# ETSI GS NFV 002 V1.1.1 (2013-10)

Group Specification

# Network Functions Virtualisation (NFV);
# Architectural Framework

***Disclaimer***

This document has been produced and approved by the Network Functions Virtualisation (NFV) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

Reference
DGS/NFV-0010

Keywords
architecture, NFV

*ETSI*

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

## *Important notice*

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, please send your comment to one of the following services:
http://portal.etsi.org/chaircor/ETSI_support.asp

## *Copyright Notification*

# Contents

# Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://ipr.etsi.org).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

# Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification Group (ISG) Network Functions Virtualisation (NFV).

# 1 Scope

The present document describes the high-level functional architectural framework and design philosophy of virtualised network functions and of the supporting infrastructure. The document also defines the scope of the NFV Industry Specification Group (ISG) activities to realize this framework.

# 2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the reference document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at http://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

## 2.1 Normative references

The following referenced documents are necessary for the application of the present document.

[1] ETSI GS NFV 003: "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV".

[2] ETSI GS NFV 004: "Network Functions Virtualisation (NFV); Virtualisation Requirements".

[3] ETSI GS NFV 001: "Network Functions Virtualisation (NFV); Use Cases".

## 2.2 Informative references

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] NFV White paper: "Network Functions Virtualisation, An Introduction, Benefits, Enablers, Challenges & Call for Action. Issue 1".

NOTE: Available at http://portal.etsi.org/NFV/NFV_White_Paper.pdf.

[i.2] ETSI TS 123 002: "Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); LTE; Network architecture (3GPP™ TS 23.002)".

NOTE: Available at http://www.3gpp.org/ftp/Specs/html-info/23002.htm.

# 3 Definitions and abbreviations

## 3.1 Definitions

For the purposes of the present document, the terms and definitions given in GS NFV 003 [1] apply.

## 3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

| | |
|---|---|
| BSS | Business Support System |
| CDN | Content Delivery Network |
| CMS | Cloud Management System |
| COTS | Commercial-Off-The-Shelf |
| DHCP | Dynamic Host Configuration Protocol |
| E2E | End-to-End |
| EMS | Element Management System |
| EPC | Evolved Packet Core |
| GW | Gateway |
| IT | Information Technology |
| LAN | Local Area Network |
| MME | Mobility Management Entity |
| NAS | Network Attached Storage |
| NAT | Network Address Translation |
| NF | Network Function |
| NFV | Network Functions Virtualisation |
| NFVI | NFV Infrastructure |
| NFVI-PoP | NFV Infrastructure Point of Presence |
| NMS | Network Management System |
| OS | operating system |
| OSS | Operations Support System |
| PGW | Packet Data Network Gateway |
| PNF | Physical Network Function |
| RGW | Residential Gateway |
| SDO | Standards Development Organization |
| SGW | Serving Gateway |
| VLAN | Virtual Local Area Network |
| VM | Virtual Machine |
| VNF | Virtualised Network Function |
| VNF-FG | VNF Forwarding Graph |
| VPLS | Virtual Private LAN Service |

# 4 Overview

## 4.1 Document Structure and Purpose

The present document is structured as follows: Clause 4 presents the main purpose, objectives and approach of NFV. Clause 5 presents the NFV framework and scope. Clause 6 introduces main concepts behind the virtualisation process and its impact on end-to-end network services. Clause 7 presents the NFV reference architectural framework. Clause 8 describes some future study items and focus areas in NFV. Finally, clause 9 concludes and proposes some recommendations to realize the NFV paradigm.

Network Functions Virtualisation is a powerful emerging technique with widespread applicability. The initial focus is on the subset of network services described in GS NFV 001 [3]. GS NFV 001 [3] further identifies a number of use cases for the virtualisation of network functions. The present document supports these use cases, although it does not explicitly map them to the NFV framework.

The purpose of the present document is to abstract the overall problem space in such a way that the requirements and aspects unique to NFV [2] are clearly identified so that the work can be scoped and organized. The resulting network architectural framework aims at positioning NFV among relevant telecommunications and IT industry stakeholders, including network operators, solution vendors, service integrators and providers, as well as serving as a reference to NFV ISG working groups.

Another purpose of the present document is to provide guidance to the industry Standards Development Organizations (SDOs) to align existing network related specifications with the NFV architectural framework outlined in the present document. Any further standardization of network functions, architecture and interfaces that are required to properly operate in a virtualised environment will be carried out in relevant SDOs. The resulting standards are expected to support the NFV high-level architectural requirements for both intra- and inter-provider domains.

## 4.2 Summary of Objectives of NFV

A detailed description of the NFV objectives is contained in [i.1]. Briefly, high-level objectives of NFV are:

- Improved capital efficiencies compared with dedicated hardware implementations. This is achieved by using commercial-off-the-shelf (COTS) hardware (i.e. general purpose servers and storage devices) to provide Network Functions (NFs) through software virtualisation techniques. These network functions are referred as Virtualised Network Functions (VNFs). Sharing of hardware and reducing the number of different hardware architectures in a network also contribute to this objective.

- Improved flexibility in assigning VNFs to hardware. This aids both scalability and largely decouples functionality from location, which allows software to be located at the most appropriate places, referred to in the present document as NFVI-PoPs [1], e.g. at customers' premises, at network exchange points, in central offices, data centres, etc. This enables time of day reuse, support for test of alpha/beta and production versions, enhance resilience through virtualisation, and facilitates resource sharing.

- Rapid service innovation through software-based service deployment.

- Improved operational efficiencies resulting from common automation and operating procedures.

- Reduced power usage achieved by migrating workloads and powering down unused hardware.

- Standardized and open interfaces between virtualised network functions and the infrastructure and associated management entities so that such decoupled elements can be provided by different vendors.

## 4.3 Approach

Current networks are comprised of diverse network functions. These network functions are connected, or chained, in a certain way in order to achieve the desired overall functionality or service that the network is designed to provide. Most current network services are defined by statically combining network functions in a way that can be expressed using an NF Forwarding Graph or NF Set construct. A major change brought by NFV is that virtualisation enables additional dynamic methods rather than just static ones to construct and manage the network function graphs or sets combining these network functions. A major focus of NFV is to enable and exploit the dynamic construction and management of network function graphs or sets, and their relationships regarding their associated data, control, management, dependencies and other attributes as detailed in clause 5. For example, the term Network Function Forwarding Graph focuses on relations that express connectivity between network functions and the aspects that virtualisation introduces as detailed in clause 6.

Furthermore, future network services may be quite different to current network services. These services will be comprised of a diverse set of non-virtualised and/or virtualised network functions, the latter supported by virtualised computing and network infrastructure, requiring interoperability among legacy and NFV-based network domains. The overall service attributes, in particular reliability, availability, manageability, security and performance will depend on the individual (virtualised) network function attributes, as well as how these functions are connected. These attributes are not necessarily independent. It is therefore important to formulate an architecture that supports the diversity of network functions that can potentially be virtualised.

The approach taken here is to describe how the elements necessary to realize NFV can be implemented in a standardized way to enable interoperability. This allows different VNFs to be deployed over the virtualised infrastructure to support E2E network services and be applicable to diverse use cases and operator network scenarios with minimal integration effort and maximum reuse.

Virtualisation means that a network function and part of the infrastructure are implemented in software, and hence, the NFV software architecture is an important aspect of the NFV architectural framework.

# 5 NFV Framework and Scope

## 5.1 General

In non-virtualised networks, NFs are implemented as a combination of vendor specific software and hardware, often referred to as network nodes or network elements. Network Functions Virtualisation represents a step forward for the diverse stakeholders in the telecommunication network environment. As such, NFV introduces a number of differences in the way network service provisioning is realized in comparison to current practice. In summary, these differences can be listed as:

- *Decoupling software from hardware:* As the network element is no longer a collection of integrated hardware and software entities, the evolution of both are independent of each other. This enables the software to progress separately from the hardware, and vice versa.
- *Flexible network function deployment:* The detachment of software from hardware helps reassign and share the infrastructure resources, thus together, hardware and software, can perform different functions at various times. Assuming that the pool of hardware or physical resources is already in place and installed at some NFVI-PoPs, the actual network function software instantiation can become more automated. Such automation leverages the different cloud and network technologies currently available. Also, this helps network operators deploy new network services faster over the same physical platform.
- *Dynamic operation:* The decoupling of the functionality of the network function into instantiable software components provides greater flexibility to scale the actual VNF performance in a more dynamic way and with finer granularity according to the actual traffic which, for instance, the network operator needs to provision capacity.

Based upon the previous definitions and description, the present document provides an architectural framework that focuses on the aspects unique to virtualisation by:

- Outlining an architecture that supports VNF operation across different hypervisors and computing resources and which provides access to shared storage, computation, and physical/virtual networking.
- Outlining a software architecture with VNFs as building blocks to construct VNF Forwarding Graphs.
- Interfacing management and orchestration of NFV with other management systems, such as EMS, NMS, and OSS/BSS.
- Supporting a range of network services with different reliability and availability levels leveraging virtualisation techniques.
- Ensuring the virtualisation does not cause any new security threat.
- Addressing performance related issues unique to virtualisation.
- Minimizing the interworking impact between virtualised and non-virtualised network functions.
- Leveraging existing data centre technology.

A number of aspects of the network functions and network infrastructure are common to Physical and Virtualised Network Functions, and therefore are out of scope of the NFV reference architectural framework. At a high level, these are:

- The specifics of the network functions themselves, their interface protocols, including any syntax and semantics of the internally or externally stored state, as well as management functions related to the functionality performed by the NF.
- Direct control, operation and management of physical network infrastructure.
- The actual packet flow, control, operation and management of the E2E network service. These should be independent of whether specific end points, network functions and/or network infrastructure are virtualised or not.

- Implementation details of the architecture itself.

## 5.2 High-Level NFV Framework

Network Functions Virtualisation envisages the implementation of NFs as software-only entities that run over the NFV Infrastructure (NFVI). Figure 1 illustrates the high-level NFV framework. As such, three main working domains are identified in NFV:

- Virtualised Network Function, as the software implementation of a network function which is capable of running over the NFVI.
- NFV Infrastructure (NFVI), including the diversity of physical resources and how these can be virtualised. NFVI supports the execution of the VNFs.
- NFV Management and Orchestration, which covers the orchestration and lifecycle management of physical and/or software resources that support the infrastructure virtualisation, and the lifecycle management of VNFs. NFV Management and Orchestration focuses on all virtualisation-specific management tasks necessary in the NFV framework.

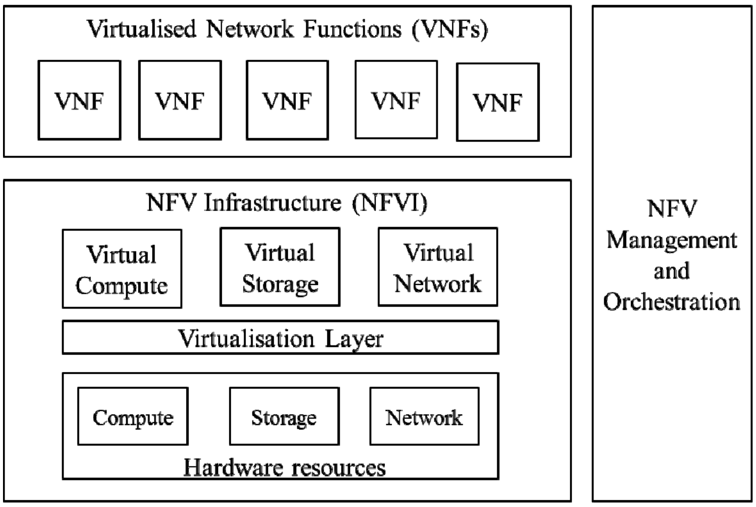

**Figure 1: High-level NFV framework**

The NFV framework enables dynamic construction and management of VNF instances and the relationships between them regarding data, control, management, dependencies and other attributes. To this end, there are at least three architectural views of VNFs that are centred around different perspectives and contexts of a VNF. These perspectives include:

- a virtualisation deployment/on-boarding perspective where the context can be a VM,
- a vendor-developed software package perspective where the context can be several inter-connected VMs and a deployment template that describes their attributes,
- an operator perspective where the context can be the operation and management of a VNF received in the form of a vendor software package.

Within each of the above contexts, at least the following relations exist between VNFs:

- VNF Forwarding Graph (VNF-FG) covers the case where network connectivity does matter, e.g. a chain of VNFs in a web server tier (e.g. firewall, NAT, load balancer), as described in use case *Service Chains (VNF Forwarding Graphs)* in GS NFV 001 [3].
- VNF Set covers the case where the connectivity between VNFs is not specified, e.g. virtualised residential gateway as described in the use case *Virtualisation of the Home Environment* in GS NFV 001 [3].

Within the present document, the specific context of a VNF and the forwarding aspects of a VNF-FG are considered. Other VNF relations are for further study.

# 6 Network Services in NFV

## 6.1 Introduction to Network Services in NFV

An end-to-end network service (e.g. mobile voice/data, Internet access, a virtual private network) can be described by an NF Forwarding Graph [1] of interconnected Network Functions (NFs) and end points.

## 6.2 Virtualisation of Functional Blocks for Network Services

A network service can be viewed architecturally as a forwarding graph of Network Functions (NFs) interconnected by supporting network infrastructure. These network functions can be implemented in a single operator network or interwork between different operator networks. The underlying network function behaviour contributes to the behaviour of the higher-level service. Hence, the network service behaviour is a combination of the behaviour of its constituent functional blocks, which can include individual NFs, NF Sets, NF Forwarding Graphs, and/or the infrastructure network.

The end points and the network functions of the network service are represented as nodes and correspond to devices, applications, and/or physical server applications. An NF Forwarding Graph can have network function nodes connected by logical links that can be unidirectional, bidirectional, multicast and/or broadcast. A simple example of a forwarding graph is a chain of network functions. An example of such an end-to-end network service can include a smartphone, a wireless network, a firewall, a load balancer and a set of CDN servers. The NFV area of activity is within the operator-owned resources. Therefore, a customer-owned device, e.g. a mobile phone is outside the scope as an operator cannot exercise its authority on it. However, virtualisation and network-hosting of customer functions is possible and is in the scope of NFV (e.g. see use cases *Virtual Network Platform as a Service (VNPaaS)* and *Virtualisation of the Home Environment* in GS NFV 001 [3]).

Figure 2 illustrates the representation of an end-to-end network service that includes a second nested NF Forwarding Graph as indicated by the network function block nodes in the middle of the figure interconnected by logical links. The end points are connected to network functions via network infrastructure (wired or wireless), resulting in a logical interface between the end point and a network function. These logical interfaces are represented in the figure with dotted lines. In Figure 2, the outer end-to-end network service is made up of End Point A, the inner NF Forwarding Graph, and End Point B, while the inner NF Forwarding Graph is composed of network functions NF1, NF2 and NF3. These are interconnected via logical links provided by the Infrastructure Network 2.

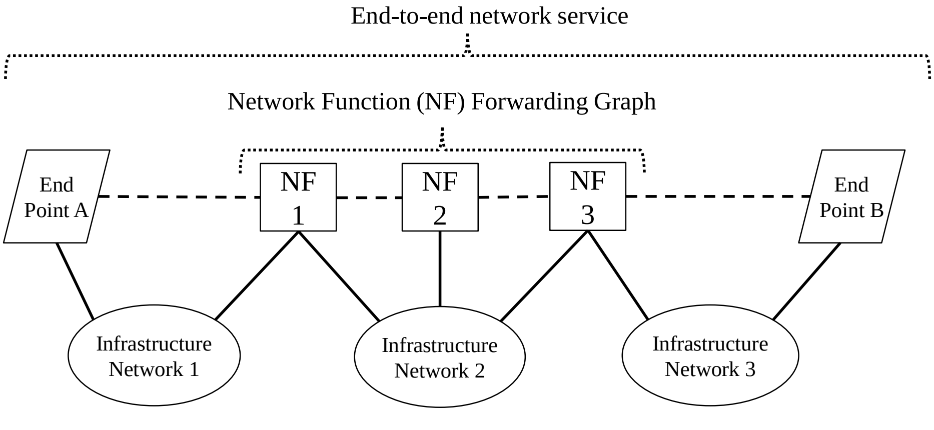

**Figure 2: Graph representation of an end-to-end network service**

Figure 3 shows an example of an end-to-end network service and the different layers that are involved in its virtualisation process. In this example, an end-to-end network service can be composed of only VNFs and two end points. The decoupling of hardware and software in network functions virtualisation is realized by a virtualisation layer. This layer abstracts hardware resources of the NFV Infrastructure.

The NFVI-PoPs includes resources for computation, storage and networking deployed by a network operator as indicated in Figure 3.

Virtualised Network Functions run on top of the virtualisation layer, which is part of the NFVI, as indicated by the arrow labelled "virtualisation". The VNF Forwarding Graph (VNF-FG) corresponding to the network function forwarding graph of Figure 2 is shown in Figure 3. The figure also depicts the case of a nested VNF-FG (i.e. VNF-FG-2) constructed from other Virtualised Network Functions (i.e. VNF-2A, VNF-2B and VNF-2C). The objective is that the interfaces between NFs and/or VNFs and the infrastructure in a multi-vendor environment be based upon accepted standards (e.g. standardized by an SDO, and/or an open de-facto standard).

NFV emphasizes the fact that the exact physical deployment of a VNF instance on the infrastructure is not visible from the E2E service perspective, with the exception of guaranteeing specific policy constraints (e.g. location awareness required to implement a virtualised CDN cache node (see the use case *Virtualisation of CDNs (vCDN)* in ETSI GS NFV 001 [3]), or to ensure redundant infrastructures are in different locations). This enables a VNF instance to be implemented on different physical resources, e.g. compute resources and hypervisors, and/or be geographically dispersed as long as its overall end-to-end service performance and other policy constraints are met. In any case, VNF instances and their supporting infrastructure need to be visible for configuration, diagnostic and troubleshooting purposes.

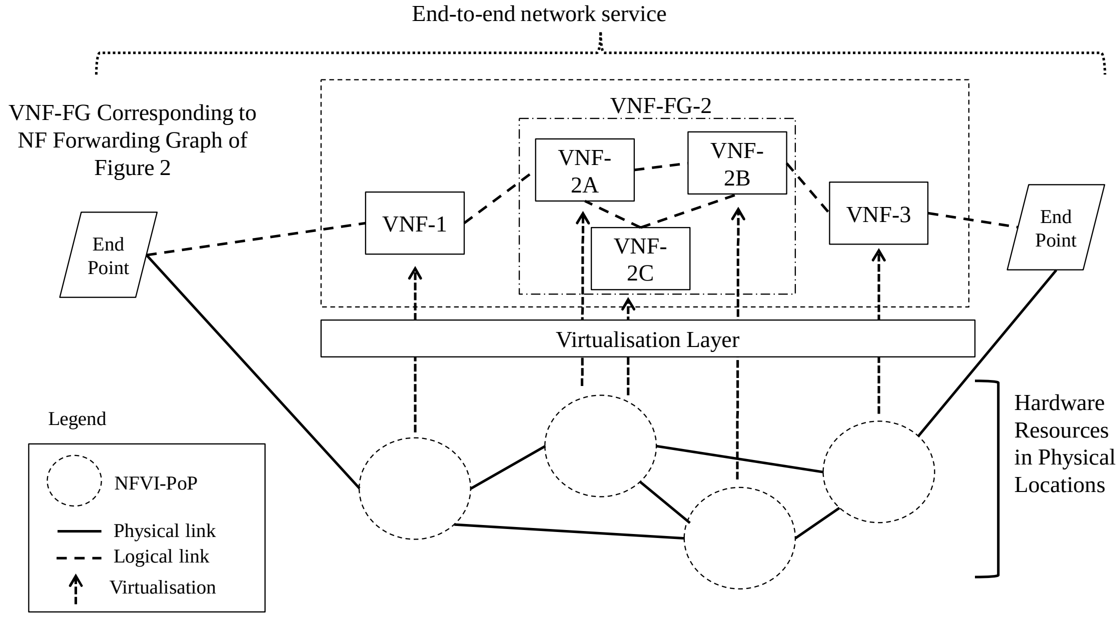

**Figure 3: Example of an end-to-end network service with VNFs and nested forwarding graphs**

## 6.3 Implications of NFV

The generic end-to-end network service model and its realization through virtualisation techniques as introduced in clause 5 require additional functionalities. GS NFV 001 [3] identifies service models applicable in the NFV context. GS NFV 004 [2] identifies requirements on the implementation of service and delivery models. The explicit mapping of these service models and requirements to the framework are not included in the present document. The NFV architectural framework addresses the following:

- The functionality that is required to be realized by the NFVI.
- The functionality that is required due to decoupling network functions into software and hardware.
- The functionality that is required for NFV-specific management and orchestration.

The required E2E network service and the delivered behaviour shall be equivalent among virtualised and non-virtualised scenarios. The focus of the present document is to describe the changes introduced by VNF, NFVI and their management and orchestration framework, and any considerations involved in providing services on top of a networking environment following the paradigm for decoupling software from hardware. As a result of the virtualisation process depicted in Figure 3, the next clause describes the high-level functional model in a greater detail.

# 7 NFV Reference Architectural Framework

## 7.1 Introduction

The NFV architectural framework focuses on the changes likely to occur in an operator's network due to the network function virtualisation process. That is, the architectural framework focuses on the new functional blocks and reference points brought by the virtualisation of an operator's networks.

The architectural framework is described at a functional level and it does not propose any specific implementation.

## 7.2 Architectural Functional Blocks

### 7.2.1 Overview of the Functional Blocks

The NFV architectural framework identifies functional blocks and the main reference points between such blocks. Some of these are already present in current deployments, whilst others might be necessary additions in order to support the virtualisation process and consequent operation. The functional blocks are:

- Virtualised Network Function (VNF).
- Element Management System (EMS).
- NFV Infrastructure, including:
- Hardware and virtualised resources, and
- Virtualisation Layer.
- Virtualised Infrastructure Manager(s).
- Orchestrator.
- VNF Manager(s).
- Service, VNF and Infrastructure Description.
- Operations and Business Support Systems (OSS/BSS).

Figure 4 shows the NFV architectural framework depicting the functional blocks and reference points in the NFV framework. The main (named) reference points and execution reference points are shown by solid lines and are in the scope of NFV. These are potential targets for standardization. The dotted reference points are available in present deployments but might need extensions for handling network function virtualisation. However, the dotted reference points are not the main focus of NFV at present. The architectural framework shown focuses on the functionalities necessary for the virtualisation and the consequent operation of an operator's network. It does not specify which network functions should be virtualised, as that is solely a decision of the owner of the network.

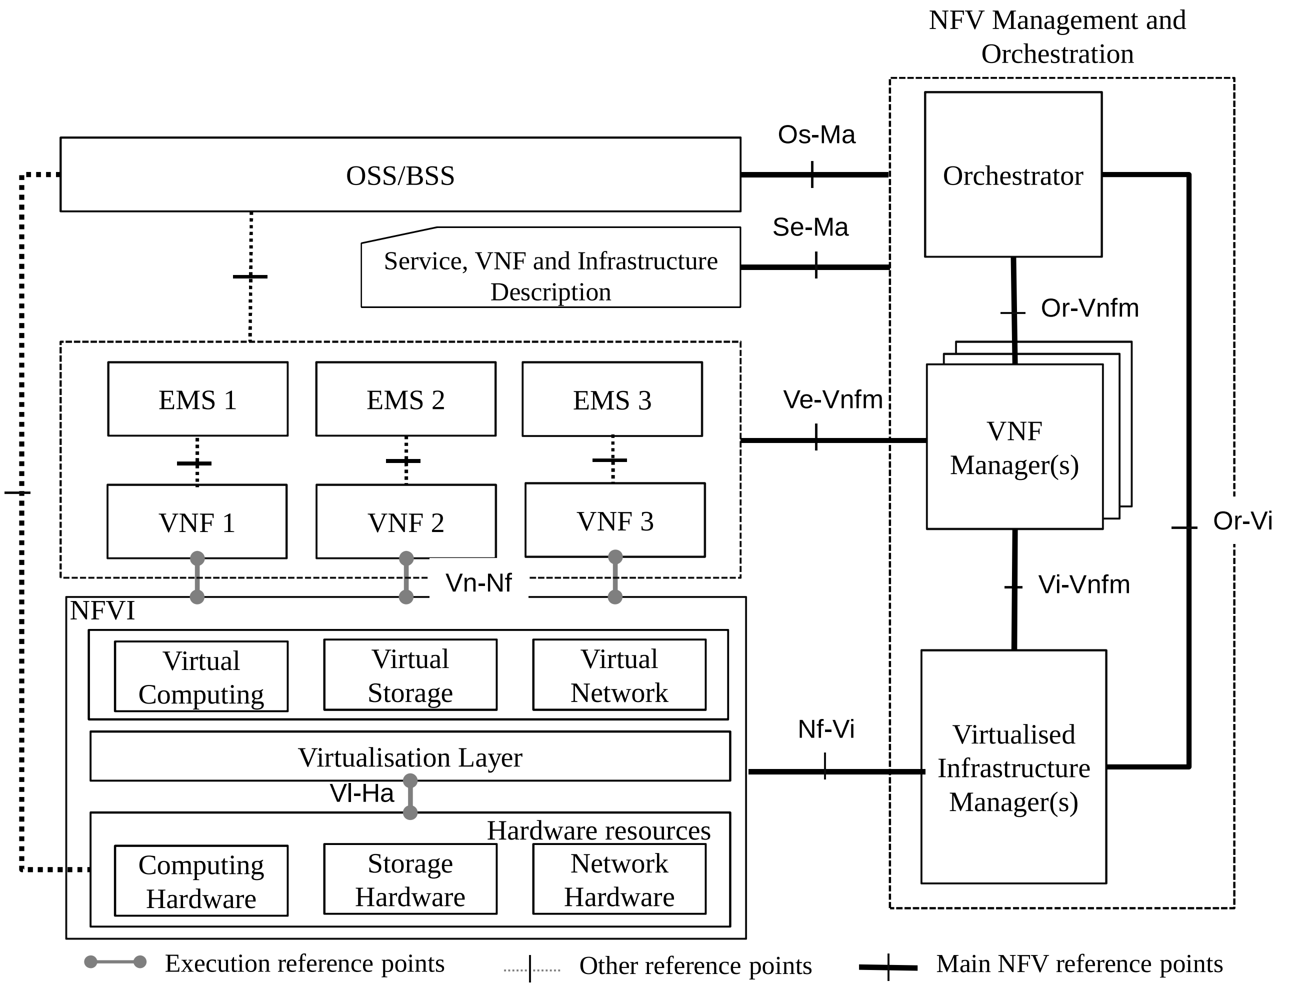

**Figure 4: NFV reference architectural framework**

The following clauses give an overview of the functional blocks in the architectural framework.

## 7.2.2 Virtualised Network Function (VNF)

A VNF is a virtualisation of a network function in a legacy non-virtualised network. Examples of NFs are 3GPP™ Evolved Packet Core (EPC) [i.2] network elements, e.g. Mobility Management Entity (MME), Serving Gateway (SGW), Packet Data Network Gateway (PGW); elements in a home network, e.g. Residential Gateway (RGW); and conventional network functions, e.g. Dynamic Host Configuration Protocol (DHCP) servers, firewalls, etc. GS NFV 001 [3] provides a list of use cases and examples of target network functions (NFs) for virtualisation.

Functional behaviour and state of a NF are largely independent of whether the NF is virtualised or not. The functional behaviour and the external operational interfaces of a PNF and a VNF are expected to be the same.

A VNF can be composed of multiple internal components. For example, one VNF can be deployed over multiple VMs, where each VM hosts a single component of the VNF. However, in other cases, the whole VNF can be deployed in a single VM as well. Detailed implementation methods are outside the scope of the present document.

## 7.2.3 Element Management System (EMS)

The Element Management System performs the typical management functionality for one or several VNFs.

## 7.2.4 NFV Infrastructure

### 7.2.4.1 NFV Infrastructure Definition

The NFV Infrastructure [1] is the totality of all hardware and software components which build up the environment in which VNFs are deployed, managed and executed. The NFV Infrastructure can span across several locations, i.e. places where NFVI-PoPs are operated. The network providing connectivity between these locations is regarded to be part of the NFV Infrastructure.

From the VNF's perspective, the virtualisation layer and the hardware resources look like a single entity providing them with desired virtualised resources.

### 7.2.4.2 Hardware Resources

In NFV, the physical hardware resources include computing, storage and network that provide processing, storage and connectivity to VNFs through the virtualisation layer (e.g. hypervisor). Computing hardware is assumed to be COTS as opposed to purpose-built hardware. Storage resources can be differentiated between shared network attached storage (NAS) and storage that resides on the server itself.

Computing and storage resources are commonly pooled. Network resources are comprised of switching functions, e.g. routers, and wired or wireless links. Also, network resources can span different domains. However, NFV differentiates only the following two types of networks:

- NFVI-PoP network: the network that interconnects the computing and storage resources contained in an NFVI-PoP. It also includes specific switching and routing devices to allow external connectivity.
- Transport network: the network that interconnects NFVI-PoPs, NFVI-PoPs to other networks owned by the same or different network operator, and NFVI-PoPs to other network appliances or terminals not contained within the NFVI-PoPs.

### 7.2.4.3 Virtualisation Layer and Virtualised Resources

The virtualisation layer abstracts the hardware resources and decouples the VNF software from the underlying hardware, thus ensuring a hardware independent lifecycle for the VNFs. In short, the virtualisation layer is responsible for:

- Abstracting and logically partitioning physical resources, commonly as a hardware abstraction layer.
- Enabling the software that implements the VNF to use the underlying virtualised infrastructure.
- Providing virtualised resources to the VNF, so that the latter can be executed.

The architectural view of NFV infrastructure and network function virtualisation is presented in Figure 4. The virtualisation layer in the middle ensures VNFs are decoupled from hardware resources and therefore, the software can be deployed on different physical hardware resources. Typically, this type of functionality is provided for computing and storage resources in the form of hypervisors and virtual machines (VMs). A VNF is envisioned to be deployed in one or several VMs.

The NFV architectural framework does not restrict itself to using any specific virtualisation layer solution. Rather, NFV expects to use virtualisation layers with standard features and open execution reference points towards VNFs and hardware (computation, network and storage). In some cases, VMs may have direct access to hardware resources (e.g. network interface cards) for better performance. Nonetheless, in NFV, VMs shall always provide standard ways of abstracting hardware resources without restricting its instantiation or dependence on specific hardware components.

The use of hypervisors is one of the present typical solutions for the deployment of VNFs. Other solutions to realize VNFs may include software running on top of a non-virtualised server by means of an operating system (OS), e.g. when hypervisor support is not available, or VNFs implemented as an application that can run on virtualised infrastructure or on bare metal. To ensure operational transparency, the operation of the VNF should be independent of its deployment scenario.

When virtualisation is used in the network resource domain, network hardware is abstracted by the virtualisation layer to realize virtualised network paths that provide connectivity between VMs of a VNF and/or between different VNF instances. Several techniques allow this, including network abstraction layers that isolate resources via virtual networks and network overlays, including Virtual Local Area Network (VLAN), Virtual Private LAN Service (VPLS), Virtual Extensible Local Area Network (VxLAN), Network Virtualisation using Generic Routing Encapsulation (NVGRE), etc. Other possible forms of virtualisation of the transport network include centralizing the control plane of the transport network and separating it from the forwarding plane, and isolating the transport medium, e.g. in optical wavelengths, etc.

### 7.2.5 Virtualised Infrastructure Manager(s)

From NFV's point of view, virtualised infrastructure management comprises the functionalities that are used to control and manage the interaction of a VNF with computing, storage and network resources under its authority, as well as their virtualisation. According to the list of hardware resources specified in the architecture, the Virtualised Infrastructure Manager performs:

- Resource management, in charge of the:
    - Inventory of software (e.g. hypervisors), computing, storage and network resources dedicated to NFV infrastructure.
    - Allocation of virtualisation enablers, e.g. VMs onto hypervisors, compute resources, storage, and relevant network connectivity.
    - Management of infrastructure resource and allocation, e.g. increase resources to VMs, improve energy efficiency, and resource reclamation.
- Operations, for:
    - Visibility into and management of the NFV infrastructure.
    - Root cause analysis of performance issues from the NFV infrastructure perspective.
    - Collection of infrastructure fault information.
    - Collection of information for capacity planning, monitoring, and optimization.

Multiple Virtualised Infrastructure Manager instances may be deployed.

### 7.2.6 Orchestrator

The Orchestrator is in charge of the orchestration and management of NFV infrastructure and software resources, and realizing network services on NFVI**.**

### 7.2.7 VNF Manager(s)

A VNF Manager is responsible for VNF lifecycle management (e.g. instantiation, update, query, scaling, termination).

Multiple VNF Managers may be deployed; a VNF Manager may be deployed for each VNF, or a VNF Manager may serve multiple VNFs.

### 7.2.8 Service, VNF and Infrastructure Description

This data-set provides information regarding the VNF deployment template, VNF Forwarding Graph, service-related information, and NFV infrastructure information models.

### 7.2.9 Operations Support Systems and Business Support Systems (OSS/BSS)

The OSS/BSS in Figure 4 refers to the OSS/BSS of an Operator.

## 7.3 Reference Points

### 7.3.1 Virtualisation Layer - Hardware Resources - (Vl-Ha)

This reference point interfaces the virtualisation layer to hardware resources to create an execution environment for VNFs, and collect relevant hardware resource state information for managing the VNFs without being dependent on any hardware platform.

### 7.3.2 VNF - NFV Infrastructure (Vn-Nf)

This reference point represents the execution environment provided by the NFVI to the VNF. It does not assume any specific control protocol. It is in the scope of NFV in order to guarantee hardware independent lifecycle, performance and portability requirements of the VNF.

### 7.3.3 Orchestrator - VNF Manager (Or-Vnfm)

This reference point is used for:

- Resource related requests, e.g. authorization, validation, reservation, allocation, by the VNF Manager(s).
- Sending configuration information to the VNF manager, so that the VNF can be configured appropriately to function within the VNF Forwarding Graph in the network service.
- Collecting state information of the VNF necessary for network service lifecycle management.

### 7.3.4 Virtualised Infrastructure Manager - VNF Manager (Vi-Vnfm)

This reference point is used for:

- Resource allocation requests by the VNF Manager.
- Virtualised hardware resource configuration and state information (e.g. events) exchange.

### 7.3.5 Orchestrator - Virtualised Infrastructure Manager (Or-Vi)

This reference point is used for:

- Resource reservation and/or allocation requests by the Orchestrator.
- Virtualised hardware resource configuration and state information (e.g. events) exchange.

### 7.3.6 NFVI - Virtualised Infrastructure Manager (Nf-Vi)

This reference point is used for:

- Specific assignment of virtualised resources in response to resource allocation requests.
- Forwarding of virtualised resources state information.
- Hardware resource configuration and state information (e.g. events) exchange.

### 7.3.7 OSS/BSS - NFV Management and Orchestration (Os-Ma)

This reference point is used for:

- Requests for network service lifecycle management.
- Requests for VNF lifecycle management.
- Forwarding of NFV related state information.

- Policy management exchanges.
- Data analytics exchanges.
- Forwarding of NFV related accounting and usage records.
- NFVI capacity and inventory information exchanges.

### 7.3.8 VNF/EMS - VNF Manager (Ve-Vnfm)

This reference point is used for:

- Requests for VNF lifecycle management.
- Exchanging configuration information.
- Exchanging state information necessary for network service lifecycle management.

### 7.3.9 Service, VNF and Infrastructure Description - NFV Management and Orchestration (Se-Ma)

This reference point is used for retrieving information regarding the VNF deployment template, VNF Forwarding Graph, service-related information, and NFV infrastructure information models. The information provided is used by NFV management and orchestration.

# 8 Study Items in NFV Reference Architectural Framework

## 8.1 Introduction

Network Functions Virtualisation represents a paradigm shift in networking. In order to realize a carrier-grade deployment of the NFV architectural framework, the following areas are recommended for further study. Specific mappings to the Requirements [2] and Use Cases [3] are not included in the present document and are left for future study.

## 8.2 Virtualisation Layering and NFVI Support

A key component in the NFV architectural framework is the virtualisation layer. As introduced in clause 7 of the present document, this layer abstracts and logically partitions physical hardware resources and anchors between the VNF and the underlying virtualised infrastructure. The primary tools to realize the virtualisation layer would be the hypervisors. The NFV architectural framework should accommodate a diverse range of hypervisors.

On top of such a virtualisation layer, the primary means of VNF deployment would be instantiating it in one or more virtual machines (VMs). Therefore, the virtualisation layer shall provide open and standard interfaces towards the hardware resources as well as the VNF deployment container, e.g. VMs, in order to ensure independence among the hardware resources, the virtualisation layer and the VNF instances. VNF portability shall be supported over such a heterogeneous virtualisation layer. Apart from presently available best practices, other performance and cost efficient methods for virtualisation layer design and VNF deployment techniques need to be investigated in order to support a diverse range of NFs found in an operator's network.

## 8.3 VNF Software Architecture

Network functions are well-defined; hence both their functional behaviour as well as their external interfaces are documented in technical specifications. In NFV, a VNF is a software package that implements such network functions. Aside from the high-level NFV architecture perspective, the VNF software architecture itself requires further study. Virtualisation presents us with an opportunity of modular and slimmer software design of the conventional monolithic NFs. A VNF can be decomposed into smaller functional modules for scalability, reusability, and/or faster response. Alternatively, multiple VNFs can be composed together to reduce management and VNF Forwarding Graph complexity. These two approaches of composition and decomposition are for further study; however, their initial descriptions are given below:

- VNF decomposition, as the process whereby a higher-level VNF is split into a set of lower-level VNFs. The NFV ISG shall provide guidelines in determining how VNF decomposition should be standardized.

- VNF composition, as the process whereby a group of lower-level VNFs is used to define a higher-level VNF.

We can take the Evolved Packet Core (EPC) [i.2] Serving Gateway (SGW) and Packet Data Network Gateway (PGW) NFs to illustrate the above. Both GW types can be provided as VNFs in their own right with the applicable properties of VNFs as described in the present document. It is also conceivable that vendors could implement the SGW and the PGW in combination, thereby creating a new combined VNF "SGW-PGW". A vendor that initially implements this combined VNF could also create independent VNFs for the SGW and the PGW. Since NFV makes no assumption on the internals relating to such "composed" or "decomposed" VNFs, a third party cannot decompose the SGW or the PGW VNFs from a combined SGW-PGW VNF.

Different management needs may arise when VNFs are composed or decomposed out of other VNFs. This may differ from when VNF Forwarding Graphs are constructed and each VNF is individually manageable, whereas in a composed VNF individual, VNF management interfaces may not be visible or in a decomposed VNF more management interfaces may be created. These issues require further study in the NFV ISG.

Furthermore, each VNF (irrespective of its hierarchy) consists of a number of software components that are manageable by corresponding management and orchestration systems. Such VNF components can only exist within the context of their "parent" VNF.

As explained above, the flexibility brought by virtualisation should be exploited to realize modular, portable, hardware independent, reusable and scalable software design of the VNFs.

Further study is required to analyse the relationship between VNF/EMS and VNF Managers (see Figure 4).

## 8.4 NFV Management and Orchestration

The decoupling of a VNF from the underlying hardware resources presents new management challenges. Such challenges include end-to-end service to end-to-end NFV network mapping, instantiating VNFs at appropriate locations to realize the intended service, allocating and scaling hardware resources to the VNFs, keeping track of VNF instances location, etc. Such decoupling also presents challenges in determining faults and correlating them for a successful recovery over the network. While designing the NFV Management and Orchestration, such challenges need to be addressed. In order to perform its task, the NFV Management and Orchestration should work with existing management systems such as OSS/BSS, hardware resource management system, CMS used as a Virtualised Infrastructure Manager, etc. and augment their ability in managing virtualisation-specific issues.

## 8.5 Performance

Performance and scalability are important since the implementation of a VNF may have a per-instance capacity that is less than that of a corresponding physical network function on dedicated hardware. Therefore, methods are needed to split the workload across many distributed/clustered VNF instances. Performance-efficient methods of deploying VNF instances on the NFV infrastructure can also be considered. It is necessary to study how to minimize performance degradation while keeping VNF instance portability on heterogeneous NFV infrastructure comprised of diverse virtualisation layer and hardware resources.

## 8.6 Reliability

NFV should provide reliability as high as that provided with equivalent non-virtualised legacy network functions, but with improved cost efficiency. Reliability could depend on specific services, (e.g. voice, health-related applications, and financial transactions) which may need higher reliability than best-effort services.

To address this issue, functions can be organized into classes or categories that have similar reliability requirements. Single-point-of failure, fault detection and recovery methods need to be investigated. The full potential of the flexibility provided by virtualisation needs to be utilized to achieve cost efficient reliability necessary for carrier-grade service availability and completion.

## 8.7 Security

Physical network functions assume a tight coupling of the NF software and hardware which, in most cases, is provided together by a single vendor. In the NFV scenario, multiple vendors are expected to be involved in the delivery and setup of different NFV elements (e.g. hardware resources, virtualisation layer, VNF, virtualised infrastructure manager, etc.). As a result, due to the virtualisation process, new security issues need to be addressed. Examples are:

The use of hypervisors may introduce additional security vulnerabilities, even though hypervisor-based virtualisation is the state of the art. Third-party certification of hypervisors should help shed light on their security properties. In general, to reduce the vulnerabilities of the hypervisors in use, it is essential to follow the best practices on hardening and patch management. To ensure that the right hypervisor is being executed calls for authenticating the hypervisor at the boot time through secure boot mechanisms.

The usage of shared storage and shared networking may also add additional dimensions of vulnerability.

The interconnectivity among NFV end-to-end architectural components (e.g. hardware resources, VNFs, and management systems), expose new interfaces that, unless protected, can create new security threats.

The execution of diverse VNFs over the NFV infrastructure can also create additional security issues, in particular if VNFs are not properly isolated from others.

# 9 Conclusions and Recommendations

The present document provides a view on virtualisation-specific functional blocks and reference points among the functional blocks necessary within an operator's network. The successful realization of the presented architectural framework depends on the detailed definition of the functional blocks and the interfaces, as well as the necessary extension of these keeping the architectural framework at the core.

The present document identifies what is in the scope of ETSI NFV ISG. Furthermore, the present document also highlights open issues that require further study that are key to achieving the NFV objective of realizing carrier-grade virtual networks.

Many of the reference points, technologies and components identified in the present document are available today. It is therefore important for NFV to identify gaps in existing standards and specifications to support the industry in filling those gaps. In this manner, the NFV architectural framework recommends NFV solutions to be assembled in a consistent and reusable manner, enabling economies of scale and faster innovation of network services.

# History

| Document history | | |
|---|---|---|
| V1.1.1 | October 2013 | Publication |
| | | |
| | | |
| | | |
| | | |