

Table 1 summarizes the requirements to help achieve these usages.

Table 1. Enhancing Services with EPA

	Enhanced with EPA			
	Accelerated SSL Security	Accelerated Floating-Point/ Integer Compute	Accelerated Encryption/ Compression for I/O	Accelerated Video Processing
What technology to specify in processors/ platforms	• Intel® Advanced Encryption Standard New Instructions (Intel® AES – NI)	• Intel® Advanced Vector Extensions 2 (Intel® AVX2)	• Intel® QuickAssist Technology (in processors and add-on accelerator cards)	• Intel® Quick Sync Video
Intel products with the technology	• Intel® Xeon® Processor E3 v3, E5 v3, and E7 v2 Families; • Select Intel® Atom™ processors	• Intel® Xeon® processor E3 v3, E5 v3 Families • Intel® Xeon® processor E7 v2 Family (supports Intel® AVX only)	• Intel® Xeon® Processor E5 2600, E5-2400 Families with Intel® Communications Chipset 89xx Series • Intel® QuickAssist Adapter 8950	• Intel® Xeon® Processor E3 v3 Family
Use this filter	• ComputeCapabilitiesFilter		• ComputeCapabilitiesFilter • pci_passthrough_filter (for PCIe-based devices)	• ComputeCapabilitiesFilter
Workload must use the specific technology	Yes			
Specify the feature in the flavor	Yes			

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 (Intel® AVX) and Intel® Advanced Vector Extensions 2 (Intel® AVX2) –** accelerates floating-point and integer operations with an expanded instruction set for Intel® Streaming SIMD Extensions.
- **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 QuickAssist technology may also be found embedded in certain accelerator PCIe* boards.
- **Intel® Trusted Execution Technology (Intel® TXT) –** supports trusted launches by measuring the hardware and software environment to verify VMs boot into known good states.
- **Intel® Node Manager –** server management technology that extends component instrumentation to the platform level and can be used to make the most of every watt consumed in the data center.
- **Data Plane Development Kit (DPDK) –** dramatically improves packet processing performance when running on Intel® Xeon® processors.
- **SR-IOV –** technology embedded in Intel® Ethernet Controllers, SR-IOV accelerates network I/O between VMs and the NIC using hardware resources dedicated to VMs, without intervention by the hypervisor.

PCIe-based Accelerators

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

Implementation

Implementing EPA is straightforward (Table 1):

1. Build the OpenStack compute node with Nova.
2. Ensure the Nova filter scheduler uses the *ComputeCapabilitiesFilter* and the *pci_passthrough_filter*.
3. Specify new servers with the hardware technologies/ capabilities you specifically want to apply to your VMs.
4. Add those servers to your infrastructure.
5. Confirm the instance for the VM utilizes the desired technologies.
6. 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.

Summary

With EPA added to OpenStack, a new level of control and configuration is available for enterprise IT administrators, Telcos, and CSPs. The benefits extend from the administrators and operators to their customers, with improved web site throughput, faster network performance, new revenue streams, and greater optimization of existing and new data center resources.

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.

Plan Your Cloud for the Future of OpenStack

OpenStack is quickly becoming the dominant paradigm for agile and efficient IT service delivery for cloud computing. As it evolves, it will continue to help enterprises, CSPs, and Telcos move forward more quickly and cost effectively. EPA is only one of recent enhancements to OpenStack. The community continues to add critical features to all OpenStack Modules, readying for upcoming releases.

Key EPA Features in Kilo Release		
Feature	Benefit	Intel Technology Supported
CPU pinning and huge pages	Improved performance for NFV workloads	–
I/O (PCIe)-based NUMA-aware scheduling	Added efficiencies and performance for NFV workloads	Intel® QuickAssist Technology
Power and Thermal Aware Scheduling (PTAS)	Greater efficiency of resources and lower cost by making scheduling decisions based on instrumented power and thermal data from servers	–

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.

[3] https://software.intel.com/en-us/articles/how-intel-avx2-improves-performance-on-server-applications

0515/DBF/HBD/PDF 330584-002US