



Cisco Systems and Intel Corporation NFV Partnership



This document details the first steps in the strategic partnership between Intel Corporation and Cisco Systems Inc. to make Network Function Virtualization (NFV) part of a flexible, open and successful transformation in service creation for their Service Provider customers.

Table of Contents

Goals of This Document.....	3
Audience	3
Challenges and Objectives.....	3
Our Partnership Today (OpenStack and OpenDaylight)	5
OpenStack and GBP	6
OpenStack and Enhanced Platform Awareness	7
NFV-specific OpenStack EPA Extensions.....	9
OpenDaylight and SFC	9
Our Partnership Today (ONP and DPDK).....	11
Cisco VPP, VNFs and Intel DPDK.....	13
Cisco Joins Intel Network Builders.....	15
Looking Ahead At The Partnership	15
Conclusion	16

Goals of This Document

The purpose of this architecture document is to clarify our perspectives on the roles of jointly supported NFV Infrastructure (NFVI) components, our shared interest in making performance improvements and best practices available through these infrastructure components to our joint customers, our work to unify testing of NFV performance and how we cooperatively pursue our common vision in elevating the discussion of NFV beyond mere virtualization. In doing so we will highlight areas where our companies currently collaborate and where we think our partnership will take us in the future. [A complete description of both company's NFV and NFVI strategy & solution is out of scope for this document, as the intention is to focus on the areas where both companies are collaborating today.](#)

Audience

This document is intended to provide architectural and technical direction to network architects, service designers and strategic planners in our Service Provider and Enterprise operator communities, our channel partners and integrators interested in truly transforming businesses through service creation via NFV.

While the virtualization aspects of NFV have the potential to spawn new service creation, the number of performance enhancement vectors available in virtualization, the current focus on these optimizations, lack of clarity around how to integrate the virtual with existing network components (brown field deployments), and a lack of direction on how to create multiple function services currently limit the impact of NFV. The direction and recommendations resulting from our joint work show consumers how to deal with the ongoing variance in performance and “level-up” their focus to service creation.

Challenges and Objectives

The current dialogue around NFV seems to focus largely on the short-term; use cases that tend to be emulation of existing services on virtual infrastructure using Commercial Off The Shelf (COTS) components and optimizations of individual aspects of that virtual infrastructure. Ultimately, while both Cisco and Intel have joint interests in the fundamentals of NFVI (including pursuit of optimal performance), and there is potential immediate savings in the short-term optimization of existing services, we agree that it is a new service creation model and architecture (on an open framework) that will deliver enduring value.

From a virtualization perspective, the Virtual Network Function (VNF) components can be located in the Guest or Host OS (multiple choices and combination of OS) as can the Virtual Forwarder¹ (possibly multiple choices – both open source and commercial) and the VNFs can be run independently or chained. When combined with resource-aware placement options (platform architecture awareness), better functionality, performance and economics (i.e. lower

¹ A virtual forwarder can provide network services to applications (VNFs) that range from a simple Forwarding Information Base (FIB) to a full network stack.

cost per unit of workload) can be achieved! However, testable deployment permutations may also increase. Many such permutations may require specific actions in/on the Guest or VNF itself to be fully optimized.

Far from a static environment, NFV architecture may be on the move from application packaging as a Virtual Machine to include containers and bare metal. The processor environment will undoubtedly continue to evolve. And the incorporation of even more hardware acceleration – both specialized and generic (NIC enhancements, FPGA incorporation, GPU investigation, etc.) is imminent.

Service providers face the quandary of deploying the most cost-effective virtualization solutions “today” while maintaining both options on future enhancements (a future-proof “confidence”) and the “open” environments they seek.

To compound the situation, systematic approaches to understanding how the entire solution stack might interact (tradeoffs) as well as repeatable, standardized test methodologies do not yet exist. In this early stage of NFV, performance claims are proliferating even as the hardware and software infrastructure of NFVI roils in a continuing and rapid cycle of improvements.

While it is currently possible to implement a highly optimized standalone VNF (like a virtual firewall) through the lens of potentially transient enhancements to specific components (e.g. custom tweaks to hypervisor environments), this approach might limit placement of the VNF, may negatively affect components of other VNF co-located on the same physical platform and may also be so vendor specific that it compromises openness. A coordinated approach to VNF placement that is platform aware (capabilities as well as resources availability), capable of supporting multi-tenancy and providing resource guarantees, open and standards-based will usher in a more open multi-vendor VNF eco-system. What is called for is simultaneous cost saving enabled by optimized NFVI coupled with business transformation through new policy and service abstractions that enable service creation.

To refocus the NFV conversation, our jointly pursued direction proposes to:

- Enable the linkage between Management and Orchestration (MANO) and NFVI, so that more compelling services can be integrated into existing environments.
- Create and support the abstractions that foster Next Generation service enablement through the use of Policy and Service Function Chaining (SFC).
- Enable optimum performance by exposing to the Orchestration and MANO the particulars of performance management and platform architectural awareness through open source contributions to NFVI projects.
- Normalize our approach and recommendations around performance optimization and testing.

This paper will reference several examples of where Intel and Cisco are specifically collaborating to evolve many of the open source and standard software ingredients that are foundational to the Intel® Open Network Platform (ONP) reference architecture including co-authoring new open standard contributions such as Network Service Header (NSH) to enable strategic capabilities such as service function chaining for the SDN/NFV market.

OpenStack and OpenDaylight are key ingredients of ONP that provide an open “hub” of functionality in the NFVI. On top of this open hub, Cisco Systems builds value-add through their service functions (VNFs), orchestration and management applications.

Our Partnership Today (OpenStack and OpenDaylight)

Cisco and Intel have been partnering in enhancing NFV Infrastructure management using two major open source platforms, OpenStack and the OpenDaylight (ODL) project.

OpenStack is a leading open-source software suite for creating private and public clouds. The functionality provided can, for the most part, be classified under similarly intentioned names such as Cloud Operating System or Virtualized Infrastructure Manager (VIM). Working within the OpenStack community, Intel and Cisco help advance the project to create an open source infrastructure for enterprise and service provider IT cloud service delivery.

Both companies align around the belief that while OpenStack is a critical component of an open NFVI, it can't be burdened with delivering the entirety of the technology required for Virtual Infrastructure Management (VIM), particularly if we are to break the current networking constructs to add abstractions that allow us to build a service plane.

Network control is rich and complex. While OpenStack is trusted with providing orchestration services, a network controller is required for providing the control layer services (more real time configuration, management functions). OpenDaylight provides a framework through which OpenStack can direct the network operation (e.g. Neutron ML2 interfaces) or the network application can directly control decisions (which is particularly useful for experimentation, before these NFVI components are synchronized). OpenDaylight is a modular system composed of many components including; a network virtualization controller, multiple services (e.g. topology), common data store, basic network services as well as northbound open APIs. It offers a very rich southbound plug-in architecture supporting a variety of popular network protocols ranging from OpenFlow and OVSDDB (for server Open vSwitch control), to NETCONF, BGP, PCEP and LISP. The Northbound (programmatic) and Southbound (implementation) interfaces are clearly defined and documented APIs through which the application developer can interface with the controller to take advantage of the rich set of protocols and services within OpenDaylight. Multiple applications, and 3rd party distributions of OpenDaylight are available that are designed to enable SDN application abstractions². As Platinum Members of OpenDaylight, Cisco and Intel are actively shaping the functional development, usability,

² Cisco Systems will be demonstrating its Cisco Open SDN Controller (commercial distribution of OpenDaylight software) at Cisco Live US San Diego.

stability, scalability and performance of ODL to accelerate its commercialization and deployment.

One of the advantages of this loose coupling between OpenStack as an orchestrator and OpenDaylight as the network controller is the use of appropriate tools to deliver new functionality in a more timely fashion with purpose made tools with specific strength areas. While it is possible to attempt the abstractions we've introduced with proprietary solutions, openness at these critical points in the architecture allows the power of a community to deal with the inevitable architectural changes and velocity of NFV better than any single vendor might. Supporting open solutions here eliminates some less desirable outcomes for the community: dilution of effort (in maintaining drivers or other hooks to multiple infrastructure solutions within the VNFs and other infrastructure components) or barrier to entry (lack of access to appropriate integration API).

To maintain this open environment, forking of open source projects should be avoided (forking essentially puts a vendor in the same position as a proprietary solution provider).

Within these confines (of using open source and avoiding forking), Intel and Cisco have also collaborated on advancing Policy interaction between the “application” and the infrastructure. Some key examples are Group Based Policy (GBP) and Service Function Chaining (SFC) in OpenDaylight.

OpenStack and GBP

Group Based Policy provides top-level abstraction in which ‘nodes’ are groups of endpoints and edges are the directional policy (contracts) between them.

Policy specifies the semantic ‘what’ or intent for network flows and object relationships. These policies can be layered and their rule sets are re-useable, supporting inheritance and the concept of redirection (graph abstraction). These characteristics reduce the time to implement application requirements, providing a contract of features and enabling service assurance. A separate white paper³ describes GBP in more detail.

³ https://wiki.openstack.org/w/images/a/aa/Group-BasedPolicyWhitePaper_v3.pdf and https://wiki.opendaylight.org/view/Service_Function_Chaining:Group_Based_Policy_Integration

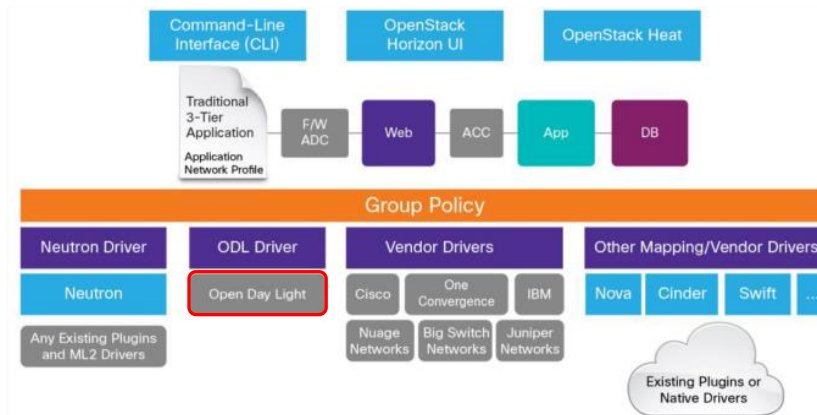


Figure 1 The architecture of Group Based Policy in OpenStack.

The architecture contains a number of plugins (more are in development) that provide policy mappings: to other OpenStack projects, specific vendor's products, Neutron specific structures and the OpenDaylight project.

The construction of GBP for OpenStack was a project generated from a blueprint that was the collaborative work of Intel, Cisco Systems and others. It continues to be an area of mutual interest and collaboration.

OpenStack and Enhanced Platform Awareness

Enhanced Platform Awareness (EPA) in OpenStack enables deployment of NFV virtual machines onto server platforms with the specific hardware and silicon capabilities that are optimal for the particular needs of the VM. This helps ensure the application is provided the best combination of features available across the spectrum of available systems in the data center. This is key to SLA adherence (including performance and functionality, high infrastructure utilization level and reduced cost per unit of workload).

Intel and Cisco are collaborating on contributions to a number of core and top level projects to enhance the server data plane for the specific needs of network applications. Combining packet processing optimizations such as DPDK with hardware acceleration capabilities on the Open vSwitch running on OpenDaylight will deliver new levels of performance and programmability on the server for example. Adding advances in OVSDB, Service Function Chaining, Group Based Policies, and Dynamic Resource Reservation will further grow the stability, scalability, security and performance of the OpenDaylight platform.

Intel's contributions to OpenStack are designed to match VNF specific requirements with the multiple hardware and software features of a modern IA server platform. It enables the VNF through the use of OpenStack to take advantage of the myriad performance optimizations of the

evolving Intel Architecture (IA), while relieving the VNF from having to explicitly configure the platform and ensure the features are engaged as desired. Current Intel contributions (to OpenStack) with this intent include:

- Trusted Compute Pools for Security and Compliance for Virtualized Servers - Intel TXT⁴ provides high value by ensuring trust in the platform through verification and attestation of launch time components.
 - **Trusted launch** is the basic verification of platform integrity, with lower risk from critical system malware and reducing support costs and data breach risks.
 - **Trusted pools** are the aggregation of multiple trusted systems and enabling platform trust status as a data point for security applications for control of workload assignment – such as restricting sensitive VMs to only run on trusted systems.

With these features Intel gives customers more visibility and control they seek for their clouds. Cisco Systems has collaborated in the development of these concepts and has demonstrated this functionality at both Intel Developers Forum and Cisco Live Milan (2015).

- Hardware-Accelerated Encryption/Decryption - Another key OpenStack software capability is the ability to recognize and take advantage of platform features. This becomes especially important in a typical datacenter having multiple generations of servers, with increasingly advanced features on the newer ones. It is very important that certain workloads only be scheduled to run on servers with these advanced features e.g. taking advantage of the crypto processing capabilities on Intel platforms. The Nova scheduler can assign crypto-processing workloads to platforms with support for Intel[®] Advanced Encryption Standard New Instructions (Intel[®] AES-NI), enabling customers to Protect Data in Motion with Up to 10x faster encryption and sensitive workloads are run on enabled platforms with Intel AES-NI. Workloads that do not have these requirements could be scheduled on older platforms. Even though support for AES in the hardware is now available on competing platforms, Intel's implementation of AES-NI still provides a performance edge for end-user benefit.
- Intelligent Workload Scheduling - INTEL[®] platform technologies that are exposed for intelligent scheduling:
 - Intel[®] AES-NI¹ for high-speed encryption/decryption.
 - Intel[®] AVX² and Intel AVX 2.0 for high-speed vector, floating point, and integer computations.
 - Intel[®] QuickAssist Technology for accelerating cryptography and data compression.
 - Intel[®] Quick Sync Video for high-speed transcoding of certain video codecs.

⁴ <http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/malware-reduction-general-technology.html>

- Intel® Xeon Phi™ coprocessors for massively parallel technical computing.

NFV-specific OpenStack EPA Extensions

From an NFV perspective, OpenStack does not need to be innately aware of the NFV applications or their functions. However, OpenStack does need the ability to provide a suitably advanced selection of tuning capabilities that enable a service provider to deploy NFV with the necessary performance and efficiency characteristics. OpenStack features and capabilities are developing rapidly. This section describes some of the features that have a particular applicability to deploying NFV workloads effectively.

CPU Feature Request - Some NFV applications may have been developed to make specific use of certain CPU instructions. For example, a VPN appliance that requires a high-performance cryptography library could be coded to leverage specific instructions such as the Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI). Best practice guidelines would recommend that software developers check for the availability of instruction set extensions via the cpuid instruction before invoking code that leverages them.

In the OpenStack Icehouse release a change was made to the Nova libvirt driver to expose all of the CPU instruction set extensions to the Nova scheduler. This correlated with a related change in libvirt to make this data available via the libvirt API. These changes made it possible to create Nova flavors that contained specific feature requests by adding these to the flavor as extra specs. In turn this enables the scheduler to place the VM only on those platforms that possess the desired feature.

NUMA Extensions - Awareness of NUMA topology in the platform was added in the OpenStack Juno release with the Virt driver guest NUMA node placement and topology extension. This feature allows the tenant to specify its desired guest NUMA configuration. The Nova scheduler was extended with the numa_topology_filter to help match guest NUMA topology requests with the available NUMA topologies of the hosts. Proper NUMA setting ensures the code is adjacent to the CPU core executing it, with significant performance implications.

SR-IOV Extensions - The OpenStack Havana release included support for SR-IOV for non-networking devices. This included the ability to allocate the PCIe device VFs to the VM for PCIe cryptographic accelerators such as Intel® QuickAssist Technology. In the OpenStack Juno release, this capability was extended to include support for network devices (e.g. NICs). By doing so, the highest performing I/O path from physical NIC to the VM was enabled.

OpenDaylight and SFC

Service Function Chaining is the key to a service layer abstraction (below the policy layer) that provides the semantic “how” for service graph traversal. With SFC Nodes are network functions (physical or virtual) and edges indicate the direction, order and sequence of the flow

of traffic through those chains. SFC allows the processing of network traffic through an ordered set of services, in a given environment, with dynamic and fine granularity (e.g. per-Flow).

This is enabled through adoption of the pending IETF SFC Work Group standard for a Network Service Header (NSH)⁵. Cisco Systems and Intel worked collaboratively on this standard (coauthoring architecture and protocol specifications) and have jointly committed to hardware-based forwarding support, including hardware offloads for header processing on Intel NICs (Network Interface Cards). SFC benefits include:

- The service path is decoupled from transport header.
- The NSH header preserves the end-to-end service path for troubleshooting and verification.
- The encapsulation may carry metadata enabling new inter service Function interactions and saving reclassification in multiple cases.

Of all of the benefits, the ability to carry metadata in-band offers the greatest benefit to NFV and service creation going forward. Metadata can be imputed (in-band measurement), a VRF context, a user context (UserID – in combination with a trustsec authentication system), an application ID (inserted by the classification step at the head of the chain) or an intermediate result (for future services that are envisioned as a series of peer process sharing a common memory space and processor).

Metadata passing is an important tool in the imagination of new services or the embellishment of existing services that are currently stuck in “emulation mode”. Without this tool, NFV functions will remain non-communicating blocks and the services created with them will remain limited.

Because OpenStack networking (Neutron) feature growth was slowing as the OpenStack project grew and given the potential complexity of service chain rendering and the amount of resulting state that may need to be managed, Cisco Systems and Intel worked together to make SFC available to VNFs directly from OpenDaylight⁶.

⁵ http://datatracker.ietf.org/doc/draft-quinn-sfc-nsh/?include_text=1 also see Paul Quinn and Jim Guichard’s article in IEEE Computer (Vol. 47 Issue 11)

<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?reload=true&arnumber=6965275>.

⁶ https://wiki.opendaylight.org/view/Service_Function_Chaining:Main

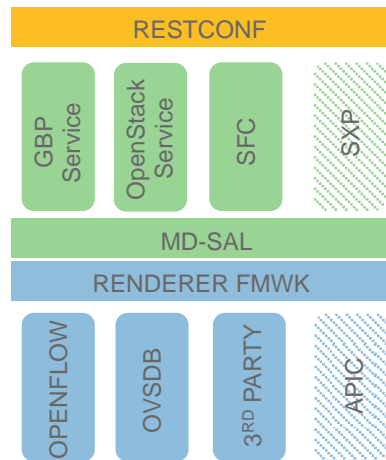


Figure 2 OpenDaylight Policy/SFC rendering architecture.

The resulting architecture in ODL is “multi-renderer”. Renderers are provided for OpenFlow and Open vSwitch (OVS) in the Lithium release, enabling SFC functionality on Intel servers and Cisco also has provided renderers for physical equipment and virtual routers as well as proposed expansions that leverage trustsec-based infrastructure (e.g. SXP – Security Group Tag Exchange Protocol⁷) in development.

Both companies are working to add SFC capabilities to OpenStack Neutron with the common interest of eventually providing a more holistic Policy solution for OpenStack.

At Mobile World Congress 2015, Intel and Cisco jointly demonstrated SFC with NSH running on the Open vSwitch optimized on DPDK with an ODL control plane. This was all running on Cisco UCS servers powered by Intel® Xeon Processors, and Intel Ethernet Technology such as the Intel Ethernet 100GbE SDI Adapter. The main point of the demonstration was the ability to create a wire-speed service chain by forwarding traffic through a series of distributed VMs using destination information in the NSH header.

The overall functionality has also been demonstrated in an ETSI PoC (in conjunction with NTT Communications) and will be jointly demonstrated (Intel, Cisco Systems, F5 Networks and Citrix Systems) at Cisco Live US (San Diego June 10, 2015) in the DevNet Zone.

Our Partnership Today (ONP and DPDK)

Intel® Open Network Platform (Intel ONP)⁸ is a reference architecture that provides engineering guidance and ecosystem enablement support to encourage widespread adoption of SDN and NFV solutions in Telco, Enterprise and Cloud. Intel ONP is one of the many ways Intel accelerates SDN/NFV adoption and deployment. It is not a commercial product, but a reference architecture that delivers proven software, tested on Intel servers to the telco, enterprise IT and cloud markets to enable SDN/NFV deployment. Intel ONP reference architecture brings

⁷ <https://datatracker.ietf.org/doc/draft-smith-kandula-sxp/>

⁸ <https://networkbuilders.intel.com/onp>

together the use of Industry Standard High Volume Servers (SHVS) based on Intel® Architecture (IA) and a robust software stack composed of open source, open standard software building blocks to enable a broad range of interoperable commercial solutions for SDN and NFV.

Intel ONP will continue to evolve as a reference architecture that contributes innovation with partners into OPNFV and other parallel consortia (e.g. OpenStack, OpenDaylight, Open vSwitch, IETF, etc...) Intel ONP will be based on OPNFV for the Telecom vertical starting with the first Arno release this year (note however, that as OPNFV is confined to use APIs of the above stacks, ONP addresses the needs of enterprise and cloud too).

Intel's Open Network Platform (ONP) projects include an initiative to accelerate packet processing with the Data Plane Development Kit (DPDK)⁹. Intel presents DPDK, Intel® QuickAssist Technology, Intel® Virtualization Technology (Intel® VT) and other recommended performance enhancing technologies in depth in their Network Builders white paper "A Path to Line-Rate-Capable NFV Deployments with Intel® Architecture and the OpenStack® Juno Release"¹⁰.

The DPDK is a set of libraries and drivers that take advantage of Intel® instruction set architecture features to accelerate packet processing on x86 platforms. The libraries include:

- 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.

Intel displays the advantages of DPDK through their contributions to OVS¹¹. Open vSwitch (OVS) is a virtual multilayer network switch. OVS with DPDK-netdev, available in OVS 2.3 in experimental mode, provides an implementation which enables better performance of the data plan when using DPDK libraries.

The DPDK performance enhancements to the server data plane for network workload processing are now integrated into OvS, ODL and OpenStack. Open vSwitch (OvS) is a leading open source

⁹ <https://01.org/packet-processing/intel%C2%AE-onp-servers>

¹⁰ https://networkbuilders.intel.com/docs/PathToLine-Rate_WP_V1.pdf

¹¹ <http://openvswitch.org/>

project for virtual switching. OvS is also supported by ODL, thus allowing ODL to configure, manage and utilize a plain OvS or DPDK enhanced OVS delivering high performance server data plane for demanding applications. The net result is that the user or the “application” can benefit from an automated DPDK configuration of the server data plane to match its performance needs.

Another joint initiative of Cisco and Intel is an emerging API allowing for hardware assisted policy controlled vSwitch processing. This set of API allows capable NIC hardware to process vSwitch tables including overlay and NSH encapsulation and decapsulation and additional policy rules (e.g. ACL expressed by GBP, for instance). It will further enhance the capabilities to handle latency and jitter sensitive workloads and lower the platform resource consumption for these network centric processing.

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.

Cisco VPP, VNFs and Intel DPDK

Cisco Vector Packet Processing (VPP) is integrated with Intel’s Data Path Development Kit (DPDK) providing optimal packet processing performance. Cisco VPP is an innovation from Cisco that is a full featured networking stack with highly optimized software forwarding engine enabling very lightweight, multi-tenanted and high performance software data plane residing on the x86 servers. Compared to a traditional scalar forwarding engine in the x86 environment, Cisco VPP allows for parallel processing of multiple packets to maximize performance, allowing

up to 10Gbps of throughput on a single CPU core. The key benefits of the Cisco VPP technology include:

- **Highly optimized packet processor for general-purpose CPUs** – VPP is a high performance packet processor stack for general purpose CPU's such as X86, PowerPC and MIPS. By using a “vectors” each comprising a large number of packet indices and processing the represented frames using a directed graph of nodes, VPP amortizes the cache operations associated with handling a single packet across many (instructions are loaded into the instruction cache (I-cache) once per vector and not once per packet thereby reducing the number of I-cache misses per packet).
- **Exploits temporal locality of application flows** – Cisco VPP leverages the temporal locality of application flow to reap the most benefit from both Instruction cache (I-cache), and Data-cache (D-cache) hits.
- **64-bit, multi-threaded, endian clean** – Cisco VPP is based on 64-bit application and is fully multi-threaded to make full use of available compute resources.
- **Industry's first user space, multi-tenanted, line rate packet forwarder** - Cisco VPP is a full-featured and optimized software packet processor with Layer 2-4 networking stack embedded into it.
- **Supports service chaining with or without meta-data headers** – Cisco VPP supports the construction of Layer 2 or Layer 3 service-chains where the forwarding of traffic can happen based on the L2 or L3 lookup. Apart from this, VPP also supports for a very flexible and fast flow classification that is used to implement Network Service Header (NSH) based service chaining.
- **Deployment Flexibility in User Space** – Cisco VPP is a highly portable technology that can be deployed in the user space as a process or within a VM over a host kernel.
- **Sticky and Stateful load balancing of traffic** –The Cisco VPP can perform both these functionalities to allow seamless and elastic scalability of the virtual appliances in data center and NFV use cases.
- **Proven innovation** – Cisco VPP technology is proven and matured to handle packet processing in the x86 architecture. Cisco Systems has been using VPP technology for many years on its flagship Carrier Routing System (CRS) and ASR9000 routing platforms.

Cisco Systems is embracing the DPDK framework across its virtualized product lines to drive performance optimization for the virtualized services with high network IO requirements (e.g. Cloud Services Router 1000v (CSR1Kv), Virtualized Adaptive Security Appliance (vASA), Virtual Mobile Packet Core software (vPC), and Cisco IOS-XR 9000v virtual router). VPP and many of these VNFs are discussed and on display in the programs and demonstrations at Cisco Live USA 2015 (San Diego).

Cisco Joins Intel Network Builders

One of our most recent steps forward in our partnership occurred when Cisco Systems joined Intel Network Builders.

With the convergence of compute, network and storage (which is the key theme behind Software Defined Infrastructure as a superset of the SDN/NFV initiatives), Intel saw the need to enable an ecosystem that can bring together all of the different technologies and competencies required to deploy commercially viable and high performance NFV solutions.

Cisco's portfolio of compute, switching, routing and storage management enable them to rapidly create turnkey NFV solutions across multiple industries and at multiple positions in the network for extremely rapid service deployment with a high degree of automation. This convergence of network and cloud is the space where customers are looking to invest new capex, and having Cisco as part of the Intel Network Builders community is a natural evolution of the Cisco-Intel partnership and brings a high degree of value for end customers.

Cisco will be demonstrating their new data center NFV-focused UCS (Unified Computing System) platform – called the Cloud Services Platform (CSP) 2100 in the Cisco booth as well as the Intel booth at Cisco Live San Diego (June 2015). The CSP 2100 is a turn-key hosting O/S and hardware platform optimized for data center NFV. The platform facilitates the quick on-boarding and lifecycle management of any Cisco or 3rd party VNF. It includes a GUI for quick out-of-the box management, REST API and NetConf support for automation, clustering for high-availability, and Intel-enabled performance enhancements. Together Cisco and Intel will be demonstrating how key ONP ingredients such as OpenDaylight, Open vSwitch, NSH and DPDK can all be used together to enable dynamic service chaining running at line rate speeds all on Cisco UCS servers.

Looking Ahead At The Partnership

As we look to the future, Intel and Cisco are working closely together to accelerate deployment of highly agile Software Defined Infrastructure in forums like OpenStack and OpenDaylight. There is much work to be done to provide the tools and capabilities for the network operators to embrace the IT cloud service delivery model and the accompanying drastic shift in operational culture. From single pane management of virtual and physical infrastructure to continuous service introduction and innovation through DevOps models, Intel and Cisco will continue working closely together to enable the path through technology introduction and trials to robust platform deployments.

We plan to continue our collaboration to create a rich, light-weight and balanced application and infrastructure interaction. On one hand, we'd like to see a complete and comprehensive

application policy driven through the NFVI stack to the infrastructure. Our joint work for adding SFC capabilities to OpenStack Neutron with the common interest of eventually providing a more holistic policy solution for OpenStack, OpenDaylight and OvS are key there. This enables the infrastructure to perform a key role in SLA enforcement and to contribute to higher utilization and lower cost per unit of workload. On the other hand, continue to work on Enhanced Platform Awareness to provide the orchestration layers data for best workload placement.

In exploring where our partnership can go, we see potential in our joint focus on NFV performance testing (Cisco Systems will soon publish a paper on their internal RFC2544-enhanced methodologies used in the development of their vMS offering¹²) in forums like OPNFV and Network Builders, cooperation around some NFV-specific management items (analytics, service assurance and resource management), continued work on security and also in the common interests of the Cisco DevNet and Intel Developer communities.

Conclusion

Cisco and Intel have partnered for over a decade not just on constructing products but also working together in open standards and open source. Examples include our partnership in the IETF for Service Function chaining with NSH and accompanying support in Intel Ethernet NICs silicon, defining an open API to accelerate OVS performance with Intel Ethernet NICs, DPDK implementation across a variety of Cisco platforms, collaboration in OpenStack and ODL to enable the evolution towards policy based infrastructure that can dynamically respond to application requirements, and OPNFV as a platform to bring together all of these innovations. Through its UCS product family, Cisco Systems is a natural partner of Intel. When it comes to NFV, the Cisco catalogue of VNFs, development of complimentary tooling for the architecture, ongoing investment, proven experience in NFV-based service creation and integration create compliments that can span the NFV solution stack. Our true partnership derives from how Cisco products consume Intel innovation and our joint focus/contributions to open source components that push the ultimate value in NFV forward – service creation.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)