
INTERNATIONAL INTERCONNECTION FORUM
FOR SERVICES OVER IP
(i3 FORUM)

(www.i3forum.org)

Source:

i3 forum keyword: NFV

A Primer on NFV
(Release 1.0) March 2016

This document provides the i3 forum's perspective on NFV, focusing on the role of and the impact to international carriers. It does not intend to duplicate other existing specifications or documents on the same issue, but to complement these documents with the perspective of the international carrier members of i3 forum.

Table of Contents

1. Scope and Objective of the Document	3
2. Acronyms	4
3. Introduction.....	5
3.1. What Is NFV?	5
3.2. What Has Enabled NFV?.....	6
4. NFV Drivers.....	7
5. NFV Market Data.....	8
5.1. Infonetics	8
5.2. SDxCentral	9
5.3. Technology Business Research Inc.....	10
6. Benefits of NFV	12
7. Challenges to Implementing NFV	13
7.1. Organization	13
7.2. Cloud Environment	14
7.2.1. Data Plane Throughput	14
7.2.2. Orchestration.....	14
7.2.3. Transcoding.....	16
7.3. OSS	16
8. NFV Standards.....	18
9. NFV In Practice	20
9.1. Network Functions Ripe for Virtualization	20
9.2. NFV Environment	22
10. NFV Use Cases & Deployments.....	23
10.1. Use Cases.....	23
10.2. Deployments	24
11. Conclusions	26

1. Scope and Objective of the Document

On Oct. 22, 2012, over a dozen Tier 1 telecom service providers released a white paper entitled “Network Functions Virtualization - An Introduction, Benefits, Enablers, Challenges & Call for Action.” This white paper laid out the evolution from proprietary hardware to commodity hardware within the network equipment industry. That paper hypothesized that telecom service providers intend to move to a model that relies on a generic, industry-standard hardware infrastructure with the network functions themselves being implemented entirely in software.

As the hypothesis has proven to be true, this i3Forum document’s first release is intended to be a primer on NFV. It describes the current strategic environment, the architectures, the drivers and the related benefits of NFV for telecom service providers the world over. The target audience for this document is intended to be engineers and product managers within international wholesalers.

After reading this document, you should have a foundational understanding of:

- 1) What NFV actually is and is intended to accomplish within telecom service providers’ networks
- 2) What is driving telecom service providers to investigate NFV more fully
- 3) What market data analysts have published regarding NFV
- 4) The benefits telecom service providers can expect to realize via an NFV-based network architecture
- 5) What is involved in preparing not only the network, but the organization as a whole, for the migration toward NFV
- 6) The standards forming around NFV deployments
- 7) The creation and management of a typical NFV deployment
- 8) Some typical NFV use cases, and early wins

The ultimate objective of the document is to provide an initial starting point for understanding and illuminating the concepts driving the adoption of network function virtualization.

2. Acronyms

CAGR	<i>Compound Annual Growth Rate</i>
CapEx	<i>Capital Expenditure</i>
CPU	<i>Central Processing Unit</i>
DSP	<i>Digital Signal Processor</i>
EOL	<i>End of Life</i>
EPC	<i>Evolved Packet Core</i>
ETSI	<i>European Telecommunication Standard Institute</i>
I/O	<i>Input/Output</i>
IMS	<i>IP Multimedia Subsystem</i>
IP	<i>Internet Protocol</i>
ISG	<i>Industry Specification Group</i>
IT	<i>Information Technology</i>
IaaS	<i>Infrastructure as a Service</i>
KVM	<i>Kernel-based Virtual Machine</i>
MANO	<i>Management and Orchestration</i>
NFV	<i>Network Function Virtualization</i>
NFVI	<i>Network Function Virtualization Infrastructure</i>
OpEx	<i>Operating Expenditure</i>
OPNFV	<i>Open Platform for NFV</i>
OSS	<i>Operations Support System</i>
OTT	<i>Over The Top</i>
PaaS	<i>Platform as a Service</i>
PoC	<i>Proof of Concept</i>
R&D	<i>Research & Development</i>
RAN	<i>Radio Access Network</i>
SaaS	<i>Software as a Service</i>
SDO	<i>Standards Development Organizations</i>
SI	<i>Systems Integrator</i>
SIP	<i>Session Initiation Protocol</i>
SDN	<i>Software Defined Networking</i>
SR-IOV	<i>Single Root I/O Virtualization</i>
TDM	<i>Time-Division Multiplexing</i>
VM	<i>Virtual Machine</i>
VNF	<i>Virtualized Network Function</i>
VPaaS	<i>Voice Platform as a Service</i>

3. Introduction

The telecom industry is embarking on a major pivot: the transition to network function virtualization. NFV promises new ways of building communications networks at a lower cost and with greater scope for innovation in network services.

Regardless of what kind of network – fixed, mobile, metro Ethernet, long distance, interconnect – or what kinds of services – data, voice, messaging, content delivery, virtual private networks – NFV offers opportunities to transform the economics of the network while at the same time accelerating the ability to design and deploy new service capabilities.

Legacy networks are becoming too expensive to maintain, and the competition from OTT providers is too fierce to ignore. Telecom service providers' networks need an ability to deliver a true, multimedia experience across a variety of different access types and devices with sustainable costs. In short, NFV represents the most significant pivot in the telecom industry since the transition from TDM to packet got under way a decade or more ago.

This paper will illustrate why NFV is one of the key topics for telecom service providers right now.

3.1. What Is NFV?

NFV is the concept of leveraging the advances in, and technology from, IT/enterprise virtualization, the application of several tweaks along the way to smooth over the differences between IT-grade and telco-grade network architectures, and applying the same to the world of networking. Put another way, NFV is all about applying IT technologies, including virtualization, cloud and data center hardware, to the problem of building networks and network services.

The end result of a successfully implemented NFV strategy is a pool of commodity servers running scalable cloud middleware (like Openstack, or VMware's Vsphere) that enables applications to be deployed without having to dedicate specific machines and resources of various types for specific tasks.

NFV provides tremendous flexibility: it enables any given pool of servers to be used in the most efficient way possible, while at the same time making it very easy to expand capacity as demand grows.




Figure 1: NFV Pictogram

Virtualization enables the deployment of new applications and new services without investing in proprietary hardware or closed operating systems with the accompanying management burdens and training costs.

3.2. What Has Enabled NFV?

Virtualization allows for efficient sharing of physical resources, similar to what one sees in data centers used for virtualizing IT and business applications. So the technologies that power NFV evolved from the support of IT workloads such as Web serving, transaction processing and databases. NFV leverages the same server and storage virtualization technologies created and perfected there. It is helpful that some network functions actually do resemble IT workloads in terms of how they use underlying compute, network fabric and storage resources. But there's a different side to the network: workloads that involve pumping packets at terabit rates. That required some additional support.

Over the last decade, Intel has enhanced each successive generation of x86 processors to improve their ability to handle network-centric workloads. The net effect of these hardware enhancements, combined with software toolkits developed by Intel and others such as the Data Plane Development Kit, is that commercial off-the-shelf (COTS) servers can now support throughputs of hundreds of gigabits per second, and many tens of millions of packets per second. COTS servers are more than capable of supporting the great majority of network functions that power today's communications services.

Moreover, the size of the worldwide data center market consumes R&D innovation and personnel development targeted at improving generic server technology. Accordingly, telecom service providers will be able to leverage the recent advances in VM orchestration to provision, scale up or down, and move network functions dynamically across servers in response to changing user demand. This provides telecom service providers with tremendous flexibility to manage network functions cost effectively.

For example, as carriers begin migrating from TDM to VoIP platforms, NFV is becoming a reality. This is driven by the fact that all commodity servers offer up standard IP interfaces and IP network stacks. The proprietary TDM interfaces and stacks (which cannot be virtualized) are no longer a factor.

4. NFV Drivers

Cost reduction, efficiency gains, innovation and fast service deployment are the key drivers for NFV. Traditional telecom service providers have invested heavily in variety of proprietary hardware appliances. Every proprietary hardware appliance consumes space, power and network ports to operate, worse yet these appliances rapidly become end of life as technology and service innovation progress. To tackle these issues, telecommunication equipment manufacturers have focused on developing NFV and software defined networking (SDN) in recent years.

SDN is a little harder to define than NFV, although NFV and SDN have much in common. They are complementary, not exclusionary. The simplest way to differentiate between the two is to understand that SDN concerns itself with the separation of the control and data planes in a network, the centralization of control of the network, and the ability to program the network on the fly. NFV, on the other hand, requires that all dedicated network appliances be migrated over to generic servers. SDN concerns itself highly with orchestration (which we'll cover later); NFV is the building block that enables orchestration.

Over the past few years, cloud computing has made significant strides to ease the way for companies to maintain and offer new services. Increased IP bandwidth, speed and reliability, along with reduction in the cost of IP connections in recent years, have prompted companies to pursue virtualization in every aspect of the business. Hardware manufacturers are on the same page, offering products with virtualization enablement. A variety of cloud service providers are popping up to provide IaaS, PaaS, SaaS and VPaaS. These providers focus on procuring hardware, network, storage and interfaces to meet companies' needs at different levels. Network equipment vendors are also joining the virtualization world by concentrating efforts on developing SDN software, allowing telecom service providers to take the most advantage of NFV. This new landscape provides a perfect opportunity for telecom service providers to re-architect their infrastructure toward NFV.

According to Infonetics:

1. NFV and SDN are necessary to the survival of telecom service providers who are required to compete against the threats of OTTs, traffic/video, mobility, and cloud, and meet the ever increasing "on-demand" expectations of users.
2. Big telecom service providers are investing resources in SDN and NFV. Therefore vendors must follow their lead/demand/requirements or be left behind to wither or die in irrelevance.
3. Telecom service providers can start small as some NFV use cases, and SDN network domains have higher priority and/or are easier to bite off as a starting project.
4. Some smaller or aggressive telecom service providers are out early with commercial deployments, which will compel the larger telecom service providers to follow.
5. The market for VNFs (hypervisor-capable software versions of network functions running on commercial servers) has been in existence for several years, and predates the terms NFV and VNF.

An obstacle hindering VoIP carriers' move toward NFV is call transcoding without making use of proprietary DSPs. Vendors in recent years, however, have improved software transcoding capabilities to close the gap with hardware-based transcoding. We will be discussing transcoding within a virtualized environment in more detail later in this document.

5. NFV Market Data

The advent of NFV and SDN represents one of the most important technology events in the telecom industry since the arrival of digital switching. Telecom service providers that choose fully to embrace NFV have the opportunity to radically transform their businesses, both to reduce their cost base, and to become far more agile in their ability to introduce new services. In short, to become “software telcos.”

5.1. Infonetics

According to Infonetics, in the five years from 2014 to 2019:

- Total service provider NFV market will grow from \$950 million to \$11.6 billion, a 65 percent CAGR
- NFV software (NFV MANO and VNFs) will grow from \$771 million to \$9.4 billion, a 65 percent CAGR
- NFV hardware (NFVI server, storage, switches) will grow from \$153 million to \$1.8 billion, a 64 percent CAGR

From a timeline perspective, Infonetics shows that, starting in 2016, NFV will start to have wider-spread commercial deployments:




Figure 2: NFV Deployment Timeline

Infonetics also predicts that:

- Software will be a much larger investment than the server, storage and switch hardware, representing about \$1 spent for every \$5 spent on software (NFV MANO, plus the VNFs).
- The carrier spend on professional services outsourced to vendors for NFV projects will grow to 14 times its small start of \$27 million in 2014, to \$387 million in 2019.




Figure 3: Infonetics Forecast of NFV Revenues

Infonetics is by no means alone in its belief that NFV will experience explosive growth over the next five years. Several other analyst groups forecast much of the same. Two additional forecasts are reproduced below for reference.

5.2. SDxCentral

SDxCentral states:

A major portion of L2/L3 spend will migrate from HW spend to a new suite of software-only networking applications. We estimate that by 2020, the market for just L2-3 networking SW apps will be \$14B. These will appear as applications running on controllers or integrated into provisioning or orchestration systems, or in some situations run partially as VNFs on NFV infrastructure.




Figure 4: SDxCentral SDN & NFV Market Size Report – 2015 Edition

5.3. Technology Business Research Inc.

TBR states:

Early adopters such as AT&T started investing in NFV and SDN in 2014, marking 2014 as the first year of tangible spend on these technologies. They are driving a significant amount of development in the NFV and SDN ecosystem and are pushing the vendor community to rapidly adapt to this new architectural approach to networks. NFV and SDN spend volume is forecasted to ramp up in 2017, at which time use cases will be more defined and the cost benefits of using the technologies will be more apparent. This will prompt holdout operators to jump on the bandwagon and aggressively pursue transformation with these technologies to avoid getting left behind.




Figure 5: Technology Business Research (TBR) NFV/SDN Market Forecast – September 2015

6. Benefits of NFV

So what do the charts actually mean? In five years' time, we can expect to see telecom service providers enjoying the following benefits of NFV:

1. **CapEx Cost Reduction**

Almost all new network build-outs will be based on highly commoditized hardware, combined with best-of-breed software. Via equipment consolidation, telecom service providers can reduce the cost of procuring expensive proprietary hardware. This means a reduction in rack space, power and network resources by sharing generic hardware and storage. With consolidation, the cost of deploying, maintaining and operating services can also be lowered.

2. **OpEx Cost Reduction**

Automation of operations management will dramatically reduce operational costs. For example, in a conventional network, the failure of a piece of hardware is typically not service-affecting, but it requires urgent action to replace the failed hardware in order to restore the proper level of fault tolerance. With NFV, the failure of a piece of hardware has no more impact than a temporary reduction in the maximum capacity of a given service, which is accommodated within planned hardware capacity headroom. There is no urgency to replace the failed hardware, and the procedure for doing so is extremely simple and completely standardized.

3. **Faster Time to Market**

Almost all new services are created by deploying new software elements in the network, or by combining existing software elements in new ways. Without dedicated appliances, it will be possible to prototype new services in a matter of days, and bring new services to market in weeks. Using standard virtualization technologies, telecom service providers can quickly load software to generic servers remotely, allowing carriers to trial new services rapidly. It will be possible to respond to feedback about new services far more quickly than currently. New services that are successful can be scaled very rapidly, while new services that fail to gain traction can be stood down without the loss of major investments.

4. **Scalability**

Virtualization allows a particular service to scale by dynamically allocating shared resources efficiently amongst servers. If needed, generic servers can be added to expand the resource pool.

5. **Customization**

The flexibility of a virtualized infrastructure enables telecom service providers to offer entirely new types of services. For example, a telecom service provider can address the needs of a customer who wants some specialized communications service by creating a virtualized service instance specifically for that customer, and customizing it accordingly. This can be accomplished far more quickly, easily and safely than attempting to satisfy specialized needs by making incremental software enhancements to some common, shared service platform based on proprietary hardware.

6. **Reduce Vendor Dependency**

With NFV, telecom service providers can easily enable different vendor platforms in the existing infrastructure to increase speed to market, new features or innovation.

7. **Ease of Management**

With consolidation of hardware infrastructure, managing, from inventory to security to hardware faults, becomes easier. In addition, evaluating new software in the QA lab could be much faster.

8. **Ease of Software Lifecycle Management**

With virtualization, it's easier to upgrade or downgrade software versions with minimal service interruption by spawning a new instance on the same physical hardware.

7. Challenges to Implementing NFV

7.1. Organization

The NFV vision of the future will only be achieved by telecom service providers that embrace NFV fully and are prepared to implement a fundamental organizational transformation to make the most of what they enable. Below is a list of potential pitfalls that could make the migration to NFV a challenging task:

1. Culture

Decades of procuring and building networks in the traditional way have created a culture that may put up strong resistance to the kinds of changes that are needed. Resistance may take many forms, including:

- Persistent disbelief that cloud-based software can ever deliver carrier-class, five-nines service availability
- Inability to trust any vendor other than the traditional suppliers to deliver telco-grade solutions
- Reluctance to acquire the new skills necessary to build and manage services in cloud-based software environment

2. Headcount

Resistance will also arise from the simple observation that, if NFV delivers major savings in operational costs, along with more and more automation, a lot fewer heads will be involved in operations.

3. Performance Evaluation

Telecom service providers must evaluate different vendors' solutions versus their existing proprietary hardware equipment. This applies in particular to transcoding evaluations, as software based transcoding still lags behind purpose-built DSPs. However, vendors are actively coming up with innovative NFV solutions for transcoding.

4. Re-architecting NFV Solution to Fit Existing Infrastructure

This requires a lot of planning, from system integration to backend office modifications. The migration will likely occur in phases to be successful. For example:

- Product introduction
- Hybrid existence between legacy and NFV
- Full NFV enablement

5. Inter-departmental Commitment

NFV will touch many different teams within the organization. It's imperative for all departments to understand the benefits and changes that are required on their part to make NFV successful. Examples of some of the challenges include:

- The capacity management team has to evaluate the new platform capabilities and different ways to size capacity.
- The technical staff has to be trained on the virtual environment and new setup steps.
- IT needs to plan for the new way of managing to, and investing in, virtualization and legacy equipment.
- The inventory system has to be updated to reflect the virtualization setup.
- Service Operations needs to understand out how NFV impacts the provisioning of new services.
- IP Management needs to re-architect the new NFV schema, possibly using higher bandwidth/speed port for NFV equipment.

6. Customer Migration Plan

Customers have to be informed about the changes and the migration plan.

7.2. Cloud Environment

It is entirely possible to start out with NFV deploying virtualized network functions just as if they were IT workloads. However, current cloud environments have significant shortcomings in three areas, which need to be addressed to some degree in order to realize the full benefits of NFV – namely data plane throughput; operations automation or “orchestration”; and transcoding.

7.2.1. Data Plane Throughput

Most IT workloads such as Web serving, transaction processing and big data analytics are not particularly input/output intensive. The capacity of most virtualized IT applications is determined primarily by CPU utilization rather than network I/O. However, in the NFV domain, data plane workloads are concerned with the routing, switching, relaying or processing of network traffic payloads. These kinds of workloads are I/O intensive, often requiring some combination of total I/O bandwidth and packets-per-second throughput that is orders of magnitude greater than is typical of IT workloads. Currently, all of the commercially available cloud computing environments impose significant limitations on the throughput of data plane workloads.

These limitations arise because virtualization introduces a layer of software between virtualized guest applications and host networking hardware. This layer of software, the vSwitch, has not (yet) been optimized for high-bandwidth or high-packet-rate applications.

The vSwitch performance issue is widely recognized by NFV industry players, and is being addressed by multiple vendors. There are two contrasting approaches to solving the problem, and we are likely to see both approaches succeeding in the market.

- The first approach, known as SR-IOV, is supported on many current Ethernet NICs. However, SR-IOV is a relatively new technology, and is currently not well supported by cloud management software solutions. But it's only a matter of time for cloud stacks to catch up and offer full support for SR-IOV.
- The second approach is to optimize the performance of the software vSwitch for high-bandwidth and high-packet-rate applications. In general, this involves re-architecting the vSwitch software so as to reduce or even eliminate packet-copying operations in the data path. Since cloud stacks already deal with the configuration of the software vSwitch, this approach to data plane performance improvement doesn't require anything new of the cloud stack.

Both of these approaches can deliver levels of data plane performance that make NFV an economically sound proposition for the great majority of virtualized network functions. SR-IOV is known to deliver data plane performance that is very close to that of bare metal. While the outlook for vSwitch software acceleration suggests substantially less impressive performance gains, it may offer some compensating benefits in terms of improved cloud flexibility.

7.2.2. Orchestration

Existing cloud management tools provide quite sophisticated operations management support in an IT environment, but arguably they lack certain functions that may be valuable in a carrier networking NFV environment, particularly in the area known as “orchestration.”

Orchestration is concerned with automation of the life-cycle management operations for virtualized network functions, and includes the following:

- **Service Instantiation**
Deployment of the VNF software components and configuration of the virtual network infrastructure so as to create an instance of a network service.
- **Service Component Health Monitoring & Repair**
Monitoring of the virtual machines that are running VNF software components, reporting on errors and failures, and performing repair operations when a component failure is detected.

- **Elastic Scaling**
Continuous monitoring of the load on VNF software components that support a given service, and dynamically increasing or reducing the population of VNF component instances in response to changing load conditions so as to maximize the usage efficiency of the NFV hardware.
- **VM Migration**
Managing the movement of VNF software instances off a given host in order to enable that host to be taken out of service—for example, for maintenance.
- **Software Upgrade**
Managing the in-service upgrade of VNF software that supports a given service.
- **License Management**
Tracking the usage of VNF software for vendor revenue assurance purposes.
- **Service Termination**
Gracefully shutting down a service when it is no longer needed.

All of these operations could, in principle, be performed manually via a cloud management console. However, automation of frequently performed life-cycle management operations will substantially reduce operational costs and is also likely to reduce the likelihood of human error during the performance of operations.

The value of automating any given aspect of NFV operations management depends to a considerable degree on the nature of the service being deployed on NFV. For example, a large-scale, multi-tenanted service such as Voice over LTE supported by IMS is instantiated only very infrequently, so automating service instantiation in this case is likely to be a low priority.

Furthermore, the load on such a service is likely to be highly predictable, so automatic elastic scaling may not be a high priority. On the other hand, monitoring the health of the service and automatically repairing failed VNF component instances is likely to be very important. By contrast, some services such as enterprise VPNs are single-tenanted, so a service instantiation operation is required for each new customer of the service. There is obvious value in automating that operation, and possibly even enabling self-service instantiation by the customer via a Web portal.




Figure 8: Detailed Orchestration Diagram

The concept of orchestration exists in IT clouds as well as in the world of NFV, but NFV demands rather more from orchestration. In the IT world, the focus is mainly on service instantiation, with some limited capabilities for monitoring and repair.

The emerging NFV industry is starting to demand enhanced orchestration solutions that address additional aspects of VNF life-cycle management such as elastic scaling, optimized placement of VNFs relative to the underlying switching fabric and more fine-grained control of the switching fabric via SDN controllers.

When setting priorities for which network functions to deploy first in virtualized form in an NFV environment, network operators would do well to assess the kinds of orchestration that may be required in order to successfully manage those virtualized network functions. Given the relative immaturity of NFV orchestration solutions, in this first stage it will make sense to focus on those VNFs that are least demanding in terms of management automation.

7.2.3. Transcoding

The proportion of telecom service providers that will require transcoding is significant. In general, there are two different approaches to transcoding, each with different capabilities. Hardware transcoding uses DSPs to perform the actual transcoding and generally provides for better density. Hardware transcoding supports some features software transcoding currently does not, notably tone generation and ptime interworking (sometimes known as transrating).

Software transcoding that runs on standard Intel CPUs is generally a better solution for small volumes. Part of the reason for that is software transcoding is highly resource intensive. In real-world side-by-side comparisons, purpose-built DSPs handle seventeen times the number of transcoded sessions of Intel CPUs.

Transcoding is a necessary evil as it's costly in terms of compute resource required. So, at the moment, telecom service providers looking for variable-scale, high-flexibility, cloud-based NFV deployments that fully embrace the Telco-in-the-Cloud must either design a hybrid solution where high-density DSPs are deployed in dedicated hardware to handle transcoding duties, or cope with the massive general-purpose CPU compute resources required to perform transcoding in software.

7.3. OSS

NFV will clearly have a major impact on the way in which network services are managed. Traditional OSS solutions are architected around the concept that network services are built on a set of appliances in which software and hardware are tightly integrated. In the NFV world, software and hardware need to be managed separately. Hardware comprises a more or less homogeneous pool of processing, switching and storage resources, while services are created entirely by deploying and configuring virtualized software elements.

An NFV environment looks very different from a traditional network environment from the point of view of the OSS.

NFV also introduces a new set of operations management requirements in the area of orchestration. Concepts such as point-and-click service instantiation, elastic scaling and automatic recovery following hardware failure just don't exist in the traditional network environment.

However, it's obvious that the OSS is going to have to deal with these concepts. It's also worth pointing out that one of the key benefits of NFV to network operators is the ability to deploy new services more quickly by eliminating the need to qualify and approve new hardware.

Traditionally, the time taken to integrate new network equipment into the OSS was a major factor in the timescales for deploying a new service. It should be much quicker to integrate new VNFs into the OSS because there is no need for the OSS to deal with any hardware management functions associated with

the VNF, but some integration effort will be required to enable the OSS to configure the new VNF and handle fault reports and performance statistics. Ideally, the OSS should be designed to make it quick and easy to integrate new VNFs, so to reflect the automatic node instantiation covering the OSS integration as well.

It's clear that an OSS solution that addresses the new realities of NFV is going to look very different from a traditional OSS. Indeed, the degree of change is so significant that network operators should seriously consider whether NFV justifies the introduction of a whole new generation of OSS. With many of the OSS solutions currently in use having their roots in network practice dating back 20 years or more, NFV might be the triggering event for an in-depth review of future OSS strategy.

8. NFV Standards

One of the more significant impacts to NFV is how the industry is handling the need for standards. As the virtualization of network elements and telco platforms was accelerating in the late 2000s, so also was the push from the involved telecom service providers and manufacturers toward a standardization process. The first objective of standardization was to create a framework that would not only include pure virtualization, but also define the requirements. The result was:

- A consistent ecosystem hosting logical nodes
- Definitions of roles and interfaces
- Encouragement of interoperability and openness
- Providing guidelines
- Addressing technical challenges
- Promoting proofs of concept and similar activities

ETSI is the main driver in terms of setting NFV standards for the telecom industry. A specific ISG called NFV was set up, and it gradually grew to include other service providers, as well as IT and network technology vendors, currently numbering more than 270 members.

Since NFV as a concept encompasses a great number of issues, ETSI NFV decided to organize the work in phases, in order to achieve gradual results in reasonable timeframes and in line with the growing awareness of the many related aspects from all stakeholders. Their white paper published in 2012, titled “Network Function Virtualization,” is essential reading for anyone who wishes to understand the topic in more depth. It is worth noting that the NFV ISG, as the name suggests, does not provide binding standards itself (ETSI’s SDOs do). Yet this group is a central point where the lack of proper standards for NFV can be pointed out and the relevant standardization groups can be involved subsequently.

The NFV ISG recently released documents related to the completion of the “Phase One” work on establishing a framework for NFV. The latest NFV documents include:

- An infrastructure overview
- An updated architectural framework describing the following aspects of the infrastructure:
 - Compute
 - Hypervisor
 - Network domains
- Management and orchestration
- Security and trust
- Resilience and service quality metrics

These new and updated documents are built on the initial Phase One work first released in late 2013. They provide the core of the ETSI documentation for NFV, on top of which current and future documents will be released describing more detailed aspects.

ETSI also announced that the NFV plans had entered “Phase Two.” Phase Two work is set to include growing interoperability across the NFV ecosystem; specifying reference points and requirements that were defined in Phase One; growing industry engagement to ensure that its NFV requirements are met; and clarifying how NFV intersects with other standards, including SDN and open-source initiatives. Since the beginning, the NFV group has promoted collaboration with external organizations, making their initiatives complementary to the specifications and standardization work done in ETSI and aiming to speed up the adoption of this technology and deal with specific matters about NFV.

Outside of ETSI, a number of organizations popped up in 2014, either based around a certain vendor’s work in the NFV space that they then opened up to others, or more diverse organizations geared toward bringing vendors together to work on specific aspects of NFV. Those organizations include:

- OpenDaylight Project
- OPNFV
- ON.Lab
- Open Stack

Many of these organizations are looking to tackle certain slices of the NFV pie. For instance, the Linux Foundation’s OPNFV program said it “will establish a carrier grade, integrated, open source reference

platform that industry peers will build together to advance the evolution of NFV and to ensure consistency, performance and interoperability among multiple open source components. Because multiple open source NFV building blocks already exist, OPNFV will work with upstream projects to coordinate continuous integration and testing while filling development gaps.”

In this initial stage, the work of OPNFV is focused specifically in the NFV-I and VIM layers, proposing solutions that make use of other open-source organizations, like OpenStack, KVM and Linux. As an advantage, many of the OPNFV members participate in NFV ISG as well.

9. NFV In Practice

9.1. Network Functions Ripe for Virtualization

In a traditional network, each distinct function was typically implemented as a specialized appliance based on proprietary hardware. Such appliances invariably include a substantial amount of software, but the software and hardware can't be separated. They are highly dependent on one another. Examples of traditional proprietary hardware-based network elements include:

- Routers of various kinds
- Deep packet inspection devices
- Content delivery network appliances
- Firewalls
- Load balancers
- Network address translators
- Session border controllers
- Mobile base station controllers
- Mobile packet gateways




Figure 6: Network Functions Can Migrate Over to Commodity Hardware

Any network function that is capable of being deployed in the network over generic Ethernet interfaces can potentially be virtualized and deployed in an NFV environment. This obviously rules out network functions that depend on specialized physical interfaces such as optical transport devices, but it leaves a very long list of possibilities.

NFV is based on the concept that network functions can be implemented entirely in software running on "industry-standard hardware." In general, industry-standard hardware is taken to mean commercial off-the-shelf servers based on Intel's x86 architecture, together with commercial off-the-shelf Ethernet switching devices. In the new paradigm the network function itself -- for example, a session border


controller function -- is delivered to the network operator as piece of pure software. This is then installed by the network operator on a standardized hardware infrastructure that is typical of a data center environment: rack-mounted or blade servers connected by Ethernet switching systems.

In general, it doesn't make financial sense to use NFV to replace existing physical network functions with virtualized equivalents, unless, for example, those physical network functions are at or near end of life. The first steps in NFV are likely to be around areas of the network where new build-outs are planned. This is going to vary from one telecom service provider to another, but some good examples of new network build-outs where NFV can make sense are:

- Evolved packet core (EPC) - mobile broadband access IP infrastructure for LTE
- IP multimedia subsystem (IMS) - SIP infrastructure for Voice over LTE
- Session border controllers (SBC) - for IP-based voice network interconnect and business voice access via SIP trunking or hosted PBX

9.2. NFV Environment

Currently, the main players in an NFV ecosystem include:



Hypervisors	Virtual Networking	Cloud Software	Public Clouds	Distro & SI	Orchestration & Automation
•KVM •Xen •VMware •Hyper-V •LXC	•Open V Switch •6Wind •Nuage (ALU) •Contrail (Juniper) •Nexus vswitch (Cisco) •VMware vswitch •Calico (Metaswitch)	•OpenStack •VMware •Apache CloudStack	•AWS •HP Helion •IBM Softlayer •Rackspace •Google	•CloudBand (ALU) •Mirantis •Red Hat •Cisco •Canonical •VMware •Windriver •HP	•Blue Planet (Cyan) •SoftLayer (IBM) •Amdocs •Cisco •Cloudband •Helion •OpenStack Heat •Juju (Canonical) •Overture

Figure 7: NFV Ecosystem Participants

But the true starting point for NFV is a cloud environment of the kind that is widely used today to support IT workloads. This includes three main ingredients:

- Commercial off-the-shelf servers
- A hypervisor such as KVM or ESXi
- A cloud management solution such as OpenStack or VMware vSphere

10. NFV Use Cases & Deployments

10.1. Use Cases

ETSI in Document “DGS/NFV-009” describes several different use cases. Here is a small description of them applicable to the interconnect market:

Use Case #1: NFV IaaS

The NFV infrastructure's function is providing an environment in which virtualized network functions can execute. The NFV-I shall provide compute capabilities comparable to an IaaS cloud computing service as a run-time execution environment, as well as support the dynamic network connectivity services that may be considered as comparable to NetworkaaS. This use case provides an approach to mapping the cloud computing service models IaaS and NaaS as elements within the network function virtualization infrastructure when it is provided as a service.

Use Case #2: Virtual Network Function as a Service (VNaaS)

There are two business models; either the service provider or the enterprise can own and operate a CPE. Virtualization of the enterprise may include:

- Virtualization of the CPE functions (vE-CPE) in the service provider cloud.
- Virtualization of the PE functions (vPE) where the virtual network services functions and core-facing PE functions can be executed in the service provider cloud.

Use Case #3: Virtual Network Platform as a Service (VNPaaS)

Enterprises and other service providers (operators) may deploy certain services based on a service catalog within the network of a hosting service provider (hosting operator). Enterprises and service operators may use predefined service templates or certain orchestration functions, or even deploy own (black-box) services.

Use Case #4: VNF Forwarding Graphs

In some VNF FGs, packets have a specific destination (e.g., a set of virtual server functions) while in others packets have no specific destination (e.g., the Internet). Many other use cases share characteristics with this VNF FG use case. Requirements, architecture and specifications on these common characteristics should meet the NFV goals for enabling migration from existing physical network functions to virtual analogues, as well as enabling implementation of new functions and arrangements not previously envisioned.

Use Case #5: Virtualization of Mobile Core Network and IMS

In the IP multimedia subsystem (IMS), which is a session control architecture to support the provisioning of multimedia services over EPC and other IP-based networks, examples of network functions include P-CSCF, S-CSCF, etc. HSS and PCRF are other 3GPP network functions, which are required in the end-to-end architecture to work in conjunction with the EPC and IMS for providing the service. Similarly, the online and offline charging systems (OCS and OFCS) are systems that capture the charging records as part of the session management. This use case aims at applying virtualization to the EPC, the IMS, and these other network functions mentioned above.

Use Case #6: Virtualization of Mobile base station

In major mobile operators' networks, multiple RAN nodes from multiple vendors are usually operated with different mobile network systems, e.g. 3G, LTE and WiMAX, in the same area. These multiple platforms expect to be consolidated into a physical base station (BS) based on IT virtualization technologies. A RAN node utilization is usually lower than its max capacity because the system is designed to cover the peak load. However, the average load is far lower, and each RAN node resource cannot be shared with other nodes. Base station (BS) virtualization can achieve sharing of resources among multiple logical RAN nodes from different systems, dynamically allocating the resource as well as reducing power consumption.

Centralized-RAN (C-RAN) technology with virtualization can leverage more efficient resource utilization among different physical BSs

Use Case #7: Virtualization of the Home Environment

NFV technology facilitates virtualization of services and functionality migration from home devices to the NFV cloud.

Use Case #8: Virtualization of CDNs (vCDN)

Specific cases are third parties like CDN providers or large content providers who ask operators to deploy their proprietary cache nodes into the ISP network

10.2. Deployments

The largest intended NFV deployment in a network was likely the announced by AT&T in 2014. In fact, AT&T stated in early 2015 that its goal is to virtualize and control over 75 percent of its network by 2020 using their new User-Defined Network Cloud architecture. AT&T has already begun to virtualize and put into production critical network functions such as domain name service (DNS), network analytics, intelligent data platform, and virtualized provider edge router, improving cycle time, elasticity, and operational efficiency.

Since then, other telecom service providers have stood up and taken note, and from a carrier perspective, virtual POPs are becoming a reality. The following is a collection of companies that have begun to embrace virtualization with some information about their deployments.




Figure 9: ATT NFV Case Study

CASE STUDY – BRITISH TELECOM



Figure 10: BT NFV Case Study

INTERCONNECT CASE STUDY: TISCALI



Figure 11: Tiscali NFV Case Study

11. Conclusions

In this paper, we've described the concept of NFV, how it came to be, and the drivers for carriers to change their infrastructure using virtualized network functions. We provided several forward-looking analyst opinions and forecasts detailing what the industry can expect. We walked through a number of benefits to be realized from (and several challenges hindering) a migration toward the virtualization of network functions.

Carriers have much to gain from leveraging private cloud technology to minimize the capital and operating costs of deploying software-centric networking and service functions. The industry as a whole is moving rapidly toward pure software implementations of network functions operating in a virtualized environment. Virtualized functions can be deployed in the cloud and will deliver some really compelling advantages over the traditional deployment model based on proprietary hardware appliances.

The industry is at the dawn of an era of dramatic change. Carriers that successfully implement NFV will find themselves in a far better position than they are in today to compete with over-the-top services and with aggressive new entrants into the network business. Never before has there been such strong pressure to evolve, and to evolve rapidly. NFV is the tool needed to enable that evolution, and therefore the onus is on the telecom service providers to embrace NFV and virtualize their networks.

To assist carriers in that endeavor, this i3Forum Working Group will be producing a second NFV paper, which will present a deeper dive designed to help you “understand the NFV ecosystem.” The next paper will cover:

- Understanding the main components necessary to migrate successfully
 - Hypervisors
 - Virtual components
 - Cloud software
 - Orchestrators
- Multi-vendor / open source implementations
 - How to put all the pieces together
 - Interop issues / concerns / gotchas
- Post mortems on PoCs
 - Issues encountered
 - Lessons learned
- Infrastructure improvements
 - How to support
 - What can be achieved
 - What are still open issues
- Impact on BSS

Look for it to be published in the second half of 2016.