White Paper

Best Practices for vCPE Deployment

May 2015

RAD

Your Network's Edge

Abstract

The virtual CPE is a prime candidate for initial commercial deployment of NFV, especially for business services. For service providers, vCPE's sweet spot lies in hardware abstraction and the ability to carry out shorter and more flexible deployment cycles for new services.

This paper presents vCPE implementation options available for service providers. It also reviews the various factors that should be considered to avoid pitfalls and ensure optimal bandwidth efficiency, security, survivability, performance, diagnostics, and QoE.

Contents

The State of NFV

The networks of tomorrow need to be agile, efficient and well-orchestrated to address the service providers' challenges of automated and speedy service delivery, along with a lucrative cost structure. The common perception of the programmable network as a "must" for communications service providers drives network functions virtualization (NFV) and software-defined networking (SDN) from theory to practice at an accelerated pace. This process involves the migration of networking functions from vendor-specific, proprietary hardware appliances to software hosted on standard compute infrastructure (NFV), as well as the replacement of distributed protocols with centralized, dynamic and programmatic management of network elements (SDN). The distributed approach to NFV (D-NFV) places virtualized networking functions (VNFs) wherever it makes the most functional and economic sense, be that at the data center, POP or the **customer edge**.

Distributed NFV

Virtualized network functions (VNFs) should be located where they will be most effective and least expensive, be that the data center, network nodes or the customer edge. Key criteria for VNF placement include:

In a little over two years, NFV has gone from a skeletal framework outlined in an introductory white paper to a collaborative endeavor in which service providers and vendors engage in clearly defined proofs of concept (PoCs) of productized solutions. Such activity is accompanied by, and contributes to, a flurry of standardization efforts by such groups as the European Telecommunications Standards Institute (ETSI)'s NFV Industry Specification Group (ISG), the Broadband Forum and the Metro Ethernet Forum (MEF). The latter recently launched its Third Network initiative to retain the best features of Carrier Ethernet 2.0 and routed IP-based networks, while using SDN and NFV to provide lifecycle service orchestration.

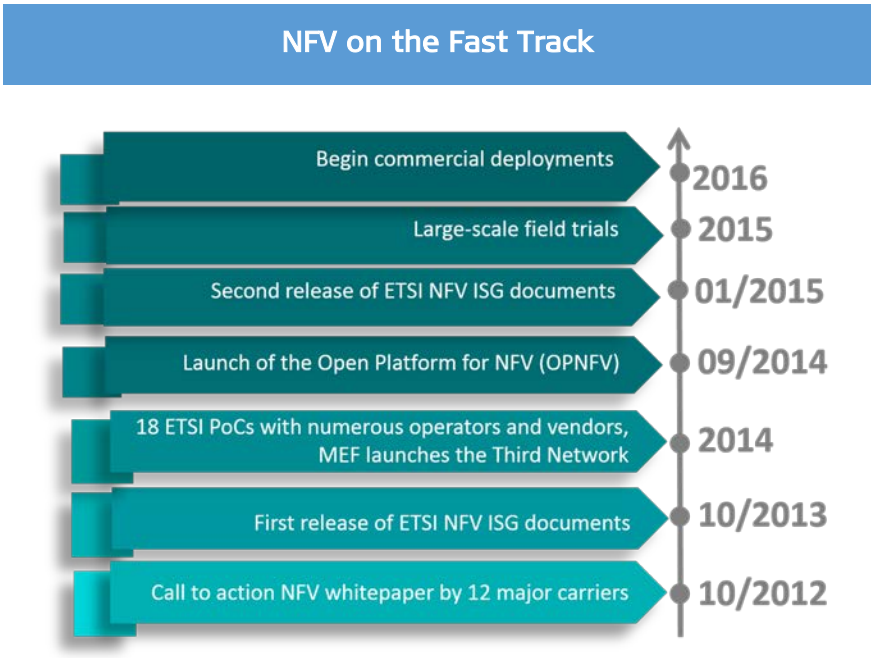

Demonstrating a speedy evolution and an overwhelming market acceptance, NFV is expected to enter a phase of extensive field trials in 2015, with initial commercial deployments anticipated from 2016 onwards.

Infonetics Research estimates the carrier SDN and NFV market to hit $11 billion by 2018, with NFV representing the lion's share of the business[1]. According to a survey it conducted among leading service providers in April 2014, 93% of respondents plan to deploy NFV in some aspect of their network[2]. The survey also listed **the virtual CPE (vCPE) for business services as the top rated use case – both for NFV deployment in general and for revenue generation**.

What is vCPE?

The virtualized CPE is a new approach to customer premises equipment whereby at least some of the networking functionality associated with conventional customer premises equipment is virtualized and, perhaps, relocated to other network locations. Note that while virtualization facilitates relocating some functionalities from the customer premises to data centers, computational power in the vCPE enables relocating other functionalities from deep in the network to the customer premises. The business vCPE (also known as the vE-CPE, for virtual enterprise customer premises equipment), is a virtualized networking appliance at the customer edge that delivers communication services to enterprises. What once has been a collection of single-purpose, hardware-based devices at each customer location (e.g., a router, load balancer, firewall, etc.) has been transformed, in

1 Carrier SDN and NFV Hardware and Software Market Size and Forecast Report, Infonectics Research, November 2014

2 SDN and NFV Strategies: Global Service Provider Survey, Infonectics Research, April 2014

the vCPE model, into virtualized appliances that can be dynamically added or dropped as needed. Physical and virtualized vCPE functionalities are divided between the customer site and the network (at the provider edge – PE – or the local/central data center) to ensure maximum flexibility and performance. The network-located functionality can also be shared among multiple users, following a multi-tenant unit model.

The ETSI NFV ISG provided a general definition for vE-CPE in an early paper listing NFV's use cases[3], while the Broadband Forum (BBF) is currently in the process of further defining a virtual business gateway (WT-328 vBG). A parallel BBF track is focusing on a network enhanced residential gateway (WT-317 NERG). Both organizations are collaborating to streamline standardization efforts and promote interoperability.

Market Drivers

It's no wonder that the vCPE is seen as the ideal candidate to test the waters of NFV and the programmable network concept. Enterprise CPEs are both CapEx- and OpEx-intensive because of their huge quantities and the complexity involved in their deployment and maintenance. These costs have a significant impact on network operations and severely limit a service provider's ability to roll out new services quickly, or make timely service modifications and upgrades. As a result, carriers are less successful in thwarting competition, shortening their time to revenue and maintaining customer satisfaction. Conventional appliance-based CPEs entail slow, expensive deployment processes, which are no longer acceptable in today's market. The NFV-enabled vCPE is, therefore, the paramount use case for service providers, promising increased revenues and lower TCO (total cost of ownership).

Increase Revenues

- Fast introduction of new services without changing customer-located equipment
- Fast response to competition
- Offer value-added services at a premium
- Improve quality of experience to reduce churn

Reduce TCO

- Simplify and reduce cost of operational processes, management and truck rolls
- Replace multiple appliances with a single device, reducing power and space requirements and improving reliability
- Migration to new cost-effective technologies

3 ETSI GS NFV 001 – Network Functions Virtualizations (NFV) Use Cases, ETSI, October 2013

Implementation Options and Factors

The vCPE architecture combines physical and virtualized entities located at the customer premises and elsewhere in the network (at the data center or local aggregation/PoP). Service providers and network operators contemplating the proper implementation architecture should consider these issues:

Minimal required functionality at the service handoff: Although virtualized CPE functions may run on virtual machines (VMs) in the cloud, some basic data forwarding and service demarcation/termination capabilities must still be located at the customer site. This means that the customer-located physical CPE (pCPE) must include at least a simple switch with essential forwarding functionality. Some scenarios call for more advanced capabilities, as described below.

To virtualize or not to virtualize: There are network functions that *may* be virtualized and relocated to the cloud, but *should* remain embedded in the CPE. Typical examples, particularly for high access rates, are data-plane functionalities, such as packet forwarding, traffic queuing and prioritization. This will affect the pCPE architecture – and its built-in hardware-based capabilities – to be deployed at the customer edge. Different physical and virtualized functions can be **service-chained** regardless of their location; however, there are often speed and performance benefits when they are performed within the same location/device, as explained below. Such function implementations can be **hybrid,** using both physical and virtualized resources. In the hybrid model, an application awareness functionality, for example, would use a DPI engine as a virtualized control plane located in the network, and hardware-based forwarding and flow marking functionalities to ensure wire-speed operation at the customer site. A similar hybrid implementation could fit routing functionality.

Where should VNFs reside: There are various views on the placement of virtualized vCPE functionalities; these can be generally narrowed down to three scenarios:

Scenario A: all virtualized functionality located in the network

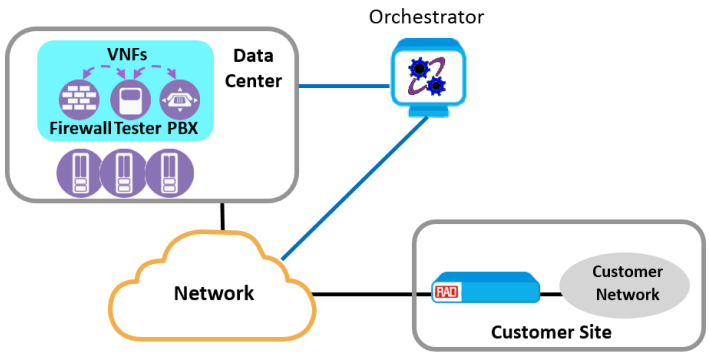

This is a straightforward scenario, using only a basic switch/router as the physical device (pCPE) at the customer premises, while all virtualized functions reside at the data center. Service providers can choose between deployment of a vendor-specific, integrated vCPE software package featuring a pre-selected set of related service applications, and service chaining of individually-sourced network functions addressing their needs. The most likely use of this approach would be in the SME services market, in which speeds and performance requirements can be fully supported by cloud vCPE implementation.

Scenario B: all virtualized functionality located at the customer edge

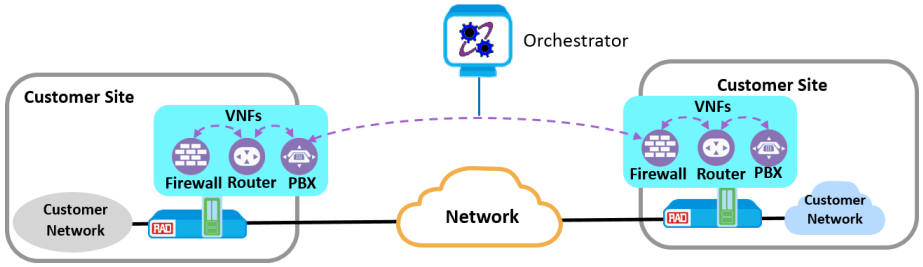

In this scenario, both types of vCPE functionalities – physical and virtualized – are placed at customer end-sites, while no VNFs are deployed in a central network location. An advantageous implementation is a network interface device (NID) with an integrated compute platform, such as an x86 server. This platform acts as the NFV infrastructure (NFVI) on which VNFs run. Alternatively, a standalone server can be colocated with the NID; however, this provides a less effective solution, as it doesn't facilitate per-application traffic handling nor can it support traffic offload for hardware-based processing.

Scenario C: virtualized functionality distributed between the network and customer sites

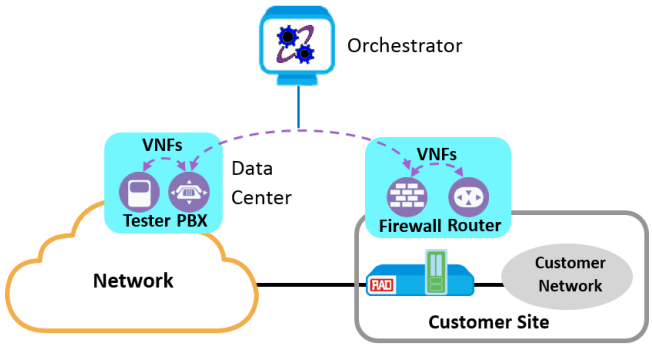

Here, virtualized functions are placed wherever their performance, cost and policy compliance are optimal – either at the customer edge or at the network. In this scenario, VNFs residing in different locations can be dynamically ordered, configured and chained to meet customer business needs.

Both scenarios B and C are based on the D-NFV model. They are particularly suited for value-added services characterized by high networking costs and stringent performance requirements (e.g., international business services). Placing at least some virtualized functionality at the customer location is a reasonable solution to avoid the bandwidth inefficiency and application performance degradation issues described in the following section.

Scenario B's *center-less* implementation may also fit early NFV deployment stages, in which the service provider wishes to run limited trials or service pilots without heavy investment in full-scale DC upgrade or network redesign.

Avoiding Pitfalls

Function relocation may have unintended effects on service quality and other aspects. Several factors that service providers need to consider include:

- **Bandwidth efficiency:** What is the bandwidth "cost" of moving functionality deeper to the network? Excess bandwidth expenditures could have a critical effect on service delivery in areas still served by relatively low-speed connections such as DSL.
- **Security:** Does moving the VNF to the network expose sensitive end-user data? For example, an encryption application located anywhere but at the customer premises doesn't provide adequate protection, as traffic interception can occur en route, in an unsecure access segment.
- **Survivability:** Critical functions must remain operative, even when the access link is down. Hosting IP PBX or router functions at the PE or data center, would result in an inability to locally place calls or deliver traffic in case of network failure.
- **Application performance:** Is the delay added by the network for DC-based functions acceptable? This is a critical factor when engineering delay-sensitive workflows, for example, for financial applications such as algo-trading. In some cases, function relocation may result in a severe drop in performance due to inadequate access link bandwidth or excessive delay. A firewall, for example, might be affected by frequent session time-outs due to packet loss and reordering – the likelihood of which rises considerably when the traffic needs to travel all the way to the data center.
- **Diagnostics and QoE:** Testing and troubleshooting applications need to accurately measure link and end-to-end service quality, as well as localize faults, starting from service handoff. When such a function resides at the data center, it cannot reliably distinguish between performance issues arising from faults at the access link, traffic handling lapses and user traffic impairment. As a result, identifying and fixing the problem becomes a lengthy and expensive process, if at all feasible – adversely affecting QoE.

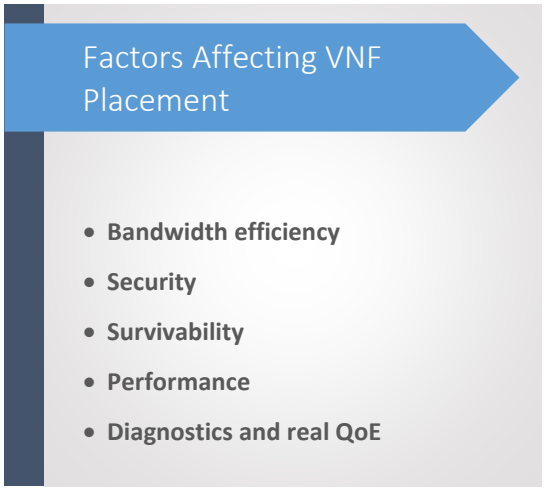

It is conceivable that, in the future, a network orchestrator could dynamically relocate functions to provide optimal QoE based on various factors, such as changing network loads. At the moment, however, this decision is part of the pre-launch NFV planning stage; it should factor in all the above to allow service providers maximum flexibility and agility while ensuring effective performance.

Future Outlook

Over the next two-to-three years, we are likely to witness a multitude of vCPE options being streamlined to fit a variety of practical use cases and address real-life challenges – both at the service and operations levels. In addition to VNF placement and end-point device capabilities, much of the industry's attention will be focused on management and orchestration. There is a clear trend towards programmability and automation in the network, making the case for the use of SDN principles when implementing vCPE solutions.

The vCPE implementation options currently being discussed typically rely on manual provisioning of VNFs alongside various proprietary virtual networking mechanisms. The next evolutionary phase will address automation of connectivity and optimization of VNF selection and placement. This evolution does not necessarily require the deployment of SDN switches at the customer premises as long as existing equipment can be configured by the new orchestration platform.

Conclusion

The virtual CPE is a prime candidate for initial commercial deployment of NFV, especially for business services. For service providers, vCPE's sweet spot lies in hardware abstraction and the ability to carry out shorter and more flexible deployment cycles for new services. When implementing the vCPE architecture, there are several options for VNF placement – in the data center, at the customer edge or a combination of both – each fitting a different scenario. When planning vCPE deployments, service providers need to consider how functionality placement affects bandwidth efficiency, security, survivability, performance, diagnostics, and QoE.

As the programmable network becomes a reality, much of the industry's focus is turning to automation and control. Over time, vCPEs are expected to transform from loosely coupled to integrated entities, with management functionalities increasingly becoming part of a dynamic control plane.

RAD's vCPE Offering

RAD's award-winning D-NFV-enabled networking devices support any vCPE implementation scenario, ensuring maximum flexibility to meet diverse needs. Part of RAD's Service Assured Access (SAA) solutions for business, mobile and wholesale service providers, these devices are designed to improve the way service providers compete. They enable service agility to minimize time to revenue, complete visibility of network performance for greater operational efficiency and better QoE to reduce churn and lower TCO.

RAD's vCPE offering includes:

- ETX-2 L2/L3 NID with a powerful x86 platform for hosting virtual network functions and applications at the customer edge
- MiNID® miniature programmable L2/L3 NID for downloading demarcation and networking applications

- RADview D-NFV orchestrator for management of virtual machines and application services in ETX-2 and MiNID
- Special VFs for central locations:
  - Complementary virtualized functionality for service demarcation and traffic monitoring, such as TWAMP for IP performance monitoring
  - VFs for VAS, such as application awareness, to enrich available vCPE packages with advanced value-added capabilities

ETX-2

L2/L3 Demarcation with D-NFV

MiNID

Miniature Programmable NID

RADview D-NFV Orchestrator

Management of Virtual Machines and Application Services

For more information about RAD's vCPE implementation options, contact us at market@rad.com.

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

Your Network's Edge