White Paper
Communications and Storage Infrastructure
Intel® Open Network Platform for Servers

Intel® Open Network Platform Server Reference Architecture: SDN and NFV for Carrier-Grade Infrastructure and Cloud Data Centers

The Intel® Open Network Platform for Servers streamlines the adoption of software-defined networking (SDN) and network function virtualization (NFV) by data-center operators as well as cloud and telecom service providers. This reference design enables highly reliable, low-latency services while delivering the CAPEX and OPEX advantages of a tailored, standards-based architecture supported by a broad ecosystem.

1 Executive Summary

Data-center operators, telecom carriers, and cloud service providers are striving to optimize their strategies for decoupling network functions from the underlying hardware, which is critical to increasing agility and lowering costs. That abstraction is central to SDN usages, where virtualized network resources are programmatically provisioned, coordinated, and controlled by management software, dynamically as changing requirements warrant. Accordingly, physical assets can be utilized more efficiently, and novel services and usage models can be supported rapidly.

Similarly, NFV capabilities replace fixed-function hardware with implementations fully in software that run more cost-effectively on general-purpose, standards-based servers. Virtualized software functions are assembled on demand to create an open-ended variety of communications capabilities, including those needed to support novel service offerings, without requiring the addition of specialized hardware. SDN and NFV reduce costs and increase business agility, fostering competitive advantages, but those goals are more widely agreed upon than the means to achieve them are.

The Intel® Open Network Platform (Intel® ONP) offers a structured, commercially viable approach to simplify and accelerate adoption of SDN and NFV, while meeting the availability, performance, and security needs of telecom and cloud service providers, as well as data-center operators. The Intel ONP server reference design (introduced in 2014), also known as "Intel ONP for Servers," complements Intel ONP for Switches (introduced in 2013) to enable effective deployment and centralized control and orchestration of virtualized servers, appliances, and network services on open standards with Intel® architecture.

Table of Contents

The reference design offers ecosystem enablement and engineering guidance, rather than a productized offering or a prescriptive set of requirements. The open approach of Intel ONP for Servers emphasizes choice among software and equipment providers, with the flexibility to integrate with existing infrastructures. This paper introduces architects and decision makers at telecom and cloud service providers, as well as conventional data centers, to Intel ONP for Servers. It consists of the following primary sections:

- **"Business Drivers Behind Intel ONP for Servers"** outlines the competitive forces that make SDN and NFV technologies compelling and that set up the value proposition for Intel ONP for Servers.
- **"Architecture for the Intel ONP Server Reference Design"** delineates the role of Intel ONP for Servers, in the context of specific hardware features and the overall infrastructure.
- **"Enabling SDN and NFV Solutions on the Intel ONP Server Reference Design"** introduces the breadth of enablement associated with Intel ONP for Servers, in terms of industry standards, systems integration, and the software ecosystem.
- **"Intel® Platform-Level Technologies for SDN and NFV"** describes the value of certain technologies developed by Intel to Intel ONP for Servers, as well as to SDN and NFV more generally.
- **"Adoption Paths for SDN and NFV on Intel Architecture"** shows the relationship between Intel ONP for Servers and other approaches taken by TEMs and OEMs to implement SDN and NFV.

2 Business Drivers Behind Intel ONP for Servers

Data-center operators must increase agility and efficiency, to support expanding service requirements with flat or shrinking budgets. Likewise, service providers face gathering intensity in the ongoing challenge to rapidly offer new services at scale that expand average revenue per user (ARPU) and enhance the user experience for end customers. In doing so, both must respond with agility to more rapidly shifting demand among offerings, changing market forces, and the needs of new geographies. At the same time, they must also strive to minimize both CAPEX and OPEX, as a competitive necessity.

The rapid proliferation of devices such as smartphones and tablets has dramatically amplified both the opportunities and the challenges associated with offering telecommunications services. Still, the Internet of Things (IoT) will dwarf the current install base of endpoint devices; Gartner forecasts that the IoT will reach 26 billion installed units by 2020, up from 0.9 billion in 2009.[1] The data-transmission services required to support this explosive growth will be unprecedented in their scope and diversity.

Intel ONP for Servers supports the efforts of many communications networks to meet this emerging demand by adopting architectural practices and standards more typically used in data centers. Commercial, off-the-shelf (COTS) systems enhance the cost-effectiveness of flexible, scalable infrastructure for the agile delivery of new services. As conceptual frameworks, SDN and NFV are central to the development of standards-based infrastructures for next-generation service delivery, and Intel ONP for Servers helps guide service providers and data-center operators in the adoption of these technologies. SDN decouples the mechanisms that determine where traffic should be directed (the control plane) from the mechanisms that actually forward packets to those destinations (the data plane). This approach enhances efficiency and data protection in the network, including in shared multi-tenant environments, with mechanisms that include the following:

- **Network virtualization and abstraction**, allowing robust mechanisms for data-traffic control among virtual entities (e.g., virtual machines), independent of physical systems.
- **Centralized network control**, for dynamic coordination and management of computing, storage, and network resources to apply them as needed to workloads.
- **Programmability of the network by external applications**, allowing granular, programmatic control over network resources, to apply them for specialized tasks.

NFV applies similar concepts to efficiently represent fixed-function network devices in software that can be deployed on standards-based servers. This approach allows for the reduction or elimination of special-purpose network hardware such as network processing units (NPUs), co-processors, application-specific integrated circuits (ASICs), and field-programmable gate arrays (FPGAs). Instead, virtualized building blocks allow for communication services to be assembled and disassembled on demand, without the use of dedicated hardware. Therefore, costs associated with wired and wireless networking equipment at both data centers and end-customer locations can potentially be reduced.

3 Architecture for the Intel ONP Server Reference Design

Intel's enablement work for SDN and NFV within this emerging model of network infrastructure is embodied in Intel ONP. The over-arching goal of Intel ONP is to reduce the cost and effort required for service providers, data-center operators, TEMs, and OEMs to adopt and deploy SDN and NFV architectures. The foundations of this effort are the Intel ONP Server Reference Design, introduced in the second half of 2014, and the Intel ONP Switch Reference Design, introduced in 2013. Intel ONP for Servers is illustrated conceptually in Figure 1.

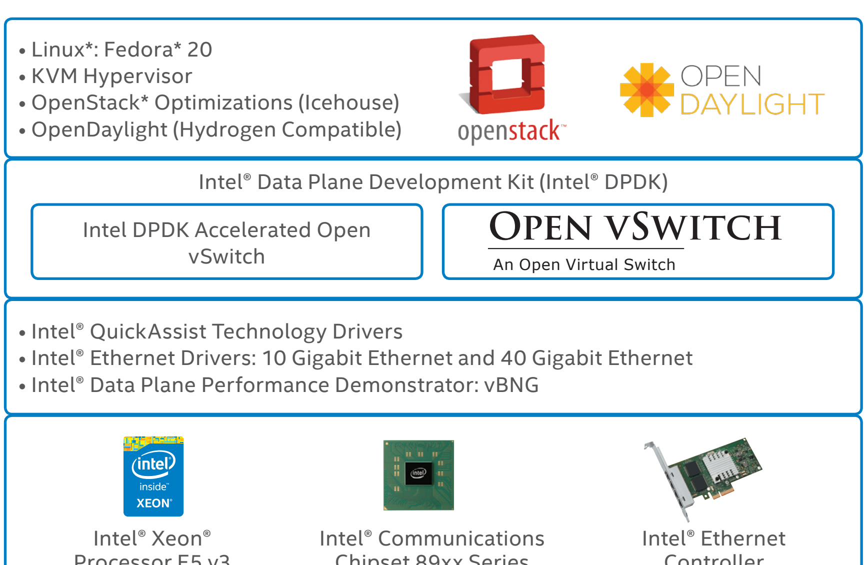

Figure 1. Intel® ONP for Servers.

Open Source Contributions and Leadership from Intel

Intel has a long-standing dedication to advancing the open-source software ecosystem, through code contributions, participation in standards bodies, and enablement for the latest hardware features. A few of the open-source projects closely associated with SDN and NFV that Intel has a particularly strong dedication to include the following:

- **OpenStack*** is an OS and a set of related tools for building and managing private, hybrid, or public clouds, including the ability to control pools of processing, storage, and networking resources.
- **Open Daylight** provides a common open source framework and platform for SDN, including a network-resource controller and utilities for management, security, and virtualization.
- **Open vSwitch** is a software-based, multi-layer virtual switch that can operate either as a soft switch running within the hypervisor or as the control stack for switching silicon.
- **OpenFlow** is a protocol that enables software on the controller layer (described below) to access the switch data plane over the network and make configuration changes in response to real-time traffic demands.
- **Intel® Data Plane Development Kit (Intel® DPDK)** is a set of libraries and drivers that support accelerated software-based packet processing.

The open-systems philosophy behind these reference designs is to offer the industry a set of engineering best practices that can be implemented using hardware and software offerings from across the ecosystem. To that end, Intel engages in both open-source and commercial software enablement, works extensively with standards organizations, and fosters expertise within the systems-integration ecosystem. These aspects of Intel ONP are described in more detail later in this paper.

Architecturally, Intel ONP is made up of three layers, which are represented in Figure 2. Starting from the bottom, the model comprises the node layer, the controller layer, and the orchestration layer, each of which corresponds to a major set of capabilities within the architecture as a whole. Each of those layers is described in more detail below.

3.1 Node Layer

The bottom layer of the Intel ONP architecture is comprised of various virtualized network nodes, such as switches, servers, appliances, and network services. Importantly, these nodes include capabilities that are already in service today as well as others that are under development or yet to be conceived, all on standards-based systems. These nodes can be managed as virtual entities that are decoupled from specific switch and server hardware, using either open-source software tools such as Open vSwitch and OpenFlow, or various proprietary alternatives.

Replacing fixed-function network appliances with virtual appliances running on standards-based Intel architecture enables software based solutions that reduce both CAPEX (through lower equipment costs) and OPEX (by simplifying the infrastructure). Moreover, it can dramatically enhance agility and support new services on existing infrastructure, for accelerated time to revenue from new service offerings, as well as greater alignment with the business as a whole for IT organizations and data-center operators.

3.2 Controller Layer

Management of network resources at the node/switch layer is addressed by the controller layer, by means of APIs, protocols, and services such as (but not limited to) those established by the OpenDaylight project, a community-based industry consortium hosted by the Linux* Foundation. The project provides code (including the controller itself) and solution recipes intended to accelerate development, commercialization, and adoption of SDN and NFV capabilities. It also helps ensure that the controller adheres to open industry standards and facilitates broad API compatibility, to foster open interoperability among solution building blocks from various providers.

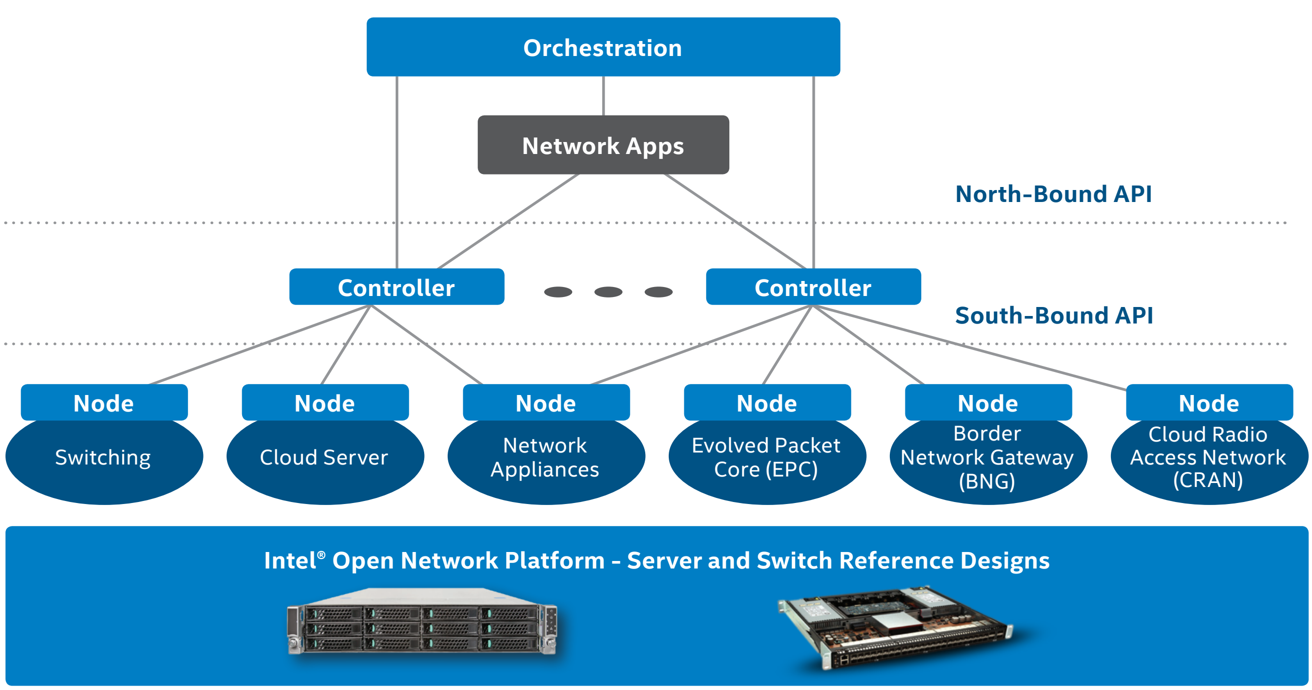

Figure 2. Three-layer Intel® Open Network Platform architecture.

The first OpenDaylight software release, named "Hydrogen," became available for public use in February 2014. The package is available in multiple forms, including a Base Edition for exploratory and academic usages, a Virtualization Edition for data centers, and a Service Provider Edition for cloud service providers and telecom carriers. In addition to Open vSwitch and OpenFlow Support, OpenDaylight facilitates the use of common security and network-management tools with SDN and NFV infrastructure. It also provides utilities for network virtualization, policy management, and DDoS protection.

3.3 Orchestration Layer

At the highest level of the infrastructure, the orchestration layer enables visibility across all physical and virtual assets. The orchestration layer exposes APIs to the controller layer that allow interrogation down to individual nodes to determine their capabilities and status, while also enabling the nodes to broadcast telematics that report on factors such as health, utilization, and availability. This enhanced platform awareness supports fine-grained matching by the orchestration layer between the needs of specific workloads and the network nodes best suited to meet those needs.

Contributions by Intel and other open-source community members to the OpenStack* project, for example, enable the Nova scheduler to assign an IPsec VPN workload to a node hosted on a server that provides hardware-accelerated encryption. In addition to accelerating performance and optimizing OPEX for particular tasks, this capability could allow providers to offer premium services to subscribers that incorporate such features. Likewise, data-center operators could support enhanced QoS and improved utilization of system resources. The potential for other such capabilities could arise from enhanced awareness at the orchestration layer of node/switch-layer hardware features such as Intel® Data Protection Technology[2], Single-Root I/O Virtualization (SR-IOV), Intel® QuickAssist Technology, and many others.

4 Enabling SDN and NFV Solutions on the Intel ONP Server Reference Design

Building on top of open-standards servers based on Intel architecture, the Intel ONP for Servers reference design reflects the value of enabling work across the ecosystem, including support and developm ent for software, systems integration, and industry standards, as represented in Figure 3. This enablement directly benefits end customers such as telecom carriers, cloud service providers, and data-center operators, with a focus on flexible options for implementation that are informed (but not constrained) by the reference design.

4.1 Software Ecosystem

Intel ONP for Servers is developed largely using open-source software components, in keeping with Intel's overall commitment to the open-source community. As alluded to elsewhere in this paper, contributions to projects and standards such as Open vSwitch, OpenFlow, OpenDaylight, and OpenStack have played a primary role in the development of the capabilities associated with the Intel ONP for Servers reference design. Open-source code developed as part of the Intel ONP initiative is promulgated directly through those communities, as well as 01.org, kernel.org, and others.

To fulfill the platform's imperative to be as flexible as possible, enablement for proprietary software (and commercial open-source distributions) is also a key aspect of the ongoing development of the Intel ONP for Servers reference design. This work is a natural extension of the ongoing co-engineering activities between Intel and major providers of OSs and solutions in areas such as virtualization, networking, management, and security.

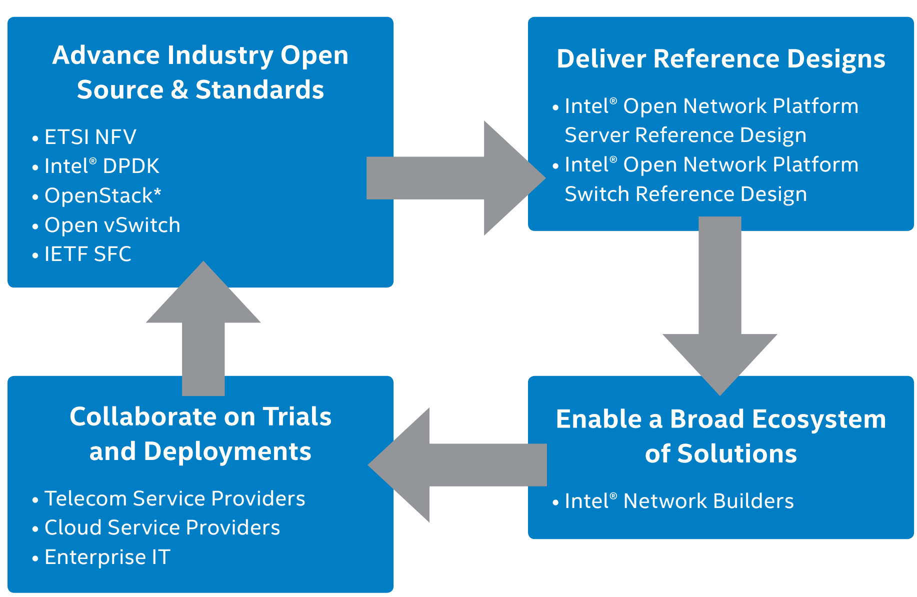

Figure 3. Ecosystem enablement for Intel ONP for Servers.

Through this broad-based approach, platform enablement encompasses most major software offerings at the node/switch, control, and orchestration layers that are relevant to service providers, data-center operators, OEMs, and TEMs. It supports vendors that prefer to tune and customize some or all of the components as they assemble commercial solutions, as well as the preference of others for "ready-made" solutions offered by external providers.

4.2 Intel® Network Builders

The Intel® Network Builders program is an ecosystem of hardware, software, middleware, and professional-services vendors that enables collaboration to bring technologies together into robust offerings. The goal of the program is to foster solutions that enhance ecosystem alignment, to reduce development effort and increase interoperability of solutions. Those goals are facilitated through member access to resources that include the following:

- **Marketing assistance** to increase visibility and connect with customers through channels such as demos, press releases, social media, and webinars.
- **Business networking** to establish relationships with industry peers and build collaborative opportunities to develop new product offerings.
- **Engineering support**, including access to a library of solution recipes and reference architectures that include proven installation and configuration instructions.

4.3 Industry Open Standards and Open Source Projects

As solutions for SDN and NFV continue to develop, new standards will emerge, existing standards must evolve, and stewardship is needed to maintain an open ecosystem. To that end, Intel maintains its standing as a member of and contributor to standards organizations that include the following:

- **European Telecommunications Standards Institute (ETSI)** produces global standards for information and communications technology, including GSM*, 3G, 4G, and DECT*.
- **NFV Open Platform for NFV (OPNFV)** is a service-provider driven initiative at the Linux foundation to create an integrated and tested open source reference implementation for NFV infrastructure.
- **Open Daylight** provides a common open source framework and platform for SDN, including a network-resource controller and utilities for management, security, and virtualization.
- **OpenStack** is an OS and a set of related tools for building and managing private, hybrid, or public clouds, including the ability to control pools of processing, storage, and networking resources.
- **Internet Engineering Task Force (IETF)** advances standards and concepts such as service function chaining (SFC), which creates complex communication services using specified sequences of virtualized network functions.
- **Open Networking Foundation (ONF)** has been established to develop and promote standards specifically related to SDN, beginning with the OpenFlow standard, which was first published in 2011.

5 Intel® Platform-Level Technologies for SDN and NFV

A number of platform-level features and capabilities enhance the performance, reliability, and security of the Intel ONP for Servers reference design, which is architected on Intel® Xeon® or Intel® Core™ processor-based systems. As mentioned elsewhere in this paper, the technologies described in this section are open source and optimized to work on open standards; therefore, they are applicable to SDN and NFV implementations outside the scope of the reference design.

5.1 Intel® Data Plane Development Kit (Intel® DPDK)

The set of software libraries comprising Intel DPDK can be used to dramatically accelerate packet processing by software-based network components, for greater throughput and scalability.

Engineering teams from Intel and Wind River collaborated to enhance the performance of Open vSwitch on large numbers of small packets by producing a code fork that supports a subset of the overall switching functionality with several modifications, notably replacing the data-plane switching logic with a new version built on top of Intel DPDK.

The resulting code, which is called the Intel DPDK Accelerated Open vSwitch, offers dramatic improvements to packet-switching throughput and is offered as a reference implementation for use with Intel ONP for Servers and other NFV solutions. Moreover, for developers at TEMs, OEMs, and other organizations that wish to extend this work or pursue similar development on their own, the Intel DPDK libraries are freely available as source code, along with use-case examples that are intended to simplify solution development. Key software-library components of Intel DPDK include the following:

- **Environment Abstraction Layer** provides access to low-level resources such as hardware, memory space, and logical cores using a generic interface that obscures the details of those resources from applications and libraries.
- **Memory Manager** allocates pools of objects in huge-page memory space. An alignment helper ensures that objects are padded, to spread them equally across DRAM channels.
- **Buffer Manager** significantly reduces the time the OS spends allocating and de-allocating buffers. The Intel DPDK pre-allocates fixed-size buffers, which are stored in memory pools.
- **Queue Manager** implements safe lockless queues (instead of using spinlocks), allowing different software components to process packets while avoiding unnecessary wait times.
- **Flow Classification** incorporates Intel® Streaming SIMD Extensions (Intel® SSE) to produce a hash based on tuple information, improving throughput by enabling packets to be placed quickly into processing flows.

5.2 Intel® QuickAssist Technology

Hardware-based acceleration services for workloads such as encryption and compression are supported by Intel QuickAssist Technology, using an open-standards approach that is well suited to use with Intel ONS for Servers and other SDN and NFV implementations. An accelerator abstraction layer provides a uniform means of communication between applications and accelerators, as well as facilitating management of acceleration resources within the control and orchestration layers of the Intel ONP architecture. Intel QuickAssist Adapters are available as PCI Express* Gen 3-compliant cards that support functionality such as the following:

- **3G and 4G LTE encryption algorithm offload** for mobile gateways and infrastructure.
- **VPN traffic acceleration**, with up to 50 Gbps crypto throughput and support for IPsec and SSL acceleration[3].
- **Compression/decompression** up to 24 Gbps throughput[3].
- **I/O virtualization** using PCI-SIG Single-Root I/O Virtualization (SR-IOV).

5.3 Intel® Virtualization Technology (Intel® VT)

Capabilities at the virtualization layer itself also play a key role in the robustness of SDN and NFV implemented on Intel architecture. Enablement by Intel for all of the major virtualization environments—including contributions to open-source projects and co-engineering with providers of proprietary offerings—provides robust support for Intel VT. The hardware assists for virtualization offered by Intel VT dramatically reduce overhead, by eliminating the need for software-based emulation of the hardware environment for each VM. As a result, Intel VT enables higher throughput and reduced latency. It also enhances data isolation between virtual machines (VMs), for greater security. Some particularly significant Intel VT features in the context of SDN and NFV include the following:

- **Extended Page Tables** accelerate memory virtualization with a hardware assist to the mapping between virtual memory spaces and the corresponding physical memory locations.
- **Intel® VT for Directed I/O** supports the creation of protection domains, which are isolated sections of physical memory to which read and write access are restricted exclusively to I/O devices that have been explicitly assigned to them.
- **PCI-SIG Single-Root I/O Virtualization (SR-IOV)** implements virtual instances of a physical network interface that can directly transfer data to assigned VMs while bypassing the hypervisor's virtual switch, dramatically increasing throughput, reducing latency, and enhancing isolation of the data stream.

6 Adoption Paths for SDN and NFV on Intel Architecture

In general, Intel ONP for Servers is a hardware/software reference architecture based on open software ingredients and supported by commercial software solutions, system integration alliances, and industry standards bodies targeting high-performing SDN and NFV solutions.

In many cases, those parties may select only a subset of the reference design's components, which they may combine with other elements and approaches to create solutions that are influenced by the Intel ONP to varying extents. Such approaches are entirely in keeping with the open nature of the architecture, which focuses on offering choices, rather than prescribing requirements. Whatever their role, and however closely or loosely their solutions follow the reference design, parties that evaluate Intel ONP for Servers have access to the same fundamental benefits:

- **Simplicity.** Intel ONP for Servers is easy to evaluate, develop, and deploy, providing a fast path to insight about design approaches and capabilities for SDN and NFV architectures.

- **Flexibility.** Adherence to standard software approaches and choice among a wide array of software options makes Intel ONP's underlying architecture readily modifiable.
- **Cost-effectiveness.** The use of industry-standard servers offers a foundation for solutions that have advantages in terms of low TCO.

6.1 Hardware Manufacturers (TEMs and OEMs)

Some TEMs and OEMs produce and distribute equipment that closely follows the Intel ONP for Servers reference design, benefiting from the engineering efficiencies of doing so. Others selectively choose only some elements, for example incorporating specific Intel software optimizations (such as the Intel DPDK Accelerated Open vSwitch) into their existing platforms that may be based on proprietary hardware and protocol stacks and may target specific SDN and NFV use cases.

Whatever option they choose, TEMs and OEMs can identify opportunities to accelerate development, reduce costs, and improve alignment with the software ecosystem.

6.2 Software Vendors and System Integrators

Software vendors and system integrators may choose to deploy Intel ONP for Servers in its native form, possibly as a proof of concept to gain better understanding of potential opportunities for developing solutions that support SDN and NFV. They may also switch out specific software building blocks to determine their level of interoperability or to benchmark various configurations against each other. Another possibility would be to use Intel ONP's core architectural approach, but with significant changes to produce a distinctive solution that uses elements of a specific software stack or that targets a particular set of use cases.

Across that spectrum of options, software vendors and system integrators can use Intel ONP for servers as a vehicle to evaluate various software for use in SDN and NFV solutions, including open-source offerings and components.

6.3 Service Providers and Data-Center Operators

End customers such as telecom carriers, cloud service providers, and data-center operators may use the Intel ONP for Servers reference design in as close to a standard form as possible. Doing so could be valuable, for example, to support usages such as a proof of concept, a test bed for pre-purchase assessment of equipment or software, or for collaborative development of industry standards.

Whether these organizations choose one of those approaches or something else entirely, the Intel ONP for Servers reference design affords them the means to rapidly determine the potential value of an SDN and NFV approach based entirely on open standards.

7 Conclusion

Network transformation through SDN and NFV is being driven forward by open-source software and open standards. Intel architecture is a common foundation of choice for innovation in this area, and the Intel ONP for Servers reference design is a key enabler for the ecosystem.

Telecom carriers, cloud service providers, and data-center operators are drawing on this flexible reference design to test and deploy SDN and NFV networking, with an eye toward accelerating adoption, mitigating risk, and reducing costs.

Learn more about the Intel® Open Network Platform:
www.intel.com/ONP

[1] Gartner, March 24 2014. www.gartner.com/newsroom/id/2688717.

[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] Configuration: [Intel® Communication Chipset DH8955 PCI Express* x16 in an Intel® Xeon processor E5 v2 platform with Intel® QuickAssist Driver/SDK 0.30; Measured by Intel]. For more information go to www.intel.com/performance.