White Paper

Distributed Network Functions Virtualization

An Introduction to D-NFV

Yuri Gittik
Head of Strategic Marketing
RAD
March 2014

Abstract

Network functions virtualization – NFV – has quickly attracted the interest of communications service providers as a means to accelerate service delivery while reducing associated costs. This paper reviews the distributed approach to NFV, discusses phased NFV deployment and introduces critical decision factors for placing such functionalities at the customer edge.

Contents

1 Introduction to NFV

Network functions virtualization, or NFV, replaces dedicated network devices with software running on general-purpose CPUs or virtual machines, operating on standard servers. In the era of NFV, adding a firewall, load-balancer, or router no longer means deploying proprietary, vendor-specific appliances. This next phase in intelligent networking draws from the successful experience in IT virtualization, as well as from technological advancements in server hardware, to bring service providers lucrative benefits:

- Faster time to market and time to revenue for business and consumer services with rapid deployment, upgrade and turnoff of network functionalities and value-added service capabilities
- Optimal placement of new network functionalities to wherever they are most effective or least expensive
- Flexible relocation of network functionalities in order to suit rapidly changing needs
- Lower expenses with commercial off-the-shelf (COTS) computational elements
- Improved economics and simplified operations by combining multiple network functionalities on a single computational platform

A strong indication of the interest NFV is arousing among service providers is the organized effort, initiated by leading carriers to provide a standardized framework for the new architecture and related technologies. This effort is being carried by the NFV Industry Specification Group (ISG) within ETSI, the European Telecommunications Standards Institute. In October 2013, the NFV ISG published its first five documents, proposing a framework to support interoperable NFV solutions. Among the topics addressed in these first specifications are NFV terminology, requirements, architectural framework, and use cases.

The rate and extent of NFV rollouts are difficult to assess at this early stage; however, 2014 is expected to abound with proof-of-concept and pilot trials, with deployments commencing in 2015 and accelerating in 2016. Analyst firm Doyle Research estimates that the market for NFV solutions will reach $5 billion by 2018, covering associated software, servers, and storage[1]. Analyst Lee Doyle referred to the figure as "best-case scenario" that depends on a range of factors, such as the various use cases and their effect on both network infrastructure and deployment time frames.

1 Forecasting the NFV Opportunity by Lee Doyle, *Light Reading* August 2013

In a RAD-sponsored survey at the Ethernet & SDN Expo in October 2013, 67 percent of respondents expressed high confidence in NFV's promised benefits. Forty percent of respondents said they expect to implement NFV capabilities in one to two years, while fully a third said such implementations would be within a year[2].

As the industry prepares itself for NFV implementations, vendors need to resolve a variety of issues. One major example is the fact that, while standard servers already populate some parts of the telecom network, proprietary ASICs or network processing units with embedded software are still widely used, primarily to ensure the required performance and scalability, making virtualization in those parts much more difficult.

Much of the current phase of NFV development focuses on mobile networks. This is only natural as the majority of those networks are in transition, with massive deployment of new infrastructure for LTE/LTE-Advanced. In addition, most of the functionalities in question, such as EPC, IMS and PCRF are already centrally located, so their implementation in data centers does not necessitate major architectural changes. On the other hand, business services for enterprises, which conventionally rely on functionality distributed throughout the network, will require a fresh view. This approach is at the center of the D-NFV architecture.

Doubtless, a key factor affecting the rate of NFV adoption is the service providers' perception of the offered benefits. It remains to be seen which of the aforementioned benefits will be the main drivers.

2 NFV and SDN

When discussing software-based networking, there is often confusion between NFV and SDN (software defined networking), leading to erroneous swapping of the two terms. In fact, NFV and SDN are independent, yet complementary and mutually beneficial technologies. Network functions can be virtualized and deployed without SDN technologies, and non-virtualized functions can be controlled by SDN. The primary distinction between the two has to do with the domain to which they apply: While NFV replaces proprietary hardware network elements (NEs) with software running on standard servers, SDN deals with the replacement of standardized networking protocols with centralized control. SDN promises to reduce the complexity of distributed networking control protocols with the simplicity of programming an omniscient controller, and hence leading to the following advantages:

2 Industry Confidence in Distributed NFV Reflected in Survey at Ethernet & SDN Expo, RAD October 2013

- Rapid development of application software in comparison to protocol standardization and development
- Simplification of network maintenance due to centralized control
- Ability to swiftly deploy, relocate and upgrade new features

For additional reading on SDN and NFV, see the "Carrier Grade Communications" blog by RAD's CTO, Dr. Yaakov Stein.

3 Distributed NFV

As previously mentioned, NFV facilitates concentration of network functions in data centers or other centralized locations. This practice is suitable for many, but not all, scenarios of interest. An alternative approach to NFV advocates distribution of virtualized network functions (VNFs) throughout the network; placing VNFs wherever they may be most effective and least expensive. The distributed NFV approach is emphasized in the recently published "Terminology" and "End-to-End Architecture" documents of the ETSI NFV ISG:

- The "Terminology" paper clearly defines a **Network Point of Presence** as "a location where a Network Function is implemented… Examples of NPOP locations include central offices, **customer premises**, mobile devices, and data centers."
- The "End-to-End Architecture" document states that "one of the NFV objectives is to ensure greater flexibility in assigning VNFs to hardware: Software to be located at the most appropriate places, e.g., at customer premises, at network PoP, in central offices or data centers."

As NFV architectures become more mature and standardized, there is an increasing emphasis on achieving maximum VNF placement flexibility. Distributed NFV promotes carrier-controlled VNFs that may reside anywhere – in data centers, in network nodes and at the customer premises.

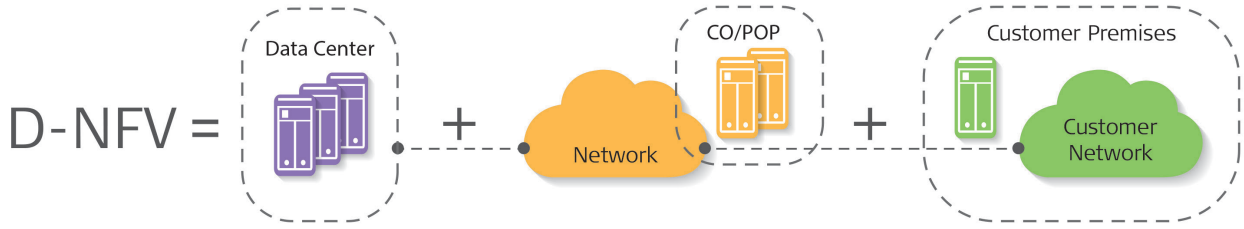

Figure 1: Distributed network functions virtualization

4 Implementing NFV at the Customer Edge

4.1 Why Some VNFs Belong at the Customer Edge

Fifty percent of the respondents in the survey referenced in Chapter 1, when asked where they think NFV will have its biggest impact, said it would be either in the data center, network edge, or customer premises. This reflects awareness to the advantages of locating NFV at various parts of the network, wherever it makes the most sense for the case at hand.

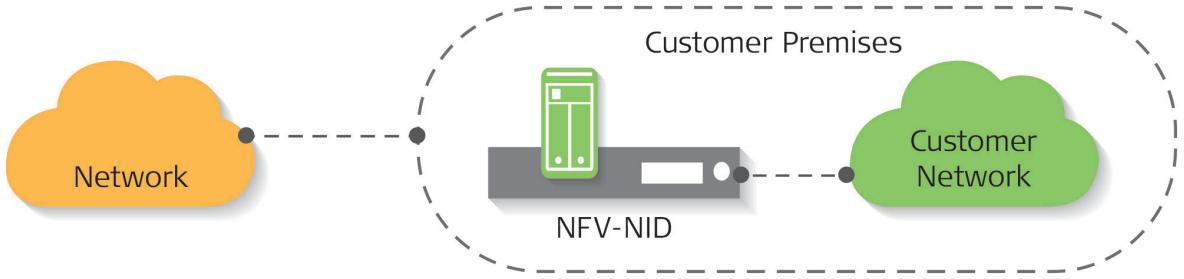

Figure 2: Distributed NFV at the customer edge

There are clear advantages for a service provider to locate certain functionalities, whether or not virtualized, at the customer premises. The reasoning behind this placement ranges from feasibility (i.e., the very possibility to implement the virtualized function), through performance to economics. For instance, diagnostic tools, IP-PBXs, traffic accelerators, NATs, and rate-limiters, are typically most effective when located as close as possible to the end-user. In fact, the customer premises are often the most suitable location for functions of traffic monitoring, QoS and security:

Loopbacks must be implemented at the customer premises as the very essence of loopback diagnostic testing is to enable verification of connectivity to that site.

End-to-End QoS/QoE Monitoring can be handled at the network edge, but may provide more accurate measurements if implemented at the customer premises.

End-to-End Security includes encryption, which most commonly necessitates implementation at the customer site. Moving encryption to the other end of an access network would leave unencrypted traffic exposed to all with access to that network. Similarly, blocking malicious traffic at the other end of an access network leaves the door open to attackers with access to that network.

Firewall functionality placement depends on the customer's corporate policy. The application of firewall rules from a remote location might create potential risks. Some customers may only consider a local firewall service due to required performance or corporate data access policy. In addition, the performance of a centrally located firewall may severely degrade (even to the point of loss of functionality) due to network constraints, such as packet delay and loss, or low connection availability.

IP-PBX centralization is frequently not acceptable, as local voice connectivity is required notwithstanding loss of WAN connectivity.

WAN Optimization of the access segment should be implemented at the customer premises, in order to maximize utilization of the access link.

The table below summarizes the benefits and considerations of network function virtualization at the customer site:

Feasibility	• Some functions must be implemented at the customer site, e.g., loopback, end-to-end security, traffic conditioning, encryption, WAN optimization
Performance	• Some functions perform better at the customer site: End-to-end QoS, QoE application monitoring • Some functions may degrade due to network constraints such as bandwidth, delay and availability, if not locally implemented at the customer site • Faster service delivery with VNF-chaining may better suit customer needs
Cost	• Higher performance and resiliency requirements may lead to a cost increase, offsetting data center economies of scale • Invest-as-you-grow flexibility becomes possible when implementing new network functions with a "Try & Buy" approach, involving limited-scale service pilots that are extended if proven successful • Cost may be cut by integrated multiple functionalities and applications per device
Policy	• Some functions need to remain close to the customer due to corporate privacy, security and access authorization policies • Regulatory restrictions (e.g., restrictions against moving data across jurisdictions) may also apply

A growing number of industry professionals share this view on distributed NFV. Tom Nolle of CIMI Corp wrote "...it makes sense to think about whether an edge device might not have some local hosting capability, and might thus serve as a platform for part of NFV deployments..."[3] Carol Wilson of Light Reading agrees: "In the race to virtualize network functions and centralize control of software-defined networks, there is good reason to consider the value of keeping some intelligence in the most distributed locations, namely the customer premises."[4]

4.2 The Economics of Customer-Premises NFV

Until recently, service providers have been overly occupied with attempts to lower their operational expenses. These days, increasing competition from OTT services, among others, their focus has widened to include the definition of new service capabilities that could lead to increase in service value and revenue-streams. In the past, the introduction of such value-added offerings was often hindered by high cost; NFV, especially when located at the customer premises, has the potential to fundamentally change this predicament. D-NFV enables the download of the desired functionality to an integrated Layer 2 and Layer 3 NID (network interface device) without truck rolls, delivery of new equipment, on-site installation, or even increase maintenance and energy costs.

The prevailing approach of centrally locating these new service functionalities would arguably provide savings in IT resources. This is due to their ability to exploit economies of scale: Real-estate expenditures are minimized by colocating as many functions as possible; cooling costs may be minimized by running functions in riverside data centers with water cooling; and placing multiple instances of a function in the same data center can vastly simplify management and maintenance. However, placing functionalities in a central location may lead to increased networking costs, the result of higher bandwidth and more stringent availability requirements, as well as the need to implement mechanisms to combat increased delay.

On the other hand, a customer-premises distributed NFV approach gives service providers the ultimate flexibility to allocate virtualization resources in the most economical way, factoring in all direct and indirect elements affecting the bottom line – costs of both IT and networking resources, potential for additional revenue and operational efficiency, etc.

3 Is Distributed NFV Teaching Us Something?, CIMI Corp September 2013

4 ESDN: RAD Rolls Out Distributed NFV Strategy, *Light Reading* October 2013

Furthermore, NFV deployment at customer sites makes new service introduction easier, by enabling a "Try & Buy" approach. Enterprise customers can easily perform D-NFV-based pilot tests of proposed services, without the complexity and cost involved in deployment of new devices. Upon successful completion of the pilot, the customer may choose to subscribe to the new service.

5 D-NFV Deployment Routes

D-NFV at the customer premises allows effective implementation via phased deployment, as illustrated below:

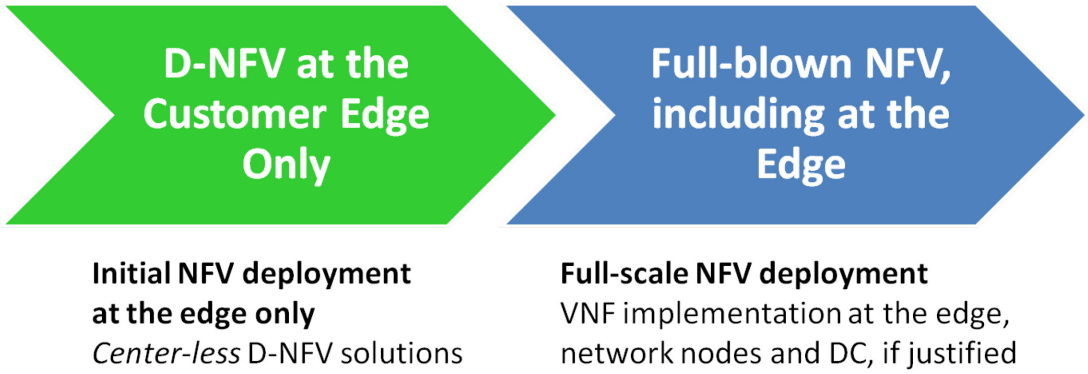

Figure 3: Phased D-NFV deployment

The initial phase involves the implementation of VNFs only at the customer edge using a NID with an integrated NFVI (NFV infrastructure) component, which can be referred to as center-less D-NFV implementation.

This allows service providers to reap the benefits of virtualization – namely, service agility and operations cost management – from day one. Furthermore, it minimizes upfront investment and mitigates the risks involved in new services introduction, as it matches the pace at which costs are incurred with that of revenue generation. In addition, it reduces the risk of early, pre-standard implementation of VNF solutions.

Another advantage offered by this approach is that it does not require (although can benefit from) sophisticated orchestration mechanisms, which are currently in a rather early stage of development in the MANO (Management and Orchestration) Working Group of the ETSI NFV ISG. This is due to the single-tenant nature of the platform. Even if several VNFs need to be chained on the platform, this can be easily handled without sophisticated orchestration. With the service and operation experience gained in the first phase of **center-less D-NFV deployment**, service providers can proceed to a full-scale D-NFV deployment adding VNFs at data centers and/or at the network edge.

Such *invest-as-you-grow* flexibility enables service providers to adopt an incremental, bottom-up approach, based on "service incubation" at a limited number of customer sites. Such an approach does not require network route modification and should not affect SLA guarantees. It may also be more economically viable, as it enables results-driven investment.

6 RAD's D-NFV Solution

RAD addresses the D-NFV approach with a novel solution that integrates an x86 server for VNF hosting with a L2/L3 NID, in a customer-located device that is controlled by the service provider. Such a NID combines all the "smart" demarcation functionality of a state-of-the-art NID (including traffic management, full end-to-end service control and monitoring, service turn-up, and diagnostic tools) with a built-in server card platform as the standard NFVI component.

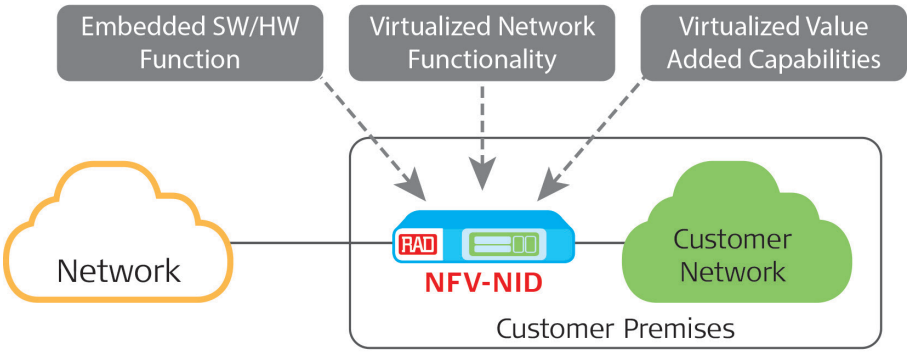

Figure 4: RAD's D-NFV solution at the customer premises

This solution joins the control of conventional network elements, VNFs and IT applications, optionally using SDN principles. The server platform enables hosting value-added service capabilities and network functionalities either directly on the x86 processor, or on virtual machines (VMs) running on the processor.

With the x86 platform in a customer premises NID, service providers can quickly and easily roll out new services and implement network capabilities without deploying new equipment. Using a single device to deploy and manage multiple functionalities and applications further reduces costs.

RAD's server platform enables hosting of any VNF that is engineered for a standard x86 CPU platform. In addition, RAD provides standard tools to facilitate this hosting (industry standard hypervisor, such as KVM, OpenStack environment for application control). As a result, RAD's solution is open to 3rd-party applications. RAD encourages application (VNF) vendors to join its ecosystem to provide a broad range of network functionalities and value-added service capabilities.

7 Summary and Outlook

The speed at which NFV enters service provider networks will be determined in large measure by the pace at which service providers can strengthen the relationship between their network and IT operations. This is true whether NFV is based on functionalities deployed in the data center or distributed out to the customer premises and other locations. To facilitate a broad implementation of NFV, large carriers will need to change their organizational structure and combine the expertise of their network, cloud, and IT teams. The situation is reminiscent of the way *Bell-heads* and *Net-heads* managed to reconcile their differences concerning legacy circuit switching and the then new packet-switched communications. Accordingly, the present revolution involves bridging a new – and even larger – gap, between *Net-heads* and *IT-heads*.

A major challenge for service providers is the need to ensure proper orchestration of the virtualized capabilities and functionalities while harmonizing IT/cloud and network resources. Gradual D-NFV deployment beginning with an initial center-less implementation minimizes the risk by mitigating complexity and reducing upfront investment.

Another challenge will undoubtedly be the assurance of service level agreements (SLAs) in a virtualized environment that involves service chaining and multiple network segments. That is why early adoption of NFV is critical; it enables earlier realization by service providers of business benefits, surmounting of operational challenges, and accelerating adoption of novel virtualization technologies in general.

D-NFV in the access network enables service providers to create much needed value for their enterprise customers. It does so by combining smart demarcation with extended service capabilities, enabled by RAD's future-proof server platform. This platform comes ready to be enhanced with new software-based functionalities for fast service creation and on-the-fly deployment of new applications.

www.rad.com

International Headquarters
RAD Data Communications Ltd.
24 Raoul Wallenberg St.
Tel Aviv 6971923 Israel
Tel: 972-3-6458181
Fax: 972-3-6498250
E-mail: market@rad.com
http://www.rad.com

North America Headquarters
RAD Data Communications Inc.
900 Corporate Drive
Mahwah, NJ 07430 USA
Tel: (201) 529-1100
Toll free: 1-800-444-7234
Fax: (201) 529-5777
E-mail: market@radusa.com
www.radusa.com