

# Distributed NFV: Ensuring that the Benefits of Virtualization Exceed the Costs

by Dr. Yuri Gittik, Head of Strategic Developments and Innovation, RAD



*Yuri Gittik is Head of Strategic Developments and Innovation at RAD. In this capacity, he is responsible for steering the innovation and leading strategic development of new RAD solutions and partnerships, with a particular focus on NFV and cloud solutions. Dr. Gittik also leads RAD's strategic cooperation with key carrier accounts around the world, as well as R&D collaboration.*

*Since coming to RAD in 1991, Dr. Gittik has held various positions in product management, product marketing and sales. In 1999 he established a new Business Development Department, which he headed until 2005, and was instrumental in leading the company in the emerging Ethernet access and mobile backhauling markets and fostering cooperation with strategic accounts and partners. From 2005 to 2012 he served as RAD's Chief Strategy Officer.*

*Prior to joining RAD, he held a variety of positions with a number of consulting firms in the fields of research and analysis.*

*Dr. Gittik holds an M.S. in Systems Engineering and a Ph.D. in Economics.*

Over the last few decades the telecommunications industry has evolved from one technology to the next at a reasonably fast pace, but nothing seems to have ever captured our collective imagination as quickly as Network Functions Virtualization (NFV).

While Software Defined Networking (SDN) has been discussed in the industry –in different approaches and flavors - for several years already, the first anyone ever heard of the term NFV was just two years ago. In October 2012 twelve service providers published what has now become a seminal white paper on NFV. In January 2013 the ETSI Network Functions Virtualization industry specifications group (ISG) was formed to steer NFV framework development.

## Riding the Virtualization Wave

To a large extent, NFV's meteoric rise is all very logical because, untypically for the telecom industry, it has come from the bottom

up, as technologies and concepts from the IT domain have now worked their way into the network domain.

In the field of IT and computation, there has been a major trend towards virtualization, by which we mean the creation of a virtual machine (VM) that acts like an independent physical computer (or other hardware device). But computers are not the only hardware devices that can be virtualized. Functionality of many (but not all) network elements can be performed by software running on a VM or a central processing unit (CPU).

Meanwhile, service providers have been grappling with a series of complicated challenges: the proliferation of bandwidth-consuming applications, growing demand for differentiated services, increased service coverage that dictates inter-carrier connectivity, LTE and LTE-A deployments with partial or full mesh connectivity, exacting customer expectations for the SLA performance they are paying for, and lower ARPU. All of this means that providers have

to find a way to offer a greater number of value-added services and demonstrate agility in addressing competition and customer needs while reducing their TCO.

Accomplishing all this is no easy task, especially financially, since the addition of new services has inevitably required the deployment of new dedicated hardware. Apart from the resultant explosion in capital expenses, there is the ongoing responsibility – and additional cost – of maintenance, which includes training and salaries for technicians, the inevitable upgrades, and electricity, not to mention rack space in cramped facilities. As CapEx and OpEx increased, competition continued to drive profit margins down, which were compromised even further by “Over-the-Top” (OTT) services. Something had to be done to break this vicious cycle, and NFV was conceived to do just that.

For want of a better analogy, NFV was designed to do for service providers what your smartphone does for you. Even though it's a single device, your smartphone runs

**“With IT seeping into networking, traditional hardware vendors are placing an increasing emphasis on software. For their part, operators will need to improve cooperation and integration of their traditionally separate IT and networking teams. Intra-industry cooperation will also reflect these changes. An NFV hardware ecosystem already has been formed by Intel, while the RAD D-NFV Alliance has been initiated for application vendors, which see D-NFV as the ideal vehicle for enabling their applications to run at the edge of the network.”**

multiple software-based applications. No less importantly, such applications are provided by a huge group of developers and are easily available (and in many cases, for free).

Similarly, the idea behind NFV was to do away with multiple networking devices and allow rapid and flexible deployment of network functionality and value-added service capabilities as software on standard compute hardware. After all, switches, routers, firewalls, among other network elements, are all good candidates for virtualization (as long as data rates are not excessive).

#### Analyst Predictions

Industry analysts disagree about the precise numbers, but while the NFV market may be embryonic today, all foresee a promising future. Mind Commerce estimates the NFV market will grow at a CAGR of 46% through 2019 to reach revenues of \$1.3 billion, up from \$203 million this year. TechNavio says CAGR will grow even faster, by 51.57% through 2018. Analysis Mason predicts revenue of \$2.4 billion in 2018. The “best-case NFV market scenario,” according to online publication Light Reading, was prepared by Doyle Research, which predicts that “the NFV market will start to gain momentum” this year and will “reach approximately \$5 billion by 2018, including software, servers, and storage.”

Meanwhile, for network architects and planners, the most important issue today is how to best accommodate the ever increasing number of demands they face. Or, to put it another way, how can they accommodate all those demands in a non-destructive, evolutionary and, most importantly, cost-justified manner? After all, in the current market, all new technologies must first and foremost prove not just their technical virtues but that their economic benefits exceed CapEx and OpEx. What matters most now is the business case, fiscal responsibility.

#### Bringing NFV to the Access

As is the (so far) prevalent case with SDN, the typical perception of NFV is that it

has to be located at the data center. But that’s not necessarily the only case. NFV can indeed be implemented in the access. This is accomplished using an approach called Distributed NFV (D-NFV), which was already envisioned in the introductory NFV White Paper and has been clearly emphasized in several already completed documents of the ETSI NFV ISG. What is important, however, is that D-NFV has now gone way beyond the discussion stage and is now being tested by several operators worldwide. Approved by the ETSI ISG and sponsored by CenturyLink, a “Multi-vendor Distributed NFV” proof of concept (PoC) was successfully demonstrated at the Big Telecom Event this summer in Chicago. Field trials and early deployment are expected in 2015 and mass deployments are to start a year later.

This approach is widely accepted in APAC as well, and leading carriers in the area are expressing interest in such innovative solutions. For instance, NTT successfully completed another ETSI-approved PoC, which paralleled the PoC mentioned above and addressed a vCPE solution with service chaining.

So what are the benefits of bringing NFV into the access? It turns out there are many.

There is no doubt that a good argument can be made for running new service functionalities at a central location, such as a data center. Admittedly, this would save on IT resource costs by leveraging economies of scale. However, placing functionalities at a central location actually increases networking resource costs as the result of stringent connection delays or availability requirements and increased bandwidth needs.

But we’re actually getting ahead of ourselves. Certain functions, even if virtualized and implemented as software, can only be conducted in the access or customer edge, or will better perform there. Just to give a few indicative examples, these may include:

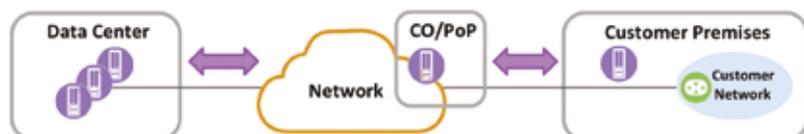
- End-to-end QoS/QoE monitoring – if not implemented at the customer site such measurements will not be sufficiently accurate
- End-to-end security – can only be implemented at the customer site if required to protect all connections up to the site
- Firewall – depending on customer policy, applying firewall rules outside the customer premises may be not acceptable or create potential risks; some customers may only consider local firewall service; another issue

## NFV: The Distributed Model



- The most prevalent approach to NFV concentrates functions in a centrally-located data center (DC) or network nodes (CO/POP)

**Distributed NFV allows functions controlled by the service provider to reside anywhere - including at the customer premises**

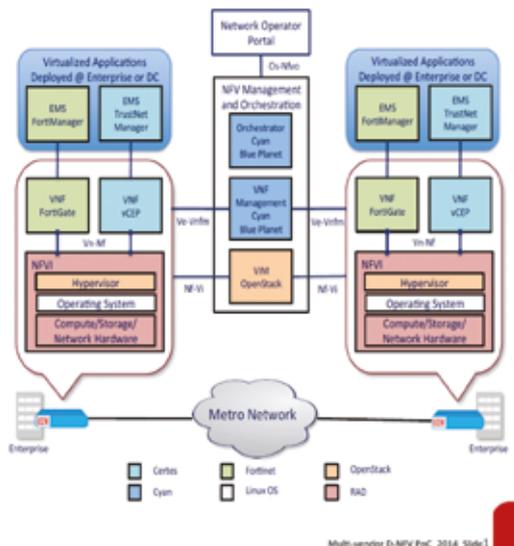


## Multi-vendor D-NFV PoC



ETSI NFV ISG-approved PoC public demo, June 2014

- Carrier sponsor: CenturyLink
- Vendor participants:
  - RAD, Cyan, Fortinet, Certes



Multi-vendor D-NFV PoC\_2014\_Slide1

could be the performance of a firewall placed too far away in the cloud

- IP-PBX – there are cases where centralization of this function is not accepted or partial only, as local voice connectivity would be required even in case of a lost WAN connection
- WAN optimization – by nature, optimization for the access segment should be implemented at the customer premises

The bottom line is that D-NFV allows virtualized functions to be located wherever they are most beneficial and least expensive. It provides true freedom in allocating virtualization resources in a most economical way, taking into consideration IT and networking resource costs, additional revenue and operations efficiency. The implications for service providers and business customers alike are obvious. Both would be best served with NFV located at all possible locations, from the data center to the network node to the customer premises, and orchestrated together as a joint pool of resources, dynamically tuned per network changes and service needs. All these considerations are of increasing importance for the priority NFV use case –virtual CPE (vCPE) for business customers.

So one final question has to be asked and answered: Do the monetary benefits exceed the costs? Can service providers rest assured that D-NFV will make the introduction of new value-added services justified economically?

Yes, they can, in particular with staged NFV deployment.

### Invest As-You-Go Flexibility

As a first phase, service providers can start deploying NFV just with center-less D-NFV approach by implementing virtualization only at the customer edge. Then, with the service and operation experience gained in this first phase, they can proceed to full-scale NFV by adding VNFs at data centers and/or the network edge, if justified. Such invest-as-you-grow flexibility enables them to adopt a phased bottom-up approach based on “service incubation” at a limited number of customer sites with incremental (and minimal upfront) investment. This would not require network route modification and should not affect service SLA guarantees. It may also be more economically viable, as it enables results-driven investment.

Moreover, new functionality can be downloaded to customer edge devices without the need to install new equipment or undertake expensive truck-rolls. This translates into far fewer on-site installations, lower maintenance and energy expenditures and better MTBF.

### Conclusion

With IT seeping into networking, traditional hardware vendors are placing an increasing emphasis on software. For their part, operators will need to improve cooperation and integration of their traditionally separate IT and networking teams. Intra-industry

cooperation will also reflect these changes. An NFV hardware ecosystem already has been formed by Intel, while the RAD D-NFV Alliance has been initiated for application vendors, which see D-NFV as the ideal vehicle for enabling their applications to run at the edge of the network.

It may only be two years since NFV first burst onto the scene, but, thanks to D-NFV, it is now poised to make a significant contribution to the virtualized future of networking in the access. ●

**Connect-World now on Facebook & Twitter**

Connect-World, the world's foremost discussion forum for leaders in the ICT industry, is now available on Facebook and Twitter.

The world's top ICT decision makers express their opinions in *Connect-World*. They use clear, non-technical, English to discuss how ICT helps shape regional and global development. The articles essentially examine the influence that ICT products and services have on the way people live and do business. With separate editions for each of the world's regions, the reports highlight the most important ICT trends and issues influencing socio-economic growth.

Connect-World is now available to follow on Twitter (<http://twitter.com/#!/ConnectWorldICT>) and Facebook <http://www.facebook.com/connectworld.ict>

Also, it is still possible, for FREE, to directly access all past and present *Connect-World* articles, ICT Industry press releases, eLetters, ICT News and more at [www.connect-world.com](http://www.connect-world.com).