Network Function Virtualization

Architecture and Framework and vCPE USE CASE

Arsalan Zahid

Capstone Project presented for Masters of Internetworks

University of Alberta

Table of Contents

Table of Figures

Chapter No 1:

Introduction to Network Function Virtualization

1.1. Introduction to NFV:

Network functions virtualization disaggregates the network functions from the network hardware. As a result, functions, such as NAT, firewalling, IPS and DNS etc. can be delivered in software and deployed on the commodity hardware instead of specialized network hardware. This gives enterprises and service providers a lot more flexibility in the way they design, deploy and manage the network services.

NFV began inside of the competitive service provider community, as they searched for approaches to cut capex and opex and accelerate the takeoff of profitable services to better utilize their systems and grow their revenues. Hardware based network systems, which are normally costly and complex to setup and manage, were restricting the SP's capacity to consolidate functionality and quickly trial new services.

Introducing virtualized network functions (VNF) into the network draws on many of the capabilities particularly the automated provisioning and management. As the VNFs are deployed, it is possible to provide services to a broader set of customer sites, adding those that previously had locally delivered value added services. The services can be delivered centrally using shared resources for and entire VPN rather than just for a single site. The end customer will see a modest reduction in power, cooling and space utilization because each site no longer requires its own appliances to deliver the services. The operator or the end customer will see a reduction in capital expenditure as it no longer has to purchase the appliances to support new sites or to provide enhanced capabilities. The operator will be able to offer enhanced variety and flexibility of services since each customer can now choose from all the available VNFs without having to order the associated appliance, install it, cable it and keep it powered within each site. The speed with which the operator can provision a service chain and start charging the end user for the service means that the time to revenue for any new service will be dramatically reduced.

Any enhanced features can be added to VNFs by the operator quickly and efficiently in a single centralized location for each customer and the operator will even have the ability to introduce a completely new VNF into a service chain to deliver an existing function without any hardware replacement on the customer premises.

In addition to serving individual end customer organizations, NFV can be used by service providers to deliver scalable functions to their fixed and mobile broadband subscribers. NFV provides a way to reduce capital expenditure on hardware appliances to deliver the services, while maintaining scalability with the ability to spin up new resources and load balance across those resources, as each function becomes capacity bound. The different measures by which a function may be bound (CPU, storage, Input/output) can all be scaled independently making for very efficient scaling.

By applying analytics and big data techniques to identify automatically when a particular VNF is becoming capacity constrained (or has excess capacity) and to automatically spin up (or down) resource to more accurately match the new load. This automated expansion or reduction of capacity means that resources need not be pre-provisioned in large chunks meaning that capital expenditure can be delayed and that resources no longer required for a particular purpose can be rapidly repurposed for other services running in VNFs. This efficient use of capital assets can save significant amounts for a large operator. Similarly, the automated nature of the expansion/reduction of capacity means that operational expenditure associated with constant monitoring of the platform can be reduced.

Within the increasingly virtualized infrastructure, service providers wanted to able to setup the network functions dynamically depending upon the requirements and did not want to be restricted to a limitation of the hardware appliance. They felt if they could disaggregate the hardware and the software it would allow them to setup the network functions quickly over the available commodity hardware.

1.2. The Promise of NFV:

The enterprises are starting to use NFV to support their business requirements and Service Providers to support the services they want to deliver within their network infrastructure. NFV facilitates the organizations to reduce the costs associated with the deployment, management and maintenance of the network infrastructures at the same time it reduces the time to deliver the service to the customers. Also, it brings flexibility and agility they require to move and scale functionality to address the challenging business requirements.

1.2.1. Reduce Capital Expenditures:

- The use of commodity servers reduces hardware costs; and increases the price performance. Also, a lot of vendors can offer commodity servers, which increases the volume and competitiveness in the market, which ultimately drives down cost.

- By setting up the network functions in software, service providers are no more compelled to depend on specific equipment to run system capacities. This implies the premium that merchants could charge for their exclusive equipment is no more material or reasonable.
- The infrastructure built on the commodity hardware can be used to setup network function redundancy, hence the need to maintain the expensive set of spare equipment is not required. The reliability is built into the software and replication is done to mitigate the failure and the workloads can be moved across redundant commodity infrastructure.
- Public cloud infrastructures have come to the benefit of the organizations especially the enterprises by allowing them to run the network functions in the cloud instead of locally in the entire Data Centers that turns the capex to opex and increase the capital efficiencies. Instead of buying the equipment upfront the organizations can rent and take advantage of the pay-as-you-grow model and avoid costly and wasteful overprovisioning.
- A side advantage of utilizing less costly equipment is that an association can conceivably cycle the equipment more frequently to enhance the general execution of the system. By updating the system each 2 to 3 years, rather than the conventional 5 to 7, an organization can proceed to adequately address the changing requests put on their system and build the worth caught all through the lifetime of those server

1.2.2. Reduce Operational Expenditures:

- Software enabled network infrastructure empowers organizations to rapidly and effortlessly move and scale usefulness to address changing needs and amplify the utility of the product equipment. A commodity server can be utilized instead of specialized hardware to provide variety of capabilities.
- With the utilization of commodity hardware, the organizations can decrease the space, power and cooling expenses of the IT infrastructure.
- The virtualized functions offer greater flexibility and less complexity[1] in management. Organizations will be able to spin up the virtual networks, take a snapshot, move the VNF and delete them like we do with the Virtual Machines today. All of this makes it very simple to move or redeploy the functionality across the organizations.

1.2.3. Accelerate Time-to-Market:

- The virtualized networks functions can be easily setup thus enabling the enterprises and service providers to rapidly deploy the services whenever they are needed. New services can be tested without experiencing much risk. The orchestration system enables the organizations to dynamically recover from the failures and decrease the risk of deploying new products from different vendors. Today, conducting a Proof of Concept requires network staging beforehand and its integration with the production environment which takes ample amount of time and effort whereas with NFV in place the POCs can be run a

1 This is context with the operations after all the services have been provisioned.

lot faster in smaller scale environments which allows the organizations to adjust and tune their offerings for a wider-scale deployments.

- NFV allows the organizations to run the VNFs on top of the existing physical network infrastructure; therefore they do not experience the time and costs of upgrading the existing systems to add new services.

1.2.4. Deliver Agility and Flexibility:

- The organizations do not need to make million dollar investment to bring in the new service for a customer instead the same service can be delivered on demand using VNFs provisioned on top of clusters of commodity servers. Similarly, the ability to easily tear down, move, scale and configure services as the requirements of customers or the business changes gives organizations the ability to bring up services anywhere in the world, any time.

Chapter No 2:

NFV Architecture and Framework

2.1 Network Service in Network Function Virtualization:

The network service is defined as the connection of network functions connected together (referred as forwarding graph of network function). It can be setup over the infrastructure of a service provider completely or it can be distributed across the infrastructure of multiple service providers. The performance of the network service depends on the performance of the network functions and the performance of the network function depends on the infrastructure.

To explain the concept of a network service further lets consider an example of a firewall. A physical firewall typically provides functions such as firewalling, DPI and NAT etc. Traffic that is requires all the aforementioned services needs to be passed through the firewall. That's a network service. With network function virtualization, now this network service can be distributed across several virtual network functions setup over an x86 platform and connected together to form forwarding graphs of network functions and traffic is passed through each VNF. Below figure represents the scenario. NF1 represents the firewalling function, NF2 DPI and NF3. Traffic from the end point A is steered across all the network functions. The physical interfaces amongst the nodes are represented as solid lines and the logical interfaces are represented as dotted lines.

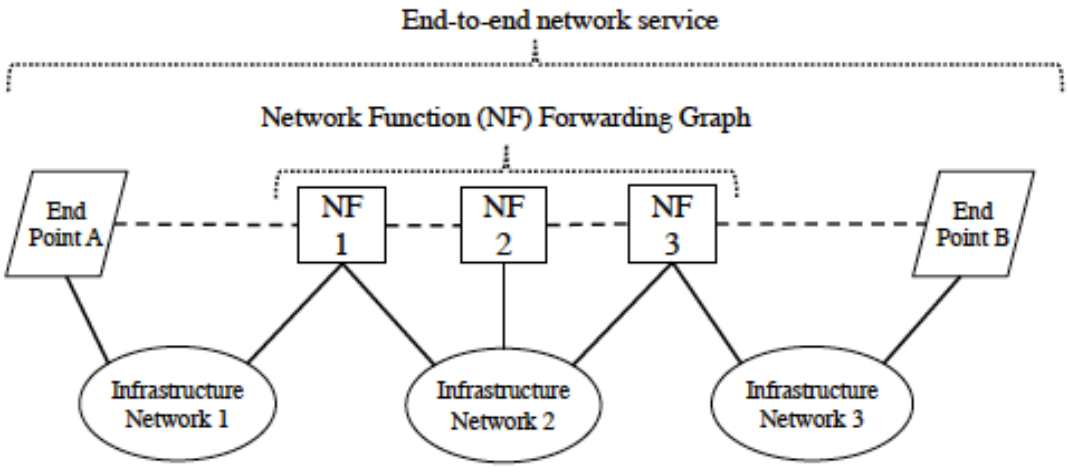

Figure 1 Example Network Service[2]

2 STANDARDS INSTITUTIONS, "GS NFV 002 - V1.2.1 - Network Functions Virtualisation (NFV); Architectural Framework," Tbd, vol. 1, pp. 1–21, 2014.

2.2. Functional Blocks of the NFV Architecture:

There are several identified functional blocks in the NFV architecture that are going to be referenced throughout the document. Following is the brief description of the critical functional blocks except the virtual network function and orchestrator that are discussed in detail in their respective sections in the following chapters.

2.2.1. Element Management:

The element management system is used for the management of the physical or virtual network functions. The interface between the EMS and the network function can be proprietary or standard. It integrates with the Network Monitoring System and the OSS/BSS (discussed later) through standard interfaces. The role of the EMS is to manage all aspects of the relevant networks elements and present the information northbound to the NMS, OSS and BSS over standard interfaces.

Example of the an Element Management System from Juniper Networks is the Junos Space Network Management Platform along with different network management applications to manage different types of network functions e.g. Junos Space management platform used with the Network Director is used to manage the wired and wireless routing and switching network functions for e.g. EX switches, MX routers and VMXs in the infrastructure. Junos Space management used with Security Director to manage all the Juniper Networks security devices e.g. SRX and virtual SRX.

It is responsible for complete fault, configuration, accounting, performance and security functions. Hence, it is one single pane of window from where the network administrators can manage the entire infrastructure.

2.2.2. Virtual Infrastructure Manager (VIM):

The virtual infrastructure manager is responsible for managing the exchanges of the VNF with the compute, storage and network. The VIM ensures the abstraction and consolidation of the resources in the physical infrastructure and creation of the virtual resources assigning them optimally across several network functions.

2.2.3. Virtual Network Function Manager:

Virtual Network Function Manager manages the lifecycle[3] (e.g. instantiation, update, query, scaling, termination) of the multiple virtual network functions. One VNF may be deployed for multiple VNFs and every VNF may be managed by a dedicated VNF.

2.3. The Components of an NFV Architecture:

Following are the components of the Network Function Virtualization Framework.

- Virtualized Network Functions (VNFs) – the software implementation of a network function.
- NFV Infrastructure (NFVI) – the physical resources (compute, storage, network) and the virtual instantiations that make up the infrastructure.
- NFV Management and Orchestration – the management and control layer that focuses on all the virtualization-specific management tasks required throughout the lifecycle of the VNF.

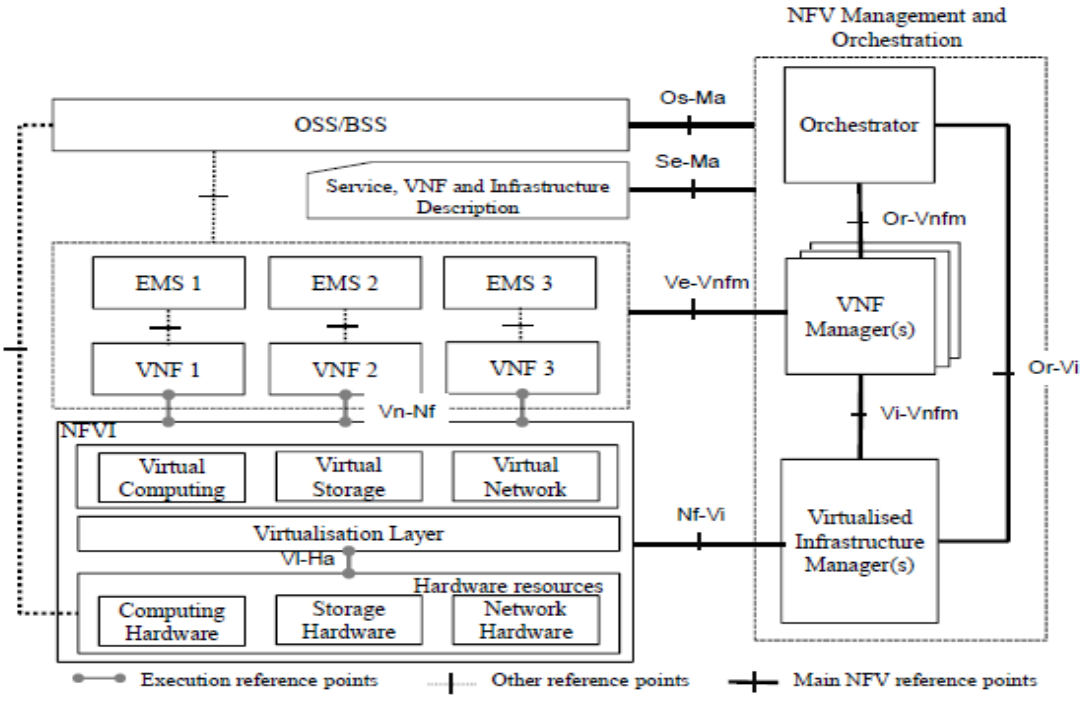

Figure 2 NFV Referential Architecture[4]

3 VNF lifecycle is discussed in detail in the Virtual Network Function Section.
4 STANDARDS INSTITUTIONS, "GS NFV 002 - V1.2.1 - Network Functions Virtualisation (NFV); Architectural Framework," Tbd, vol. 1, pp. 1–21, 2014.

Chapter No 3:

Virtual Network Function

3.1 Virtual Network Function:

Virtual Network Function is a network function, which is typically setup on the commercially available off the shelf (COTS) hardware e.g. X86 server. Below is the typical representation of a virtual network function.

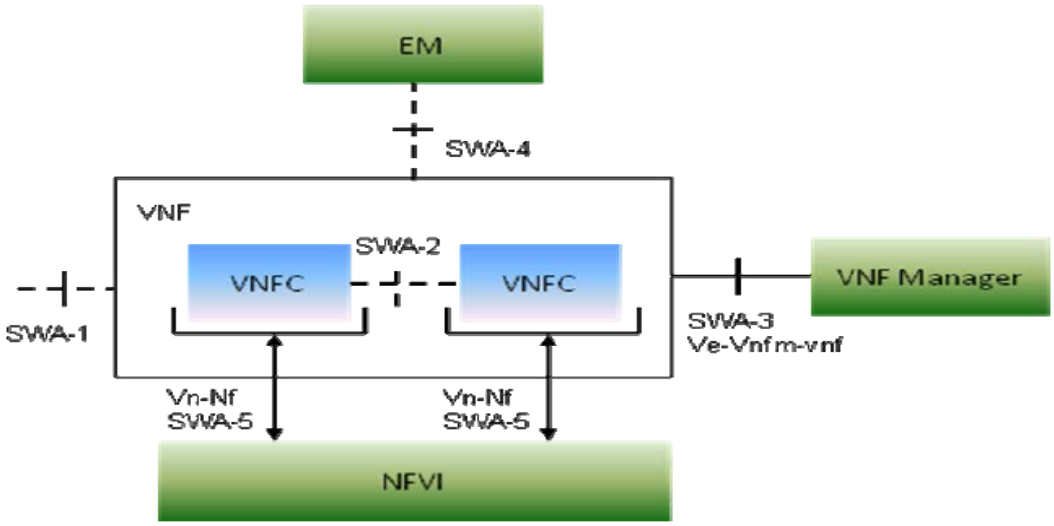

Figure 3 Virtual Network Function Architecture[5]

The VNF can be developed monolithically or modularly. Monolithic development of a VNF implements a single network device with the interfaces and functions defined by the standard organizations e.g. IETF. Modularly developed VNF breaks down the functionality of a network device into several modules where each module is responsible for a different network function. E.g. A physical network firewall typically implements the policy enforcement, IPS, NAT and DPI each being a different network function although a part of the same network firewall. Monolithic implementations tends to convert the physical form factor into the virtual form factor where one Virtual Network Function Component – VNFC implements all the network functions whereas in modular implementations each network function is implemented as a different VNFC. The interfaces between different VNFCs can be proprietary and do not need to be exposed. However, the interfaces between different VNFs have to follow standards to make VNFs from different providers interoperable with one another.

5 ETSI, "Network Functions Virtualisation (NFV); Virtual Network Functions Architecture," vol. 1, pp. 1–93, 2014.

The benefit of the modular implementations is that the innovation of different network functions becomes independent which makes scaling out a lot easier and faster. E.g. The VNF provider may decide to enhance the performance and scale of the VNF by enhancing the interface between the VNFCs. However, any change in the interface between the VNFCs do not effect outside the VNF – hence easier to scale out.

Another example of the modular implementation is a separate VNFC for the control plane and data plane. The benefit of this architecture is that it separates control operations such as routing updates and system management from packet forwarding so the router VNF can deliver superior performance and highly reliable Internet operations.

3.2. Virtual Network Function Interfaces:

There are 5 different interfaces related to the software architecture of a VNF shown in the figure

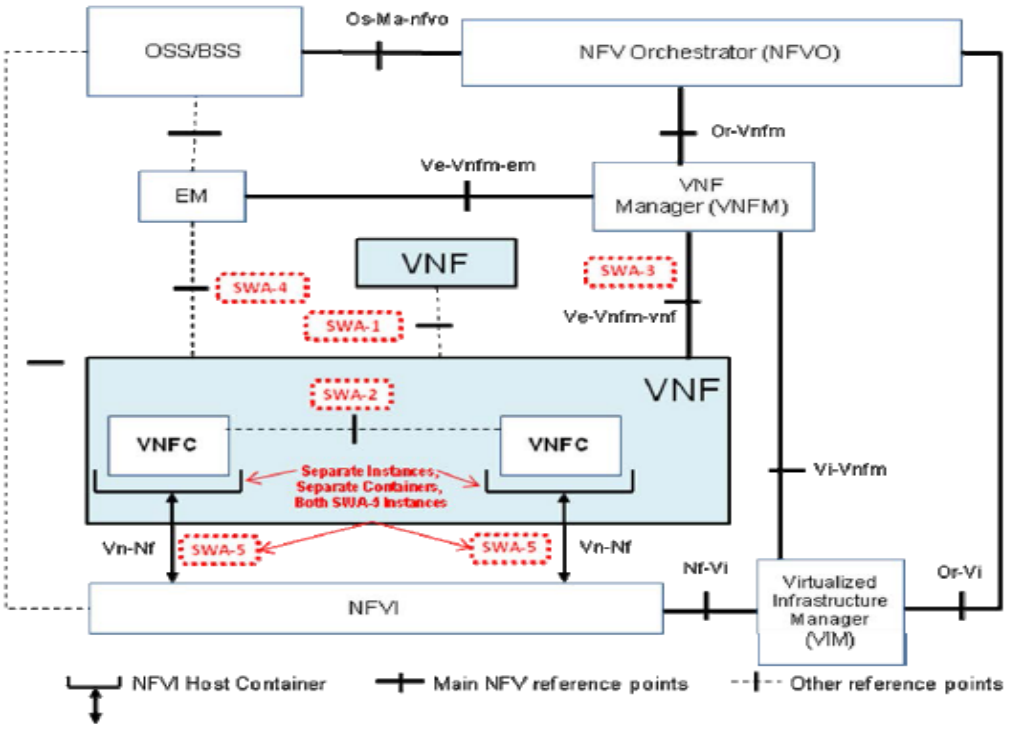

Figure 4 VNF External and Internal Interfaces[6]

6 ETSI, "Network Functions Virtualisation (NFV); Virtual Network Functions Architecture," vol. 1, pp. 1–93, 2014.

3.2.1. SWA-1 Interfaces:

The interface between two VNFs or between a VNF and PNF[7] or between a VNF and an End Point is called SWA-1 Interface. This may represent Data/Control plane interfaces of the VNF.

3.2.2. SWA-2 Interfaces:

These are proprietary interfaces defined by the VNF providers, which allow communication between different components that make up a VNF. The VNF user does not have the visibility to these proprietary interfaces but they assert performance e.g. latency, capacity requirements on the underlying virtual infrastructure. These are logical interfaces that make use of the SWA-5 interface (discussed later in this section) to gain access to the virtual resources such as connectivity services when VNFC instances are hosted across different hosts. It can be that they can be hosted on the same host. In such cases, to provide better latency methods like shared memory access can be exploited to enable message passing across multiple VNFC instances.

3.2.3. SWA-3 Interfaces:

This is the interface between the NFV management and orchestration specifically to one of the MANO[8] functional component called Virtual Network Function Manager. Virtual Network Function's instantiation and scaling is done via the management interfaces. This interface may use Layer 2 or IP connectivity.

3.2.4. SWA-4 Interfaces:

The Element Management uses the interface SWA-4 to communicate with a VNF for Fault, Configurations, Accounting, Performance and Security.

3.2.5. SWA-5 Interfaces:

This is the interface between the VNF and NFVI and it exposes the virtual resources e.g. virtual compute, network and storage to be consumed by a virtual network function or VNF service chains.

7 A PNF is a physical network function packaged in a dedicated hardware. It may be firewall, NAT, DPI, IPS, routing, switching etc.
8 MANO is management and orchestration that is used to manage and spin up/terminate the different NFV components. It is described in detail in Chapter 5.

"Generic compute functionality:

- Role - The NFVI provides an interface to access generic compute functionality.
- Interconnection attributes - These sub-interfaces have CPU-dependent attributes.

Specialized function:

- Role - A Specialized function provides the VNF with non-generic compute functionality or expanded generic compute functionality. These functions can vary from a Video or Voice Gateway card, to specialized memory and or routing application.

- Interconnection attributes - The SWA-5 specialized interface(s) are strictly dependent upon the technology being implemented and the interface types associated with that standard. Generically the interfaces should have their own identity, and mapping capabilities.

Storage:

- Role - The NFVI provides a storage interface that can support storage operations on any granularity including block, file, or object storage.
- Interconnection attributes - Physical storage in the NFVI may be provided locally or remotely.
- Storage operations - In order to access the storage capabilities of the NFVI, the VNF developer may use common storage operations. The execution environment provided by the NFVI may implement these operations using a number of different device drivers to mask technology differences.

Network I/O:

- Role - The SWA-5 Network I/O interface(s) provide the VNFC/VNF instances with network connectivity services.

- Layer 2 services (e.g. E-LAN, E-Line, E-Tree) based on the Ethernet switching network infrastructure or on BGP Ethernet VPNs (EVPN).

- Layer 3 services (directly based on the L3 infrastructure or on Layer 3 VPNs.

- Multiple interfaces for both redundancy and segmentation are normal considerations, along with hybrid Network Interface Cards (NICs) that provide multiple ports. Each virtual NIC has a driver for the supported OS/VNF.

- Interconnection attributes - Each (can be multiple) SWA-5 Network I/O interface has NIC-level attributes, and needs to be separately map-able to a network segment, and also map-able in pairs for redundancy with network functions such as Link Aggregation Groups (LAG). The Network I/O sub-interface maps to the [Vn-Nf]/VN reference point. All other sub-interfaces map to the [Vn-Nf]/VM reference point."[9]

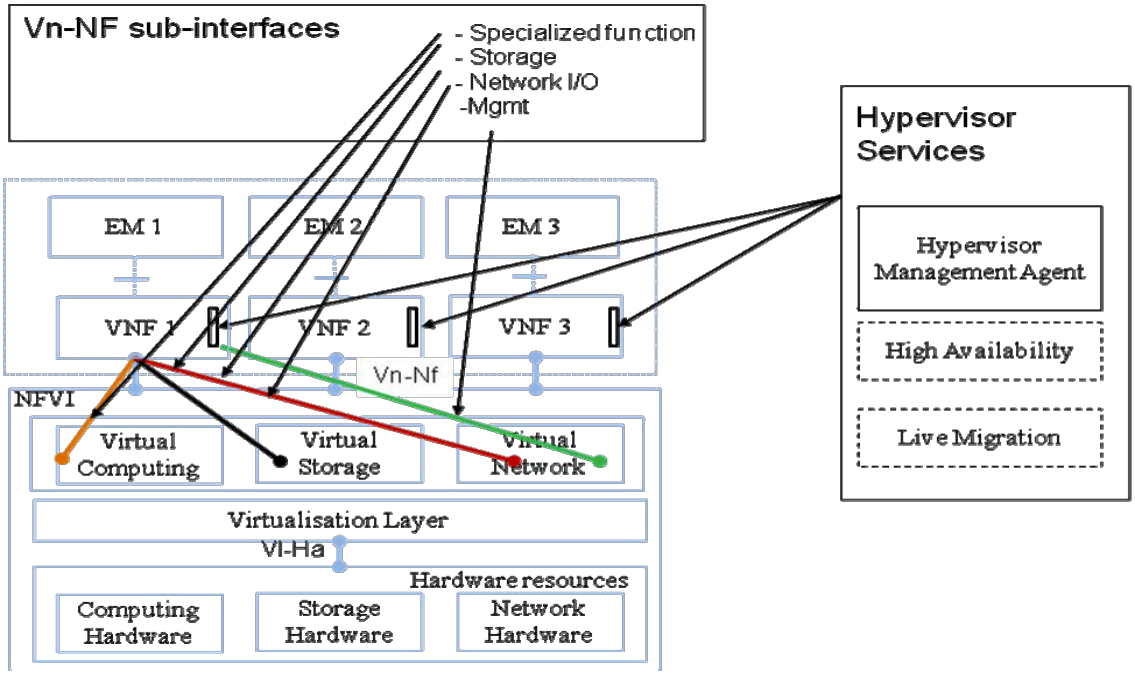

Figure 5 SWA-5 Sub-Interfaces[10]

9 ETSI, "Network Functions Virtualisation (NFV); Virtual Network Functions Architecture," vol. 1, pp. 1–93, 2014.
10 ETSI, "Network Functions Virtualisation (NFV); Virtual Network Functions Architecture," vol. 1, pp. 1–93, 2014.

3.3. Communication between VNFCs:

There are several methods by which the Virtual Network Function Components communicates with one another. The representation of the different methods is shown in the figure.

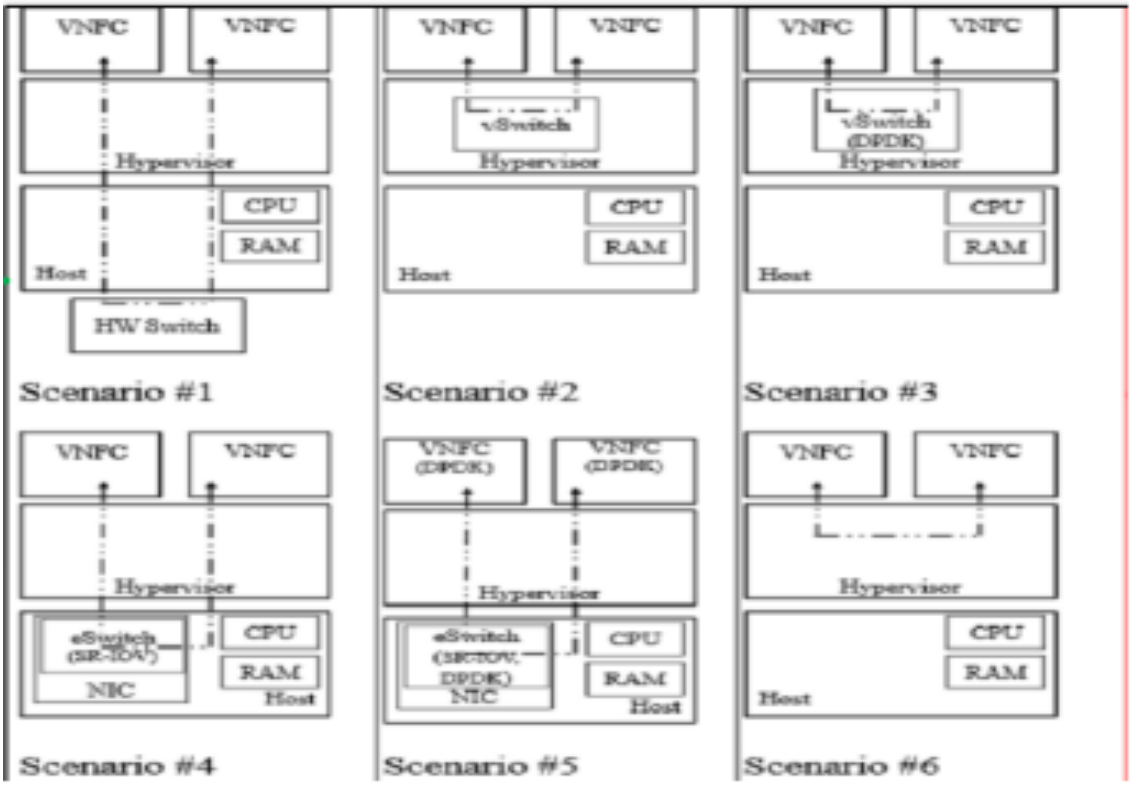

Figure 6 Methods of Communication between the VNFs[11]

The VNFs can communicate through the network hardware switch or through the software switch in the hypervisor for e.g. open virtual switch in case of VMware. Alternatively, the VNFs can also communicate via the E switch in the NIC or as simply as through the bus.

11 ETSI, "Network Functions Virtualisation (NFV); Virtual Network Functions Architecture," vol. 1, pp. 1–93, 2014.

3.4. VNF Instantiation:

To understand VNF instantiation first we need to understand the VNF descriptor. A VNF descriptor is a template of the specifications given by the VNF provider that describes requirements that can be consumed by the NFV MANO to understand the execution of the VNF lifecycle operations e.g. requirements on the operating system, interconnectivity requirements and virtual resources etc. A data model describing all these requirements is called the VNF descriptor.

VNF instantiation is the aggregation of the instantiations of the VNFCs. Once the VNFCs corresponding to a VNF are instantiated, the VNF is said to be instantiated. Now lets understand what steps are included in the instantiation of the VNFCs. VNFM consumes the VNF descriptor to request a new virtual machine and the storage resources for the VNFC. Once the VM is allocated and booted up and VNFM starts configuration of the VNF according to the VNF descriptor. Once the configuration is completed an event informing the VNFM is generated and sent to the VNFM.

Chapter No 4:

Network Function Virtualization Architecture

4.1 Network Function Virtualization Infrastructure:

NFVI is all the hardware and software resources that make up the environment. It can vary greatly across different types of organization, depending on the network's complexity and geographic distribution. The physical resources typically include compute nodes, storage and networking equipment that provides processing, storage and connectivity amongst the VNFs and between the virtual and physical infrastructure. Other supporting services, such as service catalogs, external testing and analytics will likely be built into the NFVI over time, as they become a critical part of the infrastructure and are increasingly relied on to ensure the uptime and performance of the systems.

The virtualization layer sits right on top of the hardware and abstracts the resources, so they can be logically partitioned and assigned to the VNF to perform their functions. It disaggregates the VNF software from the specialized underlying hardware, so the VNF can use the virtualized resources effectively to execute its function.

The virtualization software is not specific to any VNF instead the existing easily available technology such as a hypervisor capable of abstracting the physical resources and presenting them logically to the VNF can be leveraged to make up the virtualization layer. A VNF can also be implemented as an application in environments where hypervisor support is not available by leveraging an operating system (OS) that adds a software layer on the top of non-virtualized hardware.

Today networking systems are represented as functional blocks and the communication across different components of the network are specified by the open standard interfaces between them. This approach is often referred as system engineering. The principle behind system engineering is that the complete system specifications are cumulative of the specifications of the individual functional blocks and the interfaces between them. A representation of such a system is given below.

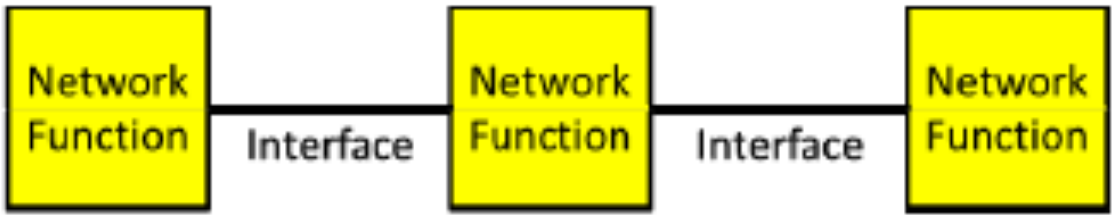

Figure 7 System Engineered - Network System Representation[12]

The objective of the NFV is to disaggregate the network function from the hardware that hosts it and creating software (VNF), which can be executed and hosted on the general-purpose hardware. VNF and NFVI requirements are therefore separately specified. Now the disaggregated system cannot be specified using the fundamentals of systems engineering. Following is the representation when the software (Network Function) is separated from the hardware it hosts it and becomes commodity hardware.

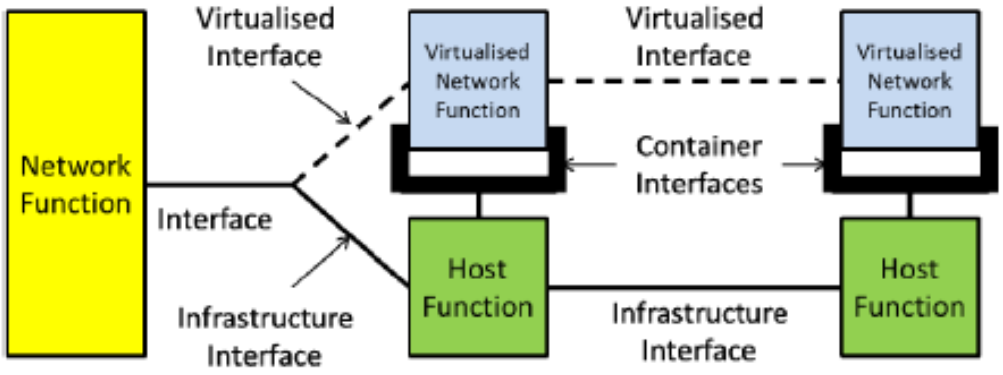

Figure 8 Post disaggregation view of Network Function[13]

Key noticeable things that arise as a result of virtualizing of the network functions:

- The functional block has been divided into the VNF and the host function (For e.g. X86 server)
- Container (For e.g. Virtual Machine) interfaces have been created between the hosted VNF and the host function.
- The interface between the two network functions has now been divided into the virtualized and the infrastructure interface.

12 ETSI, I. Overview, ETSI, and I. Overview, "GS NFV-INF 001 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure Overview," Etsi, vol. 1, pp. 1–59, 2015.
13 ETSI, I. Overview, ETSI, and I. Overview, "GS NFV-INF 001 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure Overview," Etsi, vol. 1, pp. 1–59, 2015.

- Homogeneity can been seen in the interface between the non-virtualized network functions whereas as divide amongst the infrastructure interface and virtual interface can be seen when the network function is virtualized.

The differences between the system based on the standard network function and virtualized network function are:

- The network function when virtualized does not remain autonomous i.e. it is not functional block without the host function.
- The creation of the container interface, which was not present in the standard system.

As stated above, the virtualized network function does not remain autonomous and depends on the host function for it to exist. This implies that if the host function's execution is interrupted in any way the execution of the virtual network function will be interrupted.

When the host function is configured with the Virtual Network Function, the host function appears to be the functional block executing the specifications of the Virtual Network Function. Actually, it is the host function that is the functional block but it appears to be the VNF from the outside. Therefore, the Network Function Virtualization Architecture is defined using the host functions with their related container and infrastructure interfaces and the VNFs with their related container and virtualized interfaces.

The NFV documentation uses these constructs. The NFV documentation also noted that the container relationship between host function and virtualized function could be recursive layered. Host functions with more and more specific capability can be built up by the successive configuration with layers of virtual functions.

An example related to Network Function Virtualization is discussed below.

A blade server (base host function) provides the container interface of the bare metal machine and the related BIOS. A hypervisor can be installed on top of the blade server that allows creating multiple virtual machines on top of it providing multiple virtual machine container interfaces. A virtual machine can be configured with an operating system, so now the blade server is configured with the hypervisor and an operating system that can provide several application container interfaces. On top of that any virtual network function can be setup with the capability of providing several different functions.

4.2. NFVI Domains:

Following domains are a part of NFVI.

- Compute Domain
- Hypervisor Domain
- Network Domain

We will take each domain one by one and discuss the components and interfaces of each in this section.

4.2.1. Compute Domain:

The storage and network I/O that includes the servers and capacity is included in the compute domain of the NFVI. Key noticeable thing is that the compute domain is connected with MANO since the orchestrator runs on top of the NFVI as VMs.

Figure 9 highlights the reference points in the NFV and those pertaining to the compute domain are discussed later in this section.

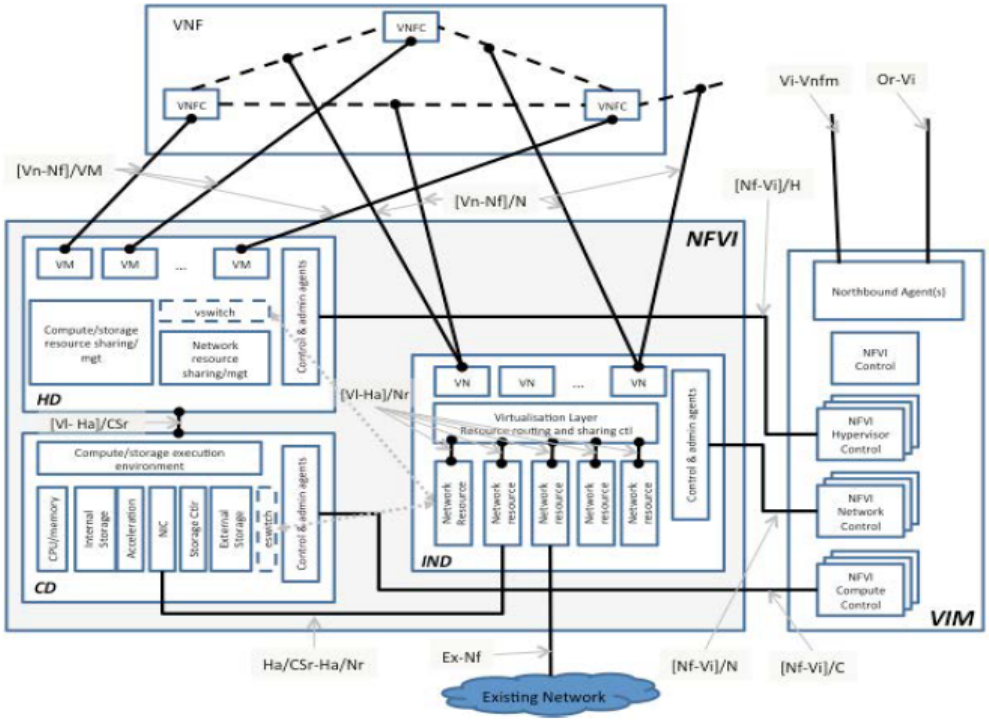

Figure 9 Reference points in NFV Architecture[14]

14 ETSI, E. E. T. S. Institute, ETSI, and E. E. T. S. Institute, "GS NFV-INF 003 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Compute Domain," Etsi, vol. 1, pp. 1–57, 2014.

4.2.1.1. NFVI Reference Points relating to the Domain:

INF Context	NFV Framework Reference Point	INF Reference Point	Reference Point Type	Correspondence with figure 2 Interfaces	Description and Comment
Internal	Vl-Ha	[Vl-Ha]/CSr	Execution Environment	12	The framework architecture [2] shows a general reference point between the infrastructure 'hardware' and the virtualisation layer. This reference point is the aspect of this framework reference point presented to hypervisors by the servers and storage of the compute domain. It is the execution environment of the server/storage.
		[Vl-Ha]/Nr	Execution Environment		The framework architecture [2] shows a general reference point between the infrastructure 'hardware' and the virtualisation layer. While the infrastructure network has 'hardware', it is often the case that networks are already layered (and therefore virtualised) and that the exact choice of network layering may vary without a direct impact on NFV. The infrastructure architecture treats this aspect of the Vi-Ha reference point as internal to the infrastructure network domain.
		Ha/CSr-Ha/Nr	Traffic Interface	14	This is the reference point between the infrastructure network domain and the servers/storage of the compute domain.
External	Nf-Vi	[Nf-Vi]/C	Management, and Orchestration Interface	11	This is the reference point between the management and orchestration agents in compute domain and the management and orchestration functions in the virtual infrastructure management (VIM). It is the part of the Nf-Vi interface relevant to the compute domain.

Figure 10 Characteristics of NFVI Reference Points - Relevant to Compute Domain[15]

4.2.1.2. Interface between VIM (Nf-Vi) and NFVI:

There are two interfaces that NFVI is concerned with. The one internal to the NFVI e.g. SWA-5 that has been in detail earlier is only meant to provide the virtual infrastructure resources to the virtual network functions. It is not meant provide virtual network function the capabilities to manage the NFVI.

The only NFV management interface between the NFVI and VIM is through Nf-VI. In the case VIM is hosted on the same infrastructure its management interface is abstracted by SWA-5 interface. Although the VIM can be a VNF hosted on the same infrastructure that it is supposed to manage, but certainly not recommended. The reason is that it puts a potential security and reliability concerns. In any event if the NFVI requires a reboot/maintenance the VIM will lose its ability to manage the NFVI. Therefore, the recommendation is to host the VIM on the completely separate NFVI.

4.2.1.3. Nf-Vi/C and [Vl-Ha]/CSr Interfaces:

The Nf-Vi/C and Vl-Ha/CSr are two interfaces external to the compute domain. VIM consumes Nf-Vi/C to manage the compute and storage and Vl-Ha/CSr is the interface

15 ETSI, E. E. T. S. Institute, ETSI, and E. E. T. S. Institute, "GS NFV-INF 003 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Compute Domain," Etsi, vol. 1, pp. 1–57, 2014.

between the compute and hypervisor domain. Hypervisor consumes this interface to gain visibility of the available resources in the compute domain.

4.2.2. Hypervisor Domain:

The hypervisor domain sits on top of the compute node and to allow several virtual machines to consume part of the physical resources assigned to it. A hypervisor can emulate every piece of hardware platform and in some cases it can emulate the instruction set of CPU making a virtual machine believe that it is running on a complete independent hardware. However, number of actual CPU cycles required to emulate virtual CPU cycles can be large enough to bring down the performance significantly.

4.2.2.1. External Interfaces of the Domain:

The interface between the hypervisor and compute domain is VI-HA-CSr as described earlier as well. Hypervisor consumes this interface to gain visibility of the available resources such as BIOS, Drivers, NICs, Accelerators and Memory etc. in the compute domain and provide the data to the Virtual Infrastructure Manager. The interface between the hypervisor and network domain is VI-HA-Nr. Hypervisor domain consumes this interface to gather the relevant metrics from the network domain and provide that data to the Virtual Infrastructure Manager.

4.2.2.2. Hypervisor to VIM (Nf-Vi-H) Interface:

The interface between the Virtual Infrastructure Manager and the hypervisor is Nf-Vi-H. Nf-Vi-H is consumed by the hypervisor in sending the information such as configurations, alerts, policies, responses and updates of the core infrastructure to the VIM. Currently, the information is sent using the proprietary packages since the standard API has yet to be defined.

4.2.3. Network Domain:

The reference point architecture of the Network Domain is given in the Figure 11 below.

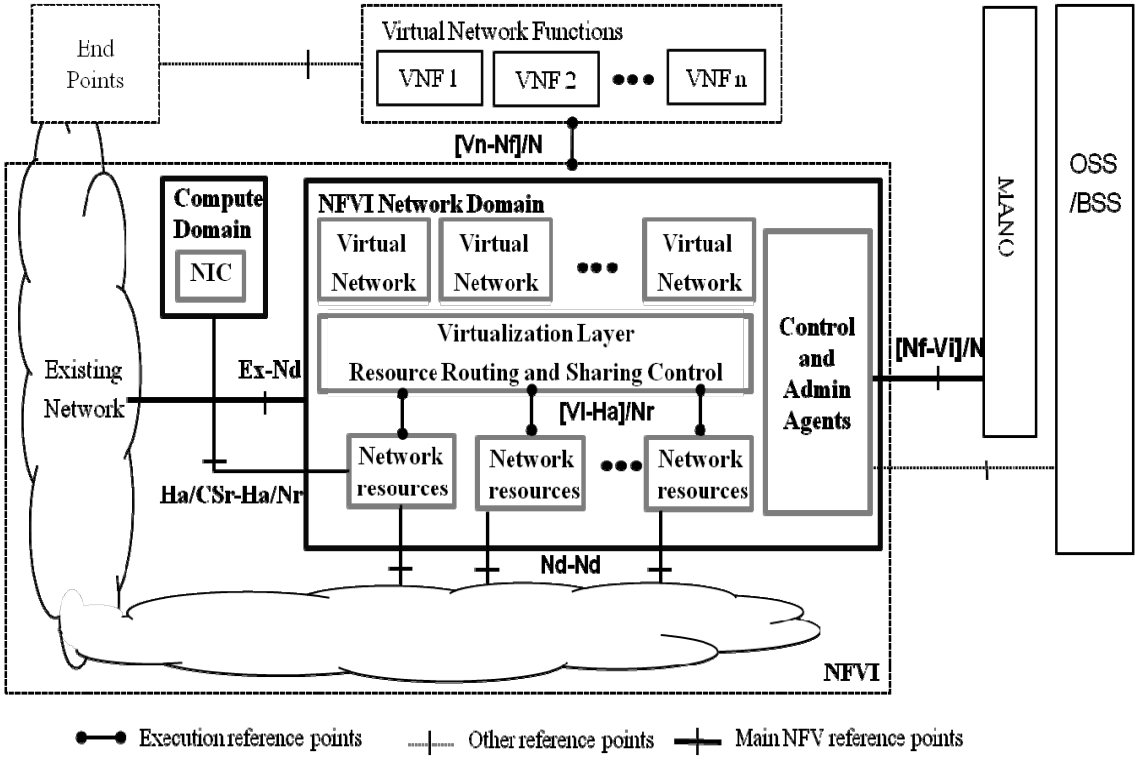

Figure 11 Network Domain Reference point architecture[16]

16 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

4.2.3.1. Characteristics of the Reference Points:

Reference Point	Description
[Vn-Nf]/N	This reference point is the virtual network (VN) container interface carrying communication between VNFC instances. Note that a single VN can support communication between more than a single pairing of VNFC instances (eg an E-LAN VN). It is the reference point over which the services of the network domain are delivered. These services may be either IP forwarding services or Ethernet private line/LAN/TREE services provided by the infrastructure. The reference point is providing services at two layers: IP forwarding services across the [Vn-Nf]/N/L3 reference point and Ethernet services, e.g. E-LINE, E-LAN, E-TREE, across the [Vn-Nf]/N/L2 reference point.
[Nf-Vi]/N	This is the reference point between the management and orchestration agents in the infrastructure network domain and the management and orchestration functions in the virtual infrastructure management (VIM). It is the part of the Nf-Vi interface relevant to the infrastructure network domain.
[Vl-Ha]/Nr	The reference point between the virtualisation layer and the network resources.
Ex-Nd	The reference point between the infrastructure network domain and external networks.
Nd-Nd	The reference point between NFVI-PoPs used to extend the virtualisation layer of a single Network Operator's NFVI over multiple geographically separated sites.
Ha/Csr-Ha/Nr	This is the reference point between the infrastructure network domain and the servers/storage of the compute domain.

Figure 12 Characteristics of NFVI reference points - relevant to Network Domain[17]

4.2.3.2. [Vn-Nf]/N Interface:

The Vn-Nf/N interfaces shall provide transparent network services to VNFs. This interface interconnects different VNFC instances to one another be it a part of the same VNF or different VNFs, VNFC instances to the storage and VNFC instances to the Physical Network Functions.

Network services are emulated entirely to expose virtual network services across the VNFs in such a way that VNFs are not aware of how the services are exposed. Specific implementations of the how these services are provided across the VNFs and across the VNFs and storage are discussed in detail below in the section.

In the event when VNIC has one-to-one correspondence with the VNFC instance port it is more difficult for a provider of the network service to make use of resources owned by a distinct NFV Infrastructure provider because this method requires either the network service provider to control the network e.g. VLAN, VXLAN to connect two VNICs while the virtual networks are owned by the infrastructure provider or it requires implementation of the virtual networks within the VNFC instances. Neither of the approach is desirable.

17 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

In order to support all use cases

- Possibility to connect several ports within a single VNFCI to a single VNIC.
- The VNF forwarding graph and the network connectivity between the VNFs of a network service enforced by infrastructure network must be independent of the VNF.
- The network service provider is able to make VNF forwarding graph decisions.
- Traffic received on the virtual NIC can be delivered to the VFCI port.

4.2.3.2.1. Layer 2 Service [Vn-Nf]/N:

The [Vn-Nf]/N layer 2 service provides connectivity of the Ethernet data frames over Ethernet switching network infrastructure or over L2 VPN services based on the IP network. VNFC instances can consume any of the methods to send Ethernet data frames across.

Figure 13 shows the Layer 2 service definition.

Service Specification		Illustrative Parameters
Control Operations	Establishment	Request • list of vNICs • (pairwise) bandwidth requirement • (pairwise) delay requirement • resiliency requirement Return • instance id • list of vNICs • (pairwise) bandwidth • (pairwise) delay • Resiliency
	Modification	Request • instance id • change to list of vNICs • change to (pairwise) bandwidth requirement • change to (pairwise) delay requirement • change to resiliency requirement Return • instance id • list of vNICs • (pairwise) bandwidth • (pairwise) delay • Resiliency
	Removal	Request • instance id Return • success
Instance Interfaces	vNIC end point	VLANiD/MAC address Physical location
Operational status		OAM parameters
Performance stats	Establishment	Virtual network provisioning latency Virtual network diversity compliance Virtual network provisioning reliability
	Operation	Packet loss Packet delay Packet delay variation Delivered throughput Network outage
Instance Functionality	Forwarding and transport	MAC forwarding

Figure 13 [Vn-Nf]/N - Layer 2 Service Definition[18]

18 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

4.2.3.2.2. [Vn-Nf]/N – Layer 2 VPN Service:

The BGP EVPN based solution provides layer 2 connectivity across the VNFC instances independent of the infrastructure be it MPLS or IP technology. Ethernet VPNs is another address family in the Multiprotocol Border Gateway Protocol that allows distributing mac-routes (mac-addresses with the next hop ip-address) across the provider edge routers providing layer 2 connectivity across different sites. With EVPN the mac learning happens in the control plane as opposed to in the data plane in the case traditional Ethernet switching which makes it extremely scalable. Mac learning across the VNFC instances and the edge routers occurs exactly the same way as for the layer 2 bridges without adding any extra complexity for the VNFC instance.

4.2.3.2.3. Vn-Nf]/N – Layer 3 Service:

Two different methods to Layer 3 services exist with distinct features as it relates to address space isolation, however there are no major differences from VNFC instance view point as both services are presented to VNFC instance as a ability to send and receive data packets.

Following differences exist in the service definition of Layer 3 compared to one presented in figure 13.

Instance Interfaces	vNIC end point	IP address, MAC address
Instance Functionality	Forwarding and transport	IP forwarding of IPv4 or IPv6 data packets only

[19]

4.2.3.2.4. [Vn-Nf]/N – Layer 3 VPN Service:

A L3 VPN service based on BGP IP VPN makes utilization of an overlay network to provide discretionary number of logically separate virtual networks, each of which has its own range of ip-address space. The ip-address across multiple virtual networks may overlap. Hence, The layer 3 VPN service allows the ip-address space segregation as opposed to L3 infrastructure discussed in the following section.

No complexity is introduced in the L3 VPN service for the VNFCI, since the overlay is transparent to the VNFCI. It simply provides the network virtualization solution through which end system is provided with the IP service. The solution disaggregates the control and forwarding plane and allows implementing the distributed forwarding on multiple devices.

19 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

4.2.3.2.5. [Vn-Nf]/N/L3 Infrastructure based VNS:

Following are some key noticeable things considering the Infrastructure based Layer 3 services.

- Traffic is exchanged across the VNFC instances using the infrastructure network.
- The layer-3 infrastructure service does not allow the ip-address space isolation, which means that each VNFC instance is assigned a unique ip-address. The address can be a private address or a public address or a mix of both.
- Security policy can be implemented to prevent communication across different VNFCI.

4.2.3.3. [NF-Vi]/N

4.2.3.3.1. Nature of the Interface

The interface to the virtualized infrastructure manager dynamically provisions the infrastructure connectivity services. Infrastructure connectivity services require orchestration and management of both infrastructure network resources (Ethernet switches, routers, Firewalls), compute resources (e.g. NICs) and hypervisor resources (e.g. vSwitches or vRouters) in order to provide the desired infrastructure connectivity services to the VNF.

It is the role of the VIM to create compatible configurations across the multiple domains in order to construct the infrastructure network services. The scope of a particular infrastructure network service is likely to go beyond that on anyone VIM. The NFV orchestrator provides co-ordination of configuration across all the VIMs associated with an infrastructure network service and will also co-ordinate with existing infrastructure, which are not managed by VIM.

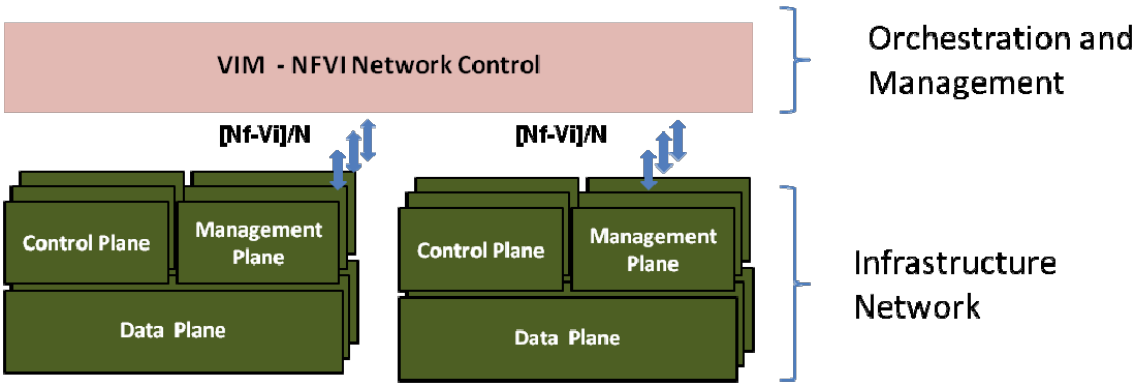

Figure 14 Orchestration and Management within infrastructure management domain[20]

20 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

Following are different approaches to infrastructure network management and control.

- One method is that it may interface with the equipment directly as shown in figure 14 above. This method does not scale well. In environments where scalability is required usually the approach illustrated in figure 15 is adopted.
- Another method is that the virtual infrastructure manager may interface with a network controller that manages the equipment using standard interfaces. The control and management is completely or may be partially centralized and the controller may only provide the abstract view of the network domain to the virtual infrastructure manager as shown in figure 15 below.
- Third method is the combination of the first two where the VIM interfaces with both network controller and equipment.

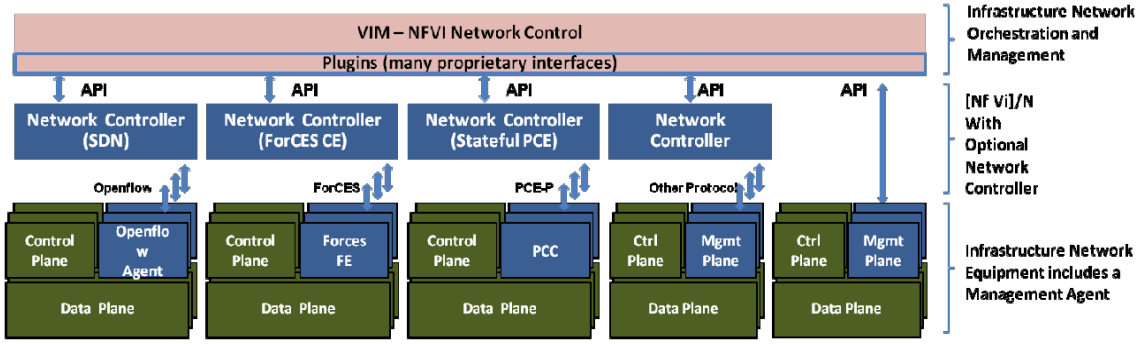

Figure 15 Infrastructure Network Domain Plane[21]

4.2.3.4. Ex-Nd:

Since, the overlay networking is used in the L2 and L3 VPN service, therefore, gateway functionality is required for encapsulation and de-encapsulation between the overlay and external network. One thing to note here is that the layer 3 infrastructure based service does not use the overlay networking, so the VNFC instances connected by the infrastructure service do not require the gateway function. The following section address the gateway functions required to allow communication between PNFs and VNFC instances that are connected through the L2 and L3 VPN based service.

4.2.3.4.1. Nature of the Interface

Connectivity between the different VNFs or between the VNF and PNF is provided using this interface. The connectivity is transparent to the VNFs.

21 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

Figure 16 shows a simple representation of this interface.

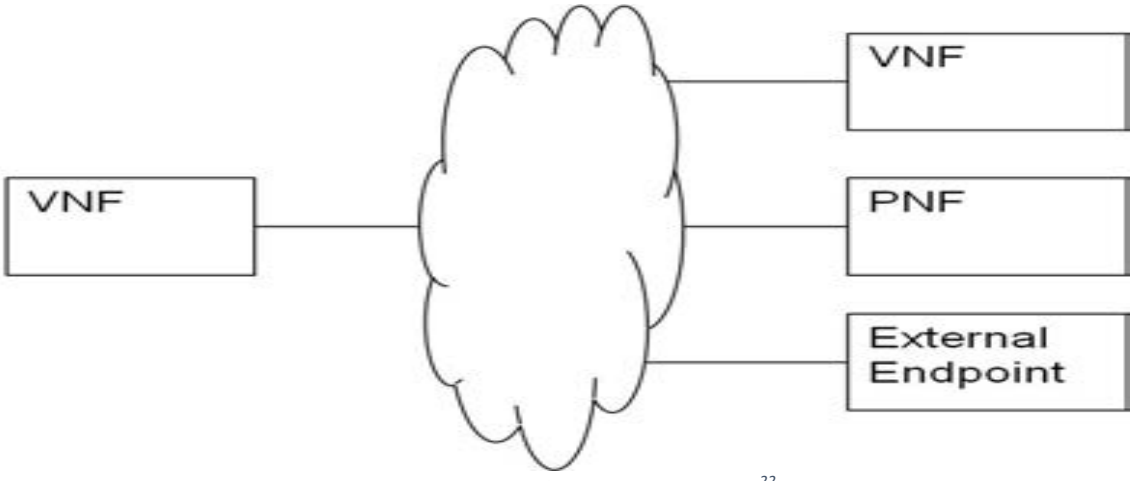

Figure 16 Inter VNF/VNF-PNF Connectivity[22]

A method has to exist between the virtual network and the physical network to allow the physical and virtual network co-exist within the same infrastructure and allow the communication between them. Typically, traffic leaving the virtual network is encapsulated to identify which virtual network it belongs to and then sent out towards the gateway that de-encapsulates and sends it to its destination. In the reverse direction, the reverse steps receive the packet for the virtual network de-encapsulate it and send it on the existing network.

In its most straightforward scenario, the traffic coming from the virtual network might be forwarded taking into account the physical or virtual port on which it is received at the edge of the virtual network, i.e. the gateway between the virtual network and the physical network.

In more difficult scenario, it might require deep packet inspection before it can be classified.

22 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

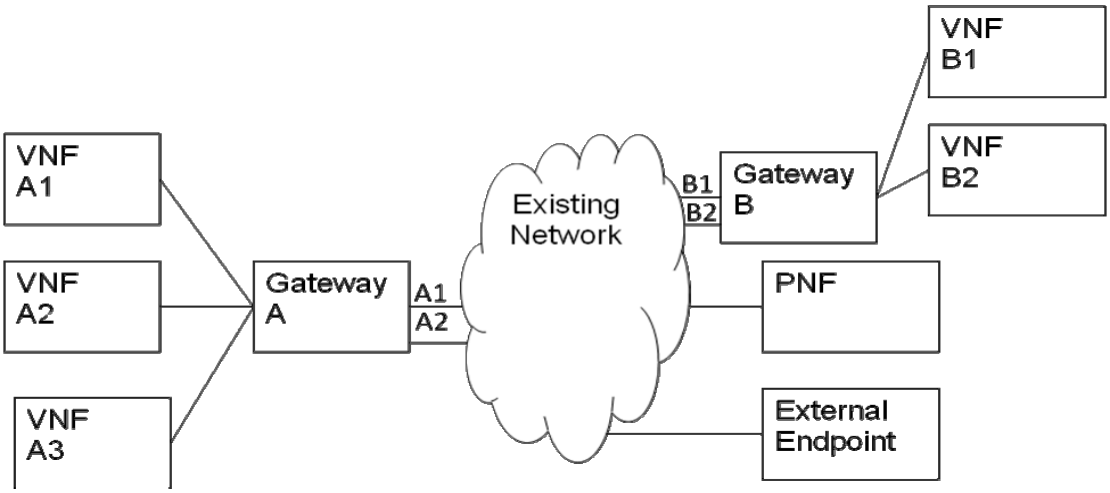

Figure 17 Gateway node[23]

As discussed earlier in the VNF section that a VNF can be designed monolithically or modularly consisting of multiple VNFC instances and the interfaces between them can be proprietary and may not be exposed to other VNFs, so it can happen that the externally exposed interface may not correspond to the VNFC instance interfaces. Several options are discussed in the following section that outlines different methods that can be used for connectivity services and gateway functions.

The simplest case of the connectivity services e.g. E-LINE service, is created between the gateway and a single VNFCI of a VNF. Note that the VNF may include other VNFCIs but for external connectivity to other NFs, a single VNFCI is responsible for that interface.

This is an example and other similar examples include two VNFs with E-LINE service to the same gateway

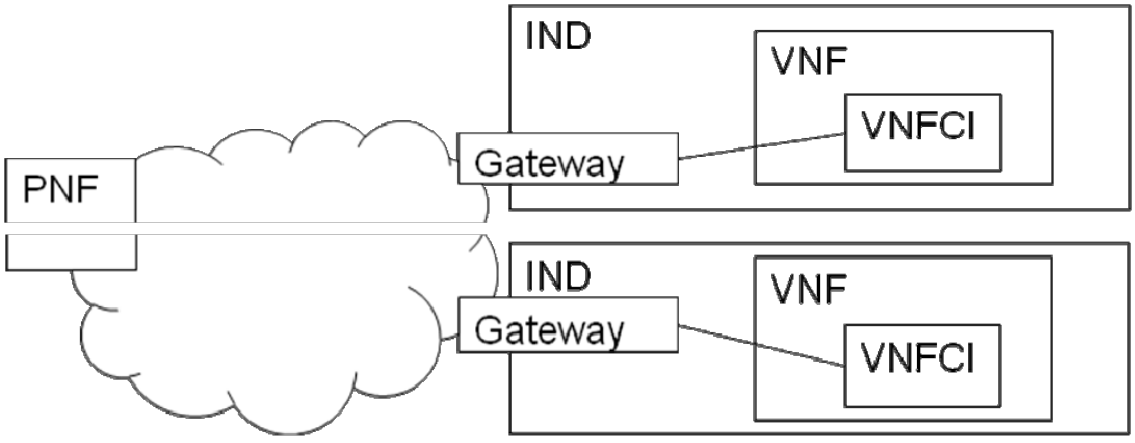

Figure 18 VNF external interfaces and VNFCI interfaces 1:1 correspondence[24]

23 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.
24 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

Another case of connectivity services is when a VNF exposes one interface to other network functions but multiple VNFCI interfaces to the gateway providing functions such as load balancing. In this case, many point-to-point services may be requested; one E-LINE between each VNFCI and the gateway. Point-to-multipoint services may be requested depending on the characteristics needed by the virtual network function. In these cases, the gateway is responsible to steer the traffic based on some protocol fields as visible on the external network.

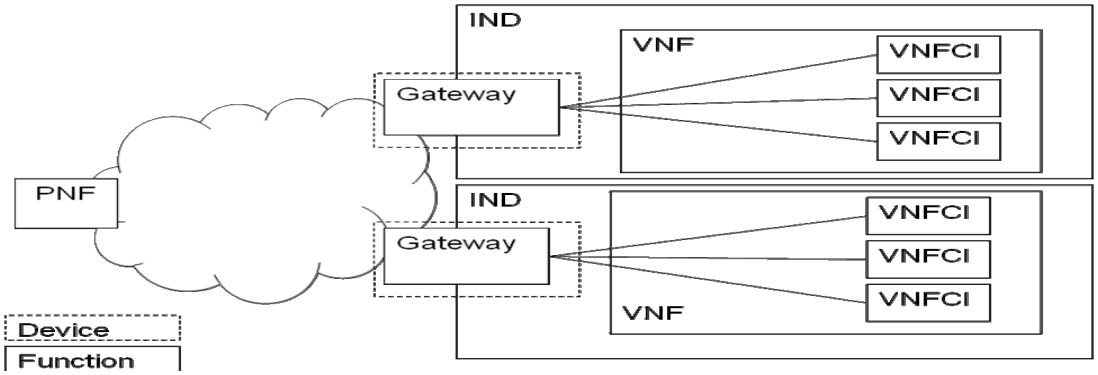

Figure 19 VNF external interfaces and VNFCI interfaces 1:N correspondence - Native Gateway[25]

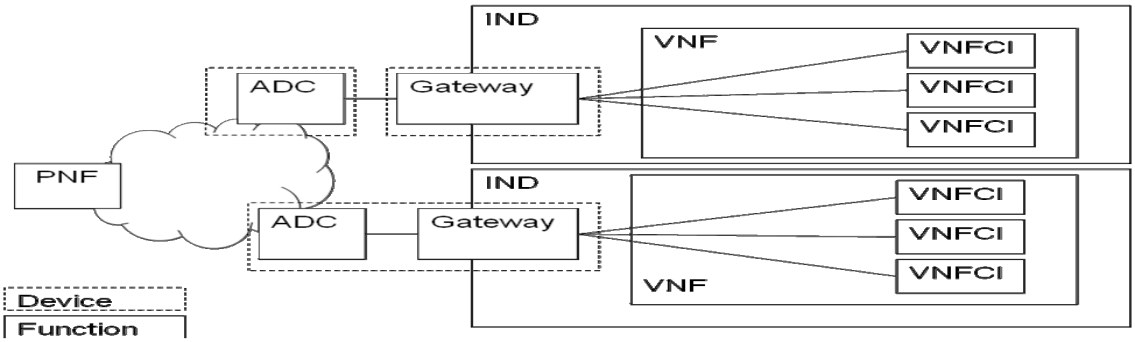

Figure 20 External interfaces and VNFCI interfaces 1:N correspondence - Extended Gateway[26]

Another example when a VNF exposes multiple interfaces to the gateway regardless of the number of VNFC instances it contains. Traffic from the VNF is mapped to the external network based on the source VNFC instance. In this case, multiple point-to-point connectivity services between the gateway and VNFCI instances may be used.

Other options based on point-to-multipoint services are also possible where it is expected that the VNFCI would select the VNIC that corresponds to the externally visible interface.

Note: "As the only virtual networks exposed to the VNFs are L2 and L3, there is no need to steer traffic towards the external network based on higher layers."[27]

25 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.
26 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.
27 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

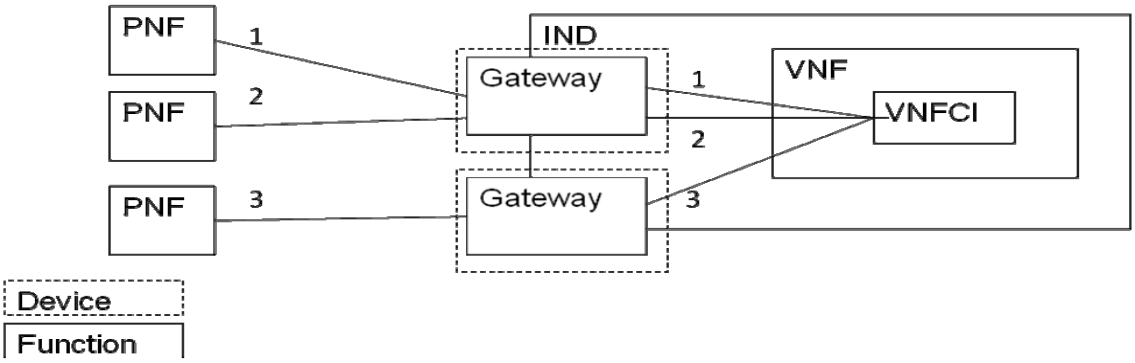

Figure 21 External interfaces and VNFCI interfaces M:1 correspondence - Native Gateway[28]

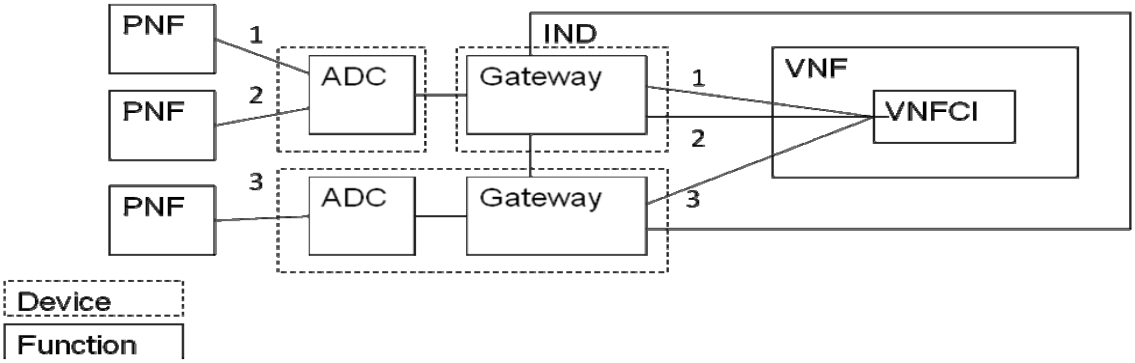

Figure 22 External interfaces and VNFCI interfaces M:1 correspondence - Extended Gateway[29]

Note: "If an application delivery controller offering load balancing functions and other traffic steering applications is included in the infrastructure, the functions that it provides that are based on layers 4-7 are considered outside the infrastructure network domain, i.e. they are not included in the virtual network service definition and the configuration is done outside the Nf-Vi reference point."[30]

4.2.3.5. Nd-N:

4.2.3.5.1. Nature of the Interface:

"The interface consists of the protocols that are exposed between the NFVI-PoPs. On-demand creation of network connectivity among NFVI-PoPs is essential to meet the dynamic nature of traffic flows produced by VNFs.

If new data plane network technologies are adopted in this interface, it may be necessary to extend protocols to provide a full set of connectivity services across domains. At the edge of each NFVI-PoP, a gateway provides the Nd-Nd data plane interface.

28 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.
29 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.
30 N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

This interface includes connectivity service virtualisation layers over carrier networks such as LANs, MANs, and WANs:

- L2 overlay models.
- L3 models.

A gateway at the Nd-Nd interface includes a number of tools and protocols including centralized control protocols, cross-domain orchestration and routing. The following list is particularly important for the Nd-Nd interface:

- Planning tools e.g. Path Computation Element (PCE), Traffic Engineering (TE) and on-demand bandwidth API.
- OAM protocols for measurement and monitoring, e.g. latency, bandwidth utilization.

- Protocols for traffic isolation, mutual authentication and authorization between different NFVI-PoPs."[31]

4.2.3.6. [Vl-Ha]/Nr:

The network hypervisor introduced on top of the network hardware resources allows creating virtual networks that allows sharing the same physical resources and controls across multiple tenants that provides abstraction of network resources at the [Vl-Ha]/Nr interface.

"Networking functions provided within the Compute Domain include:

- Physical network interface controllers (NICs).
- Virtual Ethernet Bridges (VEB) in NICs;
- Virtual Ethernet Port Aggregation (VEPA) in NICs; and
- Base L3 forwarding capabilities in the underlying kernel.

Networking functions provided within the Hypervisor Domain include:

- Virtual network interface controllers (vNICs);
- Virtual switches (vSwitches), which may be supplemented by a native or virtualized router; and
- Virtual routers (vRouters)."[32]

The connectivity between NFVI-PoPs and the transport networks is provided by the gateways. Not only that, the gateway also provides the connectivity between the virtual and

31 ETSI ISG. "ETSI GS NFV-INF 005 V1.1.1 (2014-12) Network Functions Virtualisation (NFV); Infrastructure Network Domain." 2014.
32 ETSI ISG. "ETSI GS NFV-INF 005 V1.1.1 (2014-12) Network Functions Virtualisation (NFV); Infrastructure Network Domain." 2014.

the physical networks.

4.2.3.7. Ha/CSr-Ha/Nr:

4.2.3.7.1. Nature of the Interface

This interface consists of the protocols visible between the server NIC and its adjacent Ethernet Switch or router.

Chapter No 5:

NFV Management and Orchestration

5.1 NFV Management and Orchestration:

The NFV Management and Orchestration Architectural (NFV-MANO) Framework is comprised of orchestrators, VNF managers and virtualized infrastructure managers (VIMs). It is responsible for managing and maintaining the data repositories, reference points and interfaces that are used to exchange information between all the components that make up the service to ensure the ongoing orchestration of the NFVI and VNFs.

The orchestrator is in charge of orchestrating, managing and automating the end-to-end network service that is delivered by the VNF and NFVI. Typically there is a single orchestrator that oversees the realization of the NFV service.

The VNF managers are responsible for the VNF lifecycle, including instantiation, updates, queries, scaling and termination. Multiple VNF managers can be deployed, depending on the environment a manager may be required for each VNF or may be in charge of several VNFs.

VIMs are used to control and manage the interaction of a VNF with the underlying computing, storage and network resources under its authority. VIMs are often a part of the virtualization layer, versus a separate solution, so the way an organization supports the virtualization layer (via hypervisor, OS or application) is likely going to be the way they support VIM functionality.

VIMs provide visibility into the underlying infrastructure and handle resource management, including the:

- Inventory of resources available to the NFVI
- Allocation of virtualization enablers
- Ongoing management of infrastructure resources and allocation shifts to optimize utilization and efficiency

Virtual Infrastructure Manager, NFV Orchestrator and Virtual Network Function Manager are the NFV-MANO functional blocks. We will discuss the reference points relating to NFV-MANO in the NFV reference architecture given earlier in the following section.

5.2. NFV-MANO reference points

Following are the reference points between NFV-MANO functional blocks and other functional blocks.

5.2.1. Os-Ma-nfvo:

This reference point is used to exchange information between the NFV orchestrator and OSS/BSS. It supports the management of network service descriptor and virtual network function packages along with the lifecycle management of network service instance that includes network service instantiation, updating and querying the network service instance. Network service scaling and termination is also supported. The support is included for management the lifecycle of virtual network function. It can also do the policy enforcement for NFIs, VNFCIs and NFVI resources for the purpose of authorization, resource reservation etc. It also reports the resources consumed by the VNFC instances.

5.2.2. Ve-Vnfm-em:

The information between the Element Management and the VNF manager is exchange over this reference point. It supports instantiation of the virtual network function, updating the configuration of the VNF, querying the VNF for retrieving the run-time information. VNF instance scaling out/in and scaling up/down is also supported. Not only this it also sends out the events and configurations from the Element Management to the Virtual Network Function Manager.

5.2.3. Ve-Vnfm-vnf:

The exchanges between the virtual network function and virtual network function manager is sent across this reference point. It supports instantiation of the virtual network function, updating the configuration of the VNF, querying the VNF for retrieving the run-time information. VNF instance scaling out/in and scaling up/down is also supported. It sends out the events and configuration from the VNF to VNFM. It also supports sending the heartbeats to the VNF to identify if the VNF is alive and functional.

5.2.4. Nf-Vi:

The exchanges between the VIM and NFV Infrastructure are sent across this reference point. It supports the complete lifecycle management of the VMs such as allocation of the physical resources to VM, VM migrations, VM terminations and setup/remove connectivity between the VMs and forwarding performance metrics to VIM.

5.2.5. Or-Vnfm:

The information between the NFV Orchestrator and VNF manager is exchanged over this reference point. It supports NFVI resource allocation, resource authorization, resource validation and release for the virtual network function It also supports instantiation of the

virtual network function, updating the configuration of the VNF, querying the VNF for retrieving the run-time information. VNF instance scaling out/in and scaling up/down is also supported. It sends out events about the VNF that can potentially impact the network service instance.

5.2.6. Or-Vi:

The information between the NFV Orchestrator and VIM is exchanged over this reference point. It supports NFVI resource allocation and reservations. It also supports the addition, deletion and updating the VNF operating system image. It sends out the performance metrics of the NFVI to the NFV orchestrator.

5.2.7. Vi-Vnfm:

The information between the VNF manager and VIM is exchanged over this reference point and it supports the NFVI resources reservation information retrieval and the NFVI resource allocation. It sends out events, measurement results, and usage records related to NFVI resources consumed by a VNF.

Chapter No 6:

NFV Use Case – Virtual CPE

6.1 Virtual Network Function as a Service:

The enterprise organizations require several services at the edge of the branch. One way to address the requirement is to position a dedicated appliance for each network function and connect them inline to the traffic so that it may give the desired results. The challenge with this approach is that the cost of the dedicated hardware per network function is high which brings down the price performance significantly while increasing the ROI. Clearly, this approach seems inflexible. Another approach is to deploy an integrated services devices e.g. Cisco's ISR and Juniper's SRX etc. to the edge and exploit services like routing, firewalling and content filtering. The challenge with this approach is the scale and performance that these integrated hardware appliances offer. It may not scale for very large enterprises. This forces a lot of the enterprises to adapt to a newer model and move the network services from the dedicated appliances to the virtualized appliance setup over COTS. Two models are available. It may be hosted on the customer premises or it may reside in the providers cloud.

The service provider can setup the VNF by utilizing its own NFVI and can provide services to the enterprise. If the enterprise requires multiple services per branch then the service provider can potentially setup multiple VNFs per branch for the customer and create service chains for steering traffic from the desired functions. Making the VNF functionality available to the enterprise, as a service is synonymous to the cloud computing idea of Software as a Service. The benefits of virtualization don't stop here. The provider can also use the same infrastructure to scale its own infrastructure e.g. Provider edge etc. One key thing to notice is virtualization of the PE and virtualization of the CPE are independent of each other. Service Providers typically terminate multiple customers on a single Provider Edge whereas every customer has dedicated Customer Premises Equipment; therefore virtualizing the CPE makes more sense to be done prior to virtualizing the PE.

Figure 23 and Figure 24 depicts the difference between the legacy network edge and virtualized network edge.

Figure 23 Legacy CPE network infrastructure[33]

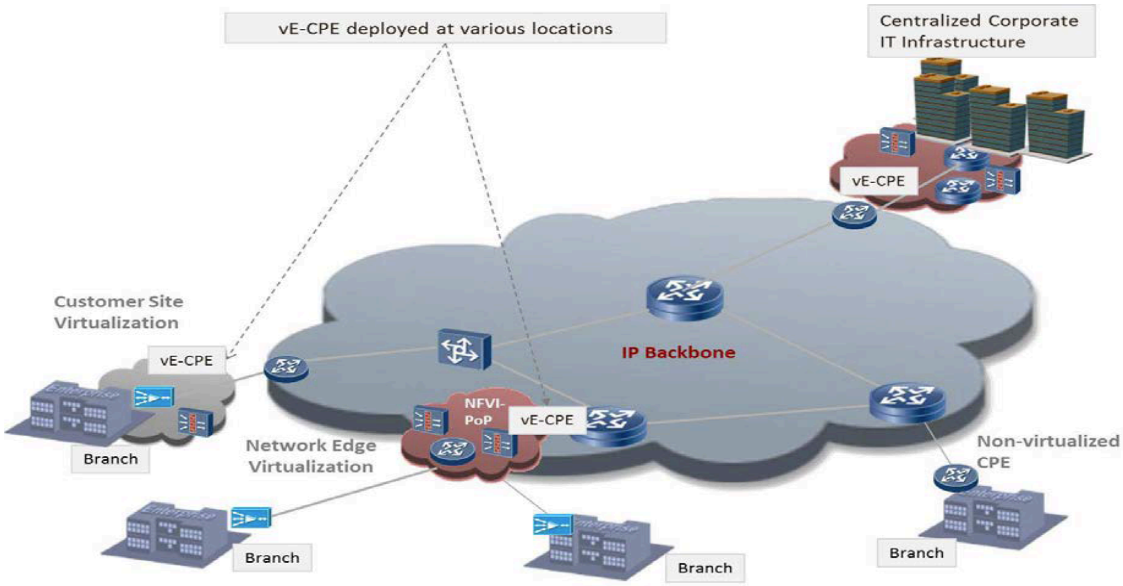

Figure 24 Virtualized infrastructure[34]

Let us analyze the impact in the enterprise traffic after virtualizing the CPE. An L2/L3 switch at the customer premises handles the local branch traffic and stiches the customers network to the service providers. Except for the local branch traffic all the traffic is steered up to the service provider's cloud. Example functionality, which can be provided by the VCPE, includes routing, VPN termination, QoS support, DPI, NG-FW, and a WOC (WAN Optimization Controller). Figure 25 contrasts the non-virtualized customer site with virtualized customer site. The dotted purple lines indicate where this VCPE functionality is located.

33 E. G. N. 001 V1.1.1, "Network Functions Virtualisation (NFV); Use Cases," IEEE Netw., vol. 1, no. 5, pp. 1–50, 2013
34 E. G. N. 001 V1.1.1, "Network Functions Virtualisation (NFV); Use Cases," IEEE Netw., vol. 1, no. 5, pp. 1–50, 2013.

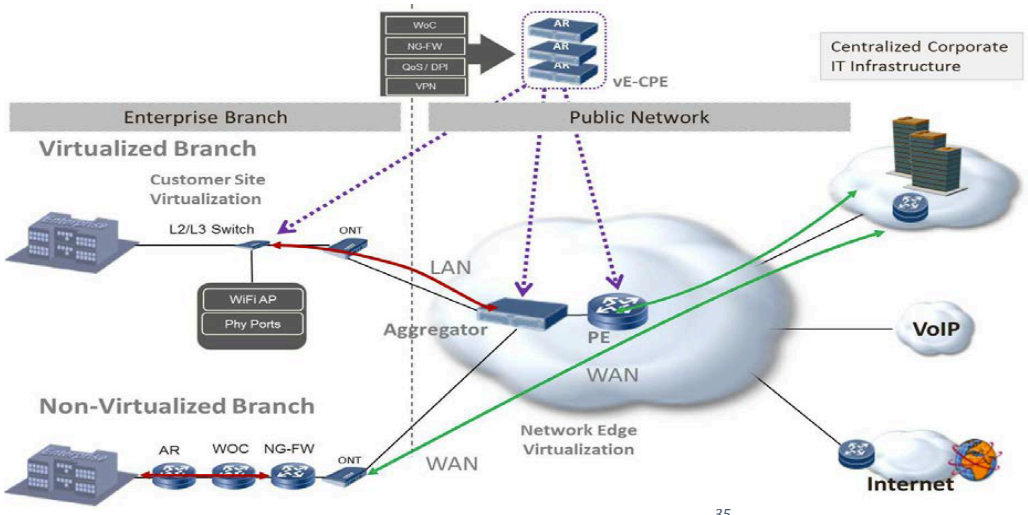

Figure 25 Cloud CPE deployment scenario[35]

6.2. Example Deployment Scenarios

6.2.1. Simple VCPE:

In the most simple use case, the L3 functionality is brought back from the CPE into a virtualized CPE running as a routing-instance on the provider edge (PE) router. Depending upon the architecture of the operator's network and the services to which the customer subscribes, the routing-instance could either be a virtual-router or a VRF (L3VPN).

An Ethernet switch remains on the customer premises with the uplink Ethernet configured to receive traffic on a particular VLAN as defined by the operator. This VLAN is configured on the PE as an interface within the routing-instance. The PE can be configured to provide either a local DHCP service or DHCP relay service (or a local DHCP server can be maintained on the customer premises). The address pools in multiple routing-instances may overlap.

If the routing-instance is part of a VRF with connectivity to a central HQ site at which all the relevant services are offered, then no more need be done to support this function. If an Internet access service is required then a physical or virtual node, which is attached to the customer routing-instance and the Internet routing-instance and which offers a NAT service, must advertise a default route into the customer routing-instance.

In this scenario, all configurations are applied manually to the PE router through the CLI or through a Selfcare Portal. There is little or no automation of the VCPE setup and only simple management of the underlying PE platform and its associated routing-instances

35 E. G. N. 001 V1.1.1, "Network Functions Virtualisation (NFV); Use Cases," IEEE Netw., vol. 1, no. 5, pp. 1–50, 2013.

6.2.2. VCPE with Service Chain:

In this use case, more of the network functions are migrated into the operator's network. In order to achieve the flexibility required to meet the array of different customer requirements without the requirement for racks and racks of appliances, it is necessary to virtualize the network functions so that they can be run as VMs either inside the PE or in a Datacenter on COTS hardware with an SDN controller.

The underlying VCPE is identical to that proposed above. This provides the basis for the entry point into the service chain. If the VNFs run locally on the PE then the routing-instances can still be either a virtual-router or a VRF. However, if the VNFs are running in the Datacenter, the routing-instance will likely always be a VRF thus providing a virtualized transport network between the PE and the Datacenter Gateways.

Within the datacenter, the SDN orchestration platform is used to build the chain of services on the COTS hardware as a series of virtual machines with a set of interfaces connected with virtualized links. It is common in such a service chaining environment to have an "east-bound", a "west-bound" and a management network interface as a minimum. These virtualized links can be constructed as virtual point-to-point links or as virtual multi-access networks. Thus, it is possible to build redundant and scalable service chains.

The diagram below shows two customer networks indicated by the two layer 2 CPE nodes on the far left. The upper customer site has two VLANs, each of which represents a unique service terminated on the PE. This could be delivered as either a single VRF with multiple incoming links and bridging between the links or as multiple VRFs where each service is strictly separated. In the diagram below, the VPN1 is attached at the DC Gateway to a load-balanced web caching service before passing through statefull firewalls (including NAT) to the Internet.

The service chain for VPN2 is not shown. Similarly, if the second customer requires a similar service to VPN1, it will have a uniquely provisioned set of VNFs providing that service with connectivity back into the Internet through the main routing instance on the DC GW.

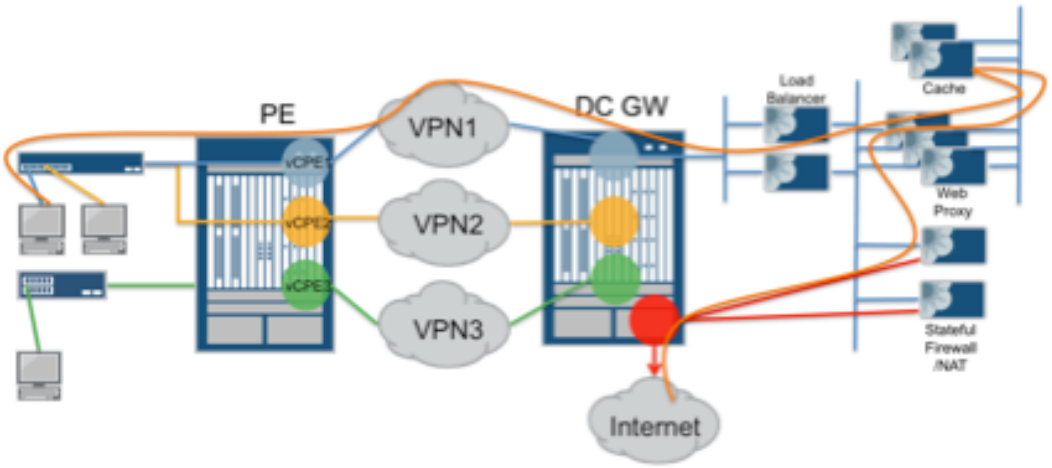

Figure 26 VCPE with service chain[36]

6.2.3. VCPE with Policy Based Service Chain:

In the previous use case, it is assumed that all traffic within a single routing-instance will follow the same path through a service chain. However, such an approach isn't always the most efficient since it is common for certain functions only to apply to particular traffic types or particular subscribers. By introducing a policy function, it's possible to first identify the traffic requiring treatment by the service chain and only forward traffic through the service chain that requires that treatment. This can significantly improve the efficiency of the virtualized network functions since they don't have to handle traffic for which they perform no function.

Some of these policies can be statically configured, for example, we might send only traffic destined to TCP port 80 to a web proxy/cache rather than loading those VMs with 100% of the traffic leaving a particular customer's network. Such a policy can be created as a firewall filter applied to the ingress interface of the customer.

In a subscriber management environment, where each user authenticates, it is also possible to introduce per-subscriber policy using the Traffic Direction Function (TDF). In this case, the subscriber's username is bound to an IP address by the BNG. Service chains in this environment tend to be more closely tied to a particular function than a particular access point, so the service chains are often shared by many subscribers with no specific link to each other except for the service to which they subscribe. In the same way as the entry point to a service chain can be associated with a particular VRF in the previous use case, in this case, traffic can be leaked into service chain specific VRFs based on policy. Thus, as an example, a

36 Contrail and IBM SCO Cloud cCPE Reference Architecture.

child using a system may have all their web traffic directed through a web proxy system, which implements a filtering service; while an adult's traffic would not be filtered using that service and would, therefore, pass direct to the Internet (or via a different service chain).

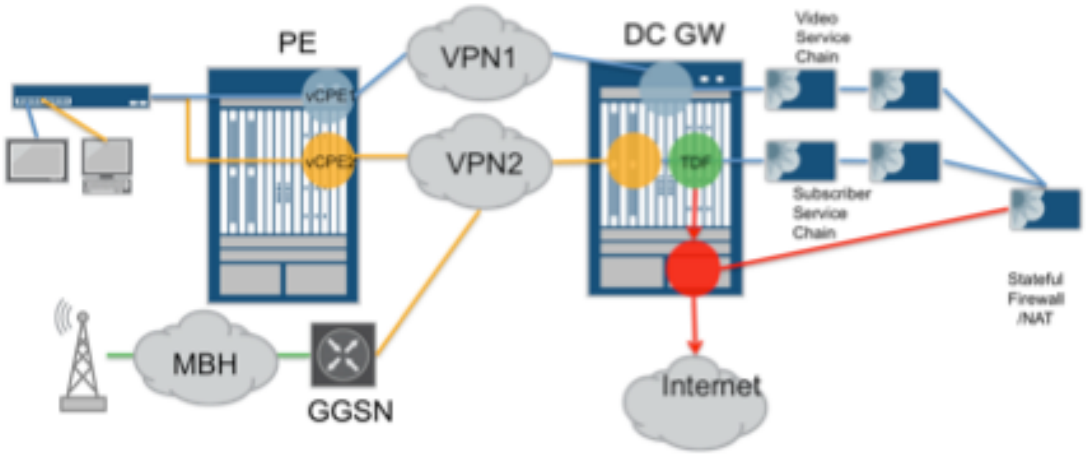

Figure 27 VCPE with policy based service chain[37]

6.3. Juniper Networks Implementation of VCPE:

6.3.1. Centralized Cloud CPE Deployment Model:

The cloud CPE model from Juniper Networks de-positions the network services from the customer premises and moves them to the service provider cloud. The new services can be orchestrated through a customer portal on demand. Contrail service orchestrator does the all the complex service orchestration, service chaining[38] and life cycle management of the VNF automatically to deliver scalable multitenant services.

Centralized Cloud CPE considerably simplifies the positioning of managed services, letting telco providers offer on-demand availability, low-risk purchasing, personalized marketplace options, and highly differentiated services. Centralizing capital asset investments helps service providers quickly improve efficiency and ROI metrics.

37 Contrail and IBM SCO Cloud cCPE Reference Architecture.
38 Service Chaining is the connectivity across multiple VNFs to provide a network service

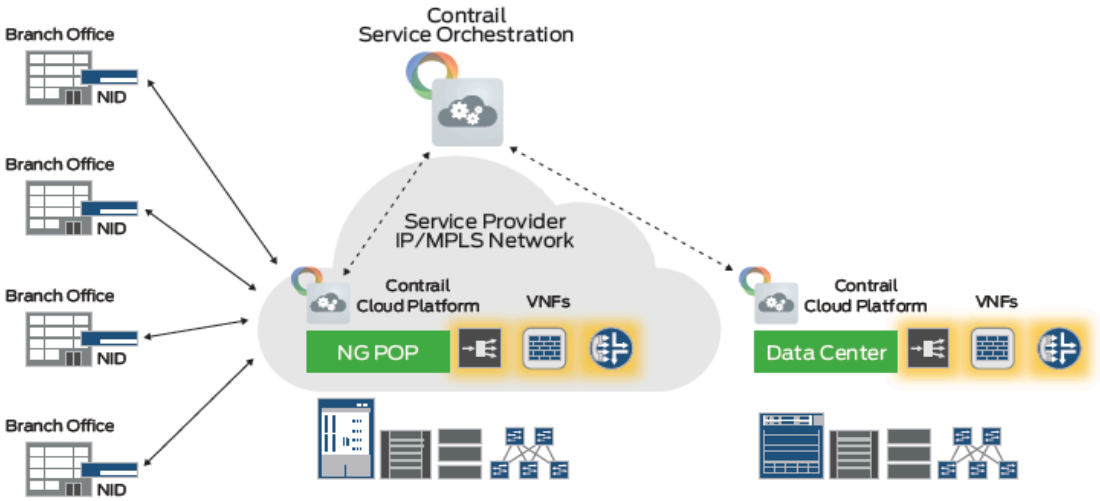

Figure 28 Cloud CPE model - Juniper Networks[39]

6.3.2. Distributed Cloud CPE Deployment Model:

Distributed cloud CPE model from Juniper Networks has x86 based NFX platform that allows setting up to 8 VNFs currently on top of single hardware. Currently the available VNFs from Juniper Networks are VSRX for security and VMX for routing. It also promises to integrate with the third party VNFs as well. In this model the NFX is positioned in every branch and the VNFs can be orchestrated by using the orchestration and management software from Juniper Networks (Service Orchestrator or Service Maestro). It gives the flexibility of creating services dynamically, which are delivered instantaneously.

By setting up VSRX on top of NFX platform, the customer can benefit from the edge security and protect the business-sensitive environment. Working together, the NFX Series and vSRX simultaneously virtualize IP routing for site survivability, meeting the strict demands of always-on application availability.

This flexibility eliminates traditional service silos and enables enterprise customers to scale their network services while addressing market seasonality and business requirements with a virtualized, software-driven approach.

39 "Production-Ready NFV Solution - Juniper Networks." [Online]. Available: https://www.juniper.net/us/en/solutions/nfv/. [Accessed: 01-Mar-2016].

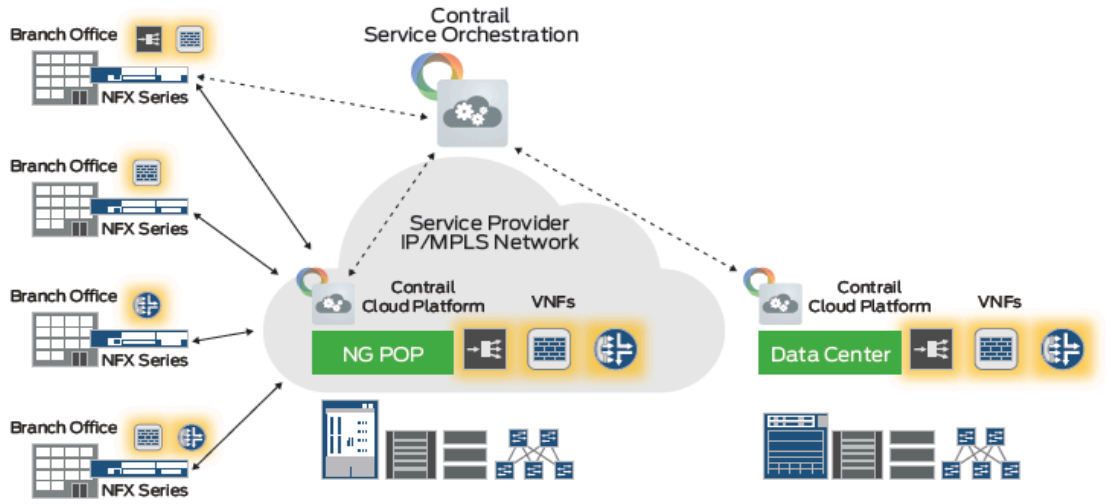

Figure 29 Distributed CPE model - Juniper Networks40

6.3.3. Solution Overview:

Contrail service orchestration is the network management and orchestration platform by Juniper Networks. It has several components that work together to provide the end-to-end NFV-MANO functionality. Below figure represents the entire solution and the functionality of each component.

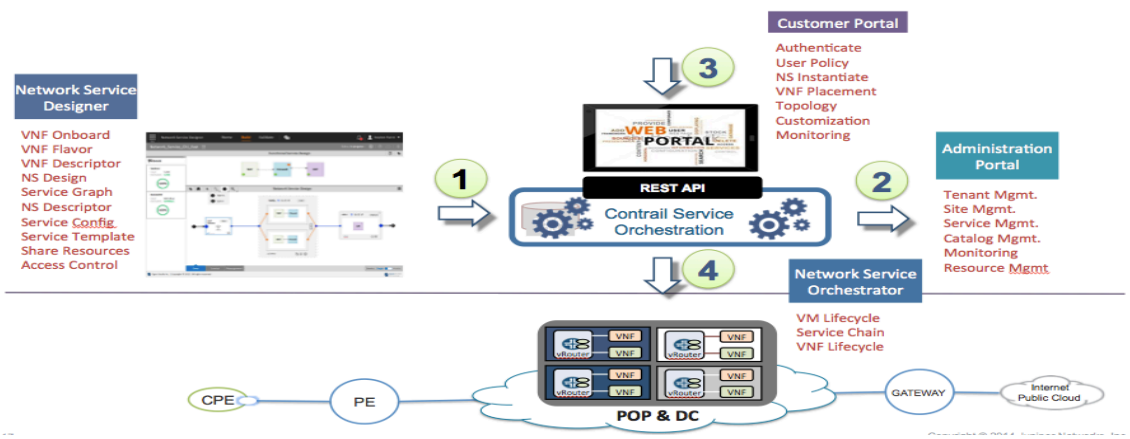

Figure 30 End-to-End virtual CPE solution - Juniper Networks[41]

40 "Production-Ready NFV Solution - Juniper Networks." [Online]. Available: https://www.juniper.net/us/en/solutions/nfv/. [Accessed: 01-Mar-2016].

41 Cloud CPE Solution Enablement Presentation Playbook – Juniper Networ

List of abbreviations:

API: Application Programming Interface
BGP: Border Gateway Protocol
COTS: Commercial off the Shelf
DHCP: Dynamic Host Configuration Protocol
DNS: Domain Name Server
IPS: Intrusion Prevention System
NAT: Network Address Translation
DPI: Deep Packet Inspection
EM: Element Management
EVPN: Ethernet VPN
FCAPS: Fault, Configuration, Accounting, Performance, and Security
E-Line: Ethernet Line
ETSI: European Telecommunications Standards Institute
NFVI-PoP: Network Functions Virtualisation Infrastructure Point of Presence
VM: Virtual Machine
PNF: Physical Network Function
VNF: Virtual Network Function
SWA Software Architecture
VNFI: Virtual Network Function Instance
VNFC: Virtual Network Function Component
NFVI: Network Function Virtualization Infrastructure
MANO: Management and Orchestration
VNFM: Virtual Network Function Manager
EMS: Element Management System
VIM: Virtual Infrastructure Manager
CPE: Customer Premises Equipment
VCPE: Virtual Customer Premises Equipment
PE: Provider Edge
VPE: Virtual Provider Edge
VEB: Virtual Ethernet Bridge
VNICs: Virtual Network Interface Card

References:

[1] M. Chiosi, S. Wright Bell Canada, J. Erfanian, B. B. Smith, B. Briscoe, A. Reid, P. Willis CableLabs, D. Clarke, C. Donley CenturyLink, M. Bugenhagen, J. Feger, J. Benitez, N. Fischbach Deutsche Telekom, K. Martiny, U. Michel DOCOMO, T. Nakamura, J. Triay Marques KDDI, K. Ogaki, T. Matsuzaki KPN, S. Zhang, A. K. de Boer, K. Ok, E. Kyoung PAIK NTT, K. Shimano, T. Shimizu Ooredoo, M. Stura Orange, B. Chatras, C. Kolias Portugal Telecom, J. Carapinha, A. S. Gamelas Telecom, D. Lee, J. Han Park Softbank, R. Wakikawa, K. Nishi, S. Matsushima Sprint, L. Laporte, F. Feisullin Swisscom, M. Brunner Telecom Italia, E. Demaria, A. Pinnola Telenor, P. Waldemar, G. Millstein Telefonica, D. López, F. Javier Ramón Salguero Telstra, D. Kirkham Turk Telekom Argela, M. Ergen, M. Ahmet Karaman, A. Ulas, E. Lokman Verizon, N. Khan, R. Morera Vodafone, S. Sabater, A. Neal Windstream, and A. Nichols, "Network Functions Virtualisation (NFV) CONTRIBUTING ORGANISATIONS & AUTHORS PUBLICATION DATE," no. 1, pp. 1–20, 2015.

[2] N. Domain, "GS NFV-INF 005 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Network Domain," vol. 1, pp. 1–53, 2014.

[3] M. Ersue and ETSI NFV MANO WG, "ETSI NFV Management and Orchestration - An Overview Virtualization as a Paradigm," *IETF 88 Proc.*, 2013.

[4] ETSI, "Network Functions Virtualisation (NFV); Virtual Network Functions Architecture," vol. 1, pp. 1–93, 2014.

[5] ETSI, "Network Function Virtualisation Updated White Paper," *Terminol. Main Concepts NFV*, no. 1, pp. 1–16, 2013.

[6] ETSI, E. E. T. S. Institute, ETSI, and E. E. T. S. Institute, "GS NFV-INF 003 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Compute Domain," *Etsi*, vol. 1, pp. 1–57, 2014.

[7] ETSI, I. Overview, ETSI, and I. Overview, "GS NFV-INF 001 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure Overview," *Etsi*, vol. 1, pp. 1–59, 2015.

[8] Etsi and J. Quittek, "GS NFV-MAN 001 - V1.1.1 - Network Functions Virtualisation (NFV); Management and Orchestration," vol. 1, pp. 1–184, 2014.

[9] European Telecommunications Standards Institute, "Network Functions Virtualisation (NFV)\,;~Terminology for Main Concepts in NFV," vol. 1, no. GS~NFV~003~–~V1.1.1, pp. 1–13, 2013.

[10] A. Introduction, C. Action, M. Chiosi, D. Clarke, P. Willis, A. Reid, J. Feger, M. Bugenhagen, W. Khan, M. Fargano, C. Cui, H. Deng, D. Telekom, and U. Michel, "Network Functions Virtualisation,An Introduction, Benefits, Enablers, Challenges & Call for Action," *Citeseer*, no. 1, pp. 1–16, 2012.

[11] J. Networks, "The NFV Service Edge," pp. 1–4.

[12] N. F. V Poc, "What is an NFV PoC?"

[13] M. Report, “SDxCentral Network Functions Virtualization Report,” 2015.

[14] A. Service, “Juniper Networks Cloud CPE Solution,” pp. 1–5.

[15] G. Specification, “GS NFV-INF 007 - V1.1.1 - Network Functions Virtualisation (NFV); Infrastructure; Methodology to describe Interfaces and Abstractions,” vol. 1, pp. 1–30, 2014.

[16] STANDARDS INSTITUTIONS, “GS NFV 002 - V1.2.1 - Network Functions Virtualisation (NFV); Architectural Framework,” *Tbd*, vol. 1, pp. 1–21, 2014.

[17] E. G. N. 001 V1.1.1, “Network Functions Virtualisation (NFV); Use Cases,” *IEEE Netw.*, vol. 1, no. 5, pp. 1–50, 2013.

[18] “Contrail and IBM SCO Cloud cCPE Reference Architecture.” .

[19] “Production-Ready NFV Solution - Juniper Networks.” [Online]. Available: https://www.juniper.net/us/en/solutions/nfv/. [Accessed: 01-Mar-2016].

[20] “Cloud Computing Patterns | Mechanisms | Virtual Infrastructure Manager.” [Online]. Available: http://cloudpatterns.org/mechanisms/virtual_infrastructure_manager. [Accessed: 01-Mar-2016].

[21] “What is Service Chaining? - Packet Pushers -.” [Online]. Available: http://packetpushers.net/service-chaining/. [Accessed: 01-Mar-2016].

[1]–[4], [4]–[7], [7]–[21]