WHY DISTRIBUTION MATTERS IN NFV

STRATEGIC WHITE PAPER | NFV INSIGHTS SERIES

Classical telecommunications networks are highly distributed. At the time they were created, bandwidth limitations together with economics left no other choice. Today bandwidth is much more available, and market competition is driving the demand for a highly automated, responsive and open infrastructure. With cloud-related technologies, such as Network Functions Virtualization (NFV) and Software Defined Networking (SDN), we are able to rethink network architectures. The IT industry, for instance, has moved towards highly centralized clouds; will this also be true for carrier networks with NFV? In this paper we examine the role of distribution in NFV networks by assessing performance, network offload, availability, and security requirements of carrier applications. We also address the requirements this places on the NFV ecosystem vendors.

About the NFV Insights Series

NFV represents a major shift in the telecommunications and networking industry. NFV applies virtualization and cloud principles to the telecommunications domain, something that appeared to be impossible until recently due to the stringent performance, availability, reliability, and security requirements in communication networks. Many service providers are now keen to implement NFV to help them gain an advantage through automation and responsiveness to deliver an enhanced customer experience and reduce operational costs. This series of whitepapers addresses some of the key technical and business challenges on the road to NFV.

Telefonica

Alcatel·Lucent

TABLE OF CONTENTS

1. TODAY'S TELECOMMUNICATION NETWORKS ARE DISTRIBUTED

Historically, service providers are coming from an age of massive distribution. In the United States of America, for instance, there are roughly 20,000 central office buildings, each of them hosting one or more telephone switches (Figure 1).

Figure 1. Map of the Central Office locations in the US

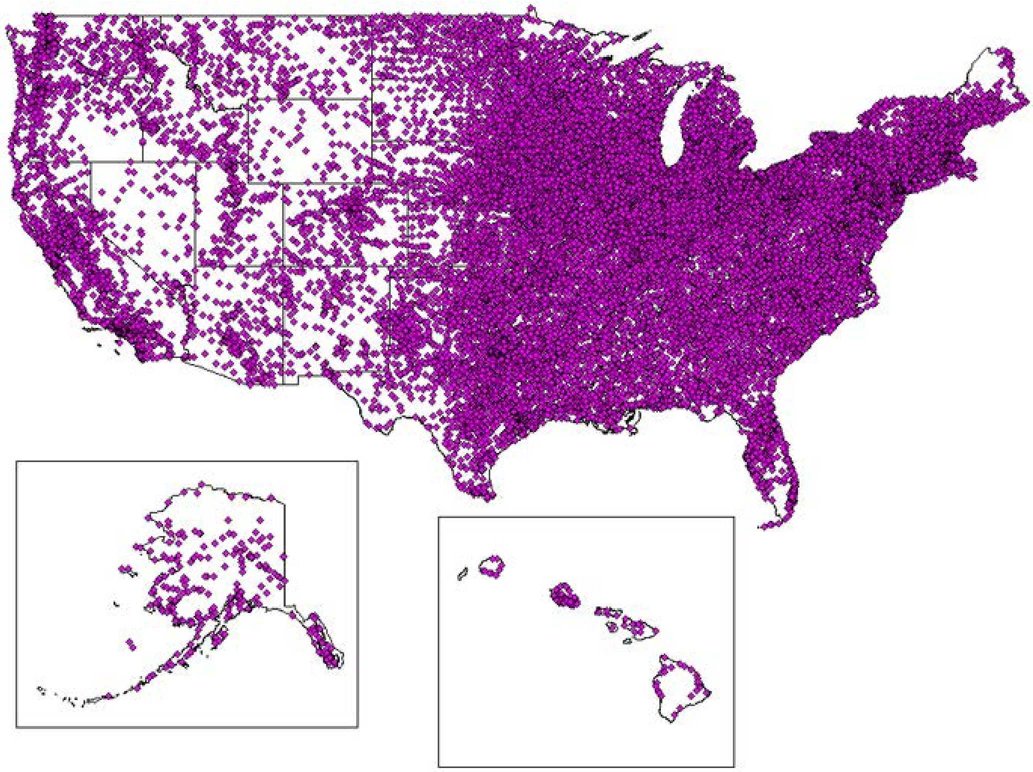

Source: Wikipedia

Decentralized, hierarchical topologies of switching centers were necessary historically, in part because of the limitations of analog voice switching technology, in part because of traffic patterns. The population was not as mobile when the telephone system was built. The majority of phone calls were local and could be handled locally. Local switching centers managed most of the load, and the long distance capacity, which was scarce and expensive, was optimized. This resulted in a hierarchal and distributed network topology (Figure 2) that offered the capillarity for reaching households and, much later, mobile base stations, while minimizing the number of international trunks and exchanges.

Figure 2. Hierarchy of telephony switching centers

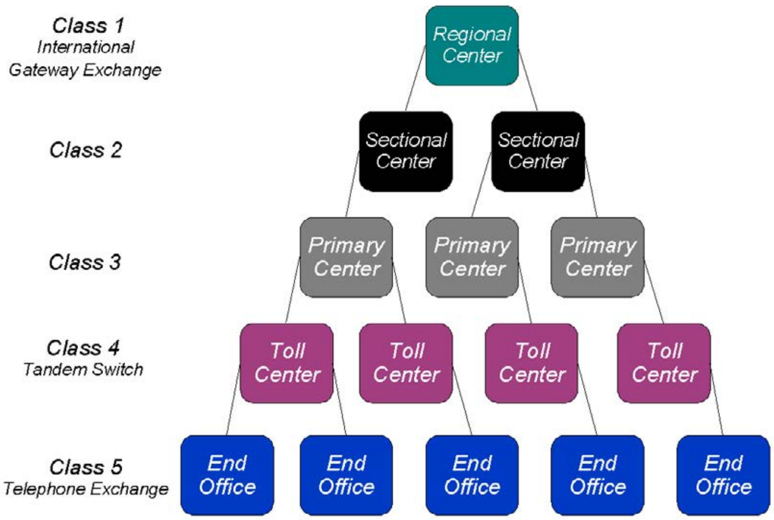

Source: Wikipedia

These old networks have served as a solid base for voice service and also, during more than two decades, for the beginnings of the IP networks that support our digital world. Nonetheless, many things have changed since the time of circuit switched telephony, and service providers are understandably asking if the decisions of that time are still valid.

2. NFV CHANGES THE GAME

Network Functions Virtualization (NFV) has marked the beginning of a new era in telecommunications networking. The virtualization of network functions on top of an industry-standard server infrastructure, typically a private carrier cloud, provides a radically new technology for building networks.

Virtual network functions (VNF) are often decomposed into multiple components that run on different virtual machines, each of which can be placed in the same or different locations. Virtualization thus brings heretofore unseen placement flexibility as network functions — and even components of network functions — are no longer tied to specific physical locations. This gives us the flexibility to distribute network functions throughout a geographic area, either in regional data centers, metro areas, neighborhoods or even on customer premises and mobile devices.

3. IT CLOUDS ARE CENTRALIZED

It is widely accepted in the IT world that a small number of warehouse-size data centers are more cost-effective than many small, widely spread data centers. First, the Internet companies that have dominated this industry do not have to build and operate local access networks. Second, bandwidth is inexpensive and high quality, making the geographical placement of data centers extremely flexible. As a result, they can choose locations based on the availability of cheap power, tax concessions, or lower labor costs. And having fewer facilities overall lowers security and general facilities management costs. The result is that they operate very large and highly centralized infrastructures to receive and process the aggregated traffic. This is opposite, as we have seen, to how traditional telephone networks have been built.

Figure 3. Nine large data center locations of Amazon Web Services

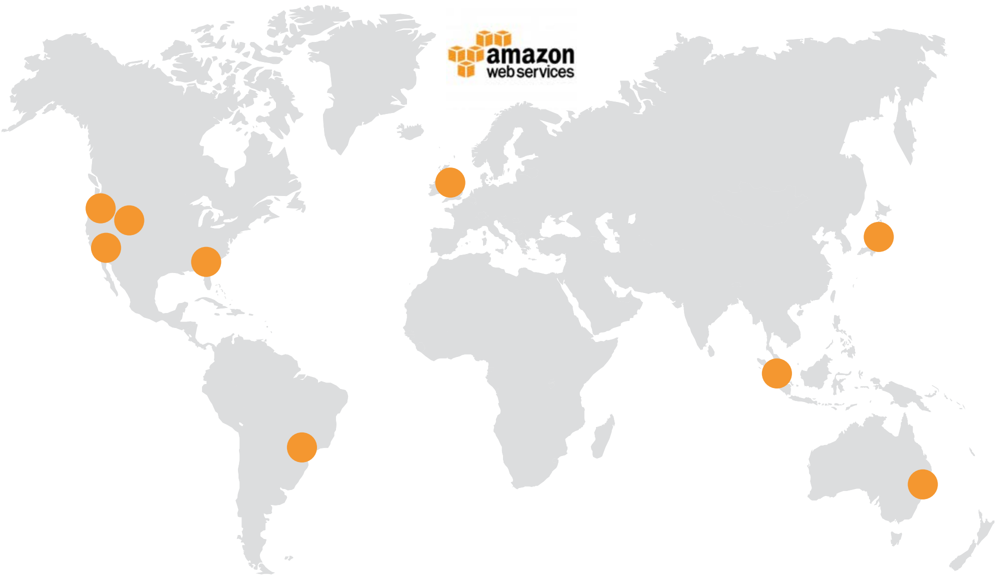

Amazon Web Services, the largest cloud provider in the world is a good example of a highly centralized architecture, serving its customers from as few as 9 locations worldwide (Figure 3). The success of Amazon and its peers shows that centralized clouds make sense for their business. However, it is also the nature of their business, which is largely web applications and certain types of content applications and transactional applications, that make this architecture optimal. The applications and services they offer can tolerate the latency created by data packets having to travel long distances between user equipment and servers.

4. NFV-BASED NETWORKS NEED DISTRIBUTION

In contrast to IT clouds, such as Amazon's, distribution still matters in NFV networks. Many carrier applications (as well as some IT applications) have needs for which a centralized architecture isn't suitable. These demands are related to network offload, low latency and jitter, availability, security and regulations.

4.1 Network offload

Whereas in the past voice traffic dominated the networks, today, video and data traffic use the majority of network capacity.[1] Although for different reasons, they also benefit from a distributed architecture. Streaming the same video content over and over again from a central source to each viewer is an inefficient use of network bandwidth. The two principal methods for overcoming this inefficiency, content distribution and multi-casting, both benefit from a hierarchical, distributed architecture. Similarly, back-hauling mobile data traffic to a central location to authenticate users and apply security is not an efficient use of network capacity. A distributed authentication function is more efficient. Peering points with other communication service providers are also an important consideration. Traffic destined to other service providers should be routed via the closest peering point to their network, not from a centralized point.

4.2 Latency and jitter

An obvious issue with centralized deployments is signal latency. In telecommunication networks, data travels at the speed of light, which is very fast but can still be noticeable, especially over longer distances. In fiber, the speed of light is about 200,000 km/s or 10 milliseconds (ms) for a roundtrip over a circuit of 1,000 km distance. And here we have to keep in mind that the fiber lengths may be significantly longer than direct line of sight. Additional latency is caused by switches, routers, and other network equipment along the way. Although a high speed router should take less than 100 microseconds (µs) to transmit a single packet (thus a roundtrip across 5 routers should be less than a millisecond), this is only the case if the switches are not highly loaded, and there is no significant switching contention and port queuing.

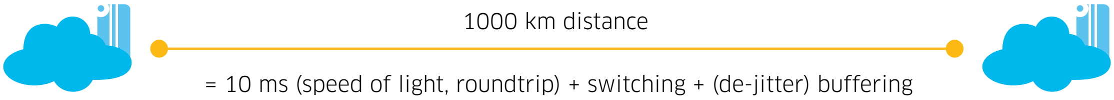
1000 km distance
= 10 ms (speed of light, roundtrip) + switching + (de-jitter) buffering

A comparatively large portion of latency is incurred in the access domain. The smallest latencies occur with fiber access, e.g., between 10 and 20 µs, however, mobile networks can cause round trip delays of several hundred milliseconds due to measures to counter noise and contention in a shared radio environment. LTE networks reduce these mobile latencies to below 100 ms, but latencies are still significant compared to fiber.

[1] Video shakes up the IP Edge

NFV introduces additional sources of latency through the virtualization layer, including virtual switches inside the servers. These latencies can multiply if data packets travel through a sequence of virtual machines, as is the case with service chaining.

For voice communication services, or any full-duplex, real-time traffic, such as video conferencing, not only absolute latency is critical. Tail latencies and delay variation (jitter) need to be contained, as well, and can add up if the signal travels across many switching points. To re-create a continuous speech flow in the presence of jitter, the signal needs to be buffered at the destination. If the jitter is high, these buffers will cause an unpleasant delay in two-way conversations.

4.3 Reliability and availability

Another reason for distribution is service reliability and availability, including disaster survivability. One type of risk is a “smoking hole” scenario or some larger geographically defined disaster. The 2004 Sumatra-Andaman tsunami, Hurricane Katrina in 2005, and the 2011 Tohoku tsunami with the Fukushima nuclear meltdown are still fresh in our memories. The latter disaster, for example, has prompted Japanese service providers to thoroughly review and change their disaster protection measures to reduce possible impact for the future, notably even in areas not directly affected by the meltdown.

Another type of risk is service provider error, for example, misconfiguration of the infrastructure or erroneous calculations. Large centralized data centers have been affected by power outages and software problems. In 2012, a power failure in an Amazon datacenter uncovered a bug in their load-balancing software which stopped applications from redirecting service traffic to other centers. These failures knocked out nationwide services, such as Netflix and Instagram. In some cases, these outages spread to distant data centers due to simultaneous recovery processes causing overload and failure even in places not originally affected.

Distribution helps to restore service quickly to users affected by an incident. These installations do not necessarily have to be small. They need to be independent and geographically far enough apart (over 1,000 km) that not all of them are affected by the same disaster.

4.4 Security and regulatory reasons

In the area of security, distribution can be both a risk but also an opportunity for improved security.[2]

Distributed networks present greater risk because they provide an enlarged attack surface with more locations and network connections between them. Once attackers have infiltrated one of the locations, they may be able to spread to other locations and even attack critical management and orchestration functions.

At the same time, distribution can also be a tool to mitigate risk. Carefully distributed NFV applications can quarantine localized attacks leaving the vast majority of nodes and users unaffected. If security measures are implemented coherently across the network, security attacks can be detected automatically, and infected elements, isolated and restored, while remaining elements continue to operate. For example, distributed denial of service (DDoS) attacks use the power of distributed resources to overwhelm service infrastructure. But a similarly distributed infrastructure can be an effective countermeasure, enabling service providers to allocate additional cloud resources dynamically to withstand attack waves.

[2] NFV Insights Series: Providing security in NFV - Challenges and opportunities

Government regulation may be an additional reason for distribution at the level of nations or groups of nations. Beyond the reliability and availability targets discussed above, critical infrastructures may have to be self-contained within national boundaries, for instance in the EU, authentication services that store personal data.

5. SERVICE EXAMPLES THAT REQUIRE DISTRIBUTION

For these sound reasons, most real-world deployments of NFV will require geographic distribution of virtualized functions for some services. In contrast to Internet companies, as we noted above, most carriers also operate access networks. Both fixed and wireless access networks have limited placement flexibility. Even in mobile networks, the trend towards using small cell radio and sensor networks to improve coverage and capacity means that network elements are placed closer to users than ever. And as we saw, latencies in mobile access networks, even LTE, dictate that some services need to be highly distributed. Other functions — call them service functions — have more flexibility in where they are placed, and other variables than proximity to users will determine what degree of distribution is optimal to provide the best service experience for the user and minimize the cost of operations.

5.1 Video services and virtual content delivery networks (vCDN)

Today, the majority of network capacity for consumer services is used for video consumption, in the form of downloads or streaming. This is also the fastest growing service and will put increasing strain on current network architectures.

Content Delivery Networks (CDNs) are the principal way to offload video traffic (including video downloads) on international and some wide area networks. Instead of streaming or downloading the same content over and over again from a central server, CDNs cache the content closer to the subscriber. The world's largest CDN, Akamai, has well over 1,000 points of presence (PoPs) in more than 80 countries. A country such as Germany with 80 million inhabitants is serviced by no more than 14 PoPs. Other CDN operators work with even fewer super-PoPs. Level 3 runs less than 30 PoPs to cover the USA and about 20 PoPs in Europe. This seems to indicate that de-centralization of CDN functions down to the range of tens of kilometers or closer is not currently happening, perhaps because the benefits of content caching at that level of granularity do not warrant the cost of the additional infrastructure. It remains to be seen if the advent of distributed carrier clouds will change that balance.

In fact, popular, high volume video streaming traffic, such as live TV and video-on-demand, can be further optimized by deeper distribution. The load that video-on-demand traffic places on the network can be further reduced by using specialized CDNs or by distributing the video platforms to cache the content even closer to the users. Live, broadcast video traffic, such as sporting events, can also be optimized through the use of multicast capabilities in the network or through specialized CDNs when this service is offered by over-the-top (OTT) providers. These techniques can provide significant savings[3] but their usage will depend on future net neutrality regulations.

[3] Video shakes up the IP Edge

5.2 Virtual radio access network (vRAN)

The physics of radio waves determine the optimal location of radio antennas. However, parts of radio base stations can be placed away from the antennas or remote radio heads. For example, the baseband processing units of base stations can be pooled for multiple radio heads. This allows service providers to better manage processing capacity and simplify the maintenance of the base band units. The vRAN is one of the most latency sensitive NFV applications identified by the ETSI NFV Industry Specification Group. Signal latencies between remote radio heads and baseband units need to be in the range of microseconds to a few milliseconds. This limits the fiber distance between them to less than 40 km.

5.3 Virtual customer premises equipment (vCPE)

Consumer and business CPE, such as DSL routers, firewalls and set-top boxes, are the most numerous network elements. Failures and maintenance in most cases involve sending technicians and adding or replacing equipment which generates important operational expenditures. In addition, introducing a new service to a large customer population that requires a different CPE can be a costly barrier to deployment.

This is why service providers are interested in virtualizing CPE functions. Two models of virtual CPE can be considered. In one model, the physical CPE is reduced to the “bare bones”: e.g., DSL modem, switch, WiFi access. All service functions, such as packet encapsulation, security/authentication, IP address assignment/dynamic host configuration protocol (DHCP), network address translation (NAT), firewall, graphical user interface (GUI) with statistics, wireless access network (WLAN) config, access network config, voice over IP (VoIP)/private branch exchange (PBX), child protection/access control, storage/network attached storage (NAS), power management and IPv6 are virtualized and moved to the network. In another model, the physical CPE contains virtualized compute, storage and network resources and effectively becomes a mini cloud node. This allows service providers to dynamically locate network service functions either at the customer premises or in the network, whichever is more advantageous.

When functions are moved out of the CPE and into the network, resilience and performance become major concerns because so many more customers are affected by simple CPE issues. In most cases, this means that subscribers are grouped into regional clusters, with each cluster having access to the resources of neighboring clusters as a form of redundancy. Management of this distribution, thus becomes very important.

6. DISTRIBUTION IMPOSES NEW REQUIREMENTS ON THE NFV ECOSYSTEM

As we have seen, the ability to distribute functions becomes mandatory for many of the new cloud-based services, but this need for distribution doesn’t come without its challenges. This is where NFV infrastructures become very different to typical IT-centralized clouds.

Highly automated NFV points-of-presence: NFV data centers and PoPs are more numerous and new ones need to be set up more frequently. This means that tools are needed to install and commission NFV PoPs efficiently. The smaller PoPs will be run in a lights-out mode without staff permanently on location. That is, virtually all management actions should be doable remotely. For example, in case of a server failure, it may not be cost-effective to send maintenance personnel to the remote location to replace it immediately. Instead, the remaining servers will host the affected virtual network functions. Ideally, replacements can be deferred until a whole rack is end-of-life and can be replaced with a new hardware generation.

Intelligent placement of virtualized network functions (VNF): The placement of VNFs needs to take into account specific performance and security requirements, as well as the availability of resources. To be able to automate VNF lifecycle management with deployment and scaling, placement needs to be automated and policy-driven. Placement algorithms should execute appropriate policies that describe both the needs of the VNFs, and the way the service provider wishes to fulfill them.

Automated networking with SDN: The geographically distributed placement of VNFs also means that the necessary network structures with the required service levels need to be created, not only inside the data center, but also across the WAN. For example, when a VNF component is placed in another data center, then all the connected networks need to be able to follow in order to reach the new location. This is where SDN can significantly simplify these changes.

A single view of the infrastructure: Managing all NFV data centers and PoPs as separate clouds would incur a high cost and be error prone. Service providers need a coherent pool of compute, storage and network resources across all locations with a single management view of the resources, the deployed VNFs and networks. In case of degradation or failure, there needs to be ways to quickly identify affected resources, including root cause analysis, and remedy the situation to provide service assurance. This single view should also include capacity usage trends to allow service providers to anticipate future needs and risks of failure or service degradation. To enable this, analytic applications will need extensive data to be collected and analyzed from the different infrastructure locations, the network and the deployed VNFs.

7. MODELS FOR DISTRIBUTION

Many service providers have taken the advent of NFV as a trigger to start rethinking their network architectures. Service providers are looking to simplify architectures and move toward a more consistent and flexible model.

For the reasons discussed above, network offload, latency, jitter, reliability, availability and security, service providers will choose a multi-tier network architecture giving them the flexibility to distribute network functions optimally (Figure 4). In reality, existing organizational structures and ownership will also influence technical architecture.

Figure 4 shows an example of a four-tier network comprising customer premises, city, regional, and core tiers. Clearly, the detailed design, including the number of tiers, will depend on the requirements of the specific service provider. The figure also illustrates some placement constraints due to the required sequence of network functions. For example, the virtual Border Network Gateway (vBNG) as an authentication and security device would always be closer to the subscriber than service functions, such as virtual IMS (vIMS). Likewise, in a mobile network (not shown in Figure 4), the Evolved Packet Core (EPC) Serving (S) and PDN (P) gateways need to be closer than Gi-LAN service functions. To provide the necessary service availability, at least the Tier 1 and Tier 2 locations need to be geographically redundant and the tiers closer to the subscriber are multi-homed into more than one location in the tier above. Peering interconnections with other service provider networks will primarily occur at the regional and core level, and these connections need to be protected with firewall and session border controllers.

Figure 4. Multi-tier architecture

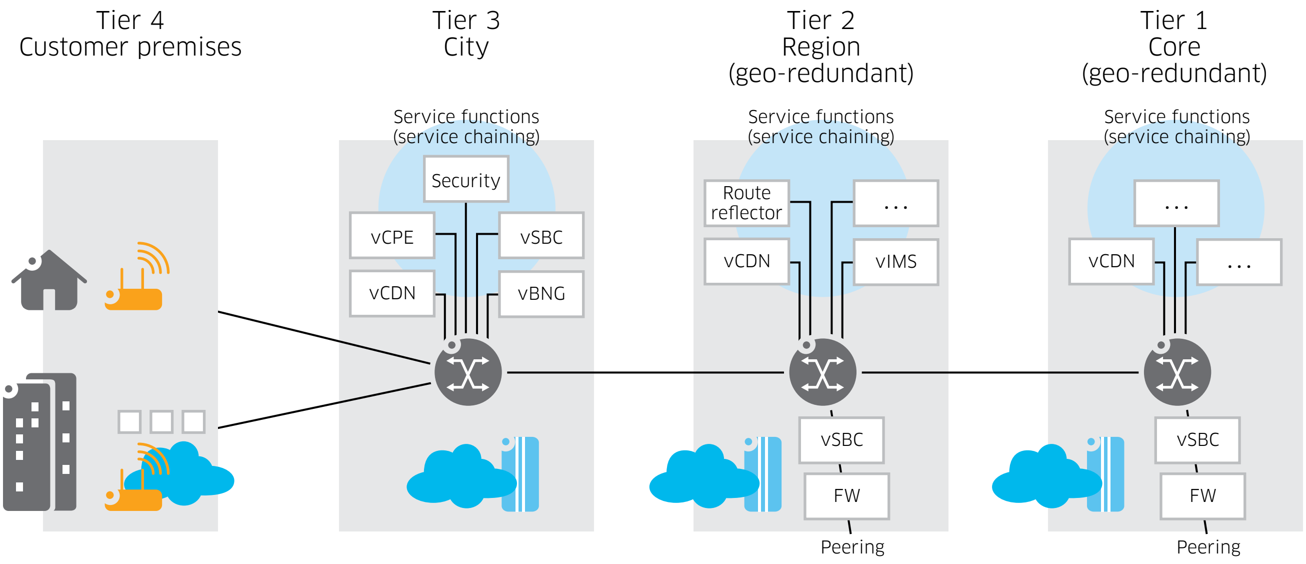

In general, functions that are data plane intensive (e.g., virtual CPE (vCPE), vBNG, virtual session border controller (vSBC), and firewall (FW)) should be distributed when transmission is an issue. Functions more related to the control plane (e.g., vIMS, Home Location Register (HLR)) can be more centralized whenever techniques for ensuring availability are in place, as discussed previously (Figure 5).

Figure 5. Distribution of data plane and control plane functions

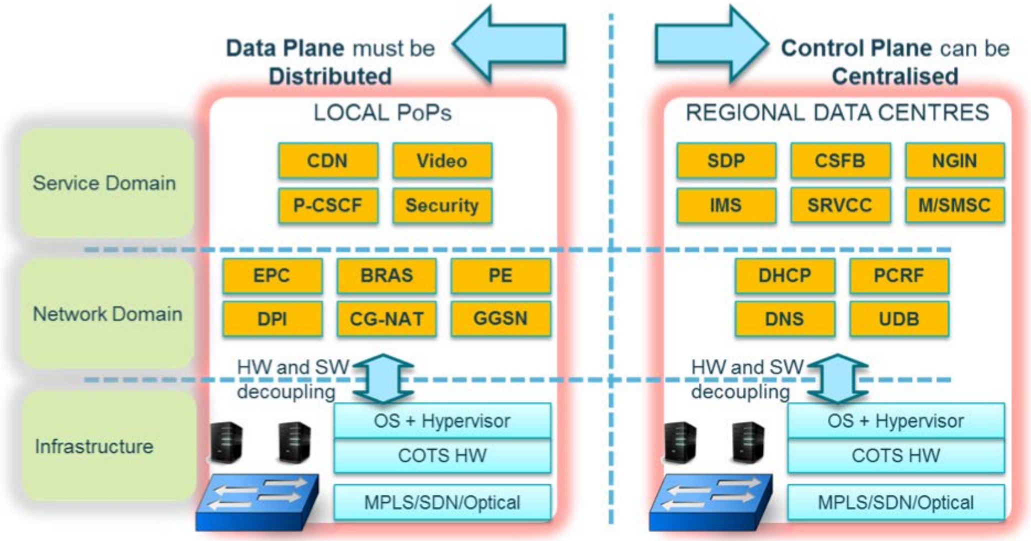

Source: Telefonica

8. CONCLUSIONS

NFV enables the separation of network functions from dedicated hardware using a private carrier cloud infrastructure. This allows service providers greater freedom in distributing service functionality throughout a network coverage area, and network functions can be shared across common hardware. However, this does not necessarily imply the degree of centralization we see in the case of IT cloud providers, such as Amazon. Telecommunications networks and services are, in many instances, different to IT workloads and so distribution in NFV becomes very important. Clearly, distributed access and aggregation networks must be in place to service geographically dispersed users, but cloud-based service functions also need to be placed at various distances between customers and subscribers and regional centers.

Instead of simply replicating a fixed network architecture, with NFV it is possible to define the position of cloud-based network functions based on policies. Knowing the requirements with respect to latency/jitter, network offload, service availability and security allows service providers to define the right placement policies. To run a distributed NFV infrastructure efficiently, a number of challenges concerning placement, networking, resource management and operations need to be addressed.

To attain the operational cost advantages promised by NFV without degrading service performance, an NFV platform with SDN tightly integrated is needed to automate policy-based placement, resource management and service assurance and, hence, manage the distributed NFV infrastructure as a coherent pool of resources.

9. GLOSSARY

BNG	Border Network Gateway
CDN	Content Delivery Network
CPE	Customer Premises Equipment
DHCP	Dynamic Host Configuration Protocol
DDoS	Distributed Denial of Service
DSL	Digital Subscriber Line
EPC	Evolved Packet Core
ETSI	European Telecommunications Standards Institute
FW	Firewall
GUI	Graphical User Interface
HLR	Home Location Register
IMS	IP Multimedia Subsystem
IP	Internet Protocol
IPv6	IP version 6
IT	Information Technology
km	Kilometer
LTE	Long-Term Evolution
μs	Microsecond
ms	Millisecond (1,000 μs)
NAT	Network Address Translation
NFV	Network Functions Virtualization
OTT	Over-the-Top
PBX	Private Branch Exchange
PoP	Point of Presence
SDN	Software-Defined Networking
TV	Television
vBNG	Virtual BNG
vCPE	Virtual CPE
vIMS	Virtual IMS
VNF	Virtual Network Function
VoIP	Voice over IP
vRAN	Virtual Radio Access Network
vSBC	Virtual Session Border Controller
WiFi	Trademark for a local area wireless technology
WLAN	Wireless Local Area Network

10. FURTHER READING

1. Maintaining Service Quality in the Cloud (http://www2.alcatel-lucent.com/techzine/maintaining-service-quality-in-the-cloud)
2. Video Shakes up the IP Edge (http://www.alcatel-lucent.com/wps/DocumentStreamerServlet?LMSG_CABINET = Docs_and_Resource_Ctr&LMSG_CONTENT_FILE = White_Papers/Video_Shakes_Up_IP_Edge_EN_Whitepaper.pdf
3. NFV Insights Series: Business Case for Moving DNS to the Cloud (http://resources.alcatel-lucent.com/?cid = 178476)
4. NFV Insights Series: Providing Security in NFV - Challenges and Opportunities (http://resources.alcatel-lucent.com/?cid = 178552)