Toward a Telco Cloud Environment for Service Functions

João Soares, Carlos Gonçalves, Bruno Parreira, Paulo Tavares, Jorge Carapinha, João Paulo Barraca, Rui L. Aguiar, and Susana Sargento

ABSTRACT

Deploying service functions, SFs, is an essential action for a network provider. However, the action of creating, modifying and removing network SFs is traditionally very costly in time and effort, requiring the acquisition and placement of specialized hardware devices and their interconnection. Fortunately, the emergence of concepts like cloud computing, SDN, and ultimately NFV is expected to raise new possibilities for the management of SFs with a positive impact in terms of agility and cost. From a telco viewpoint these concepts can help to both reduce OPEX and open the door to new business opportunities. In this article, we identify how telcos can benefit from the abovementioned paradigms, and explore some of the aspects that still need to be addressed in the NFV domain. We focus on two major aspects: enabling telco infrastructures to adopt this new paradigm, and orchestrating and managing SFs toward telco-ready cloud infrastructures. The technologies we describe enable a telco to deploy and manage SFs in a distributed cloud infrastructure. In this context, the Cloud4NFV platform is presented. Special attention is given to the way SFs are modeled toward cloud infrastructure resources. In addition, we explore the ability to perform service function chaining as one of the fundamental features in the composition of SFs. Finally, we describe a proof of concept that demonstrates how a telco can benefit from the described technologies.

INTRODUCTION

The emergence of the cloud concept, its ongoing evolution, and the opportunities it brings have led many businesses to adapt in order to get the most utility out of it. One can say that the telco sector is today one of the most active business sectors exploring the opportunities offered by the cloud. The relationship and interdependence between clouds and telecommunications can be analyzed from two distinct perspectives.

Telcos *supporting* the cloud: In a cloud environment, communication endpoints are user devices and virtual machines (VMs) that can be hosted in different physical locations according to varying conditions. Compared to traditional networking environments, network capacity requirements are no longer static, but are likely to change as the associated computing and storage resources expand and reduce. This poses a whole new set of challenges to the network, now jointly including the data center (DC) and the wide area network (WAN) segments. To provide assured levels of performance to cloud services, cloud and telco services need to be provisioned, managed, controlled, and monitored in an integrated way.

Telcos *using* the cloud: Today, the establishment, management, and composition of service functions (SFs) (e.g., router, firewall) follow a rigid, static, and time consuming process. For example, resource overprovisioning is usually necessary to cope with estimated peak demand; and a fault in a single function can disrupt an entire network, imposing the need for faster disaster recovery methods. As virtualization technologies reach maturity and are able to provide carrier-grade performance and reliability, it becomes feasible to consolidate multiple network equipment types, traditionally running on specialized hardware platforms, onto industry standard hardware, which minimizes costs, reduces time to market, and facilitates open innovation. Cloud computing, combined with software defined networking (SDN) [1] and network function virtualization (NFV) [2], promises to make SF management processes much more agile. Cloud computing represents a paradigm for information technology (IT) services, which can now be delivered in an on-demand and self-service manner. SDN brings new capabilities in terms of network automation, programmability, and agility that facilitate integration with the cloud. On the other hand, NFV, from a high-level perspective, accelerates the innovation of networks and services, allowing new operational approaches, novel services, faster service deployment (shorter time to market), increased service assurance, and stronger security.

Conceptually, an SF is a functional block responsible for a specific treatment of received packets and has well defined external interfaces [3]. An SF can be embedded in a virtual instance or directly in a physical element (the usual situation until recently). Virtual SFs offer the opportunity to compose and organize virtual SFs dynamically, opening a new set of business opportunities — and technical challenges. One of the topics that arise from the combination of

João Soares is with Ericsson Research.

Carlos Gonçalves is with NEC Laboratories.

Bruno Parreira, Paulo Teveres, João Paulo Barraca, Rui L. Aguiar, and Susana Sargento are with Instituto de Telecomunicações.

Jorge Carapinha is with Portugal Telecom Inovação e Sistemas.

João Soares and Carlos Gonçalves were with Instituto de Telecomunicações at the time this work was carried out.

0163-6804/15/$25.00 © 2015 IEEE

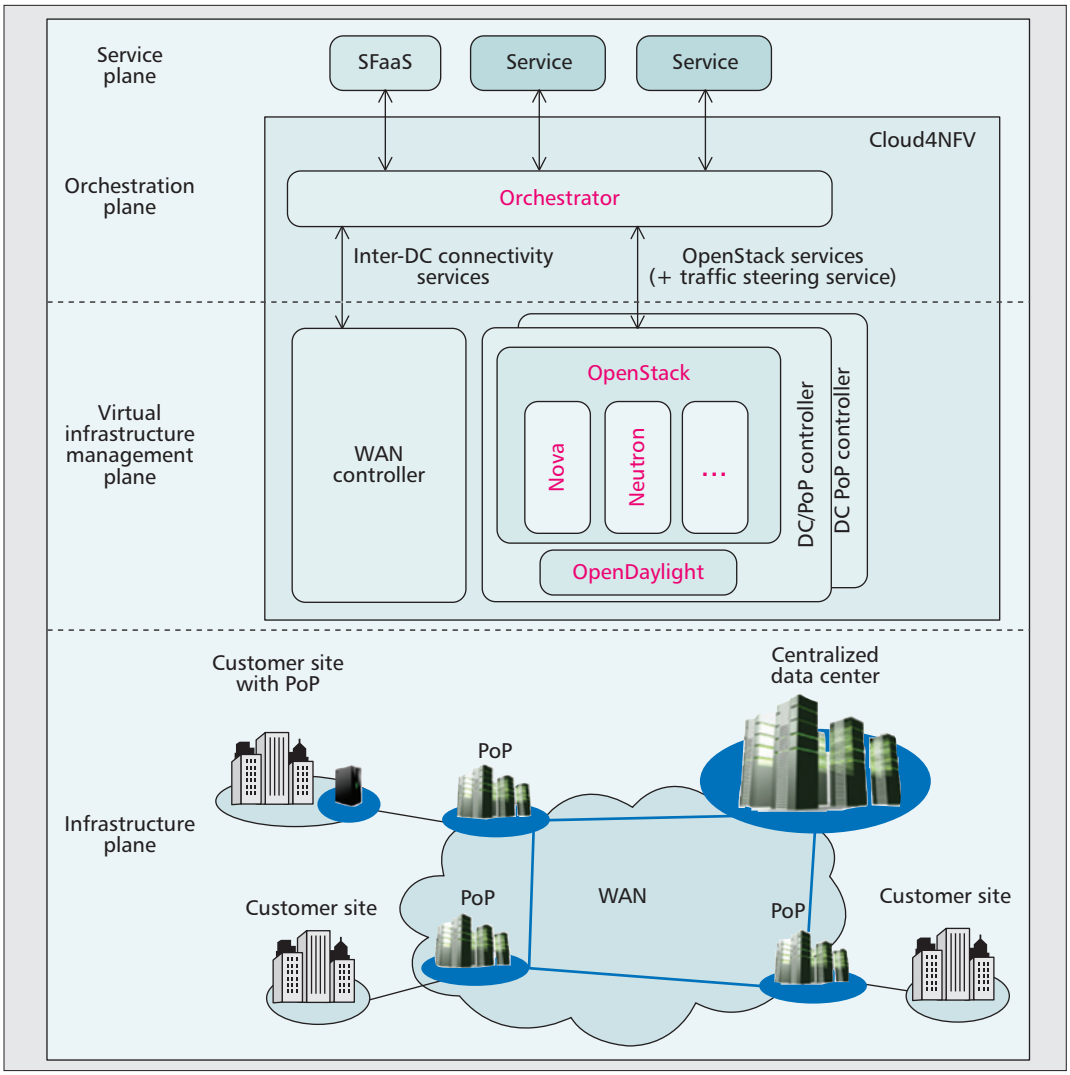

Figure 1. Cloud4NFV platform — overview.

Telcos, with their already established distributed network infrastructure and hosting centers, are ideally positioned to take the lead in this area, as they can easily create a compelling end-to-end cloud proposition that integrates their network management capabilities, adapted to a more agile and cloud service-oriented operation model.

SFs is SF chaining (SFC). SFC is loosely defined as "an ordered set of service functions that must be applied to packets and/or frames selected as a result of classification" [3]. It can be considered as a particular case of service composition. It requires the placement of SFs and the adaptation of traffic forwarding policies of the underlying network to steer packets through an ordered chain of service components. However, the lack of automatic configuration and customization capabilities increases the operational complexity.

In this article, we explore how telecom operators can take advantage of the above concepts to improve the management of SFs and potentially build new business models. First, we highlight the telcos' privileged position in this area compared to traditional cloud providers. We then present Cloud4NFV, a platform for managing SFs in a telco cloud environment. Later, we focus on SF modeling toward cloud infrastructure resources. Special attention is given to the ability to perform SFC. To emphasize possible application scenarios of the solution presented in this article, a proof of concept (POC) is then detailed. Finally, we point out future work directions and conclusions.

The Carrier Cloud Opportunity

Traditional cloud infrastructures are far from suitable for all types of businesses, especially when referring to network SFs. Most network SFs have carrier grade requirements, from guaranteed quality of service (QoS) in terms of IT resources and network connectivity, to high availability (e.g. perform detection and forecast of operational anomalies, support fault mitigation procedures such as VM migration and network replanning) and fast fault recovery through redundancy.

Telcos, with their already established distributed network infrastructure and hosting centers, are ideally positioned to take the lead in this area, as they can easily create a compelling end-to-end cloud proposition that integrates their network management capabilities, adapted to a more agile and cloud service-oriented operation model (on-demand, self-service, elastic).

We envision a near-future telco cloud infrastructure that comprises not only the traditional centralized DC domains, but also the WAN domain. In such a scenario, the telco can take

The Cloud4NFV platform builds upon Cloud, SDN and WAN technologies to allow SFs to be managed on an as-a-Service basis. The platform is targeted for Telcos to improve the management of SFs within their environment, but can also be used to build new services based on the concept of SFaaS.

	Cloud4NFV	UNIFY [4]	T-NOVA [5]	CloNe [6]	StEERING [7]	CloudBand
Distributed infrastructure	Yes	Yes (conceptually)	Yes (conceptually)	Yes	No	Yes
End-to-end service management	Yes	Yes (conceptually)	Yes (conceptually)	Yes	No	Yes
Multi-domain architecture	Yes	Yes (conceptually)	Yes (conceptually)	Yes	No	Yes
Network QoS support	Yes	Yes (conceptually)	Yes (conceptually)	Yes	No	Yes (partially)
SF management	Yes	Yes (conceptually)	Yes (conceptually)	No	Yes (partially)	Yes
Traffic steering support	Yes	Yes (conceptually)	No	No	Yes	No
SFC support	Yes	Yes (conceptually)	No	No	Yes (partially)	No

Table 1. Summary of existing approaches.

advantage of its already established distributed facilities (sometimes referred as points of presence, PoPs) to host small cloud environments. It is also possible for this distributed cloud infrastructure to extend itself into the customer site. Figure 1 depicts this scenario.

Although there are important contributions ongoing in this area, work is still required, namely when it comes to the definition of a true telco cloud platform and the details of how to model and actually realize SFCs. Table 1 presents a summary of the features supported by some existing solutions that more closely relate to the scope of this work. The information presented reflects the publically available information at the time of writing, and may meanwhile have been superseded. The recent UNIFY [4] and T-NOVA [5] projects seem to share a similar vision; however, these projects have recently started and have only provided conceptual approaches to some extent. CloNe [6] has support for the infrastructure features; however, it lacks SF management, traffic steering, and SFC. StEERING [7] supports traffic steering, and partially supports SFC and SF management ("partially" because the SFC service model and SF management features do not seem fully mature). Finally, we also consider the Alcatel Lucent CloudBand[1] solution, which supports some of the envisioned features.

Cloud4NFV Platform

The Cloud4NFV platform builds on cloud, SDN, and WAN technologies to allow SFs to be managed on an as-a-service basis. The platform is targeted for telcos to improve the management of SFs within their environment, but can also be used to build new services based on the concept of SF-as-a-Service (SFaaS), in which case SFs or bundles containing a combination of SFs can be offered as a service to customers.

Functionalities

The most relevant functionalities of Cloud4NFV are:

- Automated deployment, configuration, and life cycle management (instantiation, configuration, update, scale up/down, termination, etc.) of SFs
- Exposure of functionalities such as service deployment and provisioning, service monitoring and reconfiguration, and service teardown
- Federated management and optimization of WAN and cloud resources for accommodating SFs
- Support of SF composition through SFC

All the above mentioned functionalities are essential in the scope of an NFV platform; however, we highlight the last two due to their novelty. These two functionalities are seen as key differentiation factors from other available solutions, taking this platform closer to being fully carrier-grade compliant. The federated management and optimization of WAN and cloud resources gives the platform a broad and distributed scope. It allows the establishment of end-to-end services over a distributed physical infrastructure. The ability to perform SFC gives the platform unprecedented flexibility with respect to SF management and composition, allowing the definition and establishment of advanced services in a much more efficient and flexible way.

Architecture

Figure 1 provides an overview of the system, organized in four major planes: infrastructure plane, virtual infrastructure management (VIM) plane, orchestration plane, and service plane. The service plane handles the services that are built on Cloud4NFV, and the infrastructure plane comprises all physical resources. Special attention should be given to the VIM and

[1] CloudBand, http://www.alcatel-lucent.com/solutions/cloudband.

orchestration planes, since we consider them to be the major lever for enabling SFC. It is important to note that this architecture is aligned with the ETSI NFV architectural guidelines [8]. This fact is highlighted along the description of the platform.

Orchestrator

The orchestrator is responsible for the automated provision, management, and monitoring of SFs over the virtual infrastructure. It exposes the ability to create and delete SFs, as well as the ability to chain SFs. It relies on the VIM plane to provision the infrastructure resources where SFs run (VMs, virtual networks, etc.). Looking to the ETSI NFV reference architectural framework [8], this component considers the *orchestrator* and *VNF manager(s)* entities. The orchestrator has an interface (REST) that exposes the ability to create and delete SFs as well as to chain SFs.

Virtual Infrastructure Management Plane

The VIM plane includes the components for management of infrastructure resources. It includes cloud DC controllers (one per DC) and a WAN controller that is able to establish inter-DC connectivity services. The VIM plane can be seen as the *virtual infrastructure manager(s)* in the European Telecommunications Standards Institute (ETSI) NFV reference architectural framework [8]. However, the current ETSI specification does not take into consideration the WAN component. This is considered by ETSI to be the subject of future analysis.

Data Center Controller(s) — Although the cloud model may require, to a large extent, the redefinition of SFs and the way they are managed, SFs also require adaptation from today's cloud solutions to cope with their requirements, especially in terms of networking features. A clear evidence of this fact is the OpenStack[2] project, a reference open-source cloud management platform, which has been witnessing a tremendous evolution of its networking features in its networking project mostly known by the codename, *Neutron*. It is also important to note that Neutron provides network service logics, and relies on different backends called *drivers* to interact with different networking technologies. Among these drivers is the recent OpenDaylight[3] SDN controller. OpenDaylight is today seen as an initiative equivalent to OpenStack in the SDN domain. With this in mind, our DC controller is based on OpenStack and OpenDaylight.

OpenStack — From a networking perspective, OpenStack allows the creation and management of *networks* (L2 network segments) and *ports* (attachment points for devices connecting to networks, e.g., virtual network interface cards, vNICs, in VMs). The OpenStack community has been making a considerable effort to keep up with users' demands by introducing new Neutron network service types: L3 routing, firewall as a service (FWaaS), load balancer as a service (LBaaS), and VPN as a service (VPNaaS); however, it is infeasible (and probably unwise) in the long run to keep up with demands at this pace in a timely manner. Therefore, we argue that OpenStack should focus on offering the basic tools for network services to be orchestrated at a higher level and be deployed as VMs.

With the orchestration and composition of SFs in mind, it is easy to identify the need to fill a gap in OpenStack: steering traffic between OpenStack elements (e.g. VMs, routers). We envision a new OpenStack service abstraction that extends and relies on current OpenStack networking features, allowing traffic steering between *Neutron ports* according to classification criteria. New entities are introduced into the OpenStack Neutron data model: *port steering* and *classifier*. Both entities have a set of common OpenStack data model attributes (i.e., *id, name*, *description*, and *tenant_id*). Port steering adds to this common set a list of ports (ports attribute) and a list of classifiers (*classifiers* attribute). The former lists the sets of ports that must be targeted for classification and then steered. The classifier entity adds the following attributes: *type*, *protocol*, *port_min*, *port_max*, *src_ip* and *dst_ip*.

This functionality is very useful as it provides the means to realize, among other things, SFC, as described later. Furthermore, the primitive is seen as a foundation for future (higher-level) abstractions within OpenStack.

OpenDaylight — OpenDaylight has a module that integrates with OpenStack Neutron for the enforcement of services in the infrastructure. This module was extended in order to support and enforce the previously mentioned OpenStack traffic steering feature. It is important to highlight that this implementation relies on OpenFlow and Open vSwitch Database Management Protocol (OVSDB) for the management of network resources.

Wide Area Network Controller — The WAN controller is responsible for managing the operator network, and it exposes connectivity services to the upper layers (in this case the orchestrator). In this context, WAN services are used to support SFs (the SF is the client of the WAN service). Point-to-point and multipoint connections with guaranteed network QoS are provided. These are exposed through a service interface that, similar to cloud IaaS interfaces, is technology-agnostic. The details and mechanisms to manage the automatic establishment of connectivity services across different locations are detailed in [6].

The ability to perform SFC gives the platform unprecedented flexibility with respect to SF management and composition, allowing the definition and establishment of advanced services in a much more efficient and flexible way.

Service Function Virtualization

This section elaborates on how SFs are modeled toward virtual infrastructure resources. Figure 2 depicts the correspondent data model, and each class is detailed below.

Service function: represents an instance of a functional block responsible for a specific treatment of received packets.

Service function endpoint (SFE): represents an external interface of one SF instance that is always associated with an SF. Each SFE can have associated information regarding layer 1

[2] OpenStack, http://www.openstack.org/

[3] OpenDaylight, http://www.opendaylight.org/

(e.g., physical/virtual interface), layer 2 (e.g., medium access control, MAC, address), and/or layer 3 (e.g., IP address), or even regarding higher layers (e.g., HTTP).

From an infrastructure perspective, the resources considered to realize a SF are: *compute instance* (i.e., virtual or physical machines), *image* (disk image), *compute flavor* (hardware specification of a compute instance, i.e., CPU, memory, and root disk), *block storage* (additional disks), *port* (i.e., network interface), *network* (a network segment), and *link* (a connection between two ports from different compute instances), which has an associated *link flavor* (dedicated QoS in terms of bandwidth, delay, and jitter). An SF can be associated with multiple compute instances, while each compute instance has a single image and a single flavor, and can have multiple ports and block storages. A port can only be associated with a single network; however, it can be associated with multiple links. An SFE is directly associated with a port, but not all ports need to map to SFEs.

The network QoS, represented in the model by link and link flavor, is not considered in today's cloud infrastructure systems. However, for a carrier grade cloud this is a must, and OpenStack already has an ongoing project to support it.[4]

Figure 3 presents an example of how several SFs can be composed and organized. Furthermore, it also highlights how SFCs can be built and explored.

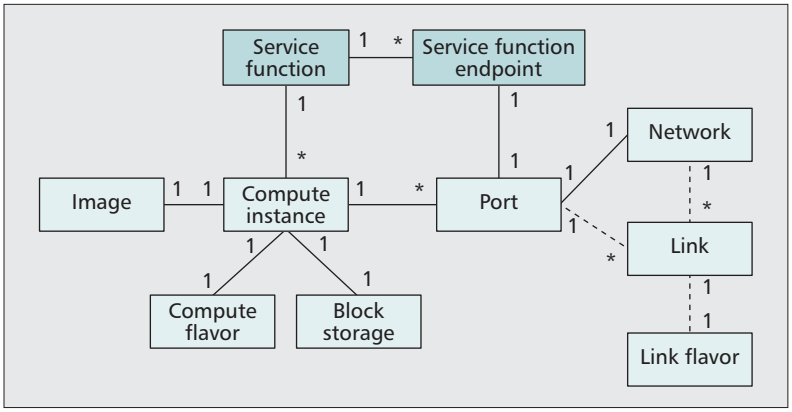

Figure 2. Service function data model toward a cloud infrastructure.

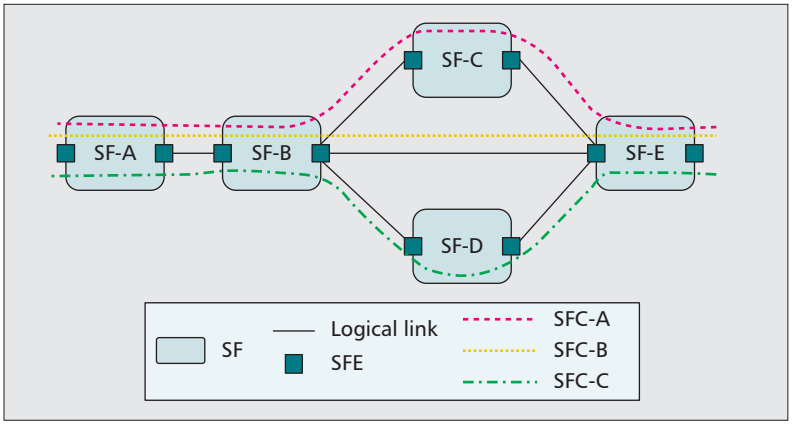

Figure 3. Service function composition — example.

Service Function Chaining

In this section we provide insights on the fundamentals and modeling aspects of SFC proposed in this work.

Fundamentals

In SFC two aspects are vital:

Classification: a policy for matching packets (e.g., HTTP traffic) used for the identification of appropriate actions (e.g., forwarding). It can be, for example, an explicit forwarding entry in a network device that forwards packets with a specific IP or MAC address into the SFC. (Re)Classification can also occur at each SF of the SFC independent from the previous SFs. In such cases, multiple classification policy entries should be allowed in an SFC system.

Traffic steering: the ability to manipulate the traffic route at the granularity of subscriber and traffic types [7]. The actual network topology should not be modified to accomplish this.

Moreover, the combination of classification and traffic steering can be done in two ways:

Tagged packet approach: classification can occur only at the initial redirection points to an SFC, if upon this classification packets are tagged. After that, packets are steered to the SFC and routed along it according to the embedded tags.

Non-tagged packet approach: classification occurs not only at the redirection points but also at each hop of the SFC. In this case, packets are not tagged and are subject to classification and steering at each SFC hop.

The consequences of following a tagged or non-tagged packet approach are felt at the VIM plane level. One of the benefits is that this choice is relatively well isolated from the higher planes. We believe the non-tagged approach to be the smoothest approach to follow due to its lower impact on SFs and virtual infrastructure management systems. The advantage of the tagged approach is that the traffic only needs to be classified and tagged (e.g., with a VLAN or other tag) once along the entire SFC. The drawback is that the SFs need to know how to handle the tags (in the simplest case, they should at least ignore them). Although we can add to the platform support for a tagged approach (e.g., classify only at one point, tag, and steer traffic according to the tag), it only makes sense if there is also support at the SF level. Hence, in this work we adopt the non-tagged packet approach.

Further aspects should be taken into account when elaborating an SFC solution, such as:

- No assumption should be done on how functions are deployed, that is, whether they are deployed on physical hardware, as one or more VMs, or any combination thereof.
- An SF can be part of multiple SFCs.
- An SF can be network-transport-independent.
- An SFC allows chaining of SFs that are in the same layer 3 subnet and of those that are not.
- Traffic must be forwarded without relying on the destination address of packets.
- Classification and steering policies should not need to be done by SFs themselves [10].

Service Function Categories

Two categories of SFs have been defined.

Active SFs: those that are in fact part of the main course of a packet, in which case two subtypes are considered:

- Functions that may drop packets or forward them, such as a firewall
- Functions that can actually change packets, such as an IPSec VPN server

Passive SFs: are considered to be out of the main course of the chain. These functions mainly inspect packets (e.g., a monitoring system or a deep packet inspection, DPI). In practice one can think of an SF in a physical device connected to a hub through a single network interface configured in promiscuous mode. Traffic is considered to be duplicated when having to reach a passive function.

These two categories are important because they impose constraints on how classification and steering can be implemented. In short, passive functions can rely on packet characteristics as packets are not modified, while active functions must be integrated at a service level because ingress and egress packets can be different (e.g., virtual private network, VPN). If an SFC has active functions that change packets, the classification may differ when passing one of these functions.

Service Function Chaining Abstraction Model

The ability to classify and steer traffic accordingly can be enough to implement low-level SFC functionality. However, it is important not to forget that the traffic steering functionality is a low-level functionality that does not explicitly express an SFC. Having in mind the considerations made so far, a base data model for SFC (that supports both tagged and non-tagged approaches) is now presented. Naturally, other SFC service abstraction proposals may appear in the future, but we consider that this model lays a strong foundation over which other service abstractions can easily be created by extending the model. Figure 4 depicts the model. Five main classes are considered: *service function chain*, *service function*, *service function endpoint*, *packet flow*, and *classifier*.

All classes have the following attributes: *id*, *name*, and *description*. The id refers to a unique identifier able to identify the class instance within the SFC system. The remaining two, name and description, are attributes that allow a human-readable characterization of the class instance. Below, we provide further detail about each class.

Service function chain: An SFC has a set of associated SFs and an attribute that defines the ordered sequence of functions (path). Since a function can have more than one SFE, the path attribute is specified by an ordered list of SFEs organized by hops. For example,

- "path= { hop={SF-A_E2, SF-B_E1}; hop={SF-B_E2, SF-D_E1}, passive={SF-C_E1}}"

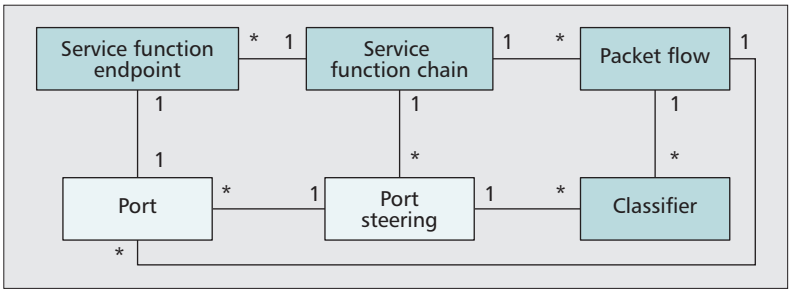

Figure 4. Service function data model toward a cloud infrastructure.

where the chain crosses SF-A, SF-B, and SF-D, and has SF-C as a passive function between SF-B and SF-D.

Classifier: A classifier represents a classification criterion applied to a packet, which determines if the packet matches that specific criterion or not. In this sense, a classifier has an attribute filter that contains the classification criteria;

- "filter={protocol='6'; port='80-90'; source_IP='192.168.10.20/32'; destination_IP='192.168.10.40/32'}" –

matches all TCP traffic using ports between 80 and 90 with source IP address 192.168.10.20 and destination IP address 192.168.10.40.

Packet flow: One *classifier* only identifies packets with a certain criterion, while a *packet flow* identifies a broader set of packets as it can aggregate packets associated with multiple classifiers. In this sense, a packet flow can have multiple classifiers, and a classifier can be associated with multiple packet flows. Moreover, a packet flow has a *source* and a *destination port*. The former identifies where the initial classification and redirection of the packet flow to the SFC takes place, while the latter identifies where packets are to be delivered after passing through the SFC. The attributes considered so far would be enough if the system realizing the SFC followed a tagged packet approach. For a non-tagged approach, an additional attribute is considered — *sfc_classifiers*. Due to the possibility of (active) SFs to modify packets, the classification initially done may not be the same along all hops of the SFC, and therefore, the *sfc_classifiers* attribute matches the classification criteria (classifiers) at each hop of the SFC.

Furthermore, the attribute *direction* is also considered to identify the direction of the SFC in which the *packet flow* must traverse; this attribute can assume one of two values: *forward* and *reverse*. We consider that multiple packet flows can be associated with a single SFC instance.

Port steering: This entity refers to the functionality presented above in OpenStack. This feature allows steering traffic between *ports*. Further details about the traffic steering functionality can be found in the OpenStack proposal,[5] for which we developed a prototype implementation.

In terms of operations, all classes are considered to allow create, read, update, and delete (CRUD) operations.

[4] OpenStack Neutron QoS support https://wiki.openstack.org/wiki/Neutron/QoS

The instantiation and configuration of SFs is done in a timescale of seconds/minutes (depending on the SF/VM and cloud infrastructure) and the SFC enforcement in a timescale of seconds. Furthermore, the user is able to control them through a dedicated SF management portal.

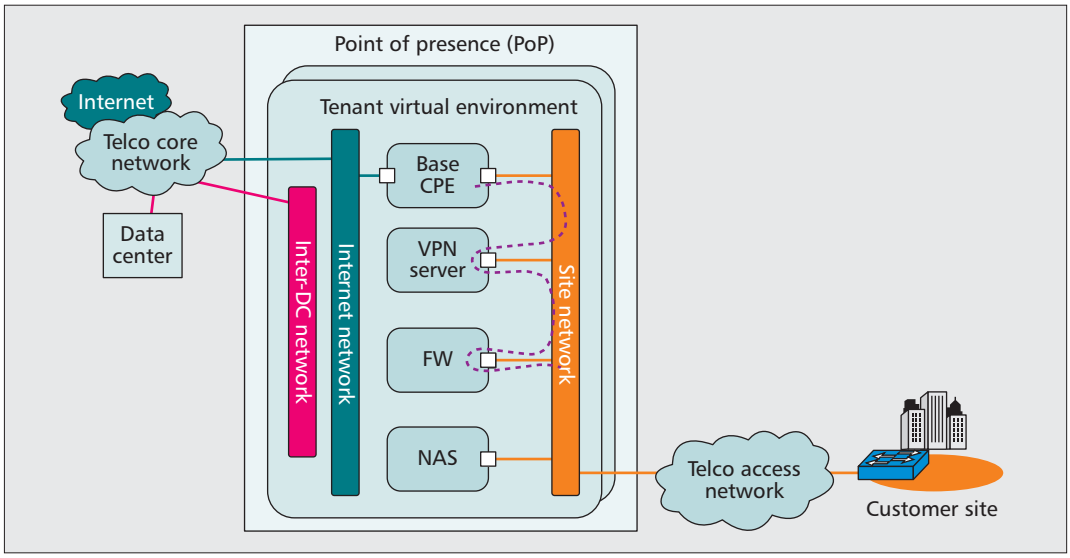

Figure 5. POC prototype setup.

Proof of Concept

A POC environment has been deployed to showcase how a telco can leverage the features described in this work. We highlight one of the most attractive use cases in the NFV scope and how it has been realized in this POC.

The testbed in place is depicted in Fig. 5, focusing on the PoP setup that is detailed later. At the core of the operator network (telco core network) there is an IP/multiprotocol label switching (MPLS) backbone composed of four provider (P) routers and four provider edge (PE) routers. The core network is managed by proprietary operations support systems (OSSs) that expose connectivity services through a service interface in a technology-agnostic manner (the WAN controller). The core network connects to two DC premises (managed by the OpenStack *IceHouse* release with traffic steering functionality), one of which represents a centralized DC and the other a PoP. Finally, the customer premises are represented by switching equipment, which is logically connected to the PoP over an access network (a simple switch-based network).

A prototype of the Cloud4NFV orchestrator, which interacts with the WAN and DC controllers, was developed using the Python language. Details regarding the orchestrator implementation (e.g., RESTful API) can be found in [11].

Customer Premises Equipment Use Case

Customer premises equipment (CPE) is often pointed out as one of the most suitable candidate SFs for virtualization [2, 12]. SFC will surely play a particularly relevant role in this case.

The CPE can be seen as a standard routing node enhanced by collection of SFs, such as network address translation (NAT), firewall (FW), voice over IP (VoIP) servers, VPN servers, network-attached storage (NAS), WAN optimization controllers (WOCs), DPI, or intrusion prevention system (IPS). These services are deployed for different scenarios, and not all traffic needs to traverse them, leaving room for optimization through SFC. It should be noted that some of the chains can even be temporary, which requires a model that enables the dynamic definition of chains.

Service Function as a Service

At the service layer we implemented a prototype of the SFaaS concept. This is exposed via a web portal. CPE functions are available in the SFaaS, and the ability to perform SFC is not exposed to the end user. The user requests CPE SFs, which already have a predetermined relation with other SFs, and associates them with one of the user's sites. The instantiation and configuration of SFs is done on a timescale of seconds/minutes (depending on the SF/VM and cloud infrastructure), and the SFC enforcement on a timescale of seconds. Furthermore, the user is able to control them through a dedicated SF management portal.

Note that the use of SFaaS requires a basic business relationship between the customer and the telco (customer sites registered and with connectivity services). In other words, the user is a customer of the telco that provides connectivity services (e.g., fiber, copper) from the client's sites (e.g., house, enterprise premises). On the site side, a layer 2 device (or a layer 3 device in bridge mode) is considered to be in place.

Currently, from a demonstration viewpoint, it is considered that the user, after having the physical connection in place, must first buy a base CPE function with routing, Dynamic Host Control Protocol (DHCP), and NAT functionalities (the POC relies on the OpenStack L3 native device). From that moment on, the user can acquire other CPE functions and services, such as Internet connection, firewall (POC relies on iptables), VPN server (POC relies on OpenVPN), and NAS (POC relies on Samba).

Prototype Setup

Figure 5 depicts the POC prototype setup with four functions as an example. Special attention is given to the setup at the PoP level.

[5] OpenStack Traffic Steering blueprint, https://review.openstack.org/#/c/92477/

Each customer has a dedicated virtual private environment in the PoP that is serving his/her site. This environment allows the creation of virtual networks and VMs (in OpenStack this is known as *tenant* or *project*). There is a point-to-point logical connection between the customer's premises (L2) device (currently we are using VLAN encapsulation to establish this connection, but others can be used). On the PoP side this logical connection is extended to a virtual network in the tenant virtual environment; in the figure, this is "site network," which has a private IP range (the OpenStack *provider network* concept is used to achieve this). Moreover, there is a virtual network shared among all tenants, which in the figure is the "Internet network" (in OpenStack this is achieved using the *external network* concept). This latter network is then connected to the core network, which provides the Internet access. Also depicted in the figure is the "inter-DC network," which provides access between the PoP and the DC over a telco VPN service in the core network (again, on the PoP side we rely on the OpenStack *provider network* concept to connect to the VPN). The processes explained so far are considered to be in place as soon as the customer establishes the basic business relationship with the telco.

All functions, when deployed upon request, are connected to the site network. When an Internet connection is requested, the base CPE is connected to the Internet network and configured to perform NAT. The figure also highlights an SFC that comprises the base CPE, VPN server, and firewall.

Future Work

Currently, the POC does not support the enforcement of network QoS in DC domains; this is only supported in the WAN connectivity services. We expect to add this support by the time OpenStack officially releases this feature. Furthermore, runtime management operations (e.g., scaling and migration of SFs) are yet to be included in the platform. On the WAN domain, we are currently adding an SDN-based network. The purpose is to have both legacy and SDN network technologies in place to better evaluate the advantages and disadvantages of each approach. Finally, we are working on exposing the ability of performing SFCs to the end user.

Conclusions

The orchestration and management of SFs is today a complex task that takes considerable time and effort. However, concepts like cloud computing, SDN, and NFV are paving the way to handling SFs in a much more flexible and agile manner. The telco will play a key role in this scenario, and we have given some insights on how that can be performed in the near future. Special attention has been given to the modeling of SFs toward cloud resources and to the combination of SFs through SFC. Finally, a platform for managing virtual SFs in a telco cloud infrastructure has been presented and a POC described that showcases how the platform and the principles here presented can be leveraged in a telco environment.

References

[1] H. Kim and N. Feamster, "Improving Network Management with Software Defined Networking," *IEEE Commun. Mag.*, vol. 51, no. 2, Feb. 2013, pp. 114–19.
[2] ETSI, "Network Functions Virtualisation (NFV): Use Cases," tech. rep. ETSI GS NFV 001 v1.1.1, Oct. 2013.
[3] P. Quinn, Kumar, and T. Nadeau, "Service Function Chaining Problem Statement," IETF Internet draft, informational, Dec. 2013.
[4] A. Császár *et al.*, "Unifying Cloud and Carrier Network: EU FP7 Project UNIFY," *IEEE/ACM 6th Int'l. Conf. Utility and Cloud Computing*, Washington, DC, 2013.
[5] G. Xilourisl *et al.*, "T-NOVA: A Marketplace for Virtualized Network Functions," *EUCNC '14*, June 2014.
[6] H. Puthalath *et al.*, "Negotiating On-Demand Connectivity between Clouds and Wide Area Networks," *IEEE CloudNet*, Paris, France, Nov. 2012.
[7] Y.Zhang *et al.*, "StEERING: A Software-Defined Networking for Inline Service Chaining," *Proc. IEEE ICNP '13*, Göttingen, Germany, Oct. 2013.
[8] ETSI, "Network Functions Virtualisation (NFV): Architectural Framework," tech. rep. ETSI GS NFV 002 v. 1.1.1, Oct. 2013.
[9] ETSI, "Network Functions Virtualisation (NFV): Terminology for Main Concepts," tech. rep. ETSI GS NFV 003 v. 1.1.1, Oct. 2013.
[10] W. John *et al.*, "Research Directions in Network Service Chaining," *2013 IEEE SDN for Future Networks and Services*, 11–13 Nov. 2013, pp. 1–7.
[11] J. Soares *et al.*, "Cloud4NFV: A Platform for Virtual Network Functions," *IEEE CloudNet*, Oct. 2014.
[12] Alcatel Lucent, "Network Functions Virtualization — Challenges and Solutions," White Paper, 2013.

Biographies

João Soares received an M.Sc. degree in electronics and telecommunications engineering from the University of Aveiro in 2009, and a Ph.D degree in 2015. He initiated his professional activity at the Institute of Telecommunications in 2009 and joined Portugal Telecom Inovação e Sistemas in 2010. He recently joined Ericsson Research as an experienced researcher in the area of cloud technologies. His interests cover the areas of cloud computing, cloud networking, SDN, and network virtualization.

Carlos Gonçalves received his Master of Science degree in computers and telematics engineering from the University of Aveiro in 2013. Currently, he is a research associate at NEC Laboratories Europe working in the areas of network function virtualization and carrier-cloud operation and management. He has participated in several open-source projects including OpenStack and Open Platform for NFV, serving as a member of the OpenStack Networking Advanced Services and Telco Working Groups.

Bruno Parreira received his M.Sc. degree in electronic and telecommunications engineering from the University of Aveiro in 2012. He started his professional activity at the Institute of Telecommunications through a research scholarship in May 2012. He has participated in several European funded projects related with cloud networking. His main interests are: cloud computing, cloud networking, software defined networking, and network function virtualization.

Paulo Tavares received his M.Sc. degree in engineering of computers and telematics from the University of Aveiro in 2011. He started his professional activity at the Institute of Telecommunications through a research scholarship in September 2011. His main interests include cloud computing, cloud networking, software defined networking, and network function virtualization.

Jorge Carapinha graduated in electronic engineering from the University of Coimbra in 1984 and got an M.Sc. degree in telecommunications from the University of Aveiro in 1998. He has been with Portugal Telecom Inovação e Sistemas (formerly CET) since 1985. Currently his main fields of interest are network virtualization, cloud networking, and software defined networking.

João Paulo Barraca received a Ph.D. degree in informatics engineering from the University of Aveiro, where he developed work focused on network management functions. He is currently an invited lecturer at the University of Aveiro

Currently, the POC does not support the enforcement of network QoS in DC domains; this is only supported in the WAN connectivity services. We expect to add this support by the time OpenStack officially releases this feature. Furthermore, runtime management operations are yet to be included in the platform.

and a researcher at the Institute of Telecommunications in areas related to programming, networking, and security. He has published more than 40 papers in the areas of networking and computer systems, and has acted as a reviewer for tens of events and journals.

RUI L. AGUIAR [SM] received a Ph.D. degree in electrical engineering in 2001 from the University of Aveiro, where he is currently a professor. He is leading a research team at the Institute of Telecommunications, and is an invited researcher at Universidade Federal de Uberlandia, Brazil. His current research interests are centered on the implementation of advanced networks and systems with special emphasis on future Internet and 5G architectures, and he is currently involved in the 5G-PPP initiative. He is a member of ACM, with more than 350 published papers. He has served as Technical and General Chair of several conferences, such as (recently) Monami '12, ISCC '14, NTMS'2014, and Mobiarch '14. He has been invited as a keynote speaker to several events, both technical and for large generalist audiences.

SUSANA SARGENTO has been with the University of Aveiro and the Institute of Telecommunications since February 2004, where she leads the Network Architectures and Protocols (NAP) group (http://nap.av.it.pt). She is also a cofounder of Veniam (www.veniam.com), a spin-off that commercializes vehicular technology. She has been involved in several national, U.S., and European projects, taking leadership of several activities in projects, such as the QoS and ad hoc networks integration activity in the FP6 IST-Daidalos Project, and the deployment of 600 nodes in a vehicular network in the EU Future Cities. Her main research interests are in the areas of future networks, more specifically routing, QoS, mobility, and cloud integration.