

Network Functions Virtualisation (NFV)

Network Operator Perspectives on Industry Progress

OBJECTIVES

The objectives for this white paper are to draw attention to the second release of ETSI NFV ISG documents scheduled to be published in January 2015, and to provide a commentary on industry progress on NFV since we published our last update one year ago.

This is a non-proprietary white paper authored by network operators who are participating in the NFV ISG. It has been produced independently of the NFV ISG; it is not an NFV ISG document and claims no endorsement by the NFV ISG.

CONTRIBUTING ORGANISATIONS & AUTHORS

AT&T:	Margaret Chiosi, Steve Wright.
Bell Canada:	Javan Erfanian, Brian Smith.
BT:	Bob Briscoe, Andy Reid, Peter Willis.
CableLabs:	Don Clarke, Chris Donley.
CenturyLink:	Michael Bugenhagen, James Feger.
China Mobile:	Chunfeng Cui, Hui Deng.
China Telecom:	Yunpeng Xie, Zhenqiang Sun.
China Unicom:	Gongying Gao, Yiqiang Hua.
Colt:	Javier Benitez, Nicolas Fischbach.
Deutsche Telekom:	Klaus Martiny, Uwe Michel.
DOCOMO:	Tetsuya Nakamura, Joan Triay Marques.
KDDI:	Kenichi Ogaki, Tetsuro Matsuzaki.
KPN:	Shuang Zhang, Alexander de Boer.
KT:	Kisang Ok, Eun Kyoung PAIK.
NTT:	Katsuhiro Shimano, Takashi Shimizu.
Ooredoo:	Marco Stura.
Orange:	Bruno Chatras, Christos Kolias.
Portugal Telecom:	Jorge Carapinha, Antonio Gamelas.
SK Telecom:	DK Lee, Jong Han Park.
Softbank:	Ryuji Wakikawa, Kazuto Nishi, Satoru Matsushima.
Sprint:	Laurent Laporte, Fred Feisullin.
Swisscom:	Markus Brunner.
Telecom Italia:	Elena Demaria, Andrea Pinnola.
Telenor:	Patrick Waldemar, Geir Millstein.
Telefonica:	Diego López, Francisco Javier Ramón Salguero.
Telstra:	Daniel Kirkham.
Turk Telekom Argela:	Mustafa Ergen, Melih Ahmet Karaman, Aydin Ulas, Erhan Lokman.
Verizon:	Naseem Khan, Raquel Morera.
Vodafone:	Susana Sabater, Adrian Neal.
Windstream:	Arthur Nichols.

PUBLICATION DATE

October 14-17, 2014 at the “SDN & OpenFlow World Congress”, Dusseldorf-Germany.

This white paper is available at the following link: http://portal.etsi.org/NFV/NFV_White_Paper3.pdf

Executive Summary

This white paper provides an update on industry progress on NFV since we published our last review in October 2013.

Since its first meeting in January 2013, the ETSI NFV Industry Specification Group (NFV ISG) has grown to 235 companies, including 34 service provider organisations. The first outputs published in October 2013 which included an NFV Architectural Framework are being widely referenced across the industry to inform product development, standardisation, and new open source initiatives such as Open Platform for NFV (OPNFV). At the same time, the NFV ISG issued a call for Proof of Concept demonstrations (PoCs) to validate NFV assumptions and to encourage growth of an open ecosystem. To date, 25 PoCs are in progress or have been completed, spanning the breadth of NFV ISG scope with the results being openly available.

In January 2015, the NFV ISG will publish a major release of documents that will be highly influential in setting the direction for NFV implementation and standardisation going forward. Drafts of these documents have been openly available on the NFV ISG portal since July. They have been brought to the attention of standards development organisations to help them frame their work to accommodate NFV concepts.

In our original white paper we envisaged a two-year timeframe for the NFV ISG to complete its work. With the upcoming release of documents in January 2015, we are satisfied that the NFV ISG has achieved its goals. It has exceeded our expectations in both fostering an open ecosystem for the emerging technology of NFV, and in the degree of influence it has had on the wider industry, including standards development organisations and open source communities. It is also evident that vendor roadmaps have been dramatically influenced by this effort, which bodes well for the emergence of competitive solutions for NFV.

The key challenge for growth of an open ecosystem is to achieve interoperability for the key interfaces identified in the NFV Architectural Framework. To ensure that momentum is not lost in achieving interoperability, we have encouraged the NFV ISG to work beyond its original two-year limit with a specific focus on addressing the barriers to interoperability. A second two-year phase of work will begin in February 2015.

We also encouraged industry and academia to participate in the NFV ecosystem by creating research programmes around NFV and to create new teaching courses to train a new generation of students to be multi-skilled in networks and software.

NFV will have a significant impact on the design of future telecommunications support systems. Using virtualisation in an up-to-date cloud environment will offer new self-managed redundancy and failover scenarios. The evolved management/support systems to handle this new environment must be highly automated, which requires new thinking on OSS/BSS that will open up opportunities to gain significant operational benefits.

Contents

1	Introduction	4
2	Overview of the NFV ISG and its Evolution	4
3	Overview of Second Release of NFV ISG Documents	5
3.1	NFV Infrastructure	5
3.2	NFV Management and Orchestration (MANO).....	6
3.3	NFV Software Architecture	8
3.4	NFV Reliability and Availability	9
3.5	NFV Performance and Portability	11
3.6	NFV Security.....	12
4	Overview of NFV Proof of Concepts	14
5	NFV Research and Education	14
6	Perspectives on Open Source	15
7	Evolving Relationship with SDN	15
8	NFV Impact on OSS/Network Management	16
9	References	18
10	Contact Information.....	19
11	Glossary.....	19

1 Introduction

In our first white paper published in October 2012^[1] we introduced the concept of Network Functions Virtualisation (NFV) and outlined the benefits and challenges for NFV technologies to be implemented and deployed. NFV will transform the way that networks are built and operated by leveraging standard IT virtualisation technology and consolidating network equipment types onto “industry standard” servers. We also issued a call for action for the industry to cooperate to address these challenges and we declared our intent to encourage growth of an open ecosystem for NFV. To provide a formal umbrella for industry cooperation, we founded the Network Functions Virtualisation Industry Specification Group (NFV ISG) under the auspices of the European Telecommunications Standards Institute (ETSI). We ensured barriers to participation were low by keeping fees minimal and allowing open membership.

Since launch in October 2012, the NFV ISG has grown to 235 companies including 34 service provider organisations and has held seven plenary meetings spanning Asia, Europe and North America. The first NFV ISG outputs were published in October 2013 including an NFV Architectural Framework which identified NFV system components and the interfaces between them.^[2] In the intervening period, the NFV ISG has become the focal point for industry progress on NFV. In October 2013 we published a joint-operator update white paper to describe and position the first NFV ISG outputs in the wider industry context.^[3]

In its first two years, the NFV ISG has moved the industry substantially forward on NFV and has provided a focal point for an unprecedented level of IT and telecommunications industry collaboration. The second release of NFV ISG documents is on schedule for publication in January 2015 with drafts already available on the NFV ISG portal.^[4]

This white paper provides an update on the NFV ISG work along with our perspectives on the way forward.

2 Overview of the NFV ISG and its Evolution

Having successfully achieved its objectives for the first two-year phase, the NFV ISG is now developing the agenda for a second phase of work which is expected to take up to two years. The structure, scope and objectives for this second phase are being discussed amongst the general NFV ISG membership and are expected to be agreed at the eighth plenary meeting in November 2014. However, the network operator community expects some major themes to be addressed, summarised as follows.

One of our most important goals is to quickly achieve interoperability for NFV solutions assembled using components arising within an open ecosystem. In support of this goal we wish to see progress towards completing the key interface specifications which are required to ensure interoperability between different implementations. Specifically, interoperability for the Orchestrator ↔ Virtual Infrastructure Manager (VIM), Orchestrator ↔ Virtual Network Function Manager (VNFM), VNFM ↔ Virtual Network Function (VNF), and VNF ↔ Network Functions Virtualisation Infrastructure (NFVI) interfaces will encourage growth of a multi-vendor implementation environment.

From an architectural framework perspective, layered architecture and inter-domain aspects are topics that have not been fully addressed to date. A layered architecture is needed for deployments

where there is a hierarchy of orchestrators which in many cases will be operated by independent organisations and in multi-tenant mode. Analysis of how these orchestrators interact is needed and the interactions clearly stated in order to facilitate smooth deployment.

Another aspect needing to be addressed is how NFV infrastructures will interact with legacy OSS/BSS and how legacy OSS/BSS systems will need to evolve in order to support NFV (e.g. requirements for the OSS/BSS ↔ Orchestrator interface and the OSS/BSS ↔ Element Manager (EM) interface). Furthermore, to facilitate real deployments, virtualised network functions must be able to co-exist with physical network functions and a migration plan developed. Some of these aspects are not within the scope of the NFV ISG and are discussed in more detail later in this paper.

Finally, as NFV is clearly gaining momentum across the industry, there should be more focus on possible service models (e.g. NFVaaS, VNFaaS, and NFVPaaS) and elaborating how NFV can facilitate operations simplification. Examples of operations simplification include smooth introduction of new network capabilities by activating new VNFs without shutting down old components, and service enhancements for specific customers by activating specific VNFs dedicated to specific customers via means of tailored forwarding graphs. We recommend these topics for further study within the NFV community.

3 Overview of Second Release of NFV ISG Documents

The contents of the second release of NFV ISG documents to be openly published January 2015 are described in this section. These documents are available in draft form on the NFV ISG portal.^[4]

3.1 NFV Infrastructure

The NFV Infrastructure (NFVI) is required to support the range of use cases and fields of application already identified by the NFV ISG while providing a stable platform for the evolution of the VNF ecosystem. The NFVI is the totality of the hardware and software components which build up the environment in which VNFs are deployed. The NFVI provides a multi-tenant infrastructure leveraging standard IT virtualisation technology that may support multiple use cases and fields of application simultaneously as shown in Figure 1.

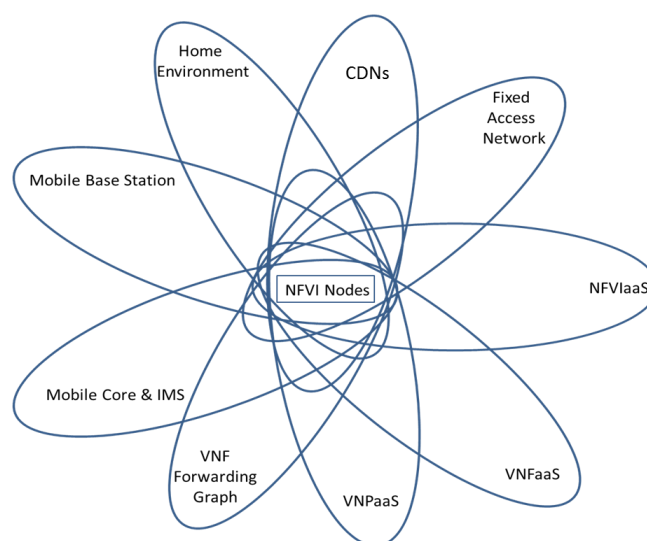


Figure 1: NFVI Supports Multiplicity of NFV Use Cases and Fields of Application

The NFVI is implemented as a distributed set of NFVI Nodes deployed in various NFVI Points of Presence as required to support the locality and latency objectives of the different use cases and fields of application. VNFs may be dynamically deployed on the NFVI on demand within the capacity limits of the NFVI Nodes. The NFVI is described in the following NFV ISG documents:

- ***NFV Infrastructure Overview***
- ***NFV Compute Domain***
- ***NFV Hypervisor Domain***
- ***NFV Infrastructure Network Domain***

The ***NFV Infrastructure Overview*** document provides an overview of the architectural principles of the NFVI and the associated interfaces; the relationships between the NFVI and cloud computing; economics and practical interoperability; key quality indicators; and partitions the NFVI into the functional domains of compute, hypervisor and infrastructure network. It also identifies a number of key challenges for the NFVI to support the required portability and performance of VNFs.

The ***NFV Compute Domain*** document identifies the functional elements of the compute domain as processors & accelerators, network interfaces and storage; the interfaces between the compute domain and other elements of the NFVI as well as those interfaces external to the NFVI that are supported directly by the compute domain. It identifies aspects of modularity and scalability that would impact implementation and deployment. The document also identifies a number of other features of the compute domain that impact various other aspects of the NFV system including management, performance, reliability and security.

The ***NFV Hypervisor Domain*** document presents the architecture of the hypervisor domain supporting the multitenant deployment and execution environment for VNFs. It focuses primarily on the use of hypervisors as an implementation technology though similar requirements would exist for implementations using Linux containers or other virtualisation technologies. The document identifies the external interfaces of the hypervisor domain as well as the functional blocks within it, including virtual switching (vSwitch) capabilities.

The ***NFV Infrastructure Network Domain*** document identifies the external interfaces of the domain as well as the functional blocks within the domain. The functional blocks of the Infrastructure network domain include virtual networks; other virtualisation layer options; network resources; control and administrative agents; as well as OAM North-South and East-West interfaces. It also considers modularity and scalability; interfaces to other NFVI domains; and other features of the networking domain that impact aspects of NFV related to performance, reliability and security.

In addition to the above documents, the NFV ISG has created an ***NFV Service Quality Metrics*** document that provides a general taxonomy for NFV service quality metrics. It identifies metrics relevant to the service quality of virtual machines, virtual network interfaces, technology components, and orchestration. It also describes use cases for service quality metrics and provides additional recommendations on aspects of measurements and service level agreements.

3.2 NFV Management and Orchestration (MANO)

NFV decouples software implementations of Network Functions from the compute, storage, and networking resources through a virtualisation layer. This decoupling requires a new set of

management and orchestration functions and creates new dependencies between them, thus requiring interoperable standardised interfaces, common information models, and the mapping of such information models to data models.

One of the key benefits of NFV identified in our first white paper is the elasticity provided by the infrastructure for capacity expansion and the roll out of new network functions. However, to take full advantage of the elasticity benefits of NFV, higher levels of automation are needed for the provisioning, configuration, and performance testing of virtualised network functions. During the past two years, the NFV ISG has studied such functions and dependencies, resulting in the creation of a Management and Orchestration Framework for NFV. The results of this study are captured in the ***NFV Management and Orchestration*** document.

The ***NFV Management and Orchestration*** document expands on the NFV Architectural Framework, further detailing the functionality of three key functional blocks: NFV Orchestrator (NFVO), VNF Manager (VNFM), and Virtualised Infrastructure Manager (VIM). The NFVO performs orchestration functions of NFVI resources across multiple VIMs and lifecycle management of network services. The VNFM performs orchestration and management functions of VNFs. The VIM performs orchestration and management functions of NFVI resources within a domain. The NFVO interacts with the OSS/BSS for provisioning, configuration, capacity management, and policy-based management. The VNFM interacts with the Element Manager (EM) and the VNF for provisioning, configuration, and fault and alarm management. The VIM interacts with the NFVI for the management and orchestration of virtualised resources. The document also describes the interworking of the NFV Management and Orchestration with network controllers when used for virtual network provisioning on the infrastructure network.

As described in the document, NFV management and orchestration functions can be grouped in three broad categories: virtualised resources, virtualised network functions, and network services. Note that in the context of the NFV ISG, a network service is constructed by chaining VNFs and/or Physical Network Functions (PNFs).

Management and Orchestration of virtualised resources encompasses all functions required to provide VNFs and Network Services with the resources they need in order to execute properly. The virtualised resources in-scope are those that can be associated with virtualisation containers and that have been catalogued and offered for consumption, consisting of those types identified within the NFVI, namely, compute, storage and network resources.

In addition to traditional Fault, Configuration, Accounting, Performance, and Security (FCAPS) Management, the NFV Management and Orchestration framework introduces a new set of management functions associated with the lifecycle management of a VNF. The NFV ISG has focused on detailing these new sets of management functions, which include, but are not limited to: on-board a VNF, instantiate a VNF, scale a VNF, update a VNF, and terminate a VNF. A difference also worth highlighting relates to fault and performance management - in a virtualised environment this is the responsibility of different functional blocks at different layers. As a result, the correlation of faults, alarms and other monitored data such as performance metrics and resource usage, and the consequent fault resolution needed to operate the service in a reliable manner, will typically be distributed.

Network Service Orchestration functions are responsible for coordinating the lifecycle of VNFs that jointly realise a Network Service. Network Service orchestration functions include on-boarding a Network Service, management of resources used by the Network Service, managing dependencies between different VNFs composing the Network Service, and managing the forwarding graphs between the VNFs. During the Network Service lifecycle, the Network Service orchestration functions may monitor Key Performance Indicators (KPIs) of a Network Service, and may report this information to support an explicit request for such operations from other functions.

Expanding on the functional blocks and reference points identified by the NFV Architectural Framework, the NFV Management and Orchestration framework defines requirements and operations on the interfaces exposed and consumed by functional blocks associated with the different management functions (e.g. VNF lifecycle management, virtualised resource management). The objective of such an approach is to expose the appropriate level of abstraction via the interfaces without limiting implementation choices of the functional blocks. The document provides an extensive description of interfaces, which is the basis for future work on standardisation and identification of gaps in existing systems and platforms. The **NFV Management and Orchestration** document annexes provide a comprehensive description of the different information flows and how such interfaces may be used.

Network Service operations and management requires common information models. To this end, the NFV ISG has defined the information elements needed in a set of descriptors used at the time of VNF on-boarding: VNF Descriptors (VNFD), Network Service Descriptors (NSD), VNF Forwarding Graph Descriptors (VNFFGD), and Virtual Link Descriptors (VLD). Such descriptors are deployment templates that describe the resource and operational requirements for each of these components. The orchestration functions take these descriptors, assign virtual resources, and perform lifecycle management operations. Example information on how existing data models can be used/expanded to support these descriptors is included in the annexes. The NFV ISG is planning to work on a standard representation in the next phase of work.

In summary, the **NFV Management and Orchestration** document describes the NFV Management and Orchestration framework with special emphasis on specifying the management and orchestration functionality, information elements and interfaces. This is all with the aim to encourage and foster a multi-vendor NFV ecosystem by enabling interoperability of the different components comprising the NFVI, VNF software, and corresponding NFV Management and Orchestration framework entities.

3.3 NFV Software Architecture

NFV is about virtualising network functions previously implemented as proprietary hardware appliances; hence, an important topic to address is how to realise virtualisation from a network function provider's perspective. Understanding the transition from a hardware-based to a software-based implementation has been an active area of study within the NFV ISG. The outcomes of this work are brought together in the **NFV Virtual Network Functions Architecture** document.

This document identifies the most common and relevant software architectural patterns that can be leveraged when decoupling the software from hardware. This has led to the specification of functional requirements with respect to management and orchestration functions, as well as

requirements towards the supporting NFV Infrastructure (NFVI). For example, a number of requirements refer to information elements that VNF descriptors and other related information models shall comprise. Other sets of requirements identify features that the NFVI should meet and enhance over and above traditional cloud technologies in order to support high-performance VNFs.

As a general approach to virtualising a network function and identifying common software design patterns, the concept of VNF Component (VNFC) has been developed. A VNFC is characterised as an internal component of a VNF that can be mapped to a single container interface to provide a sub-set of the VNF's functionality as shown in Figure 2. The relationship and decomposition of VNFs into finer granular VNFCs is detailed in a number of annexes to the **NFV Virtual Network Functions Architecture** document which cover example use cases including virtualisation of the Media Resource Function (MRF) of the IP Multimedia Subsystem (IMS), Traffic Detection Function (TDF), Enterprise Gateway, and Deep Packet Inspection (DPI) Engine.

The document lists a number of common software architectural patterns that appear as a result of such VNF componentization. Examples relate to VNFC state information, load-balancing models, scaling, etc. These patterns are used to define requirements towards management and orchestration, and NFV Infrastructure.

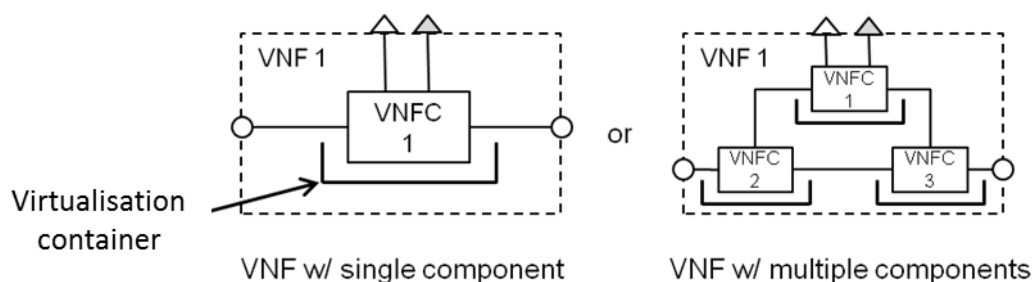


Figure 2: VNF Composition.

The VNF lifecycle is another aspect covered by the document. From a VNF's perspective, the lifecycle covers diverse procedures including scaling, configuration, operations, VNF instantiation, and VNF termination. The document also describes the pre- and post-conditions associated with each of these lifecycle procedures, and it maps such procedures to VNF instance state transitions.

In summary, the **NFV Virtual Network Functions Architecture** document defines the functions, requirements, and interfaces of VNFs with respect to the more general NFV Architectural Framework. It sets the stage for future work to bring software engineering best practices to the design of VNFs.

3.4 NFV Reliability and Availability

There are unique challenges (and opportunities) to ensuring service availability and maintaining resiliency in an NFV-based system. These challenges have been studied within the NFV ISG and the findings brought together in the **NFV Resiliency Requirements** document. This document describes the resiliency problem, use case analysis, resiliency principles, requirements, and deployment and engineering guidelines relating to NFV. The resiliency related requirements are specified for the various aspects of the NFV framework: NFVI; VNF; MANO; and, additionally cover: service availability; fault management; and, failure prevention, detection, and remediation. The use cases considered in this document include: resiliency of stateless and stateful services, network

transparency of the location of the VNFs, regression and pre-emption, spatial distribution of the VNFs, service chaining, and service continuity.

The Service Availability requirements for NFV should at a minimum be the same as those for legacy systems for the same service. In order to fulfil the requirements, the NFV components will have to provide the same or better performance in one or more of the following aspects: failure rate, detection time, restoration time, success rate of the detection and restoration, impact per failure, etc. In order to meet the Service Availability requirements, the VNF design needs to take into account factors which include commodity-grade hardware and presence of multiple software layers (i.e. hypervisor and guest Operating System).

Since NFV MANO components play an important role in maintaining service availability functions such as rapid service creation, dynamic adaptation to load, and overload prevention, they need to be highly reliable. Overall service resiliency depends on the underlying NFVI reliability as well as VNF internal resiliency.

Restoration of an end-to-end service in NFV in the events of failures and/or VNF or service relocation can be handled in one of three ways: at the service layer with the appropriate design mechanisms (e.g. IMS restoration procedures), at the VNF level (e.g. via legacy active/standby mechanism) or by the NFV MANO components. Network operators need to be able to adjust their solutions to Service Availability with various configurable parameters according to the situation and other criteria. For example, in a large-scale disaster situation, the network operator may decide to prioritise the voice call service over on-line gaming by shifting all the resources available to the voice call service; in which case, the Service Availability for the on-line gaming will be different than that in the normal situation.

The ***NFV Resiliency Requirements*** document also provides a list of faults and challenges and the relationships among them along with the mapping to the applicable NFV Architectural Framework component(s) or interface(s). These faults and challenges at the various layers of the system require detection of failures, local remediation where applicable, and notification about their occurrence to the next higher layer. In an NFV environment, the failure model and frequency of the failures are expected to be very different from the traditional network operator system because of the new challenges presented by NFV such as increased systems complexity with the addition of virtualisation, resource elasticity and VNF migration, as well as interoperability across third party software and hardware. Thus, it is essential to deploy proactive failure management approaches such as monitoring the resources usage in real-time (e.g. CPU, vCPU, virtual Memory, virtual IO, etc.), alarm correlation, and trend analysis.

A VNF implementation may choose to separate the state information from the corresponding VNF/VNFC and to store this information in a “Logical Unit” instead. A “Logical Unit” is an instantiation of Virtual Storage and has the capability of synchronizing the content automatically. This type of protection mechanism provides stateful resiliency, i.e. on-going end-to-end service sessions are maintained after the failure recovery with little or no interruption. However, in stateless end-to-end services (e.g. web/data services), there is no state information to maintain, thus it is acceptable to initialise the service at the time of the failure (e.g. by expecting end-user's retry access).

The **NFV Resiliency Requirements** document also describes deployment and engineering guidelines, which include the various redundancy schemes and other failure mitigation approaches, operator policies, and disaster recovery mechanisms for different failure scenarios. Appropriate redundancy mechanisms need to be set up depending on the criticality of the NFV MANO components. The implementation can vary: one approach could be that each component be provisioned on a separate cluster to enable extreme scalability and isolation; while another approach could be based on shared clusters (shared resources) but this can make scaling and diagnosis more complex.

In summary, the **NFV Resiliency Requirements** document describes the resiliency problem, use case analysis, resiliency principles, requirements, and deployment and engineering guidelines relating to NFV. It sets the stage for future work to bring software engineering best practices to the design of resilient NFV-based systems.

3.5 NFV Performance and Portability

The aim of NFV as articulated in the original white paper is to transform the way networks are built and operated by leveraging standard IT virtualisation technology to consolidate network equipment types onto “industry standard” servers. Hence, the NFV approach should be applicable to any data plane packet processing and control plane function in fixed and mobile network infrastructures. Given the implied restriction of using “industry standard” servers to provide a common layer to execute (potentially) different kinds of network functions, two related challenges were identified: a) would an industry standard server be capable of supporting realistic network workloads and b) if so, how could that execution capacity be abstracted to provide predictable behaviour for new network functions? In other words, would it be possible to achieve high performance with virtualised network appliances and, at the same time, make them portable between servers and hypervisors?

This effective decoupling of the capacity from the function would greatly simplify the interaction between the agents in the network ecosystem once NFV is in place. The hardware providers could be unaware of the VNFs that would end up running on their equipment in the future, while the VNF providers would still be able to provide reliable performance estimations for different hardware configurations (and hence meet Service Level Agreements), but with no awareness of the particular server where their VNF would be executed. And, of course, once this interaction model is properly defined, the operation would be greatly simplified for the service provider, paving the way for flexible resource allocation and simpler orchestration techniques.

Understanding to which extent this vision can be realised in practice (and describing how) has been an active topic of study within the NFV ISG. The results were published in July 2014 in the **NFV Performance & Portability Best Practises** document which is available on the NFV ISG portal.^[2] The analysis followed a purely practical methodology based on tests with relevant use cases and VNFs such as BNG/BRAS, DPI, C-RAN, CDN nodes, etc. Encouraging results were obtained in the most stringent scenarios, thus, not only was it proven that high performance was achievable in a fully virtualised environment (e.g. >80 Gbps per server as shown in Figure 3), but that it could be obtained in a predictable, consistent and vendor-agnostic manner, leveraging features commonly available in current state-of-the-art servers.

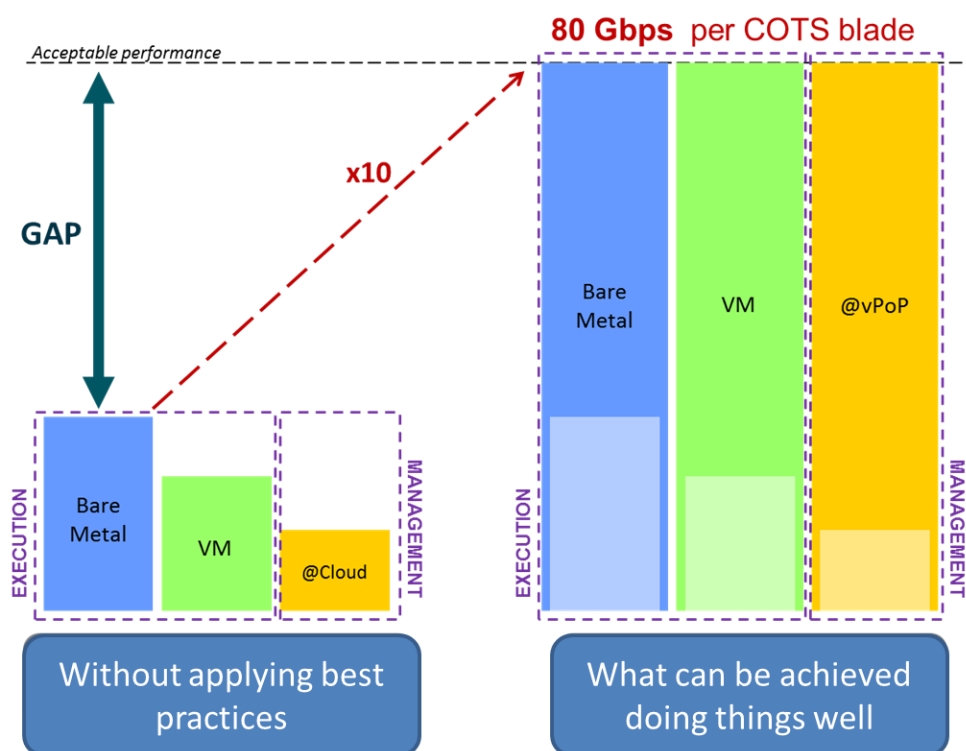


Figure 3: Performance Illustration

Key aspects such as providing the proper visibility of internal server layout (e.g. NUMA awareness), assuring that intermediate software layers in the host, including the hypervisor, do not become unexpected bottlenecks (or interact unpredictably), and, of course, making the cloud OS layer (i.e. Virtualised Infrastructure Manager - VIM) aware of all these parameters proved to be essential for optimal results in an NFV environment. The corresponding shortcomings of processors, hypervisors, and VIM environments were identified early in the work and have triggered recent developments in Linux kernel, KVM, libvirt, and OpenStack to accommodate NFV scenarios in IT-centric environments.

The **NFV Performance & Portability Best Practises** document explains how to put these principles into practice with a generic VNF deployment and should be studied in conjunction with the **NFV Management and Orchestration** document described earlier in this paper. It provides a comprehensive description of the information that is required to a) specify hardware requirements for a network function for a given performance target, and b) describe the hardware elements available in a server. These templates should allow, in a virtualised environment, each network function to be sized for a given performance target, with an appropriate definition of the underlying hardware requirements, while the corresponding hardware resources are allocated and isolated/shared in the server accordingly.

3.6 NFV Security

A large fraction of the global economy depends on the security of cloud computing and of networking. NFV uses virtualisation technology developed for cloud computing to provide networking services, so it is critical to assure the security of this new combination. The NFV ISG convened a Security Expert Group to focus on this concern. Security expertise is a scarce resource, so the original objective of the group was to initiate as little of its own work as possible in order to

prioritise advising the rest of the NFV ISG on security matters. In particular, providing security-focused review of the NFV ISG outputs and providing the NFV ISG with a feel for the size of the security problems (and opportunities) that NFV might introduce.

The group has completed a thorough assessment of potential new security concerns. The **NFV Security Problem Statement** document ^[2] focuses solely on the delta that NFV introduces on the assumption that others are addressing the security concerns of the separate component parts: virtualisation technology and networks.

The group's assessment is that NFV does introduce some potential new security problems listed in Table 1 but none are intractable, i.e. technology solutions to most of the identified concerns are already available. However, in general, processes are not yet in place to take advantage of these solutions and, once in place, they will add procedural complexity.

Topology Validation & Enforcement
Availability of Management Support Infrastructure
Secured Boot
Secure Crash
Performance Isolation
User/Tenant Authentication, Authorization and Accounting
Authenticated Time Service
Private Keys within Cloned Images
Back-Doors via Virtualised Test & Monitoring Functions
Multi-Administrator Isolation

Table 1: Summary of Potential New Security Concerns for NFV

In the few cases where solutions are not readily available, e.g. topology validation and network performance isolation, fundamental new research is not required, only engineering effort.

One area, Multi-Administrator Isolation, is still firmly in the research domain. The problem here is that, once someone is given administrator privileges over a compute platform, it is hard to prevent them from accessing the internals of a virtualised function running on their system. This makes 'separation of duties' (SoD) difficult, which is desirable security practice if sensitive functions such as lawful interception are to be virtualised. Technology is close to production that addresses this concern. In the meantime, there are many deployment scenarios for NFV where such isolation is not required or it can be arranged in other ways.

Of course, such a list can never be considered exhaustive. As NFV is increasingly used for more critical services, attackers will become more highly motivated and may uncover new concerns with NFV. Nonetheless, virtualisation technology also enhances the operator's capability to more rapidly deploy new defences against such novel attacks.

In May 2014, growing active participation enabled the security expert group to shift from problem analysis into solution mode by embarking on a guide for practitioners. The **NFV Security and Trust Guidance** document ^[4] identifies areas where evolution to NFV will require different technologies, practices, and processes for security and trust. It also gives security guidance on the environment that supports and interfaces with NFV systems and operations.

Looking forward to the next phase of work, the NFV ISG is starting the process of liaising with other standards bodies to fill the gaps identified by the security expert group, in particular those where solutions are further out, such as lawful interception. New security work on the practicalities of traffic monitoring is also being considered.

Finally, the more practical security work of the NFV ISG should not be overlooked. For example, testing is in progress of the performance of a virtualised router when subjected to a flooding attack. A problem such as this can also be seen as an opportunity: today such attack traffic is redirected to dedicated ‘scrubbing farms’ but, in an NFV future, virtualised scrubbing machines could be spun-up wherever they are needed.

4 Overview of NFV Proof of Concepts

A key goal of the collaborative NFV industry effort is to encourage implementation and growth of an NFV ecosystem. In October 2013 the NFV ISG published a framework and call for NFV Proofs of Concept (PoC),^[5] with the aim to demonstrate NFV as a viable technology and to generate practical knowledge for reference within the NFV ISG. The PoC framework enables the different actors in the industry to show their progress with NFV technologies, irrespective of company focus, company size or NFV ISG membership status (participants do not have to be members of the NFV ISG).

The PoC activity has been very successful, exceeding our most optimistic expectations in terms of participation, engagement, and technical significance. At the time of publication (October 2014), 25 multi-vendor PoCs are in progress, each sponsored by at least one service provider and spanning all elements of the NFV Architectural Framework. Over 50 vendors are participating and all use cases identified by the NFV ISG community are represented. Given the number and significance of the NFV ISG PoCs, and that more recent PoCs are leveraging findings of earlier PoC activities, we are satisfied that our goals of demonstrating feasibility of NFV technologies and encouraging growth of an NFV ecosystem are being achieved.

The next challenge for the NFV ISG PoC framework will be the journey from technical feasibility to technological maturation and interoperability, and this will be progressed as part of the next phase of NFV ISG work.

Information on the PoC Framework along with individual PoC scope and results are openly available on the NFV ISG portal.^[6]

5 NFV Research and Education

Significant industry progress has been made to encourage growth of a commercial ecosystem for NFV, but research and education are also very important for overall and long term success. In order to raise awareness around NFV Research a number of topics have been discussed. Examples include (not exhaustive and not in any particular order):

- Service chaining algorithms
- NFV orchestration algorithms
- Abstractions for carrier-grade networks and services
- Performance studies (optimisation, scheduling, portability, reliability)
- Security of NFV Infrastructure

- Impacts of data plane workloads on computer systems architectures
- Applying compositional patterns (i.e. Network Function Chains) for parallelism
- Performance monitoring and reliability of network services
- Energy-efficient NFV architectures
- Service Assurance (e.g. test & diagnostics, predictive analytics, etc.)
- New requirements on the NFV Infrastructure for supporting new types of VNFs
- NFV Infrastructure federation
- New network topologies and architectures
- Tools and simulation platforms

The rapidly increasing interest in NFV amongst the academic and industrial research communities has translated into NFV being a common keyword in many calls for papers of publications and conferences, and special tracks or workshops related to NFV are being organised.

We encourage industry and academia to participate in the NFV ecosystem by creating research programmes around NFV, and to create new teaching courses to train a new generation of students to be multi-skilled in networks and software.

6 Perspectives on Open Source

Open source software initiatives should be considered as complementary to formal standardisation processes. Since our last update where we referenced the importance of open-source for the production of open reference implementations and producing de-facto standards, open source projects such as OpenDaylight, OpenStack, etc. have created sub-groups to introduce the necessary blueprints to accommodate NFV requirements.

A key step forward is the formation of Open Platform for NFV (OPNFV) launched in September 2014 under the auspices of the Linux Foundation. The goal of OPNFV is to create an open reference platform for NFV consisting of open hardware and software interfaces based on open source projects such as those mentioned above. The reference platform will include virtualised network functions interworking with physical network functions (PNFs).

OPNFV will reference the NFV ISG as a key source for NFV requirements and OPNFV will enable vendors and users (e.g. service providers, enterprise, and cloud service providers) to work together to accelerate NFV implementation and interoperability. Hence, we recognise OPNFV as a key forum to drive open source projects to support NFV features, and to solve major technical implementation challenges including obtaining high performance using industry standard servers, as well as automating the control and management of the NFV environment .

The OPNFV effort is a welcome step, however, the NFV ISG should continue to directly engage relevant open source communities as well as OPNFV.

7 Evolving Relationship with SDN

In our first white paper^[1] we highlighted the highly synergistic nature of NFV and SDN and it is not a coincidence that these two significant industry trends emerged almost at the same time as they build on the rapid evolution of IT and cloud technologies. Perceptions of the future direction for SDN

technology seem to have evolved since the advent of NFV. Partly this is a consequence of time, but also the emergence of NFV itself has provided a stimulus to SDN.

The two key elements of SDN are the separation of the control plane from the data plane to form a domain-wide view of a network, and the ability to abstract and programmatically control network resources. Both capabilities fit nicely into the NFV paradigm and as such SDN can play a significant role in the orchestration of the NFV Infrastructure resources (both physical and virtual) enabling features such as: provisioning and configuration of network connectivity and bandwidth, automation of operations, security and policy control.

NFV creates a very dynamic network environment, driven by customers needing on-demand services and operators needing to manage utilisation and performance of services. Tenant networks will come and go, and VNFs and their connectivity will change frequently to balance load across the infrastructure. The capability to programmatically control network resources (through a centralised or distributed controller) is important in an era of continuous change. Complex network connectivity topologies may be readily built to support automated provisioning of service chains as a realisation of NFV ISG Forwarding Graphs while ensuring strong and consistent implementation of security and other policies. The SDN controller maps to the overall concept of network controller identified in the NFV architectural framework, as a component of the NFVI network domain. As such, an SDN controller can efficiently work with orchestration systems and control both physical and virtual switching, as well as provide the necessary comprehensive network monitoring. However, special attention is needed to ensure that when SDN is applied to telecommunications networks, the separation of control plane and data plane does not cause additional traffic overhead, latency, jitter, etc., as well as redevelopment of existing protocols especially for switching, routing and high availability.

SDN can also benefit from NFV-introduced concepts such as virtualised infrastructure managers and the orchestrator given that an SDN controller could run on a VM. Such an SDN controller could be part of a service chain along with other VNFs and thus rendering itself as a virtualised network function. And the SDN controller itself can be implemented as a VNF to benefit from the reliability and elasticity features brought by NFV.

Ultimately, NFV and SDN will become less distinguishable as independent topics, being subsumed into a unified software-based networking paradigm.

8 NFV Impact on OSS/Network Management

Classical OSS/BSS solutions are based on a set of interconnected applications each focusing on specific functions (e.g. inventory, supervision, performance and traffic monitoring, trouble ticketing, service configuration and activation, test and diagnostics, etc.). Several OSS and BSS sub-domains typically exist within an Operator's domain, e.g. OSS for IT Infrastructure, one or several OSS for carrier network, OSS for mobile user service management, OSS for fixed access services, etc. Some of these OSS may already be fully real-time and automated, but most are not. Current network operations models and OSS solutions are not prepared for emerging new technologies like NFV or SDN. The NFV operations of VNF on-boarding, scaling, and instantiation open the door to highly dynamic network changes. Network architecture, topology and service delivery chain can change

frequently. In combination with service delivery in a multi-vendor environment this can create challenges regarding service or application monitoring.

OSS systems will need to transform in order to accommodate NFV. Operators need to identify which parts of their OSS need to be re-designed to take into account NFV and the associated highly dynamic changes in network topology and connectivity. Furthermore, the plurality of OSS instances within an operator's domain will need to be reflected in the evolution of the NFV ISG Architectural Framework.

It should also be taken into account that NFV infrastructures will be implemented using standard IT hardware operated in a similar way to data centres, although data centre operations can be expected to evolve to accommodate the requirements of NFV. Data centre operations will be managed using optimised processes and tools with IT-centric orientation regarding operational effectiveness and efficiency, independent from NFV applications and supporting any type of application (i.e. not only VNFs).

The NFV Architectural Framework identifies a Management and Orchestration domain that includes three management components, each of which complements the current OSS functionality. The interfaces between the Management and Orchestration entity, the current OSS, and the three management components need to be standardised to reduce integration effort in a multi-vendor environment.

The introduction of NFV also has a major impact on operations. It implies reducing dramatically the complexity of Management and Operations and associated OPEX via transformation of OSS/BSS and processes towards more agility and automation. This can be achieved by introducing new operations scenarios (Fixed & Mobile Networks and associated converged OSS) and autonomic and self-management capabilities, but also with flexibility and network programmability via SDN in an optimal way. Co-existence with current networks and services as well as migration paths and scenarios need to be taken into account as well.

The goal is to strengthen and ease deployment of new services, improve customer experience, generate revenue, and reduce CAPEX /OPEX, all while keeping control of autonomic and self-management processes, programmability, virtualisation, and associated mechanisms to ensure their adoption in a smooth manner.

OSS functionalities should also be virtualised in order to make the management of the networks easier, more flexible and efficient. The harmonisation and/or standardisation of OSS interfaces is critical and new implementation strategies such as Open Source must be taken into account.

Typical vendor-specific element management strategies cannot support the highly dynamic network status changes created by NFV. To deliver the promise of NFV, OSS architectures must evolve along the following lines:

- Taking a holistic view of operations rather than the current piece-meal approach.
- Inheriting terms and definitions from standards rather than creating a separate language.
- Flexible architectures rather than static interface definitions.
- Building blocks defined by existing software rather than architecture-specific software.

- Full automation of the capacity management, optimisation and re-configuration cycle should be done by orchestration and cloud management techniques with open and multi-vendor components rather than vendor-specific management solutions.
- OSS focusing on portfolio management and end-to-end service management.

We believe that the future “telecommunications cloud” will be based largely on industry standard IT cloud technologies, while these technologies will themselves evolve to support the requirements of telecommunications networks.

We foresee that the future “telecommunications cloud” structure will need to be multi-layered as a “cloud-of-clouds”. This implies that the “telecommunications cloud” may be flexibly composed from other cloud environments. These could be private or public, national or international. Such a stacked architecture may comprise different orchestrators, different cloud operating systems, different hypervisors, etc. The “telecommunications cloud” will be highly dynamic not only for auto-scaling of applications but also for automatic infrastructure allocation. The clouds will immediately react to ad-hoc changes of parameters and constraints. An example could be to track the cheapest prices for energy costs and adapt the network topology and/or operating parameters to minimise the cost of running the network. This will require the new OSS to maintain an end-to-end view regarding all offered services.

These revolutionary and evolutionary changes will have a significant impact on the design of the future telecommunications OSS/BSS. Using virtualisation in an up-to-date cloud environment will offer new self-managed redundancy and failover scenarios. The evolved OSS to handle this new environment must be very lean and highly automated, which requires new thinking on OSS that will open up opportunities to gain significant operational benefits.

9 References

1. Original NFV White Paper, October 2012: http://portal.etsi.org/NFV/NFV_White_Paper.pdf
2. NFV ISG Published Documents: <http://www.etsi.org/technologies-clusters/technologies/nfv>
3. NFV White Paper Update, October 2013: http://portal.etsi.org/NFV/NFV_White_Paper2.pdf
4. NFV ISG Draft Documents: http://docbox.etsi.org/ISG/NFV/Open/Latest_Drafts/
5. NFV ISG PoC Framework: <http://www.etsi.org/technologies-clusters/technologies/nfv/nfv-poc>
6. NFV ISG PoCs in Progress: http://nfvwiki.etsi.org/index.php?title=PoCs_Overview

10 Contact Information

If your organisation has any comment on the contents of this white paper, please contact any of the following:

AT&T:	Steve Wright, sw3588@att.com
Bell Canada:	Javan Erfanian, javan.erfanian@bell.ca
BT:	Ian Hawkins, ian.v.hawkins@bt.com
CableLabs:	Don Clarke, d.clarke@cablelabs.com
CenturyLink:	Michael Bugenhagen, michael.k.bugenhagen@centurylink.com
China Mobile:	Robert Chen, robertchen@chinamobile.com
China Telecom:	Yunpeng Xie, xieyp@ctbri.com.cn
China Unicom:	Gongying Gao, gaogy16@chinaunicom.cn
Colt:	Javier Benitez, javier.benitez@colt.net
Deutsche Telekom:	Klaus Martiny, klaus.martiny@telekom.de
DOCOMO:	Tetsuya Nakamura, tetsuya.nakamura.cu@nttdocomo.com
KDDI:	Kenichi Ogaki, ke-oogaki@kddi.com
KPN:	Alexander de Boer, alexander.deboer@kpn.com
KT:	Kisang Ok, ksok@kt.com
NTT:	Takashi Shimizu, shimizu.takashi@lab.ntt.co.jp
Ooredoo:	Marco Stura, mstura@ooredoo.com
Orange:	Bruno Chatras, bruno.chatras@orange.com
Portugal Telecom:	Jorge Carapinha, jorgec@telecom.pt
SK Telecom:	DK Lee, dongkee.lee@sk.com
Softbank:	Satoru Matsushima, satoru.matsushima@g.softbank.co.jp
Sprint:	Laurent Laporte, laurent.laporte@sprint.com
Swisscom:	Markus Brunner, markus.brunner3@swisscom.com
Telecom Italia:	Elena Demaria, elena.demaria@telecomitalia.it
Telefonica:	Diego López, diego@tid.es
Telenor:	Patrick Waldemar, Patrick.Waldemar@telenor.com
Telstra:	Daniel Kirkham, daniel.kirkham@team.telstra.com
Turk Telekom:	Mustafa Ergen, mustafa.ergen@turktelekom.com.tr
Verizon:	Naseem Khan, naseem.a.khan@verizon.com
Vodafone:	Susana Sabater, susana.sabater@vodafone.com
Windstream:	Arthur Nichols, arthur.nichols@windstream.com

11 Glossary

BNG	Border Network Gateway
BRAS	Broadband Remote Access Server
BSS	Business Support System
CAPEX	Capital Expenses
CDN	Content Distribution Network
COTS	Commercial-off-the-Shelf
CPU	Central Processing Unit
vCPU	Virtual Central Processing Unit
C-RAN	Cloud Radio Access Network
DPI	Deep Packet Inspection

EM	Element Manager
ETSI	European Telecommunications Standards Institute (in practise ETSI has global membership)
FCAPS	Fault, Configuration, Accounting, Performance & Security
Gbps	Gigabits per second
IaaS	Infrastructure as a Service
IMS	IP Multimedia System
IO	Input / Output
ISG	Industry Specification Group.
IT	Information Technology
KPI	Key Performance Indicators
KVM	Kernel-based Virtual Machine
M&O or MANO	Management and Orchestration
MRF	Media Resource Function
NFV	Network Functions Virtualisation
NFVI	Network Functions Virtualisation Infrastructure
NFVIaaS	NFVI as a Service
NFVO	NFV Orchestrator
NFVPaaS	NFV Platform as a Service
NSD	Network Service Descriptor
NUMA	Non-Uniform Memory Access
OAM	Operations Administration and Maintenance
OPEX	Operations Expenses
OPNFV	Open Platform for NFV
OS	Operating System
OSS	Operations Support System
PaaS	Platform as a Service
PNF	Physical Network Function (typically proprietary hardware)
PoC	Proof of Concept
SaaS	Software as a Service
SDN	Software Defined Network
SoD	Separation of Duties
SDO	Standards Development Organisation
TDF	Traffic Detection Function
VIM	Virtual Infrastructure Manager
VLD	Virtual Link Descriptor
VM	Virtual Machine
VNF	Virtual Network Function
VNFaaS	VNF as a Service
VNFC	VNF Component
VNFD	VNF Descriptor
VNFFGD	VNF Forwarding Graph Descriptors
VNFM	VNF Manager
VNPaaS	Virtual Network Platform as a Service