ALTICE LABS WHITEPAPER

An NFV/SDN Enabled Service Provider: A New Generation of Digital Services

January 2015

Keywords: NFV, SDN, cloud computing, digital services, automation, OSSs

altice labs

Altice Labs
Rua Eng. José Ferreira Pinto Basto
3810-106 Aveiro – Portugal
http://www.alticelabs.com
Tel: +351 234 403 200
Fax: +351 234 424 723

Contents

Summary

Operators are facing today an enormous pressure coming from the “Internet Industry”, commonly called OTT providers. OTT players are usually agile, able to launch new services faster and cheaper than traditional operators, providing higher levels of customer experience with less resources. To cope with these challenges, operators need to do much more with much less. To address this, operators are adopting the cloud paradigm, separating the services logic (software) from the infrastructural resources, moving away from the traditional software/hardware appliance model. In this context, NFV arises as the standard to virtualize network functions, while SDN technologies come along NFV in order to improve the network flexibility and configurability. Together they are leveraging a significant mutation in traditional network and service structure.

Altice Labs, as a telecom software/hardware vendor, is currently working hard in order to introduce the SDN/NFV concepts into its portfolio of products, by collaborating with operators and relevant vendors to reach the best architectures and solutions. This whitepaper describes Altice Labs’ vision for an NFV/SDN-enabled Service Provider, focusing on the architectural impact and challenges that are still to be undertaken to fulfill this paradigm. Furthermore, it also provides an insight on Altice Labs’ portfolio evolution strategy towards NFV/SDN-enablement, as well as concrete use-cases that are being deployed in the Altice Labs NFV/SDN R&D Lab.

1. Introduction

Cloud Computing represents a new computing paradigm in which IT functionalities are provided *as-a-Service* from large data centers [1]. Technically, the Cloud is a virtualized computing environment, intrinsically providing fault-tolerance and high-availability features, as well as fully automated on-demand service delivery and dynamic scalability. The advent of a computing paradigm with these characteristics fosters *pay-per-use* based business models, in which the end-user consumes services as a utility, that is, paying only what he has actually consumed. Although this service delivery model is very attractive for the end-user, it places significant technical challenges on the service provider side since the latter has to accurately measure the *exact amount of service* that has been used in a certain period of time, and at the same time assure the infrastructural resources required to provide the contracted service grade.

Generally, there are three commonly recognized service models for delivering IT functions following the Cloud Computing paradigm: IaaS (Infrastructure as a Service), PaaS (Platform as a Service) and SaaS (Software as a Service). Prominent service providers, such as Microsoft, Google, SAP, Amazon and Rackspace, to name a few, are already successfully exploiting this computing paradigm, in both the residential and on the enterprise market segments. To achieve this maturity-level, besides virtualizing their IT functions in the first place, the above mentioned service providers had to migrate to advanced service lifecycle management solutions to fully achieve the cloud basic features – e.g. on-demand service provisioning and scalability, as well as pay-per-use charging models.

The urge to increase significantly the operational efficiency in service delivery and the agility to create new services lead the traditional communication service providers to begin the adoption of the Cloud Computing paradigm as a foundation to its service infrastructure.

The first steps were the construction of new generation data centers, based on cloud computing enabling technology / platforms, mainly as a new infrastructure for its management support systems and also as a infrastructure for new IT Services Business Lines. In this movement, aside from the inherent increase in operational efficiency, the major achievement was the significant increase in relevant knowledge in the Cloud Computing technologies.

In subsequent steps the communications industry started to consider cloud computing technologies as means to evolve its traditional service architecture. In this context, standardization bodies like ETSI and ONF started to define new network and service architectures, using cloud computing to promote the evolution from traditional networks (built over traditional network appliances providing network functions over dedicated specialized hardware), to new generation networks where the trend is to decouple network functions (software) from the supporting infrastructure and to deploy those functions in the data center. It's the rise of NFV – Network Functions Virtualization and of SDN – Software Defined Networks.

Through this movement service providers aim to dramatically increase operational efficiency, overcoming the huge CAPEX and OPEX costs imposed by network appliances, and mainly to dramatically increase agility in new services launching, aiming to compete and interoperate with OTT players in a better position.

This movement is in its beginning, network solution providers do not have mature NFV/SDN value propositions, communication service providers are conducting "proof–of-concept" initiatives in order to evaluate NFV/SDN technology maturity and gain insight to define its own NFV/SDN strategy.

Altice Labs, as the Altice R&D division, has been participating in several international R&D projects in the NVF/SDN domains, as well conducting internal exploration projects, aiming to create the necessary knowledge muscle and insight to define Altice strategy to integrate NFV/SDN in its network evolution roadmap and allowing an evolution towards an NFV/SDN Enabled Service Provider. In the

following chapters, we intend to describe the evolution vision created so far and how that vision is influencing the roadmap and positioning of our product portfolio.

2. Vision on NFV/SDN Evolution

In the last decade we witnessed the widespread use of technologies that enable virtualization of computing resources and the proliferation of commercial proposals by the major providers for creating highly flexible virtualized servers on raw computing power. This advance materialized from a paradigm in which a physical computing platform was dedicated to a single application to a paradigm where multiple virtual servers are built on a single physical computing platform. This paradigm shift increases the computational resources efficiency by enabling the share of raw computational capacity between multiple virtual machines, hence deriving evident gains in the infrastructure setup for new systems, as well as on the operation and maintenance of the infrastructure.

Nevertheless, allowing the infrastructure resources sharing is only the first step to take full advantage of the Cloud paradigm. The virtualized infrastructure model by itself does not provide the means to dynamically adjust the computing resources allocated to a specific virtual machine according to the capacity actually required by users of the system. Contrariwise, as before, regardless of whether we are talking about physical or virtual infrastructure, this is sized for peak capacity expected for the system in question, which is often hit once or twice a year. This is perhaps the biggest contributor for the operating costs of communication operators.

The evolution to a new level, in which it is common practice the dynamic adaptation of computational resources to the effective use of the system, together with the implementation of monetization models of software functions according to their actual usage, will allow significant gains for both suppliers and customers. Although this service delivery model is very attractive for the end-user, it places significant technical challenges on the service provider side since the later has to accurately measure the *exact amount of service* that has been used in a certain period of time.

Nowadays, prominent IT service providers, such as Microsoft, Google, SAP, Amazon and Rackspace, to name a few, are already *surfing* the Cloud Computing *wave* with success in the personal and on the enterprise market segments. To achieve this maturity-level, besides virtualizing their IT platforms in the first place, the abovementioned IT service providers had to migrate to advanced lifecycle management solutions to fully achieve the cloud basic features – e.g. on-demand service provisioning and scalability, as well as pay-per-use monetization models.

The Cloud Computing *wave* extends far beyond the IT field, strongly impacting the telecommunications area at a stage where these are highly pressured due to lower profit margins. Telecom operators can be proactive and *surf* this *wave* thereby taking advantage of the intrinsic advantages of the Cloud paradigm or, on the contrary, adopt a reactive approach and be overwhelmed by technological difficulties and inadequate business models.

Telco's that are willing to survive in an increasingly competitive market, the only option left is precisely to embrace the technological challenges posed by this new paradigm and adopt disruptive business models. This approach will enable them to *surf the Cloud Computing wave*, reduce operational and investment costs and enhance their portfolio with novel services. Furthermore, and more importantly, it also enables the establishment of strategic partnerships with other service providers, for example, and above all, the OTTs through the exposition of network control capabilities, which to this day are used exclusively in the operator's telecommunications services.

Following the path already adopted by IT suppliers, the evolution of telcos towards the Cloud paradigm will materialize through significant impact in their IT domain, housed in their data center infrastructure, as well as in the telecommunications network domain used to deliver services to their customers. Figure 1 reflects the Altice Labs perspective on the telcos evolution to the Cloud paradigm, reflecting the migration impact on the telco IT segment, as well as on the telco network segment.

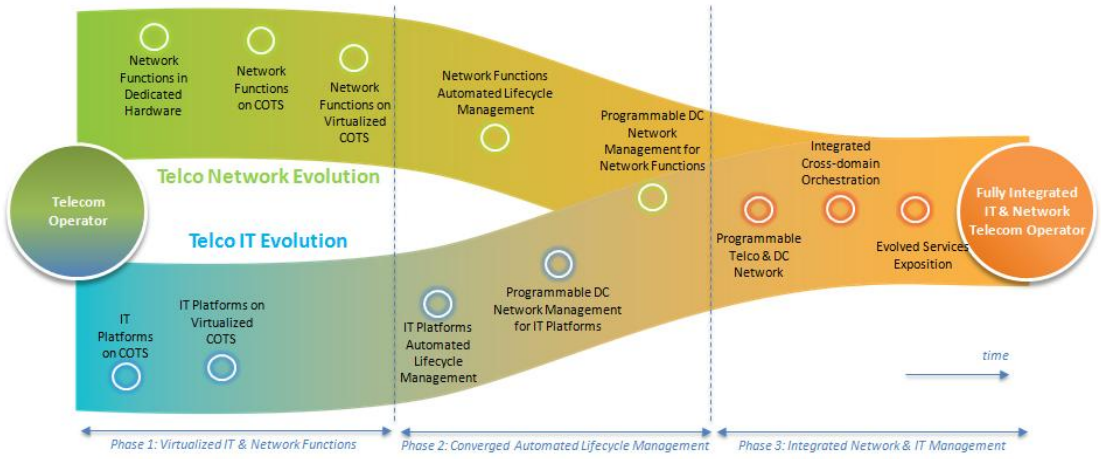

Figure 1: Evolution Path Towards a Telco Cloud Operator

The presented evolutionary path is split into three phases, as described in the following sections.

2.1. Phase 1: Virtualized IT & Network Functions

During the first phase, called *Virtualized IT & Network Functions*, there is a clear separate evolution of both telco segments (IT and Network), trying to optimize the use of the infrastructure resources. At this stage the synergies between the two movements is reduced.

Due to the massification of Cloud technology in the IT Industry and also due to the demanding non-functional requirements associated with the telco network, in the telecommunications industry the evolution to the Cloud paradigm has first happened in the IT segment. It should be noted that not all telcos are at the same evolution stage, but in general, the migration movement to the Cloud happens first in the IT field.

Taking advantage of the virtualization platforms maturity, telcos started deploying their IT platforms (e.g. OSSs, BSSs, IT services, etc.) in common hypervisors – *IT Platforms on Virtualized COTS* stage. Therefore it was possible to instantiate multiple IT platforms on the same physical server, taking advantage of all their computing power, thus reducing investment costs (CAPEX). At this phase, the management procedures for virtualizing the IT platforms (e.g. on-boarding, provisioning, monitoring, etc.) are completely manual.

In addition to the migration of their IT platforms for virtual infrastructures, telcos also started providing IT services to the enterprise market segment, such as virtual storage (SaaS), as well as development platforms (PaaS) and virtual servers (IaaS). This was an important strategic movement because it represented the beginning of the telcos involvement in a new business area, commonly occupied by traditional IT vendors.

From the network domain perspective, the evolution to the Cloud was driven by the predecessor movement in the IT domain. Despite the similarities, it should be highlighted the first stage on the network side, called *Network Functions in Dedicated Hardware*, in which the network functions acquired by telcos are instantiated in dedicated hardware, bringing enormous costs for operators. Thus, the first move in this area was the deployment of network functions in COTS, called *Network Functions on COTS*, and thereafter the migration to virtualized infrastructure, called *Network Functions on Virtualized COTS*.

In conclusion, despite the synergies between the IT and the telco network are minimal at this stage, we can already envision the beginning of a converging trend between the two domains by migrating IT and network functions to data center based virtualized infrastructure.

2.2. Phase 2: Converged Automated Lifecycle Management

In the second phase of the evolutionary process (*Converged Automated Lifecycle Management*), there is a considerable convergence of network and IT domains. As in the previous phase, the IT domain will be impacted first.

The first movement represented in Figure 1 – *IT Platforms Automated Lifecycle Management*, consists on the introduction of orchestration mechanisms to automate the lifecycle management (e.g. on-boarding, provision and supervision) of virtualized IT platforms. This automation will streamline and greatly accelerate the instantiation of these platforms.

The same trend is found in the network domain, but in terms of timing, a bit later – *Automated Lifecycle Management Network Functions*, through the use of fully automated, orchestrator-based, procedures for managing the network functions lifecycle. Although telco IT and network are areas with different requirements (mainly non-functional), the activities required to orchestrate IT and network virtualized functions lifecycle is similar. Thus, it is expected that the orchestration engines from IT and network domains converge towards a single orchestration element.

The following two movements within this phase, namely *Programmable DC Network Management for IT Platforms* and *Programmable DC Network Management for Network Functions*, bring automated and programmable mechanisms (APIs) to manage the data center network in which the virtualized IT platforms or network functions are deployed, respectively. It is the rapid growth of the Cloud phenomenon that exposes the limitations of traditional network technologies, especially (but not exclusively) within the data center. The problem is mainly related to new requirements associated with virtualization, involving dynamics, elasticity and mobility features.

Enhanced by the inclusion of programmatic mechanisms for the data center network, the scaling functionality is integrated at this stage. Adding scalability procedures, especially in the horizontal perspective, in which virtual machines are dynamically added/removed, requires dynamic and agile network reconfiguration. For this reason, adding scalability procedures should be done jointly with the network programmability functionality.

From the OSSs point of view, the existing management mechanisms will have to evolve for both domains – IT and network. More specifically, the traditional OSS systems will have to be able to interoperate with the abovementioned orchestration platforms and delegate to them the management of virtualized IT/network functions, or in alternative to integrate those orchestration functions assuming directly the lifecycle management of virtual functions.

This phase represents a strong convergence between the telco IT and network segments by automating, via technologically similar orchestrators, their lifecycle management procedures, such as on-boarding, provision, assurance and scaling. Required to support full automated management is network programmability, allowing the network to rapidly react to the changes in existing software platforms.

2.3. Phase 3: Integrated Network & IT Management

The last phase, called *Integrated Network & IT Management*, is the period during which, from the management perspective, the telco network and IT domains will fully converge.

At first, presented in Figure 1 as *Programmable Telco & DC Network*, in addition to the software defined programmability of the data center network, in the Telco network emerges the trend to separate the network control layer from the data plane layer and the ability to programmatically control (via APIs) the data layer network elements via centralized network control elements. This will simplify the legacy network management procedures and at the same time introduce new degrees of flexibility in the old plain legacy network.

In addition to easing the management procedures, separating the network intelligence from the infrastructure, and establishing open APIs, encourage innovation, in contrast to the proprietary and monolithic network solutions, dominated by a number of limited manufacturers, who now dominate the telco network.

From the orchestration platforms (IT and virtual network) point of view, given their technological similarities, they will evolve, in our view, for a common cross-domain orchestration platform for both the IT and virtualized network domains. In addition, we foresee that this cross-domain orchestrator will become part of the OSSs architecture, adding to the legacy Service & Resource Management systems the capacity to integrate virtual IT and virtual network functions lifecycle management, expanding its current footprint (Physical Network functions, supported by Legacy Network Elements).

With respect to the virtualized infrastructure (Data Center Point of Presence – DC PoP) management, it will continue to be provided by leading global players in this technical domain (e.g. OpenStack). The OSSs will interact with the virtualized infrastructure management platforms through well-known open APIs.

Although we believe that the cross-domain orchestrators will be integrated into the OSSs, it is essential that the later is prepared to deal with heterogeneous scenarios, in which 3rd party orchestrators will be made available jointly with the provided virtual network functions providers. The evolution of OSS systems, namely Resource Domain Management systems, to be cross domain (physical network functions, virtual network functions, IT functions) enabled, will boost the associated operating earnings and significantly increase the speed up capacity for creation of composite service offers, i.e., offers based on services supported by IT functions, virtual network functions and physical network functions.

The future “Service Platforms” for the Telco Industry will have this cross domain DNA, and therefore cross-domain OSSs will be mandatory to achieve the levels of automation and normalization required to be agile to launch composite services offers. This period, represented in Figure 1, is referred to as *Integrated Cross-domain Orchestration*.

The last mile in this “evolution chain” is designated *Evolved Services Exposition* period. Taking advantage of a vast network of DC PoPs scattered throughout the telco geographical area, that will come to life in the virtualization journey, as well as of the IT and network services that will be instantiated in those DC PoPs, the telcos will have all the tools to create the next generation service platform, and thus create value for customers through this new paradigm. Advanced networking services, nowadays used only for telcos internal consumption can be exposed to 3rd parties, for example OTTs, so that they can add value to their service offerings. This will allow the exploitation of new business opportunities based, for instance, on revenue-share models, thus transforming the telcos into active business partners of the OTTs. However, new business opportunities do not bind themselves only to partnerships with OTTs. On the contrary, leveraging the next generation service

platform, telcos will be able to quickly create combined IT and network services and therefore compete, in well-defined strategic areas, with OTTs, providing differentiated customer-centric services.

3.NFV/SDN-Enabled Architecture

3.1. ETSI NFV

The principle of Network Functions Virtualization (NFV) aims to transform network architectures by implementing network functions in software that can run on industry standard hardware. Furthermore, it intends to transform traditional network operations, as software can easily be moved to, or instantiated in, various locations (e.g. data centers, network nodes, end-user premises) without the need to put in place new equipment. NFV can bring many benefits, from improving operational efficiency and reducing power usage to shorter deployment/upgrade intervals and near-optimal network resource usage [2].

3.1.1. ETSI NFV Reference Architecture Description

So far we have referred several individual network functions candidate for NFV and have not looked from an integration perspective of several Virtual Network Functions (VNFs). In fact, many use cases will require, not a single network function, but a sequence of multiple network functions – a common example is a private network ingress point, where functions like NAT, firewall, DPI, load balancing, among others, are typically applied in a given order [3]. In the context of NFV, it is imperative to be able to define sequences of network functions that packets traverse. This process is commonly known as service chaining. ETSI entitled VNF Forwarding Graphs (VNFFGs) as the ability to specify complex structures, i.e. a chain, of network functions. Figure 2 illustrates a service chain, composed by several VNFs.

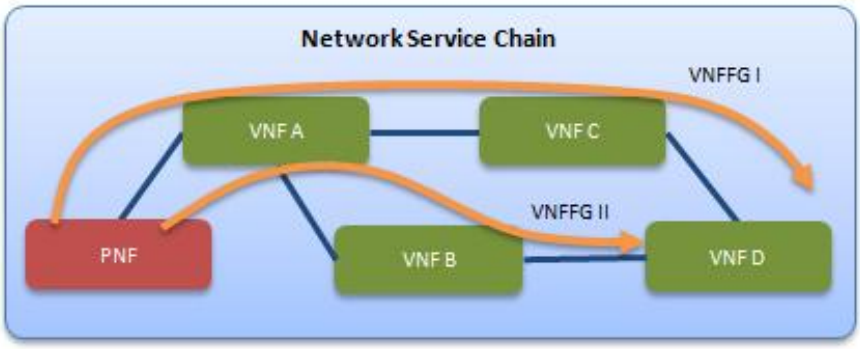

Figure 2: ETSI NFV Service Chains

Leading the leverage of NFV is ETSI with the NFV ISG Group. This group has produced what is considered today the reference architecture of an NFV framework [4], depicted in Figure 3. It focuses on the functionalities necessary for the virtualization and the consequent operation of an operator's network, identifying the main functional blocks and the main reference points between those blocks.

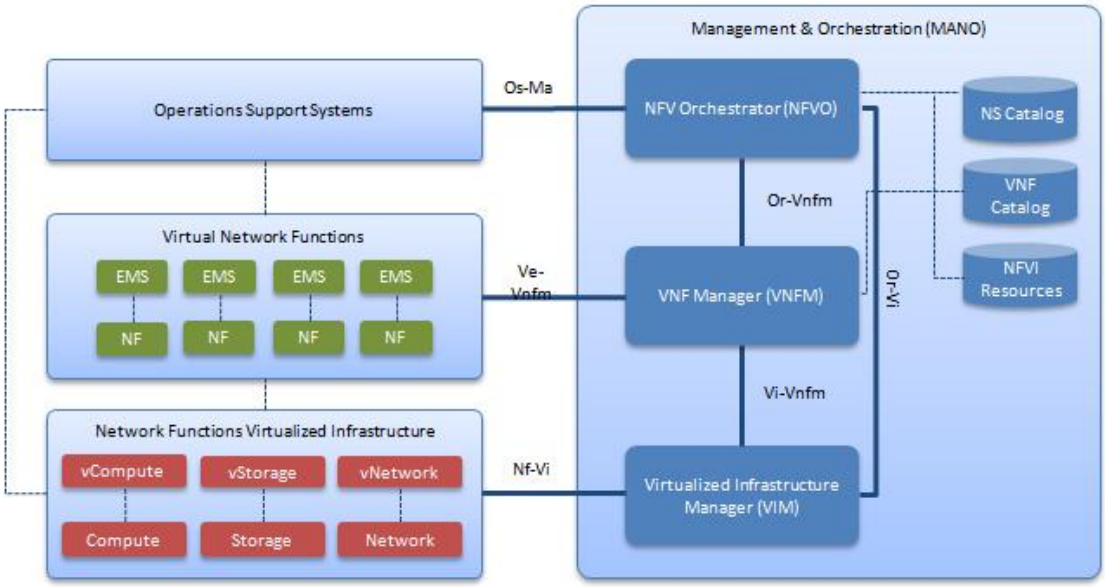

Figure 3: ETSI NFV Reference Architecture

Analyzing Figure 3, the bottom left of the picture represents the NFV Infrastructure (NFVI), which comprises all hardware and software components that support the environment in which VNFs are deployed, managed and executed. This infrastructure provides the necessary virtualized resources to the VNFs and can physically span several locations. Looking at Figure 4, it is possible to see a NFVI that comprises: a centralized data center, PoPs and also the costumer site when it has embedded on-site infrastructure to support NFV. The network providing connectivity between these locations is regarded to be part of the NFVI.

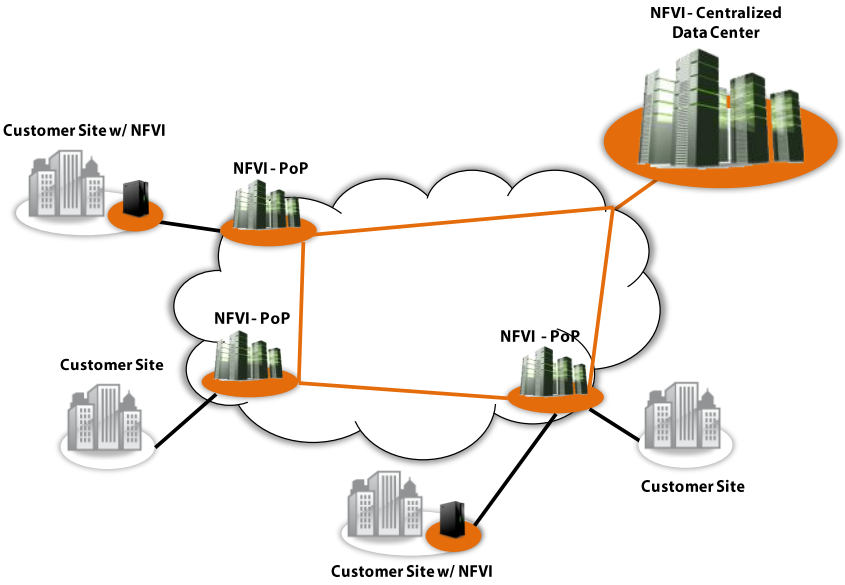

Figure 4: Distributed NFV Infrastructure (NFVI)

On the middle left side of Figure 3 are the VNFs, which use the resources provided by the NFVI. On the right side are the Management and Orchestration (MANO) elements. The Virtual Infrastructure Manager (VIM) is responsible for controlling/managing the NFVI resources (e.g. OpenStack). Note

that multiple VIM instances may be deployed. A VNF Manager (VNFM) is responsible for the lifecycle management of VNF instances (instantiation, configuration, update, scale up/down, termination, etc.). Finally, the Orchestrator (NFVO) is responsible for the orchestration and management of the virtualized infrastructure and software resources, and realize network services on the NFVI.

On the top left corner are the OSS/BSS of an Operator and on the right side are the i. Network Service, ii. VNF and iii. NFVI repositories.

3.2. ONF SDN

Networks need to keep up with the agility and rapid evolution that we see in Cloud-based applications today and this requires a fresh technological approach. Software Defined Networking (SDN) brings new capabilities in terms of network automation and programmability that facilitates the integration with the Cloud [5]. Making use of the SDN feedback loop, network control plane decisions can be made not only based on traffic engineering rules but also in response to dynamic conditions (e.g. network performance, application use trends, user behavior, congestion events, network malfunction) [6].

SDN is based on three fundamental ideas:

- Decoupling of the control and data planes;
- Abstraction of the network infrastructure resources;
- Programmability of the network via open APIs.

From the point of view of the control plane, one of the potential benefits offered by the decoupling from the data plane is the possibility to get a global perspective of the network resources and make decisions with much greater flexibility and speed compared to traditional networks. This aspect becomes especially (but not exclusively) important in the Cloud data center environment, where creation, migration and disposal of virtual machines occur on a very frequent basis. In SDN the network intelligence is logically centralized. However, this does not imply that the control of a network will rely on a single SDN controller. Multiple controllers can be considered depending on the scenario (e.g. a small data center may rely on a single controller, whereas an operator network will most certainly rely on multiple controllers). The way these controllers will interact among each other is still a subject under study.

Another advantage of SDN is a more granular end-to-end view of services with the ability to apply comprehensive and wide ranging policies, thus enabling a better quality of service and experience while improving efficiency.

Another important SDN characteristic is network resource abstraction. For the network manager, the use of a standard interface between the controller and the network elements creates an abstraction layer above the network physical substrate. The independence from the specific characteristics of the network infrastructure reduces vendor lock-in, allowing the control of network elements from different vendors transparently. The network manager only needs to worry about the supported API versions of the network devices, which need to be consistent within the network infrastructure. It should also be noted that the abstraction layer favors the creation of a programmable and automated network environment, which increases network reliability and security.

The network elements present in the network infrastructure are also impacted by SDN. Since these elements no longer needed to support control plane functions, they can be replaced by simpler elements, which only have to perform transport functions. This brings some advantages from the maintenance and acquisition costs point of view while reducing the number of causes for failures.

Finally, easy and agile adaptation of the network, following requirements raised by applications' dynamics and business requirements, is another major advantage of SDN. Programmability is the key

enabler here. Most importantly, open APIs guarantee independence from specific vendors or proprietary solutions.

3.2.1. ONF SDN Reference Architecture Description

The SDN architecture is divided in three layers, similar to the architecture found in computers, also divided in three layers, hardware, operating system and applications. Figure 5 shows the SDN three layer architecture.

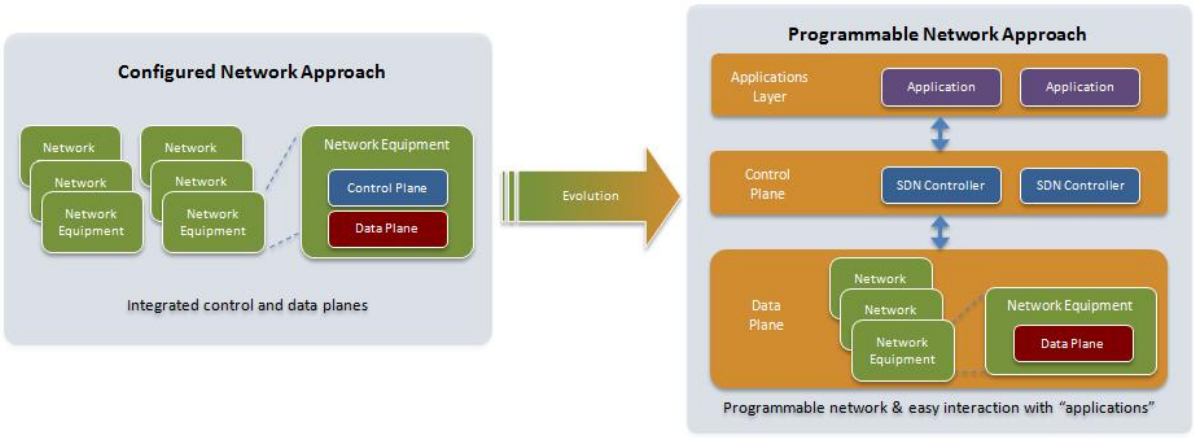

Figure 5: ONF SDN Reference Architecture

The bottom layer, which can be seen as corresponding to the hardware layer found in computers, includes the software or hardware based network devices that perform the data forwarding functions. The Control Layer represents the network operating system; the software located in this layer uses the lower layer resources to build L2 to L7 network services. This layer provides an important abstraction of the network infrastructure for the upper layer, which frees the Application Layer from the implementation details of the managed services. The upper layer is where network services translate their requirements into abstracted network resources for an optimal service delivery. This architecture allows the coexistence of different virtual network infrastructures, which can be optimized for delivery of specific services; the virtual network infrastructure is realized inside the applications.

Communication between layers is accomplished with each layer using the lower layer API. The control layer plays a pivotal role in the architecture and defines two basic interfaces – Northbound (with the Application layer) and Southbound (with the Infrastructure layer). As to the former, a standard API is still missing, but relevant initiatives in this area (e.g. ONF NBI-WG, OpenDay Light Consortium) have been looking at this specific issue and the first results are expected soon. With regard to the southbound interface, OpenFlow is the only standard defined so far, although there are other alternatives. In SDN, OpenFlow is the first standard communications interface defined between the control and forwarding layers of an SDN architecture. OpenFlow is a protocol that specifies basic primitives that can be used by an external software application to program the forwarding plane of network devices.

3.3. Foreseen Architecture

Nowadays, OSSs are responsible for managing and orchestrating the physical network functions lifecycle procedures. Therefore, since next generation network services will certainly be composed by

virtual and physical functions, it is a quite natural evolution path to enhance the Next Generation OSSs with the capacity to manage and orchestrate IT & network virtual functions. This will enable the composition and orchestration of combined virtualized IT & network functions, as well the orchestration of physical network functions in a singular management platform.

Nevertheless, the challenges to orchestrate virtual functions is much more complex when compared with physical functions. Virtual functions can be dynamically deployed / on-boarded, provisioned, started, paused, stopped, whereas the management procedures over physical functions are much more limited and static. The instantiation of a virtual network service is expected to be a fully automated procedure, without human intervention, while the instantiation of a physical service sometimes requires the explicit intervention of human resources on the field.

3.3.1. Cross-domain Telco Functions Orchestration

Figure 6 presents the foreseen telco architecture, focused on the cross-domain telco functions orchestration, corresponding to the vision that was described in section 2. Starting from the bottom of the figure, a set of functions is represented that can be part of the orchestration procedures – *Orchestrated Functions* layer. Namely, there are Virtual Network Functions (VNFs), from the core and the access network segment, as well as IT virtual Functions (ITFs), distributed over multiple DC PoP infrastructures. Traditional legacy Physical Network Functions (PNFs) deployed over non-virtualized infrastructure are also included.

The second layer is the *Functions Orchestration*, which is responsible for the cross-domain orchestration of all the available functions (VNFs, PNFs and ITFs), either they belong to the access or to the core network segments. Therefore this layer has an holistic view of all the functions and can dynamically deploy, provision, scale, monitor and terminate services made of virtual and physical supported functions, also known as service chains, according to the incoming requests from the business management layer (not represented in the figure).

The vast majority of services related with the IT domain can be centralized in a small number of DCs. Contrariwise, services that involve the networking domain, due to their demanding requirements, such as low latency and high throughput, can not be instantiated in a single centralized data center. For example, centralizing in a single data center control plane functions that are delay-sensitive, such as the Policy Management functions, can jeopardize the correct delivery of the service. Another example, now in the data plane perspective, is the Deep Packet Inspection (DPI) function, which has to filter and analyze data packets in real-time. The placement of these types of functions shall be thoroughly planned by the telco, so that data packets do not have to traverse the entire network towards a centralized data center. Therefore, instead of having a single data center infrastructure, the virtualized infrastructure of the future will be composed by several small data centers, also known as Points of Presence (PoP), geographically distributed throughout the telco network – *PoP-grid* virtualized infrastructure, that have to be dynamically interconnected according to the functions requirements. Thus, it is of the upmost importance that the cross-domain orchestrator is able to identify the most appropriate location to position virtualized elements according to i) the characteristics of the service (and constituent functions), ii) the SLA and iii) the available infrastructure resources in each PoP and WAN segment that interconnects them.

Finally, the upper layer, called *Services Exposition*, depicts the next generation service platform that will be in place to compose and deliver combined services to distinct market segments, such as, *Enterprise and Personal Customers Segment* and *3rd Party service Providers.*

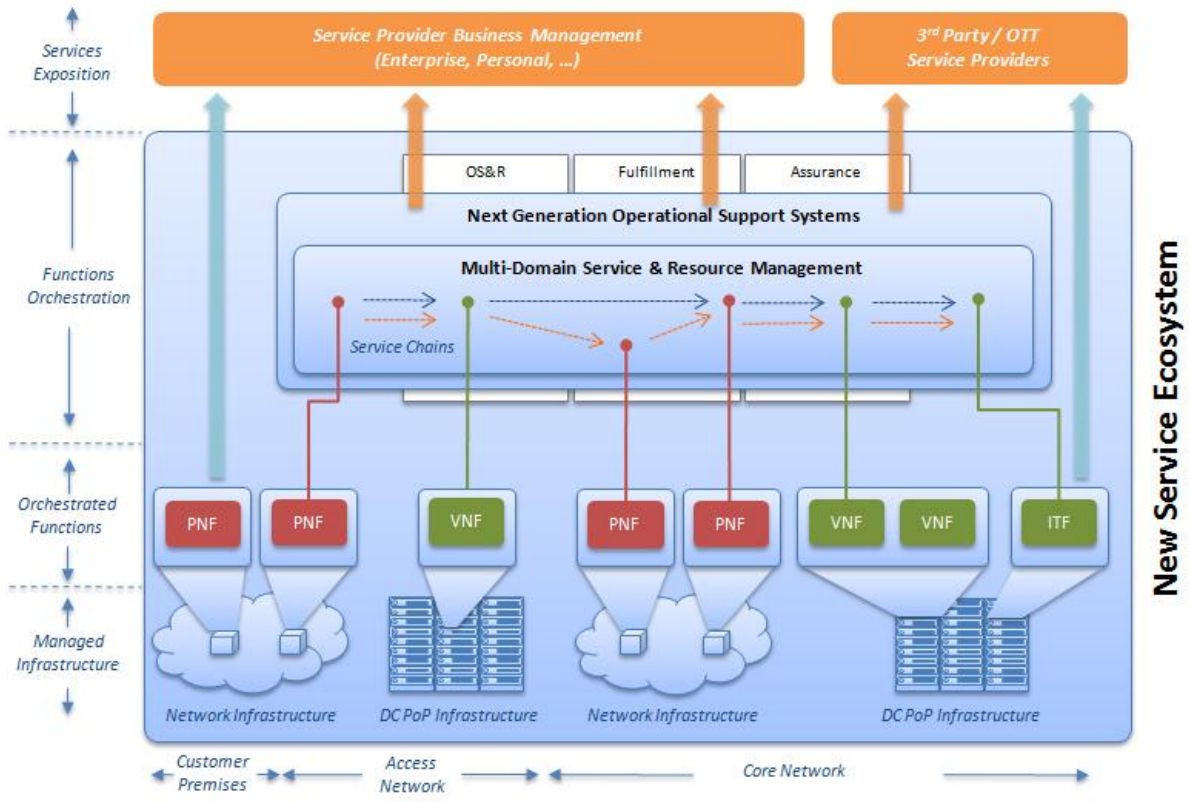

Figure 6: Cross-domain Service Functions Orchestration

3.3.2. Virtualized Infrastructure Management

With regard to the Virtualized Infrastructure Management (VIM), there are already a large number of experienced players with advanced solutions in this area. Therefore, in our view, the management of the virtualized infrastructure will be under the responsibility of these players. OSSs, as presented in Figure 7, will "consume" virtualized infrastructure management services provided through well-defined APIs. Furthermore, given the multiplicity of virtualized network functions (VNFs) that will be made available, it is essential that the cross-domain orchestrator is able to abstract and interact with various VIM providers.

The VIM interacts directly with the DC PoP for coLmputational resources provisioning, whereas the DC PoP network resources configuration is made through the SDN Controller platform, as illustrated in Figure 7. Therefore, from the OSSs perspective, the interaction with the SDN Controller might not be required in virtualized scenarios. On the other hand, for legacy network functions, since there is no middleware management platform, we foresee that the interface with the SDN Controller will be made by the OSSs.

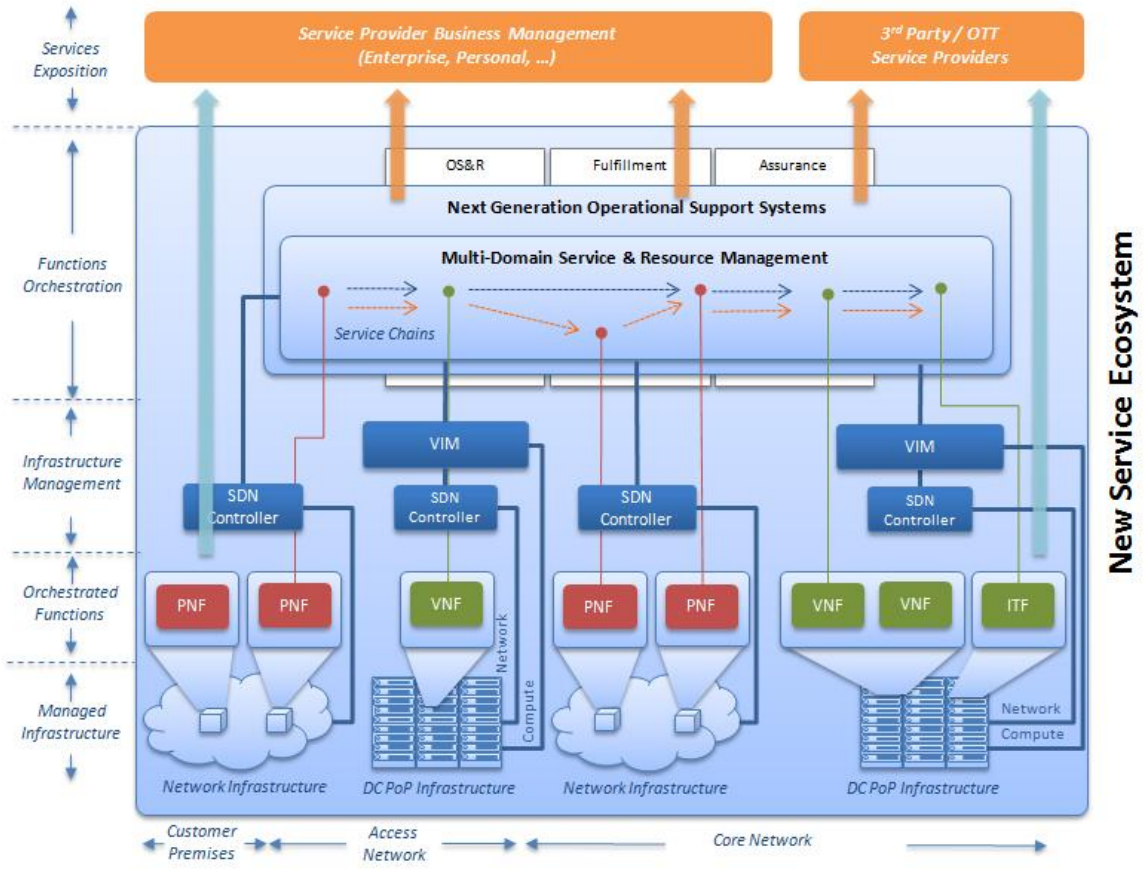

Figure 7: Virtualized Infrastructure Management

3.3.3. Federated Telco Functions Orchestration

Although we argue that telcos will benefit in having a centralized cross-domain management for physical and virtual functions spread across the access and the core network segments (cross-domain orchestrator), it is highly probable that multiple 3rd party orchestrators will proliferate. For example, VNF providers may position themselves as domain-specific orchestration platform providers as well. To address such scenarios, the telco OSSs must be able to interact with heterogeneous orchestrators, beyond its own orchestration environment.

Figure 8 illustrates this type of scenario, in which a federation of heterogeneous service functions orchestrator providers is represented. In the context of a specific service, the cross-domain OSSs will be responsible to explicitly manage some physical and virtual functions and to delegate the management of other functions to external domain specific management systems (orchestration platforms).

This scenario is well known today in the physical legacy network management, where OSSs interoperate with OMC platforms provided by network systems manufacturers and not with network equipments directly. This scenario will remain in this new network paradigm, being that the traditional OMC platform providers will, in many cases, evolve to encapsulate all required functions to manage VNFs provided by that manufacturer. This orchestrating layering will not be a desirable situation but will for sure be a reality.

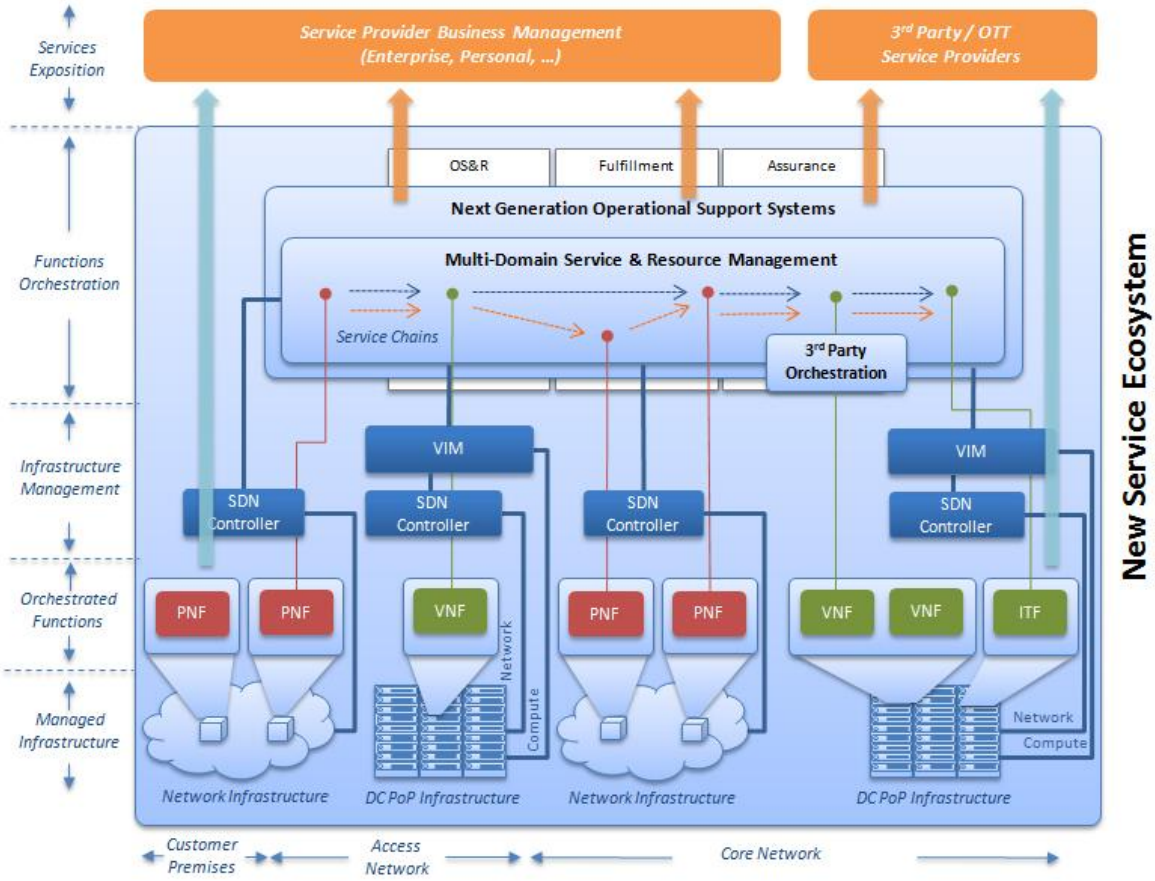

Figure 8: Federated Service Functions Orchestration

4. Solutions for NFV/SDN

4.1. Operational Support Systems

It seems likely that the virtualization of network functions, its decoupling from the hardware and the network that supports them, and even the intelligence and programmability of the network itself provided by SDN technologies will simplify some of the actual activities of OSS processes. For example, it seems less important to allocate and keep track of every individual network resource used by a customer service, because the future highly intelligent and programmable networks will handle that better. In some aspects this is similar to what happened when we changed from SDH or ATM networks to IP/MPLS with its agile and sophisticated control plane.

On the other hand, we don´t expect a scenario where the network becomes a cloud with a software management system on top, presenting itself to the OSS as a software interface. We are considering the scenario where network functions are spread across a network of datacenters and micro datacenters that are organized in a way that optimizes cost, customer experience, resilience, network load, etc. OSS will have to cope with this heterogeneous reality composed by several datacenters, SDN network segments, physical network functions that will not be suitable for virtualization, and legacy networks.

As expressed by Figure 9, to provide a service to a customer, several domains will have to work together in an organized way that will have to be orchestrated by an entity that has an holistic view of all ecosystem. So it will be up to the OSS to design services composed by the combination of resources provisioned over legacy networks, with network physical functions, and network virtual functions distributed over several datacenters. The way these different components are chosen and combined will be determinant to optimize operational costs and investment.

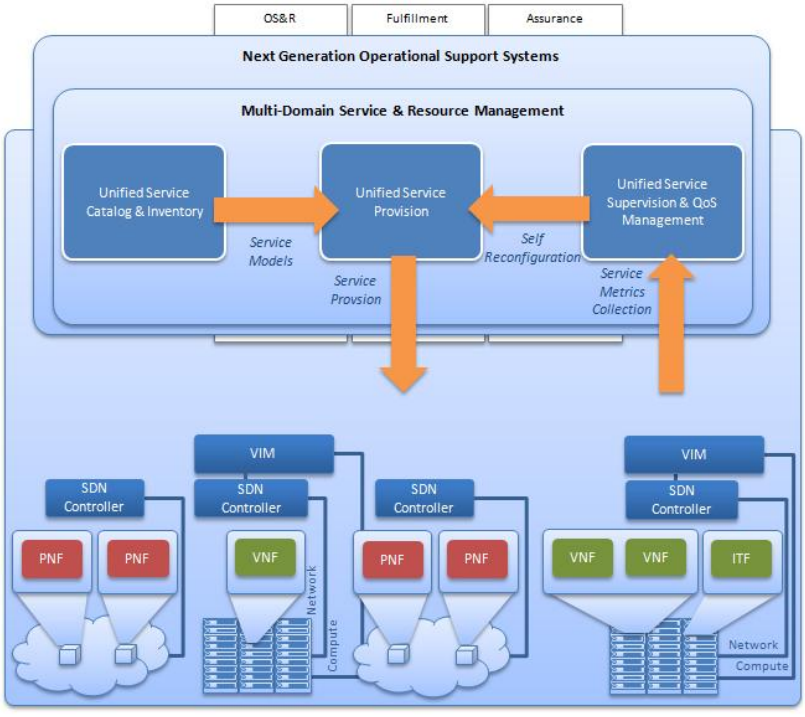

Figure 9: Multi Domain Operations Management

The service provision scenario, using OSS to configure network functions (either physical or virtual) and IT functions, does not cover it all. To take full advantage of new NFV technologies, a NFV/SDN enabled Service Provision Platform must be able to dynamically change the configuration, location and chaining of services. This will enable the roaming of services, allowing the customer to take his services with him wherever he goes. This will also enable the operator to reorganize its network on the fly, as its customers start and stop consuming the services provided. For example, when a premium customer starts using its service, we can move its network functions to the datacenter that is closer to its location, if necessary we can move other (less privileged) clients to other (cheaper) datacenters. Or if a malfunction is detected in the data center infrastructure, it is necessary to reallocate service functions to maintain service availability at contracted levels.

The only way to manage this incredible flexibility is to control it at the OSS layer, were we can have the complete vision of the network, and of course this will put pressure on the OSS ecosystem to act in real time. Real time is not a new thing to OSS. We have already experimented and deployed several OSS based real time solutions, supporting sell reconfiguration use cases in support of network self-healing and dynamic load management scenarios. The basic principle here is to collect in near real time events from the network, apply analytics algorithms to identify the need for reconfiguration and when needed to request the reconfiguration and/or reallocation of service functions and changes in the service chaining's over several distinct technical domains. As illustrated in Figure 9, those are the requirements to be supported in near real time by a collaborative work of the new NFV/SDN enabled Service Supervision and Service Provision platforms.

Another time related matter that must be taken into account, is the time (and effort) an operator has to spend when it wants to add a new service to its portfolio. NFV is expected to increase tremendously the pace at which new services are created and deployed. We must make sure that OSS layer is not a constraint. For this, an new NFV/SDN enabled Service Catalog&Inventory platform must be able to map customer facing services into the supporting technical capabilities, being either physical network functions, virtual network functions or even an IT function, able define the functions chaining and able to define the appropriate policies (e.g., for self-reconfiguration, self-healing, …) to apply.

In this context, the new NFV enabled OSS must be a bridge between the legacy network modeling using TMForum SID, using CFSs and RFSs, and describing services in a very high level way, and the world of NVF and SDN talking NETCONF, YANG and OpenFlow, with very concrete and detailed service definitions. Those are the challenges for the NFV/SDN enabled Service Catalog&Inventory.

All this potential for creating new services, and to shape them to meet the needs of every individual customer is pointless if we do not provide a fast, easy, and user-friendly channel to the customer to subscribe and to configure its services. Self-care portals provided by network or datacenter solution providers are usually not the best way to achieve that. They are confined to one technological domain, and are usually not multi-platform and difficult to integrate with corporative customer portals. Integrating their configuration APIs directly represents the same problem. Once again, an operator can take advantage of the abstraction provided by the OSS mediation layer, allowing it to configure services that are supported on several network and data center domains. These processes can be built by integrating with the BSS layer for operations that have impact on the customer account, and directly on the OSS for operations that only change network configurations that are not relevant for business processes, allowing flexibility, real time response, and agility on supporting new services.

4.2. Network Control Functions

Network Control Solution and Service Platforms can leverage immensely from this NFV/SDN movement as they will become more suitable to answer the challenges raised by nowadays telecommunications market, namely: (a) ever increasing capacity needs mainly due to the ever growth

in data traffic (by end-user or machine-to-machine applications) and (b) better time-to-market, always important for any operator that whishes to differentiate from the competition.

Nevertheless, the traditional approach to this kind of systems does not immediately fit into the requirements brought by the NFV/SDN architectures. Thus, the main challenges are:

1. Virtualization
2. Modular Architecture
3. Elasticity

With these 3 main characteristics, platforms are ready to cope with the novel architectural model we envision has a clear focus on platform automation and agility. By platform automation we understand the capacity to set up an instance of a platform for a particular customer, without any human intervention, and in short period of time (few minutes). By agility, we understand the capability of a platform to be flexible, growing when an increasing performance is required, and shrinking when the current performance is not needed anymore, in order to free resources to other purposes (again, in a few minutes).

According to our view, a platform must appear to the outside world always the same way. Changes on the level of performance must occur on-the-fly (with the platform running) and transparently to external entities (e.g. customers), without any service disruption.

In this model, the scaling procedures are fully managed internally by the platform, simplifying the 3rd entities that interact with the platform. Other solutions move the scaling burden to 3rd parties, which must accommodate the dynamicity on the number of platform components.

In order to fit with the agility requirement, the internal platforms' architecture must be organized in a way it facilitates the scaling changes. Figure 10 depicts this model.

In this architecture, the platform is organized into multiple layers of functionality, which together build up the overall service or control functionality. This model is very well-known from the past, and each layer is commonly designated as tier. Platforms may have multiple tiers, depending on his particular issues. The main concept is that each tier can be scaled-in or out, theoretically, till infinite, according to the performance requirements for the platform at any moment. This model provides a very high flexibility to make platforms growth and shrink.

Although the number of tiers is variable, 3-tiered architectures are very common, fitting many times to (a) front-end, (b) core function and (c) database tiers; however, again, this is not fixed. In that case, it is the core layer, typically the heaviest one, which is more suitable to be targeted for scaling. Of course, the others may also be scaled if needed.

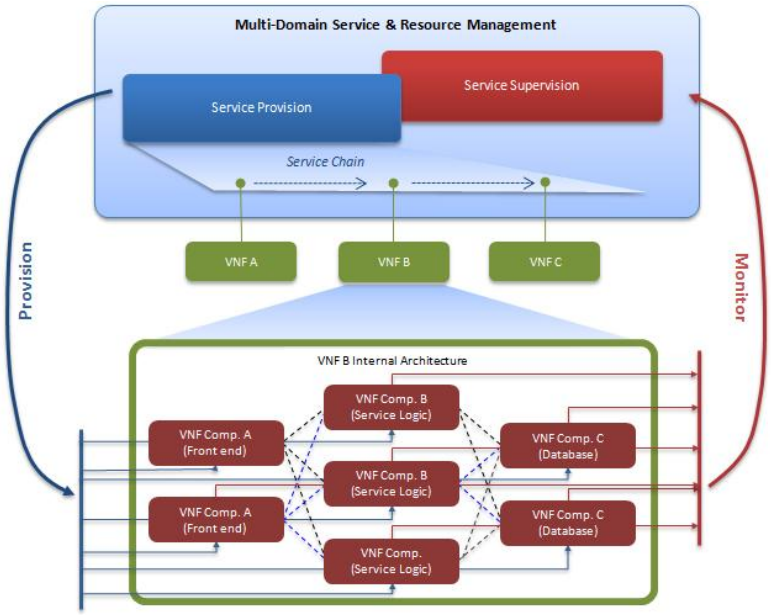

Figure 10: Virtualization Architecture

Anytime a new component is added (scale-out), in any tier, the new components must be deployed and configured. Also the surrounding environment (tiers) may need to be reconfigured, in order to adapt to the new dynamic platform architecture. This job is performed by the management/orchestration layer, which decides whether a scale-in or scale-out must take place, based essentially on the monitoring of the platform (service or infrastructure metrics).

To easy the management/orchestration task, it is important that tier components can be as stateless as possible. This way, especially on scale-in procedures, there is no need to migrate the state of the components to be removed, a task that could be really hard to implement.

At Altice Labs, we have products that work in this functional area namely: ip-Raft (PCRF and ANDSF); ipTiller (AAA); O2CS (OCS); WMS (IVR and MRF); and SEC (Enterprise Convergent Service). As mentioned before, all these systems benefit from NFV/SDN architectures and are currently being evolved towards this new network paradigm.

We recognize several advantages and opportunities for those systems:

- ip-Raft and ip-Tiller – These products (PCRF, ANDSF and AAA) can benefit from the envisioned architecture as they will be able to more easily cope with the needs to increase capacity as data usage continues to grow; improve the time-to-market for new and more complex configurations (e.g.: sponsor connectivity, IoT, Wi-Fi offload, etc.) and reduce the operational costs of the overall solution. Besides playing its usual role controlling individual accesses, the PCRF and AAA can also be used in the control of the new Data Center ecosystem.
- O2CS – As the previous set of functions, the OCS does also provide a better adaptability for increasing capacity needs; better time-to-market and reduced operational costs when ready for this new environment. The OCS can also play an important role in monetizing the usage of the new Data Centers.
- WMS and SEC – On the service platform side of this area, solutions (IVR, Enterprise Convergent Service) will benefit from reduced operational costs and agility to cope with increased demand. This evolution also provides these solutions a better suit for them to be offered through SaaS models.

4.3. Network Access Functions

The Altice Labs GPON portfolio presents the broadest and most scalable solution on the market today offering network service providers a flexible and cost effective approach for passive optical network architectures. It is able to serve multiplay services over fiber access in terms of retail as well as wholesale clients.

Initially based on compliance with the ITU-T G.984.x (GPON) recommendation, the Altice Labs solution is nowadays evolving towards next generation PON architectures as defined by the ITU-T G.987.x (XG-PON1), ITU-T G.988.x (OMCI) and ITU-T G.989.x (NG-PON2) recommendations, both for the central office and the customer premises equipments. The ONT equipment solutions are BBF.247 certified for multi-vendor OLT interoperability.

Provisioning of broadband services on the Altice Labs GPON solution can be made from an Element Management System using a SNMP or XML interface, or it can be pulled from an Auto Configuration Server (ACS) through the TR-069 ONT Management and Control Interface (OMCI).

In the most common configuration of GPON, the OLT performs advanced layer 2 switching features, including multicast IGMP snooping/proxying and DHCP relaying, while the ONT supports voice, video and data services, with fixed and wireless interfaces and built-in routing features.

In order to allow for a rapid deployment of services, layer 2 tunneling features are supported so that the customer broadcast domain can be extended to towards the operator core network along with the CPE intelligence, enabling operational cost reductions due to the usage of a simpler CPE configuration with an increased lifecycle.

The figure below shows the access network architecture for a GPON scenario where tunneling is used for the centralization and virtualization of layer 3 network functions.

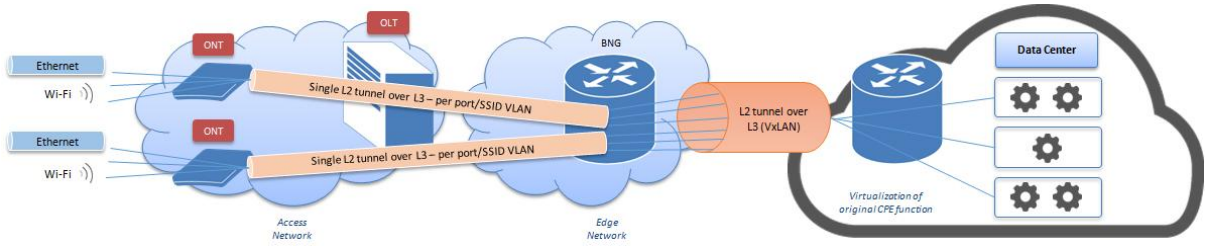

Figure 11: GPON-based Network Access Functions

The figure does not include the service provider support networks and nodes, responsible for the network establishment and management, such as the DCN network and DHCP server and ACS nodes. The objective is to highlight that the ONT equipments encapsulate the customer packets from all physical (Ethernet) and logical (Wi-Fi SSID) interfaces into a single layer 2 tunnel which is transported through the service provider L2/L3 transport network.

The customer interfaces can be identified in the tunnel by VLAN tags inserted by the ONT. It is possible to envision that RADIUS attributes used in customer 802.1x authentication may be used to establish a finer grain of VLAN tagging by end customer and not only by ONT interface. The ONT performs the role of 802.1x Port Authentication Entity (PAE), while the OLT is completely transparent to the tunnel contents.

Technologies used in tunneling can be soft-GRE, currently already supported in the Altice Labs ONT equipments, or VxLAN in the future, to identify a higher number of flows than the 4096 supported by VLAN tagging.

The tunneled traffic arrives at the data center carrying VLAN or VxLAN tag information that allows for the typical ONT layer 3 functions such as NAT to be transferred from the ONT to a virtualization infrastructure.

This architecture enables the service provider to develop and roll out innovative services with a fast time to market because now services became independent of the access network node feature support.

The whole GPON configuration can be pulled off from the ACS, in what constitutes a simple and scalable provisioning process. GPON provisioning will also not depend on the services supported by virtualization, reducing the operational burden.

5.vHGW Use-case

The virtualization of functions that are hosted today within costumers' premises can be of extreme value, both to operators and consumers. From the operators' perspective it can considerably reduce operational costs (OPEX), including the introduction of new services and their maintenance, due to the massive and distributed deployment. The operational gains can also affect the consumer as he can easily outsource the operation of these functions. Moreover, gains in terms of capital expenditure (CAPEX) can also be realized as traditional costumer equipment, such as IP routers, is amongst the most capital-intensive portions of service provider infrastructure.

On the enterprise market segment we are referring to functions such as: Enterprise Customer Premises Equipment (CPE) / Enterprise Access Router (AR); Enterprise Firewall WAN Optimization Controller (WOC); Security appliances (e.g. Intrusion Prevention System). On the residential market segment we are referring to the home environment that typically includes the Home Gateway (HGW) and the Set Top Box (STB). Both types of devices host a set of network functions, such as:

- HGW - DHCP server, NAT, PPPoE client, Firewall, Parental control, port mapping, VPN Server, etc;
- STB - DLNA media server, Media cache, VOD client, etc.

The use-case herein described, illustrated in Figure 12, focus on the virtualization of the HGW (vHGW).

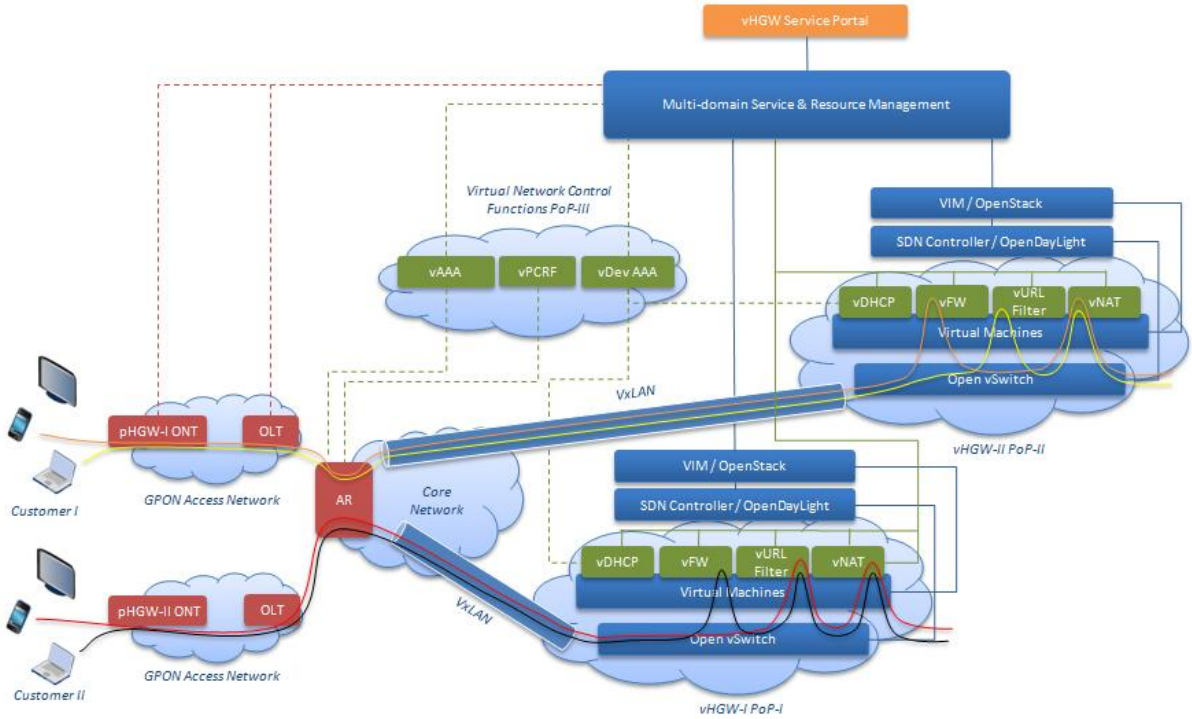

Figure 12: vHGW Use-case Architecture

Since this use-case impacts multiple segments of the telco, the described architecture is split in the following domains: i. customer, ii. access, iii. core, iv. DC PoPs and v. OSSs.

5.1. Customer, Access and Core Network Segments

On the *customer* side, there is the home network access technologies (e.g. Wi-Fi and Ethernet), as well as the end-user devices (e.g. smartphone, laptop, desktop PC and TV).

Interconnecting the customer premises and the telco *access network* is the physical HGW (pHGW). This element resembles a L2 network bridge, containing the required basic network functions to support the home environment. Additionally, it also includes the GPON Optical Network Termination (ONT), interconnected to the GPON Optical Line Terminal (OLT) aggregator.

GPON-based access networks are connected to the telco *core network* through an Access Router (AR), which splits the switching from the routing domain. The latter provides access to the telco *DC PoPs-grid* through a VxLAN-based tunneling mechanism to extend the customers broadcast domain towards the DC PoP.

5.2. vHGW Points-of-Presence

In this case, assuming that customer I and II premises are geographically spread, the customers vHGWs are deployed in different DC (*vHGW-I PoP-I* and *vHGW-II PoP-II*). These PoPs include the HGW VNFs that are moved from the pHGW to the vHGW, namely, virtual Firewall – vFW, virtual URL Filtering / Parental Control – vURL Filter, virtual DHCP – vDHCP and virtual NAT – vNAT. The role of the vDHCP within the DC PoP is to assign and manage IP addresses to the home devices.

Having the HGW network functions virtualized enables the selection of specific services to each home device and user (e.g. father laptop, son smartphone, etc.). This is achieved by chaining a specific set of VNFs, in what is commonly known in the ETSI NFV terminology, as a service chain. Data packets are routed through the VNFs that compose the service chain by Open vSwitch. The PoP Virtualized Infrastructure Management (VIM) is realized through OpenStack. In what concerns the PoP network resources programmability, an SDN Controller, based on OpenDayLight, is in place.

5.3. vNetwork Control Functions Point-of-Presence

The virtual network control functions from the telco are deployed in another DC PoP, represented in Figure 12 as *Virtual Network Control Functions PoP-III*.

The vAAA is a RADIUS-based authentication, authorization and accounting server for the pHGW. It is responsible for notifying the multi-domain Service & Resource Manager (SRM) when the residential user connects the pHGW. This will trigger the vHGW virtual environment provisioning (launch VMs and VNFs) in the PoP, as well as the creation of the VxLAN tunnel towards the DC PoP.

The virtualized Device AAA – vDev AAA, authenticates/authorizes customer devices joining the network (triggered by the vDHCP), and persists all the information related with these (per HGW). It is also the entity that notifies the SRM that a new device is attached and a service chain must be configured in the PoP Open vSwitch.

The virtualized Policy and Charging Rules Function (vPCRF) represented in Figure 12, is used to limit and control the HGWs QoS, based on a set of rules configured during the product acquisition.

5.4. Operational Support Systems (OSSs)

The SRM is responsible for orchestrating all the entities in the use-case. It manages the service chains lifecycle procedures (on-boarding, instantiation, monitoring, scaling and teardown) by orchestrating:

- Physical network elements (pHGW ONT, OLT, AR);
- Virtualized network functions from the vHGW PoPs (vDHCP, vFW, vURL Filter and vNAT);
- Virtualized network functions from the Network Control Functions PoP (vHGW AAA, vDev AAA and vPCRF);
- PoPs virtualized infrastructure managers (VIMs) – compute and network.

In detail, when a new pHGW is connected (detected through the vAAA), the SRM instantiates the computational environment in the PoP through OpenStack, installs and configures the vHGW VNFs and enforces the service chains in Open vSwitch through OpenStack/OpenDayLight. It also orchestrates the re-allocation of service chains to devices, triggered by the vHGW Service Portal, as illustrated in Figure 12. The web portal is used by the residential users to manage the devices, as well as to associate a service chain to each device.

6.Conclusions

More than a trend or a hype, NFV and SDN paradigms will for sure be adopted by the telecommunications industry. We can go further and say that NFV and SDN adoption will be major transformation forces, driving the evolution of traditional Communication Service Providers to a new generation of Digital Service Providers, emerging over a new generation of Service Platforms.

The new generation Service Platforms will have in its core a “network of Data Centers”, where the vast majority of service functionality will reside. Being the access networks a valuable asset inherited by the Telco operators, this new generation Service Platform will be complemented with managed access connectivity being this a considerable differentiating factor when comparing with Internet players.

The new generation Service Platforms will boost the speed, agility and flexibility for new digital services creation and, at the same time, will furnish the means to dynamically control, in near real time, the configuration and topology of the service functions, managing effectively service loads, service quality levels and service affecting anomalies. In the end, the ultimate purpose is to achieve unmatched levels of operational efficiency and augment significantly the end customer experience.

From the Operations Support point of view, the adoption of NFV and SDN will bring a new set of requirements, related mainly with new needs for dynamicity and real time management of new service functions and sustaining infrastructure. Those requirements do not demand new process domains but will impose significant changes in the current Operations Support, Fulfillment and Assurance processes. Traditional OSS suppliers must accommodate those requirements in the roadmap of its OSS platforms or they will become obsolete and unable to fit in the new Digital Services Generation.

In the past years the communications industry over discussed about “Becoming a Dumb Pipe” vs. “Competing with the Internet Players (OTT’s)”. This subject is no longer relevant. The question is not anymore about becoming a connectivity provider in the lower end of the value chain, is about evolving of disappearing. Ignoring NFV/SDN is a fast track for being out of business. Adopting NFV/SDN may be the way to enter the Digital Services game, a game mastered by the Internet giants.

7.References

[1] NIST, “The NIST Definition of Cloud Computing,” National Institute of Standards and Technology, 2011.

[2] European Telecommunications Standards Institute (ETSI), “Network Functions Virtualisation, An Introduction, Benefits, Enablers, Challenges & Call for Action. Issue I,” 2012.

[3] European Telecommunications Standards Institute (ETSI), “Network Functions Virtualisation (NFV); Use Cases,” 2013.

[4] European Telecommunications Standards Institute (ETSI), “Network Functions Virtualisation (NFV); Architectural Framework,” 2013.

[5] Open Networking Foundation, “Software-Defined Networking: The New Norm for Networks,” April 2012.

[6] Open Networking Foundation, “Open Networking Foundation,” [Online]. Available: https://www.opennetworking.org/.

8. Acronyms

AAA	Authentication, Authorization and Accounting
ANDSF	Access Network Discovery and Selection Function
AR	Access Router
BSS	Business Support Systems
CAPEX	Capital Expenditure
CPE	Customer Premises Equipment
DC	Data Center
EPC	Evolved Packet Core
GPON	Gigabit Passive Optical Network
GRE	Generic Routing Encapsulation
IMS	IP Multimedia Subsystem
IP	Internet Protocol
ITF	Virtual IT Function
MANO	Management & Orchestration
NFV	Network Functions Virtualization
NFVI	Network Functions Virtualization Infrastructure
NFVO	Network Functions Virtualization Orchestrator
OLT	Optical Line Terminal
ONT	Optical Network Termination
OPEX	Operational Expenditure
OSS	Operations Support Systems
PCRF	Policy & Charging Rules Function
PNF	Physical Network Function
PoP	Point of Presence
SDN	Software-Defined Networking
VIM	Virtualized Infrastructure Management
vHGW	Virtual Home Gateway
VNFM	Virtual Network Functions Manager
SRM	Multi-domain Service & Resource Manager
STB	Set Top Box
VLAN	Virtual Local Area Network
VNF	Virtual Network Function
VxLAN	Virtual Extensible Local Area Network
WOC	WAN Optimization Controller

Rua Eng. José Ferreira Pinto Basto
3810-106 Aveiro
Portugal

Tel.: +351 234 403 200
Fax: +351 234 424 723

www.alticelabs.com