Learn About Network Functions Virtualization

Network virtualization—the abstraction of network functions that allows for them to be installed, controlled, and manipulated by software running on standardized compute nodes—is fundamentally redefining the economics of networking.

Network Functions Virtualization (NFV) incorporates cloud and virtualization technologies to drive rapid development of new network services with elastic scale and automation. These emerging technologies are often grouped together as NFV and software-defined networking (SDN), and bring increased agility in delivering network services with improved capital efficiency by removing bottlenecks imposed by manual processes, and allowing new services to be deployed on demand. NFV allows service providers to deliver services faster and more cost-effectively, and to leverage automation so that they can readily adapt to customers' shifting needs for scale and agility.

This *Learn About* provides an overview of NFV, an evolving and transformational technology that is revolutionizing traditional networking, both at an architectural level and within the individual elements of a network.

NFV and the ESTI Model

Key to the success of NFV is the work that the European Telecommunications Standards Institute (ETSI) has done to develop an NFV architecture that focuses on describing the operation of virtualized networks, optimizes performance, and drives for an open, portable architecture. The NFV architecture is open for cross-vendor coordination and interoperability, while maintaining compatibility with existing hardware-based network architectures. Figure 1 shows a simplified view of the ETSI NFV architecture.

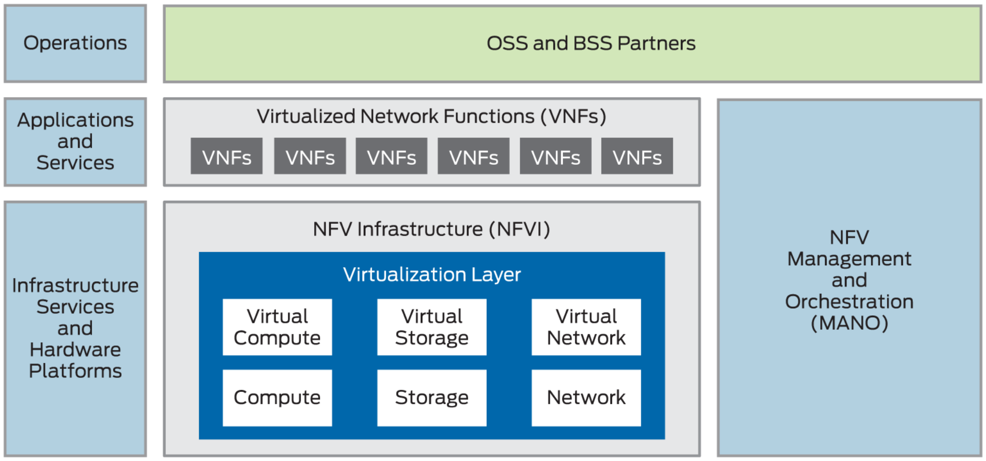
Operations
OSS and BSS Partners
Applications and Services
Virtualized Network Functions (VNFs)
VNFs
VNFs
VNFs
VNFs
VNFs
VNFs
Infrastructure Services and Hardware Platforms
NFV Infrastructure (NFVI)
Virtualization Layer
Virtual Compute
Virtual Storage
Virtual Network
Compute
Storage
Network
NFV Management and Orchestration (MANO)

Figure 1 *ESTI Architecture*

The major components of this architecture include the NFV infrastructure (NFVI), VNFs, VNF managers, and management and orchestration (MANO). Some of the critical goals of an open NFV architecture include:

- Decomposition of physical network elements into virtualized network functions (VNFs) to allow operators to choose the best implementations from a range of vendors.
- Portability of VNFs to different hardware platforms and hypervisors.
- Rapid software-based service delivery.
- Standard, open interfaces for improved interoperability in multi-vendor NFV solutions.
- Use of low-cost, commercial off-the-shelf (COTS) hardware.

Virtualization Building Blocks

The modular architecture of NFV is what allows service providers to automate at every level. Let's further explore the fundamental components of that architecture.

To begin, NFV stands on the basic building blocks common to a virtualized data center: virtualized compute, storage, and networking. The compute is built from standardized x86 servers and uses hypervisors as the foundation for virtual machines and containers for smaller footprint and faster deployment of microservices. Virtualized storage is built from physical storage that can be split into virtualized storage resources for multiple virtual machines (VM). Finally, virtualized networks use software to define an overlay network that provides a level of flexibility and agility over the underlying physical network. Data center virtualization depends on cloud orchestration and an automation framework to manage the network operation and services deployment.

Hardware

The industry is moving to generic hardware that can be reused across a range of functions for network, security, and applications, and the goal for NFV is to deploy x86 servers with general purpose processors that provide the highest degree of flexibility with regard to the workloads of each VNF. These COTS servers provide plentiful central processing unit (CPU) cycles with large RAM and disk storage that can support many VNF instances in a single physical server. COTS servers are cost effective, have shorter lead times, and can be repurposed as needed. Select a physical server that supports hardware-based virtualization technologies, such as Intel's VT-x and AMD's AM-v. While there is a goal to reach a single hardware model for NFV, different demands for the hosted VNF may require consideration of the most appropriate amount of storage and RAM within each server. For high performance workloads, you can introduce combined virtual CPUs in a VNF to achieve higher throughput.

SR-IOV

Along with high storage and RAM, COTS servers also supply the physical network interface controllers (NICs). In these servers, not all NICs are created equal, and the physical hardware server should include NICs that support Single Root I/O Virtualization (SR-IOV) and Peripheral Component Interface (PCI) pass-through to bypass the virtualization layer and provide virtual networking speeds that are nearly as fast as the physical NIC. With SR-IOV, you can share a single physical NIC across multiple VMs and give those VMs direct access to the physical NIC. However, not all NICs support SR-IOV, so you should balance the benefits and the cost of an SR-IOV-capable NIC. See the SR-IOV Primer for full details.

DPDK

When choosing server hardware and the VNF application, remember that VNFs can be more data-plane intensive because of packet handling. Choose VNFs that support the Data Plane Development Kit (DPDK). DPDK is open source software that provides libraries and drivers for fast packet processing. It allows an application to bypass the kernel and Linux network stack and directly access packets on the NIC. In addition, DPDK uses fewer CPU cycles per packet and less memory context switching to speed up the overall time it takes to process each packet. See the DPDK documentation for complete details.

Virtualization and Bare Metal

Multiple options are available for hosting your VNFs. You can use virtualization technologies such as VMs or containers, or you can host your VNF on a bare-metal server (BMS). A BMS includes no host OS and thus speeds up the VNF performance. Alternatively, you could install a hypervisor on the BMS to allow you to host multiple VNFs as VMs.

Virtual Machines

VMs became popular in the virtualized data center as a way to allow multiple operating systems and applications to run on the same physical hardware. Many VMs can run on the same host, sharing its resources. A VM has its own operating system image, which can be independent from that of the host OS or other VMs running on the same host. Further, VMs improve the overall physical hardware utilization, significantly lowering costs. The virtualization layer, or hypervisor, that runs VMs abstracts the hardware from the VM and allows for mobility and portability of VMs across hardware platforms. Common hypervisor options include Linux KVM and VMWare’s ESXi family. And VMs can provide the modular building blocks to host VNFs such as firewalls and routers.

Containers

Containers are an emerging VNF option that allow you to run considerably more VNFs on a single physical host server than you can with VMs. They do not include their own operating system, but depend more heavily on the host OS. This allows containers to better optimize resource usage in terms of the memory, storage, and CPU footprint than the equivalent function in a VM. For example, a Linux server with ten VMs will include ten guest OSs, one in each VM. The VNF would then run on top of the guest OS. With containers, the VNF has access to the host OS, so that same Linux server can include more containers than VMs, bringing even more economies of scale to the physical hardware. Docker is a prominent container technology in the Linux space. See http://www.docker.com for more information. (The remainder of this *Learn About* uses the term VNF to describe a network function that has been virtualized using either a VM or container.)

Automation

The desire to automate the orchestration and management of network, storage, and compute resources is a key driver of development for NFV and SDN. This is particularly important since NFV increases the combination of elements that need to be managed. Imagine again the scenario that includes one physical server with 10 VMs or hundreds of containers. It becomes increasingly difficult for someone to manage at that scale, from fulfillment, installation, and configuration to monitoring, optimization, add/drop/changes, and fault isolation. If manual operations were required, this concept would never scale. The complexity of managing a heterogeneous environment drives the need for open, standards-based approaches to access and control each network and server component. Automation and cloud orchestration help solve that problem by programmatically controlling and monitoring – and in some instances, repairing or replacing – networking components without direct human involvement. With automation, you can rapidly spin up or destroy VNFs (as VMs or containers) to elastically scale your network functions to match dynamic demand.

Automation isn't just the process of migrating physical tasks into computer managed workflows; it's also a culture that links infrastructure services and product development with the best operational practices. This culture is often known as DevOps. See DevOps for Communications Service Providers for more details.

Some prominent automation technologies used in NFV solutions include:

- Heat Templates – Heat templates are human-readable text files that describe the infrastructure for cloud applications, such as networks, servers, floating IP addresses, and the like, and can be used to manage the entire life cycle of that application with cloud systems such as OpenStack and OpenContrail.
- REST APIs – Representational State Transfer (REST) application programming interfaces (APIs) enable you to securely connect to systems, execute remote procedure calls (RPCs), and use a variety of formatting and display options, including JavaScript Object Notation (JSON). REST APIs integrate the provisioning, configuration, and analytics of cloud orchestrators to north-bound Operations Support Systems and Business Support Systems (OSS/BSS) and to south-bound components of the full NFV solution.
- Yang and NETCONF – NETCONF is an IETF-defined XML management protocol that client applications use to request and change configuration information on networking devices. NETCONF uses the Yang data-modeling language that describes the desired network configurations. You can use NETCONF and Yang data models to design, configure, and monitor VNFs and physical network functions in your NFV environment.
- XMPP – An extensible Messaging and Presence Protocol (XMPP) is an open XML-based protocol used to communicate between networking components. It is increasingly being used for distributed control and data plane communications in SDN and NFV architectures.
- Ansible, Chef, and Puppet – These storage management (Chef) and server management (Puppet and Ansible) tools work with scripting tools to simplify automation through the use of templates or playbooks to create a standard, consistent, and repeatable deployment and orchestration environment.
- OpenStack – OpenStack is an open-source cloud operating system that automates the management of the compute, storage, and networking components of a cloud environment.

Networking in a Virtualized Environment

With the growth of virtualization in networks, considerable traffic now travels along service chains that are comprised of multiple VNFs. This east-west traffic makes use of overlay networks, which are virtual networks that use tunneling protocols to traverse the physical network infrastructure (underlay network). Overlay networks abstract a physical network from a virtual network and allow the virtual network to be programmatically provisioned in seconds. The use of an overlay network decouples the network configuration changes for virtual services from the configuration of the underlying physical network. This overlay network supports both Layer 2 and Layer 3 transport between VMs and containers. A number of tunneling protocols can be used to create these overlays – for example, Virtual Extensible LAN (VXLAN) is the most common overlay type. Other overlay network types include Multiprotocol Label Switching (MPLS), Network Virtualization using Generic Routing Encapsulation (NVGRE), and Stateless Transport Tunneling (STT).

Since virtualization is seldom complete in the network environment, there is still a need to connect BMSs and physical network devices into this overlay network. Choose physical network devices that support these overlay networks to interconnect between BMSs and the virtualized networks.

NFVI and VIM

The NFVI building block provides the virtualization layer (hypervisors or container management systems such as Docker), and the physical compute, storage, and networking components that host the VNFs. NFVI is managed through the NFVI infrastructure manager (VIM), which controls the allocation of resources (virtualized compute, storage, and networking) for the VNFs. OpenStack is an example of an open source VIM, controlling the physical and virtual resources that the VNFs use. VMWare is an example of a commercial VIM.

VNFs and VNF Managers

A VNF is a software-based application that provides one or more network services. Examples of VNFs include routers, firewalls, and intrusion prevention systems (IPS). These VNFs use the virtualized infrastructure provided by the NFVI to connect into the overall network and provide programmable, scalable network services.

VNF Managers support the lifecycle of VNF instances (instantiation, monitoring, and teardown) as well as the onboarding and management for each revision of a VNF software image.

MANO

Management and Orchestration (MANO) provides the overarching management and orchestration of the VNFs in the NFV architecture. MANO instantiates the network services through the automation, provisioning, and coordination of workflows to the VIM and VNF Managers that instantiate the VNFs and overlay networking service chains. MANO connects the NFV architecture with the existing OSS/BSS.

Juniper NFV Solutions

In keeping with the participatory nature of the NFV ecosystem, Juniper Networks provides an open NFV solution that is compatible with the ESTI NFV architecture, as shown in Figure 2.

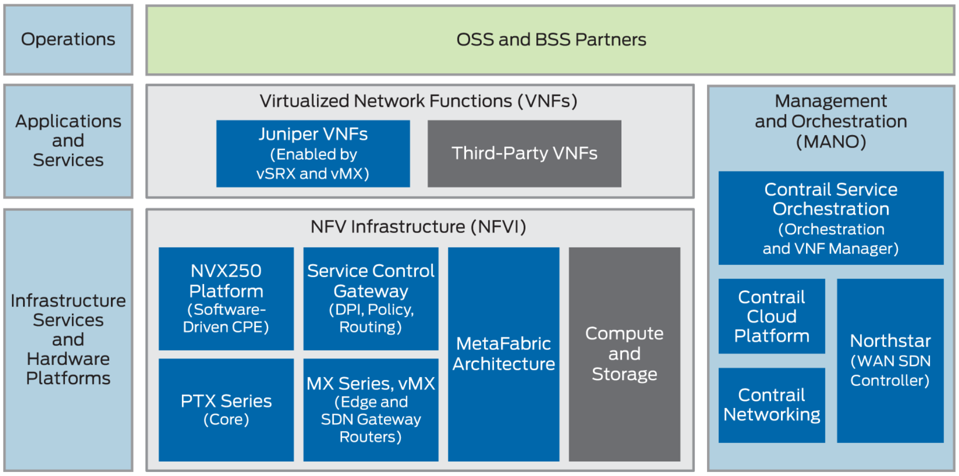
Operations
OSS and BSS Partners
Applications and Services
Virtualized Network Functions (VNFs)
Juniper VNFs (Enabled by vSRX and vMX)
Third-Party VNFs
Management and Orchestration (MANO)
Contrail Service Orchestration (Orchestration and VNF Manager)
Infrastructure Services and Hardware Platforms
NFV Infrastructure (NFVI)
NVX250 Platform (Software-Driven CPE)
Service Control Gateway (DPI, Policy, Routing)
MetaFabric Architecture
Compute and Storage
PTX Series (Core)
MX Series, vMX (Edge and SDN Gateway Routers)
Contrail Cloud Platform
Contrail Networking
Northstar (WAN SDN Controller)

Figure 2 *Juniper NFV Architecture*

Figure 2 shows the Juniper Networks products that make up the Juniper NFV solution in blue, and the building blocks in grey show the open third-party components that can be integrated into this solution. The Juniper NFV solution enables lower total cost of ownership (TCO) through automation and higher resource efficiency, while improving visibility and simplified operations with virtualization. The Juniper NFV solution uses an open platform that can be leveraged in a DevOps environment. The Juniper NFV solution incorporates:

- A programmable cloud reference architecture for MANO that leverages Contrail for a turnkey management and orchestration platform.
- MetaFabric™ architecture to provide an agile and intelligent foundation for cloud-based data centers that simplifies management of physical infrastructure and creates a more efficient pool of connected compute resources.
- Intelligent Services Edge to give visibility and control in SDNs and link the physical network and its elements into the NFVI without requiring rearchitecture of the existing physical transport network.
- VNFs enabled through vSRX and vMX products.
- Customer premises equipment to securely extend VNFs for functions that are best placed closer to the end users with the NFX250 Network Services Platform.

Physical Underlays for NFVI Components

Juniper's approach to NFV provides a range of options for NFVI placement that can expand the footprint and therefore improve the end-user experience of the services delivered using NFV. The Juniper NFV architecture uses the PTX Series and MX Series for core and gateway router functionality, along with the following products to provide the physical structure for the underlay network and to host VNFs:

- NFX250 Network Services Platform – A secure, automated, software-driven CPE platform that delivers virtualized network and security services on demand. As an integral part of Juniper's fully automated Cloud CPE solution suite for NFV, the NFX250 supports multiple Juniper and third-party VNFs on a single device. The NFX250 provides built-in security with vSRX, and dynamic, policy-based routing. See http://www.juniper.net/us/en/products-services/routing/nfx250/.
- Service Control Gateway (SCG) – Implemented on Juniper Networks MX Series 3D Universal Edge Routers, SCG utilizes subscriber awareness, deep packet inspection (DPI), and policy management to determine traffic treatment on a per-subscriber and per-application basis, enabling highly customizable and differentiated services at scale. Working with Juniper Networks Contrail Cloud Platform, the SCG can steer traffic into complex service chains hosted within the NFVI or on existing physical appliances. This granular understanding of traffic flows provides a rich set of data to analytics engines and back-office systems to permit real-time charging and end-user engagement at the application and content level. See http://www.juniper.net/assets/us/en/local/pdf/datasheets/1000540-en.pdf.
- The Juniper MetaFabric architecture is an approach to building cloud-ready data centers. See http://www.juniper.net/assets/us/en/local/pdf/reference-architectures/8030012-en.pdf.

VNFs – vMX and vSRX

When introducing services hosted in a virtual environment it is essential to have the ability to quickly migrate physical infrastructure and services into a VNF. Unfortunately, there are many external costs that make this migration expensive and risky. By delivering virtual implementations of the flagship routing and security functions, Juniper's NFV architecture ensures that existing systems interfaces, service and network architectures, operational tools and workflows, and software qualification processes can all be reused instantly to deploy these capabilities in the cloud. Juniper virtualized two of its flagship products as VNFs:

- vSRX – A virtual security appliance that provides security and networking services in virtualized private or public cloud environments. It runs as a VM on x86 servers that support virtualization, and provides firewall, robust networking, and advanced security services at Layers 4 through 7, with automated lifecycle management capabilities. See http://www.juniper.net/us/en/products-services/security/srx-series/vsrx/.
- vMX – A virtual MX Series 3D Universal Edge Router that is a full-featured, carrier-grade router with complete control, forwarding, and management planes. It runs Junos OS and supports vTrio packet handling and forwarding by compiling the programmable Junos Trio chipset microcode for x86 chipsets. See http://www.juniper.net/us/en/products-services/routing/mx-series/vmx/.

MANO

Juniper's MANO solution includes Contrail Cloud, Contrail Networking, Contrail Services Orchestration, and Northstar as the SDN WAN controller.

Contrail Cloud

Juniper's Contrail Cloud Platform is a complete turnkey cloud-management platform that is hardened and integrated from open source technologies including OpenStack cloud-management platform, Open Contrail, Ceph distributed storage system, and Puppet server management. The Contrail Cloud Platform automates the orchestration of compute, storage, and networking resources to create and scale open, intelligent, and reliable OpenStack clouds that seamlessly merge and hybridize through highly intelligent and secure networks. See http://www.juniper.net/us/en/products-services/sdn/contrail/contrail-cloud/.

Contrail Networking

Juniper's Contrail Networking automates network resource provisioning and orchestration to dynamically create highly scalable virtual networks and to chain a rich set of Juniper or third-party VNFs and physical network functions (PNFs) to form differentiated service chains on demand. Contrail Networking can be integrated with an existing cloud orchestration system such as OpenStack or CloudStack. See http://www.juniper.net/assets/us/en/local/pdf/datasheets/1000521-en.pdf.

Contrail Services Orchestration

Juniper's Contrail Service Orchestration (CSO) is a comprehensive management and orchestration platform that delivers virtualized network services built on an open framework. CSO consists of the following components:

- VNF Manager, which creates VNF instances and manages their life cycles.
- Network Service Designer, which enables design, creation, management, and configuration of network services.
- Network Service Orchestrator, which is responsible for ETSI-compliant management of the life cycle of network service instances. This application includes RESTful APIs that you can use to create and manage network service catalogs.
- Cloud CPE Tenant, Site and Service Manager, and its auxiliary component, Identity and Access Manager, which manage Cloud CPE customers and map each customer's network services to the appropriate gateway resources. These applications provide a northbound RESTful API that you can connect to OSS/BSS.
  See http://www.juniper.net/us/en/products-services/sdn/contrail/contrail-service-orchestration/.

NorthStar

In closing, Juniper's NorthStar Controller is an SDN WAN network controller that enables granular visibility and control of IP/MPLS tunnels in large service provider and enterprise networks. Network operators can use the NorthStar Controller to optimize their network infrastructure through proactive monitoring, planning, and explicit routing of large traffic loads dynamically based on user-defined constraints.
See http://www.juniper.net/us/en/products-services/sdn/northstar-network-controller/.

References and Resources

- Juniper Networks provides a carrier-grade NFV solution that combines a turnkey software cloud management system and a supporting reference architecture based on virtualization and programmability. See http://www.juniper.net/us/en/dm/big-nfv-idea/, and, http://www.juniper.net/us/en/solutions/nfv/.
- Juniper Networks provides a reliable open SDN solution that combines open software systems for improved business agility through network virtualization, orchestration, and automation. See http://www.juniper.net/us/en/products-services/sdn/.
- VXLAN is a tunneling protocol that encapsulates Layer 2 Ethernet frames in Layer 3 UDP packets, enabling you to create virtualized Layer 2 subnets, or segments, that span physical Layer 3 networks. See http://www.juniper.net/techpubs/en_US/learn-about/LA_VXLANinDCs.pdf.
- ETSI is an international organization that develops the NFV standards and shares NFV implementation experiences. See http://www.etsi.org/technologies-clusters/technologies/nfv.
- Docker is a Linux container technology that packages an application and the application dependencies into a portable software package. http://docker.com.
- The following white paper describes how to build large IP fabrics using the Juniper Networks QFX5100 line of switches. It covers why you would want to build a Clos IP fabric and how to build, configure, and verify Clos IP networks using the QFX5100 series. Example configurations are included, enabling you to architect and configure your Clos IP network based on the cited examples. See http://www.juniper.net/us/en/local/pdf/whitepapers/2000565-en.pdf.
- This Juniper Cloud CPE paper examines the market opportunity for cloud-based managed services and reveals how Juniper helps service providers transform product development cycles and service delivery and, ultimately, increase their revenues. See http://www.juniper.net/assets/us/en/local/pdf/solutionbriefs/3510561-en.pdf.

Learn About Network Functions Virtualization

by Sandra McCann and Helen Shaw

A group of emerging cloud and virtualization technologies are evolving to provide increased agility when delivering network services, promising to drive rapid development of new network services with elastic scale and automation.

Get ready for NFV, where service providers can deliver new services not only faster but with lower OPEX and CAPEX than ever before. Learn all about this evolving and transformational technology that is revolutionizing traditional networking.

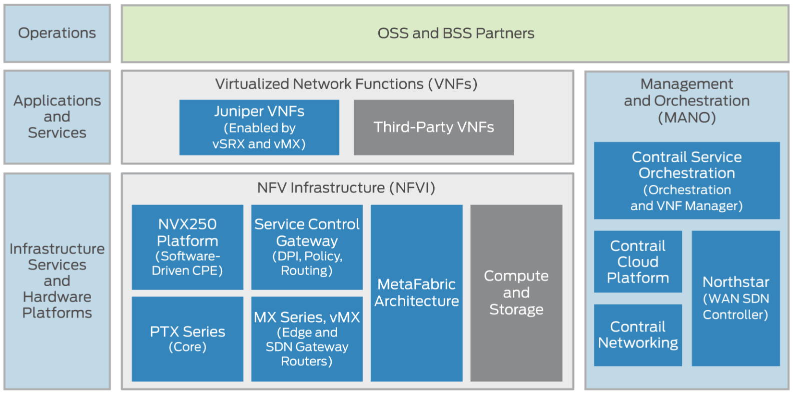
Operations
OSS and BSS Partners
Applications and Services
Virtualized Network Functions (VNFs)
Juniper VNFs (Enabled by vSRX and vMX)
Third-Party VNFs
Management and Orchestration (MANO)
Contrail Service Orchestration (Orchestration and VNF Manager)
Infrastructure Services and Hardware Platforms
NFV Infrastructure (NFVI)
NVX250 Platform (Software-Driven CPE)
Service Control Gateway (DPI, Policy, Routing)
MetaFabric Architecture
Compute and Storage
PTX Series (Core)
MX Series, vMX (Edge and SDN Gateway Routers)
Contrail Cloud Platform
Contrail Networking
Northstar (WAN SDN Controller)

About the Authors:
Sandra McCann and Helen Shaw are both technical writers with significant experience in the networking industry. They work for Juniper Networks.

Version History: First Edition, June 2016 2 3 4 5 6 7 8 9

For more information go to the TechLibrary at:
www.juniper.net/documentation