Progressive Introduction of Network Softwarization in Operational Telecom Networks: Advances at Architectural, Service and Transport Levels

by

Luis Miguel Contreras Murillo

A dissertation submitted by in partial fulfillment of the requirements for the degree of Doctor of Philosophy in

Telematics Engineering

Universidad Carlos III de Madrid

Advisor(s):

Dr. Carlos Jesús Bernardos Cano

April 2021

ACKNOWLEDGEMENTS

This Thesis is the last step in an amazing journey that started long time ago. I would like to show my gratitude to Dr. Carlos J. Bernardos for his advice, direction and enormous patience and friendship for driving me to the end. I am also grateful to Dr. Ignacio Soto who also accompanied me in the initial steps of this journey.

Furthermore, I am in debt with all the colleagues I have had the opportunity to work with along these years in my professional life at Alcatel (now Nokia), Orange and Telefónica, and also in my lecturing and research activities at the Universidad Carlos III de Madrid. Listing all of them would be extremely long, and probably not fair since my memory could fail. Some of them are now retired, others are just starting their career. Some of them are yet working on the companies where we met, others moved as I did. From all of them I have learnt, and all of them have influenced my way of thinking and my way of facing the life: being always positive, and always trying to extract the best of our time, since time passes so fast.

I have being also enriched by the interaction with so many people from other organizations (vendors, service providers, research institutions, universities, etc.) as result of the industrial and innovation activities carried out along these years, helping me to grow intellectually and professionally.

Lastly, I would like to thank the continuous and encouraging push from my parents, Diego and Pilar, on reaching this goal, and the exemplary paths always shown by my sister and brother, Elvira and Manuel. Thanks also to the rest of the crew (nephews, parents-, sisters- and brothers-in-law) who always showed and provided support on this task.

Finally, my deepest and loved gratitude to my wife Yolanda and my daughters Celia and Elena for all the time I took from them for concluding this work. *Os quiero.*

Luis M. Contreras,

Madrid, April 2021.

CONTENT

FIGURES

TABLES

ABSTRACT

Technological paradigms such as Software Defined Networking, Network Function Virtualization and Network Slicing are altogether offering new ways of providing services. This process is widely known as Network Softwarization, where traditional operational networks adopt capabilities and mechanisms inherit form the computing world, such as programmability, virtualization and multi-tenancy.

This adoption brings a number of challenges, both from the technological and operational perspectives. On the other hand, they provide an unprecedented flexibility opening opportunities to developing new services and new ways of exploiting and consuming telecom networks.

This Thesis first overviews the implications of the progressive introduction of network softwarization in operational networks for later on detail some advances at different levels, namely architectural, service and transport levels. It is done through specific exemplary use cases and evolution scenarios, with the goal of illustrating both new possibilities and existing gaps for the ongoing transition towards an advanced future mode of operation.

This is performed from the perspective of a telecom operator, paying special attention on how to integrate all these paradigms into operational networks for assisting on their evolution targeting new, more sophisticated service demands.

1 INTRODUCTION

This Thesis presents a number of contributions in related areas of network softwarization, accompanying the evolution of the involved technological paradigms as being researched and introduced in operational telecom networks.

1.1 Context and motivation

The technological paradigms of Software Defined Networking (SDN) and Network Functions Virtualization (NFV) have enabled the dynamic and flexible management and control of both services and networks. Ultimately, Network Slicing, which sits on top of the other two, has emerged as the paradigm permitting the provision of tailored end-to-end logical networks to external customers on top of a common and shared physical network infrastructure.

These three paradigms are encompassed on the overall concept of network softwarization, which determine a clear evolution and transformation of the traditional way of operating telecom networks.

The motivation of the work in this Thesis has been to analyze the impacts of all these transformations as well as to explore specific and realistic operational situations to perform such evolution and transformation in a smooth and sustainable manner.

1.2 Objectives and contributions

Four objectives were targeted during the elaboration of the Thesis:

Objective 1. To understand *how the network softwarization approach can impact telecom networks* from the perspective of a telecom operator, considering both the technical aspects and the realization of novel service offerings.

Objective 2. To *define novel architectural models* that could support such transition towards software-driven telecom networks, addressing novel situations enabled by these new technological paradigms.

Objective 3. To *assess the feasibility of services* leveraging on the aforementioned new paradigms, understanding issues and gaps for fully enabling them.

Objective 4. To develop *mechanisms for introducing network softwarization aspects at transport level*, in order to fully exploit the expected advantages of this approach.

These four targets are presented here as a collection of advances at architectural, service and transport network levels.

1.3 Thesis organization

This section gives a brief outline of how the thesis is organized and which topics are presented and discussed in the next chapters.

Chapter 2. This chapter provides a general overview of the technical paradigms of SDN and NFV which form the baseline of the advances described in this Thesis. It introduces in more detail the concept of Network Slicing, leveraging on the other two, and provides an analysis

of the challenges of all of them in telecom operators' networks. This chapter addresses Objective 1.

Chapter 3. This chapter focuses on the advances proposed at architectural level, describing a number of proposals and studies, including multi-domain interconnection and federation of infrastructures from multiple providers, collaborative separation of service and transport programmable concerns, determination of proper service edge, service blocking in federated infrastructure scenarios and efficiency gains in shared virtual network function solutions. This chapter focus on Objective 2.

Chapter 4. This chapter considers advances proposed at service level, reporting results for two concrete services: virtualized roaming solution in multi-domain infrastructures, and programmability of multicast proxies supporting multiple upstream interfaces. This chapter targets Objective 3.

Chapter 5. This chapter covers advances proposed at transport level, describing the cases of programmability of wireless transport networks and the handling of isolation in transport network slices. This chapter is related to Objective 4.

Chapter 6. This chapter provides the conclusions of this Thesis and identifies main lines of future work.

2 NETWORK SOFTWARIZATION

This chapter provides an overview of the fundamental technology paradigms that are transforming the telecom operational networks. All these paradigms, namely Software Defined Networking (SDN), Network Functions Virtualization (NFV) and lately Network Slicing, are all part of the overarching concept of Network Softwarization.

Essentially, network softwarization reflects the trend in which the networks are gradually being driven by software, in aspects such as programmability, virtualization, automation, negotiation, and even, decision. The gradual transition towards the full network softwarization permits a dynamic and flexible configuration of both services and infrastructure.

These new capabilities are enabled by the increasing introduction of programmability [O1][O2] and virtualization [O3][O4] within operational networks. Nowadays established techniques such as SDN and NFV have brought these capabilities to traditional telecom networks, dependent until recently on static provisioning methods and monolithic functional equipment. That softwarization process enables a number of important features. On one hand, it allows automation by the introduction of complementary techniques that could permit triggered corrective actions in a closed loop manner. Secondly, it facilitates a rapid service provisioning through easy reconfiguration of network connectivity and fast deployment of network functions across distributed computing facilities, adapting the service in every moment to the specific circumstances of the network, thus interacting gracefully.

This section briefly overviews the concepts of SDN and NFV for later on providing more details on Network Slicing which is built on top of the other two.

2.1 Software defined networking

SDN departed from the idea of separation of control and data plane in the forwarding devices, grouping the control capabilities in a logically centralized entity known as controller. That radical proposition was the base of an industrial and research movement around the ideas of network programmability as mechanism for flexible and dynamic configurations of networks, as compared with traditional ways of configuration, either based on manual procedures or leveraging on per-vendor Network Management Systems (NMS). SDN architectural principles have been promoted by industrial fora, being the Open Networking Foundation (ONF) the first one producing prominent architectural designs, as in [O5] and [O6].

By enabling programmability methods in the forwarding elements, it becomes possible, in principle, to arbitrarily configure the behavior or logic of the network elements, such as switching, forwarding, routing, filtering, etc., in an on-demand fashion. The concept results even more powerful with the consideration of creating solutions for vendor-neutral network programmability.

The key component retaining the control logic is the SDN controller. The controller instructs the network elements in order to accomplish specific, coordinated actions. For doing so, it keeps an overall view of the network topology on top of which different algorithms and

computation methods could be applied for determining optimal and sophisticated configurations appropriate to a given service scenario.

Figure 2-1 represents the architectural concept of SDN. The SDN Controller is the central piece interacting with the forwarding network elements for creating the delivery paths end-to-end, using specific protocols and models at its South Bound Interface (SBI). At the same time, the SDN Controller exposes a North Bound Interface (NBI) to external applications enabling sophisticated control functions that can take advantage of the overarching and comprehensive view of the network and its status as offered by the controller.

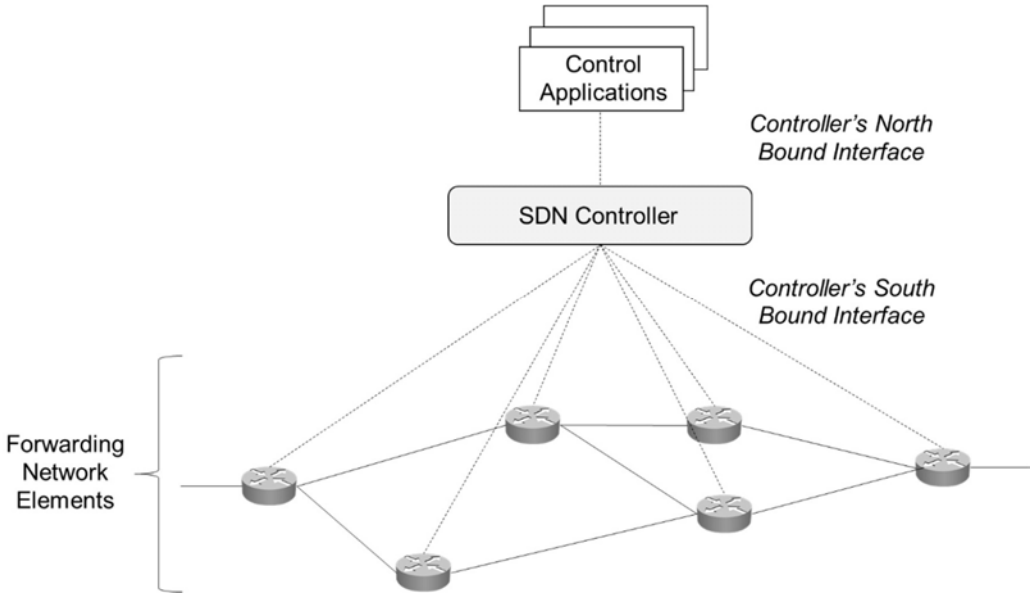

Figure 2-1. Architectural concept of SDN

Despite there were different precursor propositions in the direction of separating control and data plane (e.g., ForCES [O7]) the Openflow protocol (latest specification in [O8]) was determinant for the advances associated to SDN. Using OpenFlow as open and standard protocol at the SBI to interact with the forwarding nodes, a controller can dictate specific rules to simple network elements. These rules, using protocol specific control messages, define specific actions per traffic flow, where decisions about forwarding, re-routing, or packet modification, dropping and policing are taking according to specific flow characteristics.

The OpenFlow protocol was the one fostering the evolution that telecom industry is facing now. However, the SDN approach have been evolving along the time with different propositions nowadays.

The direction taken by the industry is less radical than the original approach in the sense that the objective is to retain and exploit control capabilities existing on the network elements while complementing the overall control of the network through central controllers. Then the trend is towards a programmatic interaction with the network rather than a full and complete

programmability of the behavior of the network elements. For simplicity, network programmability is used here with this sense.

This approach provides a more powerful scenario, where synergies between centralized and distributed control can be exploited. Essentially, the network elements accomplish forwarding and control functions which are configured dynamically from controllers with an overall view of the network status and tutelary relation with respect to the nodes. This permits a more lightweight behavior of network nodes while at the same time capturing the advantages of a programmable network.

2.2 Network function virtualization

NFV departed from the idea of separating the software implementation of specific network functions from the dedicated nodes were traditionally those functions were running, in a monolithic and tightly integrated mode. This followed the successful approach experienced in the Information Technologies (IT) industry with the cloud computing model.

The main goal of the NFV approach is to enable the decoupling of software and hardware for network function nodes allowing the software implementation of the function to run on top of commercial off-the-shelf (COTS) servers. Thus, the network function becomes a virtualized network function (VNF) that can be deployed as virtual machines (VMs) over hypervisors, which are totally transparent to the actual hardware below. From an industrial perspective, this technical movement has been promoted by the European Telecommunication Standardization Institute (ETSI) NFV Industry Specification Group (ISG) starting from the first specification on management and orchestration of VNFs [O9]. Figure 2-2 introduces the ETSI NFV architecture.

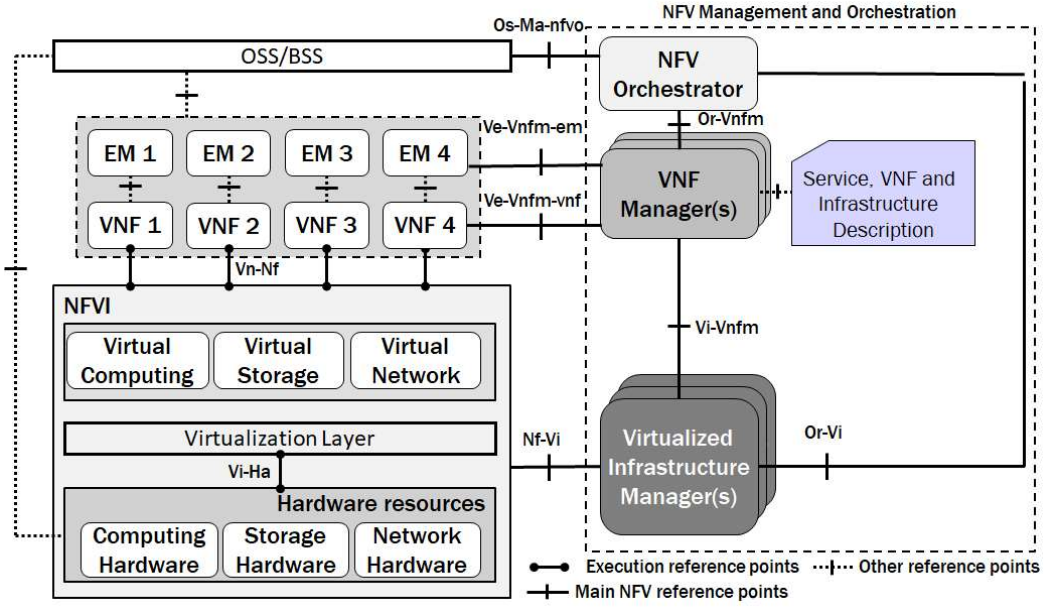

Figure 2-2. ETSI NFV architecture

The NFV Orchestrator (NFVO) is a functional block that manages a Network Service (NS) lifecycle. It coordinates the Virtualized Network Function (VNF) lifecycle (supported by the VNF Manager - VNFM-), and the resources available at the NFV Infrastructure (NFVI) level to ensure an optimized allocation of the necessary resources and connectivity to provide the requested virtual network functionality. The VNFMs are functional blocks responsible for the lifecycle management of VNF instances (e.g. instance instantiation, modification and termination). Finally, the Virtual Infrastructure Manager (VIM) is a functional block that is responsible for controlling and managing the NFVI computing, storage and network resources.

The original approach of running VMs on hypervisors have evolved in time by incorporating more lightweight schemas such as containers and unikernels. With this latest approach, a VNF can be executed over any hardware platform with a virtualization solution in place (i.e., hypervisor, containers, etc.) as it provides a unified interface to access virtual computing, storage, and network elements. Importantly, this can be done in a dynamic and flexible way, through appropriate management and orchestration capabilities, as defined in [O9].

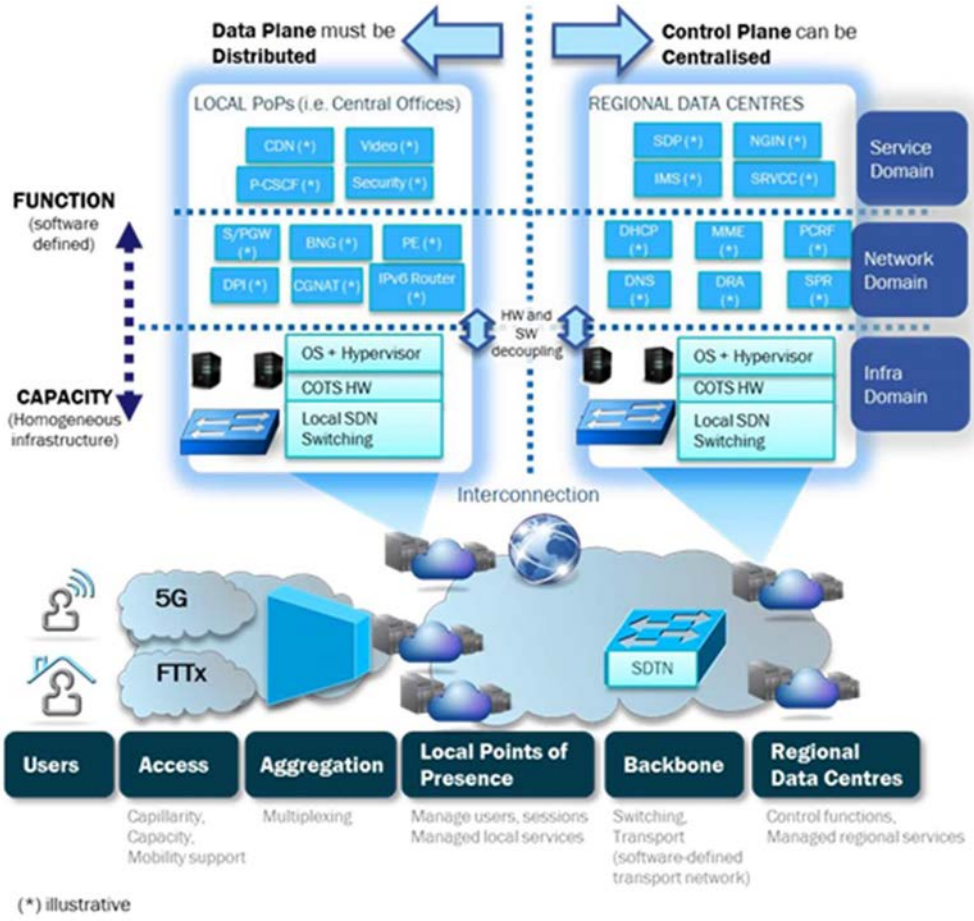

Figure 2-3. End-to-end virtualization [O10]

Figure 2-3 [O10] represents the overall idea behind network virtualization. It predicates the separation of functionality (the processing of the information with specific purposes) from the capacity (the resources necessary for such processing as dictated by the function). Depending on the particular characteristics of the functions, it could be adequate to instantiate them in different topological places in the network, according to characteristics such as data intensity (i.e., need of transferring higher volumes of traffic), latency sensitiveness (i.e., need for fast delivery of the service), or processing needs (i.e., need for high number of processors to accomplish the specific function).

2.3 Network slicing

Network slicing is a paradigm through which different virtual resource elements of a common shared infrastructure (in both connectivity and compute substrates) become allocated to a specific customer who perceived the resulting slice as a fully dedicated, self-contained network for it. The resources are virtualized through a process of abstraction of the actual physical lower-level elements, providing a great flexibility and independence of the specific network element being used along the customer service lifetime, which permits exercise advanced actions such as scalability, reliability, protection, relocation, etc. All of them represent an incredibly asset for a novel way of service provisioning with respect the present mode of operation in telecom networks.

Thus, network slicing will be provided by an integrated SDN and NFV architecture [O11], exploiting the commented advantages and permitting a flexible consumption of the network capabilities.

The possibility of dynamically instantiating slices through automation enables the provision of slices in an on-demand fashion, dealing to the concept of Slice-as-a-Service (SlaaS). The final objective of a customer when requesting a slice is to dispose of a complete logical end-to-end network on which to deploy the vertical service with full guarantees. The SlaaS approach is a versatile tool for trading tailored network capabilities with external third parties such as vertical customers, opening up new opportunities for telecom operators. The network is then transformed into a production system merging both business and operation domains [O12]. On one hand, different business models can be formed around the offering of network slices: Business-to-Business (B2B), Business-to-Consumer (B2C) and Business-to-Business-to-Consumer (B2B2C), depending on the kind and variety of stakeholders involved in the business part. On the other hand, distinct operational implications (life cycles, service objects, and slice scales) have to be taken into account derived from the service scenario where the requested slice applies.

A critical point on the overall provision of a slice is to allow control of the allocated abstract resources to the customer (i.e., the possibility of programming them by the customer). Without such control, the slice is simply made available but cannot be reconfigured by the customer, leading to a kind of static network. On the contrary, if control capabilities are enabled for the customer, the network can be then flexibly managed, for example by reconfiguring forwarding paths adapting them to changing conditions of traffic within the slice.

Different kind of slices can be considered from the operator perspective in that sense [A1], as follows:

- *Internal slices*, defined as the slices in which the operator retains the total control and management capabilities. These kind of slices will be typically devoted to provider's internal services.
- *External slices*, offered to vertical customers which perceived them as dedicated networks, but which in reality run on top of shared infrastructure. For external slices is yet possible to distinguish:
    - Slices managed by the operator, where the operator performs the control and management of the slice and the vertical customer simply runs the service on top of the capabilities and resources offered by the operator.
    - Slices managed by the vertical customer, where the customer actually has control of the resources and functions allocated. The level of control could be limited to a set of operations and/or configuration actions, but in any case, the vertical has the possibility of govern the slice behavior to some extent.

The latter case is the more challenging one. The referred control capabilities in that case should be enabled with care, since different actions from distinct customers could collide. The particular actions from each customer will be performed on their specific slice, that is, on the abstracted resources previously allocated. However, those actions could affect to a common physical resource shared among slices. Thus, if contradicting configuring actions happened, one customer could negatively impact on the slice of another customer.

The way of avoiding such an impact is the provision of isolation capabilities among slices of customers requiring higher degrees of control. That isolation can be achieved at different levels (e.g., control and forwarding plane) and usually implies a strict and dedicated allocation of resources per customer. Isolation is further discussed in Section 5.2.

From the provisioning point of view, this implies that in the case of external slices managed by the operator, different vertical customers with similar service needs can be accommodated in the same slice (e.g., customers requiring a generic enhanced Mobile Broadband – eMBB – service) if that slice is properly dimensioned, while in the case of external slices managed by the vertical customer the slices should be essentially dedicated per customer.

Figure 2-4 graphically represents this distinction, showing the different responsibilities in each case.

At the time of identifying solutions for the implementation of slices, all of these characteristics have to be taken into account, since there are different technical implications related to them, as will be considered later on in this Thesis.

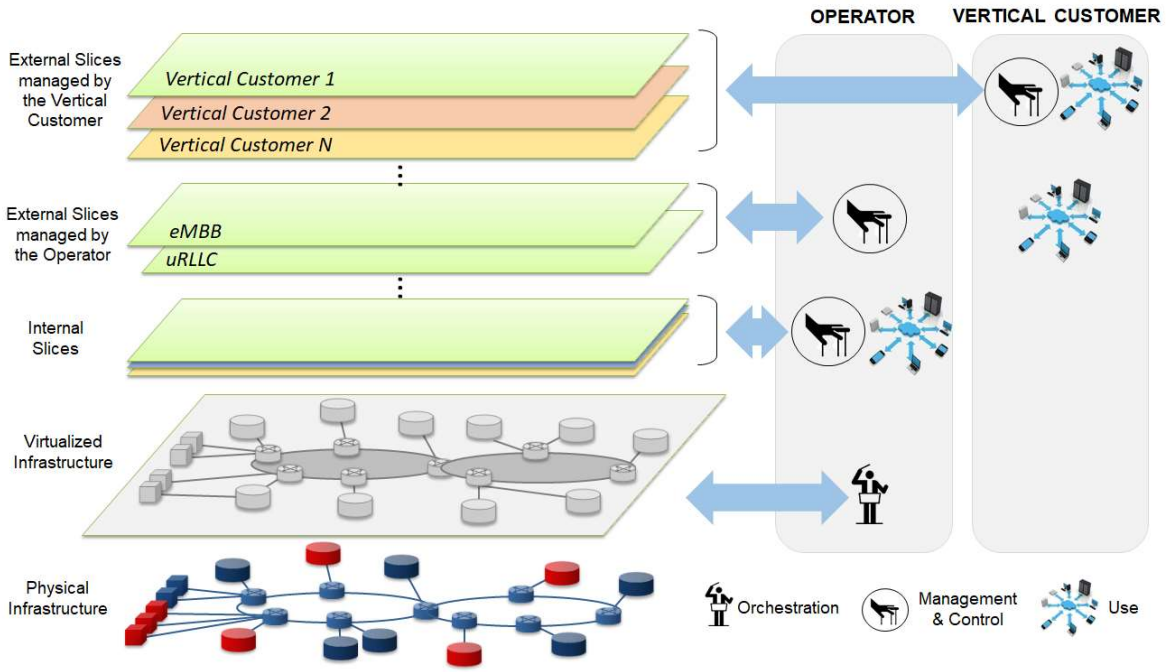

Figure 2-4. Types of network slices according to management and control levels of responsibility

2.4 Analysis of challenges due to SDN, NFV and network slicing in operational networks

The introduction of these three concepts in telecom operational networks, that is SDN, NFV and Network Slicing, brings a number of challenges to be faced in order to make all of them work. While SDN and NFV mainly affect technology aspects, Network Slicing impacts on new forms of service provision.

2.4.1 Challenges due to SDN and NFV as technological paradigms

This section provides an overview of the challenges faced by the operators adopting SDN and NFV in production networks. The following sections cover three main dimensions of those challenges: operational, organizational, and business issues.

2.4.1.1 Operational challenges

This section covers the challenges that relate to operator activities concerned with the building and operation of networks.

2.4.1.1.1 Network planning

The appearance of both NFV and SDN change the design rules and common practices traditionally used in telecom networks.

The dynamic invocation of network functions enabled by NFV changes the traffic patterns, the traffic engineering (TE) requirements and the need for quality of service assurance across the network. The traffic patterns are certainly different from those previously experienced, becoming less predictable; such a change in the observed traffic patterns complicate the network planning and operation tasks. In response to this, traditional planning approaches like overprovisioning of capacity are not an economical option. Over-dimensioning network links to accommodate traffic where the peak load varies widely and changes frequently creates a significant Capital Expenditure (CapEx) inefficiency. An alternative is that new, on-demand, transport control and traffic management mechanisms are developed to permit network adaptation dynamically in line with an offered load.

The traffic loads in conventional networks mean that overprovisioning, while not optimal, has been tolerable becoming a common practice. SDN and NFV enable new services that, in response to increasingly diverse customer and application needs, generate greater and more variable traffic loads. It is in this context that SDN and NFV are among the key catalysts for introduction of smart and dynamic traffic engineering into operator networks.

To face this challenge, network programmability, as key feature of SDN, permits new ways of resource optimization by implementing sophisticated traffic engineering algorithms to go beyond the capabilities of contemporary distributed shortest path routing. In addition, multilayer coordination can help to rationalize the usage of technological diverse resources for a common purpose. This new way of planning and operating networks requires a comprehensive view of the network resources and planning tools capable of handling these multilayer problems. An optimal planning process and tool chain then becomes multi-dimensional and multi-layered and has an important impact on operational cost and flexibility. In this respect, it is expected that the introduction of Artificial Intelligence (AI) and Machine Learning (ML) techniques could foster developments in this area, assisting on both the planning and the operation of the network.

2.4.1.1.2 Network deployment

The traditional cycles for the deployment of new network elements (NEs) in existing networks are long. A new NE or technology is first tested extensively to ensure compatibility with already deployed systems. Once validated, new equipment can be introduced and integrated in the network. Conventional product homologation is done simultaneously and in an integrated way for both the control and forwarding plane; they are, after all, present in the same NE in the traditional model.

The new situation, with centralization of control plane functionality in SDN and separation of function from device in NFV, means that the homologation of new products to be deployed in a network should change and adapt to this new reality. For instance, in the SDN case, control plane interworking e.g., of SDN controllers, must be assessed not only with the rest of the control elements in the network (controllers, orchestrators etc.), but also with all the different forwarding devices that they will control. In the NFV case, a new network function should be checked against the infrastructure that will support it (hypervisors, containers, servers, etc.) in addition to testing interworking with complimentary functions present in the

network. Finally, if the operator decides to deploy conventional and virtual versions of the same function, both implementations should be tested in a coordinated fashion.

With respect to commissioning, NFV definitely changes the way in which network functions are deployed in the network. Leveraging on available computing resources, deployed in data centers (DC), the setup and running of these functions in the network is accelerated. Site survey, cooling, cablings, and the rest of preparation for the technical environment are no longer necessary, avoiding site preparation delays. However, the DC investment must be done in advance, when there is not yet a clear demand and here is therefore some uncertainty and risk of overinvesting. The modular approaches for building data centers can mitigate this risk to some extent. Telecom network infrastructure is evolving towards the support of distributed cloud-computing services [A2]. Some applications are bandwidth intensive and/or latency sensitive which favors the distribution and placement of resources close to the points of traffic consumption, while others have more intense compute resource requirements and in consequence can benefit from large scale resource concentration. Therefore, large data centers will co-exist with micro-DCs placed in selected core locations to accelerate content delivery, reduce core network traffic, and ensure lower latency.

Multiple users (either internal or external to the operator) will use the same infrastructure that should be prepared to absorb their changing demands while satisfying committed SLAs.

Resource and energy usage are important aspects of the efficiency at which a network operates. Since service availability differs according to customer and application needs it is possible to use SDN and NFV to orchestrate a mix of service availabilities that trade efficiency (cost) against degree of protection (robustness).

Finally, the criticality of the DCs employed in the new model requires the implementation of strict and broadly based security measures.

2.4.1.1.3 Support systems

Both SDN and NFV rely extensively on software. This fact becomes clear in SDN since all the control capabilities are implemented in software programs running on the (logically) centralized controller. In the case of NFV, in addition to the software realization of the network functions themselves, capabilities for both control (e.g., for service chaining) and orchestration (e.g., of data center resources) are needed for deploying the virtualized functions in the network.

All these software components have to be integrated smoothly with the Operation Support Systems (OSS) and Business Support Systems (BSS) systems of running networks. This integration should cover not only the Fault, Configuration, Accounting, Performance, and Security (FCAPS) framework [O13], but also the inventory of the network. This is important because the decoupling of functions from devices alters the current way of performing service inventory in addition to increasing its dynamism. The conventional relationship between service and network device that previously existed now disappears. Whereas the inventory of the services was directly inferred from the supporting network device, the advent of NFV now breaks that binding.

In addition, the processes of assurance and fulfilment may now occur at different times. For instance, configuration and activation (including forwarding rules) of virtual network functions might be done weeks after the instantiation of such virtual network functions to a given network infrastructure (for instance, because some threshold motivates scaling out the network function).

It is almost certain that any deployment of SDN and/or NFV includes Open Source (OS) software. Indeed, many Tier-1 network operators are involved in or sponsoring organizations dedicated to producing software artifacts precisely for this purpose. This is a fundamental change from the traditional mode of operation where vendor-proprietary software is the norm. The use of OS software offers advantages but presents significant challenges as well. One frequently quoted advantage is the ability for operators to develop features on their own, more quickly than their vendors, thereby achieving shorter time to market for new services or capabilities. For many operators this involves a potentially significant change to their organizational skill set; they must become SW developers with all the testing, integration, quality assurance, maintenance and other SW lifecycle activities that entails. The OS software complements this transition by providing many pre-built components of SDN and NFV infrastructure but challenges it because there is a need to manage the relationship with the OS community that develops these components and, unlike the proprietary vendor solutions being used in the Present Mode of Operation (PMO), there is no telephone hotline line to call when something goes wrong.

How the OS developments can complement or to some extent substitute the existing core OSS and BSS systems is not yet clear. Most likely these solutions will be integrated through standard interfaces to provide certain capabilities of management and control over the programmatic infrastructure and the virtualized functions, but yet relay on conventional systems for certain other functions (e.g., billing).

2.4.1.1.4 Network operation

Multi-service networks today are composed of a variety of transport technologies. End-to-end path provisioning requires control and management capabilities to be present for all the technologies employed on the path. Traditionally there has not been possible an integrated way of operating this diversity of technologies in a common way. A uniform control and management capability across multiple technologies and network layers would tremendously simplify the operation of new networks. An initiative in this direction is further elaborated in Section 5.2.1

Furthermore, for each technology or network layer, a number of vendors are typically present in a single network. This implies different implementations of the same concepts and standards, which are not always fully compatible because of specification gaps and proprietary differentiators.

In this situation the network programmability concept of SDN can, when used with suitable models and abstractions, provide common ways of operation independently of the technology or vendor, thus simplifying network operations. The standardization effort in this case has to

focus on defining open interfaces, often referred to as Application Programming Interfaces (APIs), and appropriate data models to allow interoperability.

Controlling and operating a network using logically centralized control as envisaged in SDN, allows increasing the level of automation of network operations. Self-learning and self-healing capabilities can be expected to be developed for facilitating the operation and maintenance of the networks based on those controllers, minimizing manual intervention.

In this increasingly automated environment it is extremely important that the control elements behave predictably, deterministically and are free from malfunction. Again, the development AI and ML techniques as complement of programmable control is seen as one of the next steps in technological development.

Furthermore, the centralization of control decisions could impact operational reliability if proper resiliency mechanisms are not employed. For instance, there are stringent synchronization requirements for databases containing network and connection status that must be met in order to avoid data inconsistency.

The emergence of NFV also presents a fundamental change in the paradigm of network operations. While the network function itself remains the same, the supporting infrastructure radically changes, transforming the operation from being network oriented towards becoming cloud-oriented. This brings new concepts to the operation of a telecom network such as multi-tenancy, workload migration and the virtual binding of physically separate elements.

In the new model, the operations now include computing and storage as well as networking. Problem tracing, troubleshooting and backup activation all now change, requiring the production of new contingency plans and new guidelines for network operation. As an example, alarm handling and correlation has to be revisited to reconcile the existing monolithic and the new virtual approaches under a common operational model. The transition to this new mode of operation has to consider failures in the IT infrastructure itself separately from failures of the specific network functions.

2.4.1.1.5 Service provisioning

Service provisioning is also impacted. Programmable transport capabilities and function instantiation together break the service creation determinism. If the service is actually decoupled from the transport network then a new level of abstraction has to be defined, in this case at service level, to allow the operation of the service, either at its creation (fulfillment stage) or during maintenance (assurance stage).

The capability to dynamically instantiate services makes it necessary to ensure a proper mapping between service requirements and network capabilities. Mechanisms for exposing such capabilities are required, and procedures for negotiating service SLAs need to be defined to ensure that the service delivery is not affected by the underlying network, independently of where the service is instantiated. These mechanisms and procedures should include information relative to service verification, maintenance, and accounting.

Common APIs, programmatic interfaces and information and data models should be developed and standardized to facilitate the proper abstractions required at both service, resource, and device level.

2.4.1.1.6 Investment protection and migration

The introduction of SDN and NFV into existing networks is being done progressively or selectively, i.e., in some parts of the network but not others. This is not just a matter of SDN/NFV product availability. Current assets in telecom networks need to be fully amortized for investment protection. In consequence these two new approaches have to coexist with conventional networks and technologies for a long period. There is then a need for interworking between conventional and SDN/NFV systems, without impacting service during the migration period. That interworking has to be implemented for both the control and forwarding planes. During this migration phase the network should behave as it does now while, at the same time, allowing incremental deployment of the new possibilities due to SDN and NFV.

2.4.1.2 Organizational challenges

This section analyzes the impacts of the introduction of SDN and NFV on operator's organization.

2.4.1.2.1 Departmental organization

A typical operator's technical organization is structured in departments (e.g., network architecture, engineering and planning) which are divided into technological silos (e.g., service platforms, IP, transmission and radio). The advent of SDN and NFV is shaking this structure since both the departmental frontiers and the technical boundaries become blurred. This traditional structure has to be re-adapted to the new cross-technical and cross-functional reality brought about by both innovations.

For instance, the technical teams deal with multiple technologies like IT and optical transport resources simultaneously. As another example, an engineering decision on concentrating several functions on the same server can produce impacts on the planned network capacity reaching the hosting data center and the service architecture previously defined.

2.4.1.2.2 Personnel skills and know how

As mentioned before both SDN and NFV heavily relay in software. The work with these technologies requires professional profiles that include familiarity with IT technologies as a complement (not a substitute) of the current telco-oriented skills widely present in the operator's staff. Multidisciplinary teams are required to work together to accomplish all the operational functions described in the previous section.

Aligning these employees' diverse experiences and skills to focus on how things affect customer's services requires special attention and coordination.

2.4.1.2.3 Partnership ecosystem

Traditionally operators have relied on equipment and system vendors to create and develop their networks. Monolithic boxes and solutions hindered the development of a broad partnering ecosystem. This situation is now changing

NFV and SDN have the potential to change the vendor and supplier community with which an operator engages. Both technologies have the potential to enable new market entrants and both have spawned extensive standards development and Open Source software activities. New specializations emerge in the ecosystem; for instance, some companies being specialized in control plane software and others building the (white box) hardware it runs on, and yet others specialized in testing and validating network functions and their integration from a vendor neutral point of view.

All of this increases the number and dynamicity of relationships that have to be managed between an operator and the vendor ecosystem, which requires an additional effort for partner management from the operator side.

2.4.1.3 Business challenges

This section introduces challenges present in the business arena, some of which are related to the commercial sustainability of telecom operators.

2.4.1.3.1 Analytics

Information about the habits and network resources usage of customers is important for the development of new and advanced services. Detailed information can be obtained today, but the collection of that information is based on the deterministic steering of traffic towards the collection functions build on top of monolithic boxes.

NFV changes such determinism because of the dynamic instantiation of functions. In such situation the traffic pattern changes as the functions act as traffic attractors. While the service plane is logically maintained, the transport plane usage can change. In consequence, the collection of information has to consider both planes to present a consistent view; data plane usage is a function of both what services the customer uses and where the network chooses to instantiate those services, and both of them change over time. This consideration applies to billing and to accounting processes as well.

In addition, Big Data analytics for processing the data collected from computing, network, and services can help to optimize the usage of the network and DC resources in a predictive manner. Dynamic traffic engineering or virtualized network function reconfiguration are examples of actions that may benefit from the utilization of Big Data techniques.

2.4.1.3.2 Customization

SDN enables programmability of the network, which allows or simplifies, the provisioning of services both by the operator or triggered externally by the customers, i.e. on-demand network service consumption. To support such flexibility the transport network needs to be

dynamic, allowing reconfiguration of network elements and changing their behavior without impacting other services in place.

The ability to program and reprogram the network requires the development of standard APIs both toward the customer, whose applications should not have to change because of modification in the network, and toward the data plane which needs to remain independent of flux in the higher (software) layers of the network.

From the perspective of NFV the possibility of instantiating functions on demand provides an agile and personalized way of managing services. However since the underlying infrastructure is shared, both isolation and security are important considerations.

In both cases, billing capabilities that incorporate actual resource usage are required.

2.4.1.3.3 Network and IT equipment lifetime

The equipment deployed in current networks has specific lifecycles defined by vendors according to their roadmaps and product development strategies. Two main milestones characterize the time when equipment replacement becomes a concern: the End-of-Sale (EoS) and the End-of-Life (EoL). The former indicates the time when new units of a given equipment cannot be acquired directly from the vendor, while the latter refers to the date from where the equipment is no longer supported by the vendor. Normally network equipment is exploited by the operators long after the EoL date.

SDN can help to extend the lifetime of hardware assets in networks by allowing (re)programming and this can be used in combination with relocating the equipment according to throughput and functional demands.

In contrast, commodity IT equipment usually offers shorter lifecycles with increasing performance in terms of compute processing, storage capacity, power consumption, space, etc. Newer hardware is better able to accommodate network functions in terms of processing and number of functions supported. NFV support will be a driver to renew IT equipment as the trend to virtualize network functions continues.

2.4.1.3.4 Procurement

SDN and NFV results in an atomization of the components – both hardware and software – that must be acquired by network operators. Where previously one single product with specific hardware supported a number of features, in the future mode of operation there is a component integration approach, where the desired operational features are the result of combining a collection of elements (i.e., atoms) with necessary integration. As we noted earlier SDN and NFV initiatives increase the number of potential vendors of these system components (or atoms).

Conventional monolithic NEs are easily compared in terms of throughput, number of ports, computing and switching capabilities, etc., and operators have significant experience with this. In comparison the benchmarking of software systems is more difficult and is dependent of the supporting hardware, so new evaluation methods are necessary.

The solution pricing models have to be revisited, as well. In principle SDN offers NE simplification potentially lowering the price for accomplishing the forwarding functions. At the same time the new centralized control plane has costs associated with it that should be considered in the overall picture. The promise is that the combined costs of control and forwarding plane grow at less than the linear (with capacity) rate experienced until now. Similar considerations apply to NFV. When doing cost vs. benefit analysis care must be taken to make sure that virtualized functions actually are functionally equivalent to their legacy NE based counterparts. Apples must be compared to apples and not to oranges.

The contractual guarantees required from the vendors also need to be re-defined. Identifying, assigning responsibility for, and replacing failed hardware units is relatively simple in conventional solutions. However identifying failure is more difficult in the software arena, as it is assigning final responsibility, especially when the software in these systems can come from multiple vendors, contain Open source components and, indeed, code written by the operators themselves. Management of the anticipated short development cycles in these complex environments also poses a new challenge which needs to be addressed. Continues Integration / Continuous Development (CI/CD) emerges as the way of handling the SW development cycles.

Finally the management of spare parts is also impacted. The commoditization of the hardware can reduce the cost of elements in the network however, those elements have to be fully interchangeable and, for example, support in service replacement and upgrade. What is really required is commodity hardware built to rigid specifications and benchmarking; simply specifying processor architecture for example, is not enough.

2.4.1.3.5 Capabilities for sharing network infrastructures

The sharing of network infrastructure between operators is now common practice as a means of reducing investment while still allowing the provisioning of service over a broad footprint.

Once again SDN and NFV might simplify the traditional mode of operation. Currently the sharing agreements on a common infrastructure require manual configurations to route traffic appropriately towards the infrastructure of each of operator participating in the agreement. This static configuration of resources lacks the flexibility to quickly react to changing demands. SDN and NFV can play a significant role in providing agile mechanisms for sharing, bringing to the network infrastructure the idea of multi-tenancy that exists today in data centers (this is further elaborated in Section 3.2). To make that possible new business models for infrastructure and capacity sharing and usage should be investigated. Regulation can play also an important role, perhaps imposing additional requirements on the kind of interfaces offered among operators for interchanging control and management information.

2.4.1.3.6 Innovation and experimentation

It seems beyond any doubt that the flexibility offered by both NFV and SDN fosters service innovation and may also help to shorten the current time-to-market cycles. The network is turned into a moldable environment where new concepts and ideas can be tested more rapidly than before (see the exemplary services described in Chapter 4). The time between service

inception and service deployment is reduced, and dependencies on monolithic boxes and closed solutions disappear. In the same manner, it is possible to revert to a previous service easily, then lowering the impacts in case of failures or misbehavior of new features and functionalities, with the benefit of avoiding investment in specialized equipment that can finally demonstrate not to be useful.

Concomitantly it is possible to easily revert to a previous service, which lowers the impact of failure or misbehavior of new features and functionalities. This also reduces speed the introduction of new services since it reduces the possibility of assets with fixed functionality being stranded by changes in service requirements. This openness and flexibility also helps operators to test self-developed service prototypes easily on their own infrastructure (in controlled or in production environments) with full production solutions being developed later by partner vendors. The challenge for an operator in this context resides in maintaining innovative teams (both technical and commercial) to develop differentiating services.

2.4.1.4 Summary

Table 2-1 summarizes all the challenges identified with respect SDN and NFV as technological paradigms that are changing the mode of operation of telecom networks.

2.4.2 Challenges due to network slicing as new service model

Through the instantiation of distinct network slices, the operator is able to provide completely different services in a dynamic and isolated manner despite that all of them actually run over the same physical infrastructure. Thus, the operators can initiate a transition from the existing design choices that rely upon multi-service networks deployed over one architecture conceived to fit all kinds of services [O14], towards the approach of getting logical networks defined per service.

Network slices are intended to behave like entirely independent networks. This implies that there should be mechanisms for properly isolating the slices (if required by the slice definition), thereby avoiding interference between one slice and the others. Furthermore, in some cases, there are situations requiring slices to be interconnected for composing a service, raising the need for interworking (stitching) of such independent logical constructs. In order to be able of managing such complexity induced by the logical slice partitioning, the operator needs to have technical means to take into account the specific requirements of each slice (for instance in terms of throughput, delay, and jitter), as well as suitable inter-slice information exchange mechanisms to avoid congestion issues and arbitrate resource needs, permanent or occasional.

It is clearly anticipated that the mechanisms for creating and operating the network slices will be supported by the softwarization trend enabled by the programmability of the network resources and the virtualization of the network functions, that is, leveraging on SDN and NFV techniques. Apart from the elasticity provided by the virtualization of network functions when dynamically deploying such network functions, flexible steering mechanisms (i.e., Service Function Chaining) are needed to shepherd the traffic flows according to the expected service behavior.

Table 2-1. Summary of identified challenges.

Dimension	***Area***	***Challenge***
Operational	Network Planning	• Comprehensive network resource planning • Multilayer planning tools
	Network Deployment	• Increased testing and product homologation • Risk of overinvestment in data center infrastructures • Criticality of data centers and need for highly secure infrastructures
	Supportive Systems	• Integration with existing OSS/BSS systems • Relies on Open Source developments
	Network Operation	• Standardization of open interfaces that facilitate uniform control and management across technologies and vendors • Control plane resiliency • Network Automation • Decoupling of service from transport
	Service Provisioning	• Mapping between service requirements and network capabilities • Common APIs and information and data models
	Investment Protection and migration	• Coexistence with legacy networks
Organizational	Departmental Organization	• Cross-technical and cross-functional reorganization of departments
	Skills and Know How	• Multidisciplinary teams
	Partnership Ecosystem	• Larger number of partners requiring more coordination and management effort
Business	Analytics	• Correlation of decoupled service and transport indicators • Big Data analytics for predictive actions
	Customization	• Standard interfaces for network service and resource consumption • Isolation and security • Proper billing mechanisms
	Network and IT Equipment Lifetime	• Re-programmability of equipment to extend the service lifetime
	Procurement	• Management of a larger number of vendors • Definition of new quotation models to compare prices and solutions respect to conventional products • Definition of new guarantees
	Capabilities for Sharing Network Infrastructures	• New business models for infrastructure and capacity sharing • Impacts of regulation
	Innovation and Experimentation	• Creation and maintenance of innovative teams

In order to address these requirements, network programmability and its integration with the function virtualization described above, are considered as the baseline pieces to approach network slicing [O15] playing an instrumental role in the control, management, and operation of future networks.

Network slicing, as a combination of network, computing and storage resource allocation, is intended to enable value creation for vertical segments that lack physical network infrastructures, or which can complement their own resources with the ones offered by the operators. The business objective is to improve and evolve the production capabilities of those industries. Network slicing can thus be seen as the evolution of current wholesale services, by improving the way network resources are allocated and consumed in an agile and ultimately efficient manner.

2.4.2.1 Requirements for Accomplishing Business Objectives

A number of key requirements can be deemed critical for accomplishing the business objectives as previously described. This sub-section covers some of them.

2.4.2.1.1 SLA Management

The assurance of Service Level Agreements (SLAs) becomes a key aspect of the provisioning of services on top of 5G networks. Acknowledging the relevance of this observation for operator's internal services, this is even more evident when considering vertical customers, who base their production totally or partially on the negotiated network slice.

Any distortion of the negotiated SLA associated to the network slice can impact not only the technical behavior of the service offered to the end-user of the vertical but also its reputation or business leadership. Even more importantly, such distortion may have legal consequences of any kind due to the incapacity of honoring the contracted service expectations (e.g., in terms of latency, bandwidth, or availability) through the network slice based on the nature of the service. A concrete case could be the result of injured people in a factory or in an assisted driving service because of the misbehavior of the network slice supporting the corresponding service. SLAs then have consequences beyond what sales agreements with strict associated penalties imply.

The negotiated terms of an SLA define a number of service indicators to be satisfied. These service indicators are translated into network indicators (to be understood in a broad sense, that is, including computing, networking and storage resources) that are used as inputs for the configuration and the orchestration of resources to form the slice. The deployment of a vertical service observes interdependencies in the different classes of resources to be provisioned and configured, their location, and their forecasted availability. Furthermore, depending on the class of slice provided (e.g., provider- or tenant-managed, as described in Section 2.3), the vertical customer could have direct control of those resources, introducing uncertainty with regards to the future behavior of the assets provided, which may raise some inconsistency issues.

Mechanisms for a timely analysis of the real network situation (including each particular slice) and quick reaction to fulfil an SLA are then required.

2.4.2.1.2 High Customization of the Slice

The vertical industries are different in nature, as described before. This implies that operators should deal with a large variety of needs and requirements to accommodate the vertical services, introducing an inevitable high customization at the time of provisioning network slices, tailored to each of the requested services.

The variety of services can be categorized according to different dimensions:

- Major type of requested service: the major types of services to be supported by 5G networks are enhanced Mobile Broadband (eMBB), ultra-Reliable Low Latency (uRLLC), and massive Machine Type Communications (mMTC). Furthermore, some additional types could be considered, such as pure connectivity services (as an evolution of existing corporate communication services, for example), or even NFVI-as-a-Service services (as an evolution of the Infrastructure-as-a-Service – IaaS – kind of services available in cloud environments). Finally, the same vertical customer can subscribe to services pertaining to more than one of these major types, and then require some interconnect of the provisioned slices to offer a smooth service experience for the vertical customer.
- Particular type of requested service: the major type of services described above can be seen as some kind of macroscopic description of the main characteristics of the service. Beyond such description, particular services require specific conditions that eventually motivate a distinction at the time of provisioning the required slice. For instance, latency requirements of uRLLC services can differ allowing for more relaxed conditions in some cases [A3].
- Vertical customer differentiation: obviously, a competitive industrial market will look for differentiation as a form of a competitive advantage compared to similar vertical competitors. Then, specific requirements coming from different customers pertaining to the same industrial sector at the time of requesting a network slice can also be expected.
- Location: it is also clear that different vertical customers will also have distinct geographical footprints, hence the need for extending the network slice coverage in a different and customer-specific way.
- Vertical service evolution: finally, the business of each vertical customer can evolve at a different pace, requiring upgrading and scaling resources at different speeds and with distinct constraints, including the need for invoking various kinds of slice instances over time.

As a consequence, the operator must endow itself with automation mechanisms as a means to simplify and foster the provisioning of the slices. These mechanisms reside in the orchestration and management artifacts used to handle the resources (either physical or virtual) constituent of the slices. Programmability and virtualization (of functions and infrastructure) are essential to reduce time-to-market, which basically is the same that reducing the time-to-revenue. All are essential to guarantee chances of capturing the market

of 5G vertical services too, since conventional forms of service provision do not sufficiently scale to support this high customization (in terms of flexibility, speed, granularity, etc.).

2.4.2.1.3 Service Segregation

The critical importance of ensuring isolation among the slices provided and supported on top of the same infrastructure has already been mentioned.

The need for service segregation has been satisfied in the past by different means. A very basic option consisted in physically separating network infrastructures, each one being dedicated to enterprise, fixed residential, and mobile services, respectively. This design is clearly neither cost-efficient nor sustainable anymore.

Alternatively, logical separation has been also performed by means of the deployment of overlay solutions, e.g., in the form of VPNs. However, traditional tools and mechanisms for VPN provisioning are neither flexible nor agile, as previously pointed out in the high customization discussion.

Once again, it is required to evolve towards scenarios where dedicated resources (including network functions) can be automatically allocated according to the needs of the service, ensuring isolation, on top of the same infrastructure. Section 5.2 further elaborates on this topic.

2.4.2.2 Slicing challenges for network operators

The fundamental aspect of network slicing is that each slice will behave as if it were an independent network. Taking this into account, this sub-section describes a list of key challenges from a provider's perspective.

2.4.2.2.1 Scalability

Scalability in the context of network slicing exhibits two dimensions: the scalability of the slice itself, and the overall scalability, as a function of the total number of slices.

The first dimension is related to resource allocation and accounting (including resources necessary for satisfying protection and availability), directly related with the kind of service requested and negotiated with the customers. Furthermore, the deployed slices can scale up or down over time, depending on several factors that include commercial success and optimization, e.g., following seasonal demand.

This scalability dimension requires orchestration mechanisms that dynamically add or remove assets to the slice in a consistent manner, based upon either computing or networking capabilities (or a combination thereof). Additionally, on the customer's side, these new resources have to be accounted as part of the provided solution. If the customer is responsible for the control and the management of the provided resources, then those that have been introduced (to scale up) or removed (to scale down) have to be added or removed according to the decisions applied by the control and management functions of the customer. Charging should also be adapted dynamically according to the consumed capabilities at any given time.

The second scalability dimension refers to the global scalability that the operator can support in terms of quantity and types of slices to be orchestrated and managed. A too much fine-grained offering of slices can provoke an unmanageable number of artifacts to be orchestrated by the provider, making them unpractical. Some kind of aggregation or grouping is needed for achieving tractability. Scalability can be much impacted by the number of external tenant-managed slices offered.

In order to partition network resources in a scalable manner, it is required to clearly define to what extent slice customers can be served or not by an existing slice. A proper application of different SLAs with the translation of service parameters into network ones (including compute needs) is essential to understand to what extent a new demand can be accommodated with existing slices, from the service point of view. In addition, if the customer requires the responsibility of control and management capabilities, then the customer-specific "individualization" of the slice is a must.

2.4.2.2.2 Arbitration

In order to resolve conflicts and to ensure negotiated service levels, the provider needs to incorporate some arbitration mechanisms to allow an efficient usage of resources (including functions), preventing on the one hand resource over-dimensioning, and on the other hand service degradation or disruption. These mechanisms have to be in place not only among the different slices (internal and external, as presented in Section 2.3) that are being deployed over the same infrastructure, but also within the individual slices themselves, since the relationship between a customer and a slice is not necessarily 1:1 (unlike the external, provider-managed slice case).

Arbitration needs to be applied not only for slice creation or customer activation, but also (and more importantly) when scaling and/or failure events happen, so resources are properly (re-)assigned according to the applicable SLAs. Such inter-slice arbitration may negatively influence the performance of other slices that share the same infrastructure, hence the critical importance of adequate arbitration solutions. Note that SLAs should enclose technical clauses which govern slice availability and the expected maximum duration before getting the slice running as expected. Such clauses need also to be taken into account.

The role of arbitration is to some extent equivalent to the role that existing Quality of Service (QoS) mechanisms play on current networks. QoS is primarily effective in situations of network congestion, when the availability of resources becomes compromised and their scarcity has to be managed to minimize any kind of impact on the service delivery. Similarly, arbitration needs to be in place when events in the network limit the availability of the resources that compose the different slices supported by the network, or their internal components.

This capability of arbitration can collide with the requirement of slice isolation. That is, the arbitration of resources should maintain the principle of isolation among slices, to avoid any kind of degradation or interference from one slice to another, especially when such isolation is imposed by the associated SLA.

The arbitration is foreseen as an internal capability of the operator, transparent to the customer who can influence the decisions and arbitration criteria only through the negotiated SLAs.

Finally, it can be expected that some prioritization could happen in the event of massive failure or outage. The criteria for establishing priorities could be diverse, ranging from the type of service to be ensured, the customers to be protected, the percentage of slice affected, the critical SLAs to be guaranteed, the associated penalties, etc. Commercial, regulation and security aspects could motivate distinct options to be taken into consideration by each operator.

2.4.2.2.3 Slice planning and dimensioning

Over-dimensioning has often been the default design principle to deploy and operate networks for avoiding any kind of congestion. Through slicing, the location of the traffic sources and destinations becomes much less predictable, if predictable at all. This is especially relevant for the case of external tenant-managed slices, where the final decision of where to deploy traffic sources and destinations (as well as some intermediate service functions that could alter or modify the traffic profile) lies in the hands of the customer.

Two different time scales can be distinguished during these processes: microscopic and macroscopic planning and dimensioning. Microscopic planning and dimensioning refers to the process applied to each individual slice. On the contrary, macroscopic planning refers to the global process applied to the overall infrastructure, including resources and functions.

In the microscopic case, a primary source of information for the dimensioning process is the SLA agreed with the customer, which specifies the expectations for the service to be provided. A number of network and compute KPIs are derived from the service parameters expressed in the SLA. These serve as inputs for determining and tweaking the required resources for honoring the requested service (as well as for identifying engineering parameters of the slice, like redundancy, etc.). This for sure is needed for data plane related resources, but also for control plane resources, including associated licenses or processing needs (e.g., for database dimensioning purposes). A certain level of over-allocation/overbooking can be expected in order to mitigate issues related to unexpected demands, failures, etc. Additionally, some excess could be derived from the resource quantization process (e.g., bandwidth granularity allocated in Gbps units, or licenses allocated in blocks of thousands of users).

The planning in the microscopic case requires some (traffic) forecast inputs from the slice tenant, either internal or external. The formalization of those inputs could vary, from being part of the negotiated SLA (renewed periodically) up to a different/specific administrative process. Adherence to the planned resources derives from a commercial obligation for the operator towards the customer. Then, an assessment of the allocated capabilities should be performed as well. The concept of in-operation network planning [O16] can be also assumed as valid for the network slicing case as a means to adjust any kind of resources to the observed demand evolution.

At a macroscopic level, the dimensioning process takes into consideration all the slice demands, the ones already in place plus the slices that are being instantiated. A punctual need coming from an aggregated demand could require borrowing resources from other slices (e.g., capabilities allocated for resource protection purposes) in order to satisfy the overall demand while the operator builds the necessary infrastructure for all the demands. Furthermore, overbooking is certainly an option when the probability of congestion is actually low or the scarcity of resources does not allow for an immediate upgrade of the slice in terms of resources.

The planning procedures also take into consideration the forecasted evolution of the existing slices plus the evolution of the (marketing and business) plans provided by the commercial and service units. The multi-tenancy approach facilitated by the slicing concept is fundamental to define rational investments for serving the expected demands, but also the programmability and virtualization techniques are solicited to make a better use of the available resources by reconfiguring services in a flexible manner.

Proper planning, dimensioning, and enforcement are needed to make the transition towards this new form of service sustainable, starting with an appropriate data collection on resource usage, especially the virtual ones which need abstract models for usage reporting.

2.4.2.2.4 Multi-domain

Multi-domain can be interpreted in different ways, since the notion of "domain" can reflect different concepts. For slice provisioning, two meanings are considered here: technological domain, taking into account the applicability of slicing to different technologies, and administrative domain, where more than one provider is involved in the provisioning of network slice(s) to a customer.

Slice-based services necessarily require the integration of different technology domains. The complete end-to-end nature of slices involves distinct computing environments and transport technologies (data switching, optical, etc.), and linking them requires a consistent orchestration approach. For those services making use of radio access technologies, the slice concept has also to be extended to the Radio Access Network (RAN) and the Core Network (CN), with their own slicing specificities [O17].

The deployment of slices in the networking technological domains is necessarily different since the forms of traffic segregation are specific to each of the technologies. Furthermore, a combination of logical (e.g., VLAN tag-based in Ethernet systems, label-based in MPLS, etc.) and physical (e.g., lambda-based in optical, slot-based in Flexible Ethernet, etc.) separation methods can be expected. Complementary computing capabilities have to be added with their own form of separation (e.g., hypervisor-based or container-based). The form of combining and integrating them for a single slice requires the combined action of multiple controllers, including an overarching system that maintains a global, systemic view of all the resources, including those that are available.

The interaction with different technologies is clearly fundamental for slice provisioning. Notwithstanding, vertical customers could require no restriction for the slice requested in

terms of coverage, service capability, resource constraints, geographical footprint, etc., avoiding any potential limitation of the network provider with whom they have a commercial relationship as their privileged provider. This leads to the necessity of enabling multi-domain slicing, which implies functional and commercial interfaces to be normalized for the sake of massive adoption.

Both business and technical implications can be deemed necessary for such multi-operator slice provisioning context. From the business side, the following implications can be listed:

- Coordination models: there is a need for business coordination to define how multiple stakeholders interact for the provisioning of a multi-domain slice in order to trade low-level resources and elementary services combined and orchestrated to deploy slices end-to-end.
- Inter-provider SLAs: each provider in each administrative domain should have its own SLA assessment capabilities internally, including interfaces with SLA aggregation components that automate the multi-domain aggregation process.
- Pricing schemes: bilateral negotiation can be expected as a regular mechanism to establish pricing agreements. Even for simple pricing formulas, the values of the parameters under consideration should be dynamically adapted, e.g., according to demand or resource/service availability.
- Service specification and customer facing advertisement: each provider may consider not only its own capabilities in its domain but also slice offerings and capabilities available in neighboring domains. Therefore, catalogue synchronization is required to be performed across domains.

From a technical standpoint, implications include:

- Multi-domain orchestrator: from the provider-to-provider's viewpoint, only certain entities within each domain should interact with each other for handling the inter-domain activities in order to keep the slice provisioning consistent end-to-end. Those entities can be identified as multi-domain orchestrators. This kind of orchestrator is in charge of abstracting the underlying infrastructure in its domain before announcing (to neighboring providers) what capabilities and functions the operator can provide.
- Slice decomposition: The vertical customer will request a slice to a provider, the latter becoming the origin provider for that customer. The origin provider should then incorporate sufficient logic for decomposing the slices across the different domains.
- Discovery of domains: although manual configuration can be used, automatic procedures are desirable for speeding up service provisioning in the network softwarized era.
- Common abstraction models: a common understanding of the description of the resources (i.e., network, compute and storage) and the capabilities per domain is needed.
- Standard interfaces, protocols, and APIs: these will be required for remote control and management of functions and slices in other domains.

2.4.2.2.5 Orchestration and Control of the Slices

The request for network slices that is sent by customers requires the orchestration of resources in order to address the said request. SDN and NFV techniques can be used by an operator to orchestrate slices [O11]. Such orchestration needs to have full control and visibility of the nodes, topology, functions, and capabilities (such as bandwidth or compute power) to make decisions. Example situations can be found in [A4] and [A5].

End-to-end slices are a service management issue. The enablers of management and orchestration are the usage of open and standard interfaces for interacting with the nodes and systems, as well as the definition of normalized models for service and devices. An overview of network management and orchestration can be found in [A6].

The result of the orchestration process is the allocation of the resources, as well as the management of their lifecycle, including service assurance and fulfillment. Then, the constituent blocks of the slice are controlled by the operator.

Different customers can require distinct levels of control for the resources they have requested to the provider. Extreme cases can be customers that do not require any capability of control and management of the allocated assets (just pure communication service), and on the other end, customers requiring full control of their assets. A gradual level of control can be found in between.

Then, the operator should provide configuration and administration capabilities to the customers according to the levels of control that they request. These capabilities could come by simply exposing some interfaces for that required control actions (e.g., APIs), up to granting direct access to the resources (e.g., IP address to access the element console). The more abstracted way, the less invasive for the operator.

3GPP [O18] defines a number of management functions needed to manage network slices to support communication services. These functions are:

- *Communication Service Management Function* (CSMF), which is responsible for translating the communication service-related requirements into network slice-related requirements.
- *Network Slice Management Function* (NSMF), which is responsible for the management and the orchestration of an instance of a network slice.
- *Network Slice Subnet Management Function* (NSSMF), which performs the same task as the NSMF, but at a sub-instance level.

Figure 2-5 represents the relation among this three slice management functions.

Both NSMF and NSSMF can be considered as functionally similar. For instance, ETSI NFV has mapped these management functions within the ETSI NFV orchestration framework in [O19] as part of the broader OSS/BSS components. This is further commented in Section 3.2.1.

Regarding the programmable control of the slices, two levels of SDN control can be considered for a given slice. On one hand, there is the tenant's controller that permits to

configure the involved functions. On the other hand, there is the infrastructure controller that is used to program the underlying infrastructure resources to provide end-to-end connectivity. These two levels of control facilitate interplay of actions at both service function and connectivity infrastructure levels, enabling flexibility on the slice instantiation and provisioning, assuming cross-layer coordination.

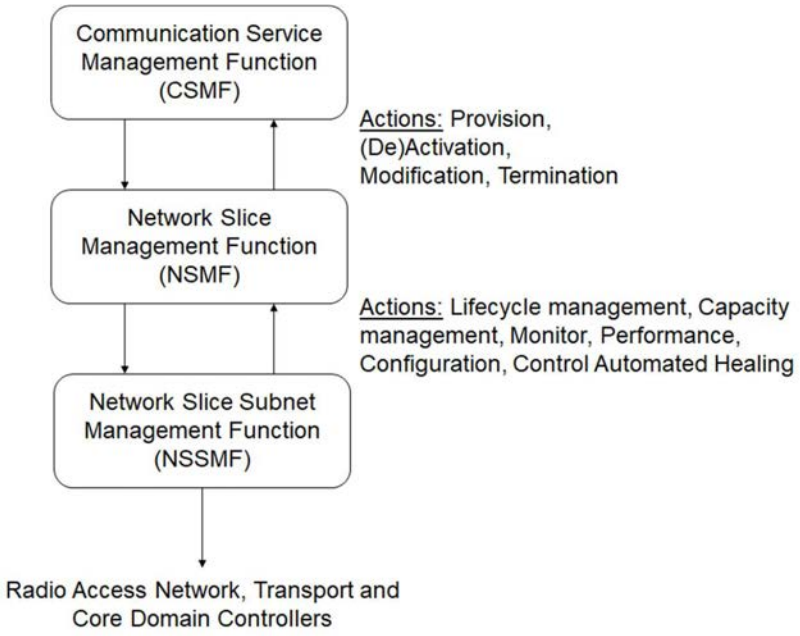

Figure 2-5. 3GPP slice management functions

In the NFV framework, such a double control level is proposed in [O24] that describes the usage of SDN in NFV environments. This is further elaborated in Section 3.1.6.

2.4.2.2.6 Slice Operation

The operation of slices, once instantiated, share many common aspects with respect to conventional networks. However, some new mechanisms and artifacts will be needed. This sub-section considers monitoring and maintenance as relevant aspects of the operation of a network, identifying how the application of maintenance and operational procedures related to a slice can raise new requirements.

Monitoring and performance information is essential for a healthy operation of a network. Monitoring data are usually associated to specific resources, either network ones (e.g., packet errors) or compute ones (e.g., CPU load). This can be complemented with indicators of the service functions that compose the service (e.g., number of active users of the service). Some mechanisms have to be defined in order to properly display and abstract the information for each slice tenant (or user). To this respect, external slices have a higher degree of complexity

since the information to be exposed, and the constraints to access it, have to be defined or even better, agreed between the provider and the customers.

At the time of creating a slice, a number of resources are allocated to a given customer. As a consequence, the monitoring information associated to the allocated resources has to be extracted in order to ensure proper operation of the slice. All data or part of them should be presented to the customer in order to provide the necessary information. It is important to preserve privacy as the information related to a given slice must not be leaked into other slices. The monitoring and performance-related information would also serve as a reference for assessing the compliance of what has been delivered with what has been agreed in the SLA.

Then, the monitoring information has to be properly processed to be provided to the customer (possibly according to a specific format and semantic). Only the information strictly associated to the customer's resources has to be exposed, which implies some filtering on the global indicators. Furthermore, slices can be supported on top of virtual resources. This means that the physical resources can change over time while the virtual ones allocated for the slice appear as unchanged, even though they may have been, for example, migrated from one virtual machine to another. This implies that the monitoring information could be originated from different sources, thereby requiring dynamic aggregation and association during the slice lifetime.

All the received information has to be processed and correlated in the same manner as in today's legacy networks. This implies the replication of operational procedures per slice, hence raising a scalability issue. Two approaches could be followed: *(1)* initial processing of all the indicators and further personalization per slice; or *(2)* initial separation per slice of the corresponding indicators for further individual processing and correlation. The first approach seems to be more scalable and practical, but requires more coordination and integration from the OSS/BSS point of view.

2.4.2.2.7 Slice Marketplace

The interaction with the vertical customers is a critical feature for understanding the needs of the service to be provided. From that interaction, the provider obtains the necessary information for creating (or reusing) a slice and mapping the customer service to the allocated resources.

The level of detail in the request sent by the customer impacts the functionality required on the provider's side to address such request. The customer requests could be expressed in terms of service or resources. The former implies that the request identifies the characteristics of the service to be delivered, without any further details about what is needed in terms of resources to be allocated (also known as intent-based requests). In contrast, the latter implies that the slice request details all the resources identified as needed by the customer. Clearly, the provider should be able to translate the service semantics into resource semantics for the actual allocation of the slice resources, based upon a computation logic that can take into account the outcomes of the possible negotiation between the customer and the provider as

input data, but also the network planning policies, the status of the network, its resources, the location of the end-user, whether the end-user is in motion or not, etc.

Different kinds of customers can have different levels of know-how and skills, thereby determining what approach to follow. From the provider's perspective, service semantic requests can allow to reach a broader market, the one formed by vertical customers neither specialized nor skilled in the knowledge of what communication services are needed for their specific business. Considering the type of slices shown in Section 2.3, most probably the external, customer-managed slices will be requested through the resource semantics approach, while the external, provider-managed slices will be requested through service semantics. The support of the different slice classes requires a consistent set of abstractions either at service or resource levels to allow the aforementioned semantics to be expressed consistently. As a consequence, proper abstractions and templates have to be defined to ensure the provisioning of a service portfolio providing a consistent view of the network and its resources, as well as their integration with the internal network management and orchestration systems.

2.4.2.2.8 Security

In any shared infrastructure, security is a key element to guarantee proper operation to each user. Slice customers must be appropriately authenticated, their rights enforced by authorization mechanisms, and the operations they perform accounted for, so that further auditing can be applied in case of any problem. This becomes crucial when considering the external, customer-managed slices, since the possibility of altering the behavior or status of the allocated functions and resources increases.

Each vertical customer offer services to its end-users thanks to the slices that have been deployed (similar to current Mobile Virtual Network Operator – MVNO - service offerings). This means that the vertical customer has also the responsibility of enforcing security measures in order to protect the service (and the end-users, indirectly), which affects the security exposure of the allocated resources. The security measures of every vertical customer are multiple and diverse, thereby requiring the provider to harden the allocated resources in order to prevent whatever issue (e.g., attacks generated from within a slice managed by a tenant). Even if the provider can influence the security implemented by the vertical customer, the control is not total. Generic guidelines and best (current) practices should be defined and updated according to new threats and security problems as they are met.

Another key issue is the privacy of customer data, as well as the privacy of the end-users' data making use of the service offered by the vertical customer. All this information has to be properly stored and encrypted (whenever applicable) in order to prevent any exposure of such data to the provider in general, or to other customers who use the provider's infrastructure. This is essential when some of the functions and resources can be shared between different customers in the same slice or across slices.

Beyond this, measures have to be in place to proactively detect and mitigate active security attacks, thereby avoiding that a security breach that affects one slice does not propagate into the infrastructure and other slices.

2.4.2.2.9 Slice Aging

The same dynamicity for the allocation of slices to tenants, leveraging SDN and NFV techniques, applies as well to the slices' lifecycle. One of the promises of automation is the possibility of invoking and deploying services faster than today, thereby overcoming the currently weak service agility. This facility in the creation of services opens new business opportunities since targeted services, scoping events and situations of short duration, can be made available easily. The duration of the existence of a slice can be certainly variable, so differentiating between short versus long is relative: "long" slices are those created as a semi-permanent service, while hourly, daily, or even seasonal slices can be considered as "short"-aged ones.

Furthermore, the frequency at which a given slice is requested can be another timing parameter to be taken into consideration. Depending on the frequency of slice requests, some resources can be freed or should be kept as booked until a new forthcoming service request shows up (that is, the resources may remain unavailable for any other slice request, except for temporary needs, for example).

This dynamic situation (motivated by allocating and freeing resources), implies that whatever the decision or the action made by the operator with respect to using some resources, the operator should consider not only the resource view and status at the time of the specific request but also over a broader timeline, since any decision at the moment of the request can negatively influence future decisions (here again AI and ML techniques can help).

Slices based on calendaring considerations (e.g., day/night operation) need some guarantees, thus motivating a certain level of resource booking. The matching of the resource availability with the time duration request introduces more constraints as far as slice resource allocation is concerned.

Whatever the duration of it, the creation and operation of a slice requires a non-negligible number of administrative and technical registers. Administrative notifications, data and billing records, systems configurations, etc., are proportional to the number of slices.

A mix of long- and short-lived slices co-existing on top of the same infrastructure should be expected. This impacts on providers in several manners, from resource planning (slice demand forecast including traffic and resources to be consumed) to security (data preservation per tenant).

2.4.2.2.10 Slice Isolation

This main requirement for isolation constitutes the essential feature any network service provider has to support in order to deliver slices to its customers. The degree of isolation achieved is critical in determining the ability of a certain provider to address the different classes of slices discussed above, and how they can be requested and used by its customers.

The isolation can be applied at different levels, such as control plane and data plane isolation, as well as resource and function isolation. Even the sharing of constituent elements within each of these planes can be allowed, the allocated capabilities have to be segregated in order to avoid any kind of misbehavior induced by any other customer in the system.

Data plane isolation can be achieved by several means, and with different degrees. The encapsulation of data in different tunnels, one per vertical customer, can be a primary measure for achieving such isolation in a shared environment. More extreme situations, like strict allocation of assets (as enabled e.g., by the concept of calendar slots in Flexible Ethernet [O20]) allow the exclusive allocation of a specific number of resources to slices. It is also possible to find several options with a higher or lower level of isolation.

For the control plane, the separation can be achieved even by supporting such separation in the actual control plane capabilities or alternatively by replicating control plane capabilities dedicated to specific customers per slice. The first approach requires the control plane element to implement such isolation mechanisms, e.g., via the creation of different virtual spaces per customer. The second approach naturally grants isolation since the different replicas act as independent control elements.

A pairing among the form of managing the data plane and the control plane should be defined. For instance, sharing the data plane capabilities while dedicating control plane ones can be problematic. Multiple replicas of control elements (i.e., hard isolation at the control plane level) acting on the same data plane elements that isolate traffic by means of encapsulation in an overlay model context (i.e., soft isolation at the data plane level) can lead to inconsistencies, hence jeopardizing the isolation.

Function isolation basically consists of the instantiation of separated service function instances per vertical customer versus the sharing of a given service function instance to be used by multiple verticals. An intuitive example could be a firewall. Implications like the number of compute capabilities associated to each option, the connectivity with the rest of the service functions per each customer, the number of licenses to be consumed, etc., have to be taken into consideration. Section 3.5.2 explores efficiencies due to the sharing of virtual functions.

As for resource isolation, the implications reside in the level of partitioning that the provider is willing to or can implement. Resource isolation can apply to compute nodes, ranging from the dedication of specific compute resources like the bare-metal approach versus the sharing of computing capabilities by means of hypervisors or containers. It can also apply to transport resources, like the allocation of specific lambdas in optical nodes to specific slices versus the accommodation of traffic of different customers carried by the same lambda. In this respect, Section 5.2 goes further in the description and proposes an index to indicate isolation feasibility at transport level.

2.5 Summary and outlook

This chapter addresses the Objective 1 of this Thesis: *to understand how the network softwarization approach can impact telecom networks from the perspective of a telecom*

operator, considering both the technical aspects and the realization of novel service offerings.

The chapter has overviewed aspects related to network softwarization, and includes a number of original contributions produced in the following manner:

- Advances on the network slicing topic published in [A1], with further elaboration in a journal paper under submission [A7].
- Analysis of operators' challenges related to SDN and NFV, which have been published as a journal paper [A8].
- Analysis of operators' challenges related to network slicing, which have been published as book chapter in [A9], recently re-printed in [A10].

3 ADVANCES AT ARCHITECTURAL LEVEL

This chapter provides a number of insights on a variety of contributions to advances on network architecture related to the work in this Thesis.

Specifically, the following aspects at architectural level have been the main subjects of research:

- Architecture describing cooperation in the programmability of service and transport concerns.
- Interconnection of multi-provider infrastructures for service orchestration, considering multi-domain NFV-enabled carrier networks, options of federation of MEC environments, and multi-domain slicing integrating fronthaul / backhaul aggregation networks.
- Decoupling of physical from service edge, with an architecture assisting to determine the proper infrastructure to deploy services at the edge.
- Analysis of service blocking probabilities in multi-domain service provisioning, considering compute and network capabilities in different administrative domains.
- Analysis of efficiency gains when sharing virtualized delivery points in a CDN-as-a-Service slicing scenario.

The following sub-sections provide further details on each of these lines of work.

3.1 Cooperation among service and transport concerns

This section describes an original approach called Cooperating Layered Architecture for Software-Defined Networking (CLAS), wherein the SDN control functions associated with transport are differentiated from those related to the SDN control of the services, in such a way that each of both concerns can be provided and maintained independently, following their own evolution path as well.

As introduced in Section 2.1, SDN advocates for the separation of the control plane from the data plane in the network nodes by introducing abstraction among both planes, allowing the control logic on the SDN Controller, to be centralized; one or multiple controllers may be deployed. A programmatic interface is then defined between a forwarding entity (at the network node) and the control entity. Through that interface, the control entity instructs the nodes involved in the forwarding plane and modifies their traffic-forwarding behavior accordingly. Support for additional capabilities (e.g., performance monitoring, fault management, etc.) could be expected through this kind of programmatic interface [O12].

Most of the intelligence is moved to this kind of centralized functional entity. Typically, such an entity is seen as a compendium of interacting control functions in a vertical, tightly integrated fashion. Such approach of considering an omnipotent control entity governing the overall aspects of a network, especially both the transport network and the services (supported on top of it) simultaneously, presents a number of issues as described next:

- From a provider perspective, where different departments usually are responsible for handling service and connectivity (i.e., transport capabilities for the service on top),

the mentioned approach offers unclear responsibilities for complete service provision and delivery.
- Complex reuse of functions for the provision of services.
- Closed, monolithic control architectures.
- Difficult interoperability and interchangeability of functional components.
- Blurred business boundaries among providers, especially in situations where one provider provides only connectivity while another provider offers a more sophisticated service on top of that connectivity.
- Complex service/network diagnosis and troubleshooting, particularly to determine which layer is responsible for a failure.

The relocation of the control functions from a number of distributed network nodes to another centralized entity conceptually places together a number of control capabilities with different purposes. As a consequence, the existing SDN solutions do not provide a clear separation between services and transport control.

Both service and transport are concerns of different nature, and that is reflected on the SDN capabilities required for each of them. On the one hand, in the case of transport, actions focused on programming the network to handle the connectivity or forwarding of data between distant nodes are needed. On the other hand, for services, actions devoted to programming the functions or services that process (or manipulate) such data are also needed.

Here, the separation between service and transport follows the distinction provided by [O21]. Despite such differentiation, close cooperation between the service and transport layers (or strata as referred in [O21]), and the associated components per layer, is necessary to provide efficient usage of the resources.

Operator networks support multiple services (e.g., Voice over IP, IPTV, critical mission applications, etc.) on a variety of transport technologies. The provision and delivery of a given service independent of the underlying transport capabilities require a separation of the service-related functionalities and an abstraction of the transport network to hide the specifics of the underlying data transfer techniques.

Such separation can facilitate configuration flexibility and adaptability from the point of view of either the services or the transport network. Multiple services can be provided on top of a common transport infrastructure; similarly, different technologies can accommodate the connectivity requirements of a certain service. Close coordination among these elements is required for consistent service delivery (inter-layer cooperation). Key aspects to consider in this coordination are: *(i)* to expose transport capabilities to services; *(ii)* to capture transport requirements of services; *(iii)* to notify service intelligence of underlying transport events, for example, to adjust a service decision-making process with underlying transport events; and *(iv)* to instruct the underlying transport capabilities to accommodate new requirements, etc.

An example is guaranteeing some Quality-of-Service (QoS) levels. Different QoS-based offerings could be present at both the service and transport layers. Vertical mechanisms for

linking both service and transport QoS mechanisms should be in place to provide quality guarantees to the end user.

Figure 3-1 shows the CLAS architecture. It is based on the functional separation in the Next Generation Network (NGN) architecture defined in [O21], where two strata of functionality are defined. CLAS adopts the same structured model described in [O21] but applies it to the objectives of programmability through SDN [O12]. These strata are the Service Stratum, comprising the service-related functions, and the Transport Stratum, covering the transport-related ones. The functions of each of these layers are further grouped into the control, management, and user (or data) planes. The cooperation between the two layers is expected to be implemented through standard interfaces.

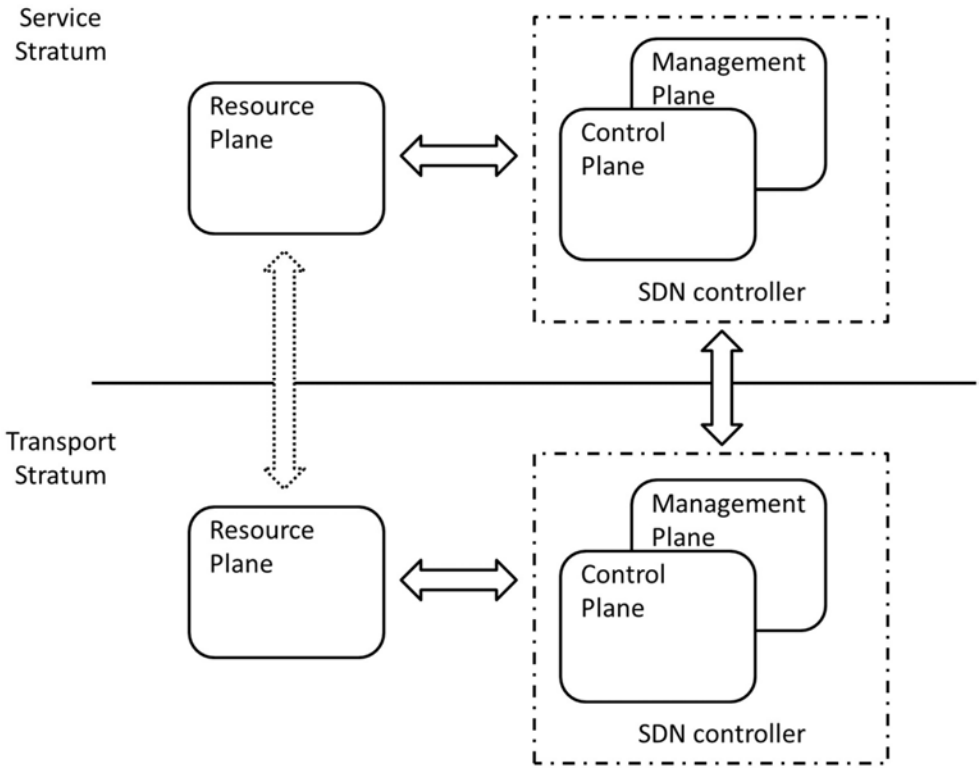

Figure 3-1. Cooperating Layered Architecture for SDN (CLAS)

In the CLAS architecture, both the control and management functions are considered to be performed by one or a set of SDN controllers (due to, for example, scalability, reliability), providing the overall SDN control in such a way that separated SDN controllers are present in the Service and Transport Strata. Management functions are considered to be part of SDN to allow for effective operation in a service provider ecosystem [O12], although some initial propositions did not consider such management as part of the SDN environment [O5][O6].

Furthermore, the generic user- or data-plane functions included in the NGN architecture are referred here as resource-plane functions. The resource plane in each stratum is controlled by the corresponding SDN Controller through a standard interface.

Both SDN controllers cooperate in the provision and delivery of services. There is a hierarchy in which the Service SDN Controller makes requests of the Transport SDN Controller for the provision of transport capabilities. The Service SDN Controller acts as a client of the Transport SDN Controller.

Furthermore, the Transport SDN Controller interacts with the Service SDN Controller to inform it about events in the transport network that can motivate actions in the service layer.

Despite not being shown in Figure 3-1, the resource planes of each stratum could be connected. This will depend on the kind of service provided. Furthermore, the Service Stratum could offer an interface for applications to expose network service capabilities to those applications or customers.

3.1.1 Functional Strata

As mentioned before, there is a functional split that separates transport-related functions from service-related functions. Both strata cooperate for consistent service delivery. Consistency is determined and characterized by the service layer.

3.1.1.1 Transport Stratum

The Transport Stratum comprises the functions focused on the transfer of data between the communication endpoints (e.g., between end-user devices, between two service gateways, etc.). The data-forwarding nodes are controlled and managed by the Transport SDN component.

The control plane in the SDN Controller is in charge of instructing the forwarding devices to build the end-to-end data path for each communication or to ensure that the forwarding service is properly established. Forwarding may not rely solely on the preconfigured entries; means can be enabled so that involved nodes can dynamically build routing and forwarding paths (this would require that the nodes retain some of the control and management capabilities for enabling this). Finally, the management plane performs management functions (i.e., FCAPS) on those devices, like fault or performance management, as part of the Transport Stratum capabilities.

3.1.1.2 Service stratum

The Service Stratum contains the functions related to the provision of services and the capabilities offered to external applications. The resource plane consists of the resources involved in the service delivery, such as computing resources, registries, databases, etc.

The control plane is in charge of controlling and configuring those resources as well as interacting with the control plane of the Transport Stratum in client mode to request transport capabilities for a given service. In the same way, the management plane implements management actions on the service-related resources and interacts with the management plane in the Transport Stratum to ensure management cooperation between layers.

3.1.1.3 Recursiveness

Recursive layering could happen in some scenarios in which the Transport Stratum is itself structured in Service and Transport Strata. This could be the case in the provision of a transport service complemented with advanced capabilities in addition to the pure data transport (e.g., maintenance of a given SLA).

Recursiveness has also been discussed in [O5] as a way of reaching scalability and modularity, where each higher level can provide greater abstraction capabilities. Additionally, recursiveness can allow some multi-domain scenarios where single or multiple administrative domains are involved, such as those described in Section 3.1.5.3.

3.1.2 Plane separation

The CLAS architecture leverages plane separation. Three different planes are considered for each stratum. The communication among these three planes (with the corresponding plane in other strata) is assumed to be based on open, standard interfaces.

- Control plane. The control plane logically centralizes the control functions of each stratum and directly controls the corresponding resources. The work in [O22] introduces the role of the control plane in an SDN architecture. This plane is part of an SDN Controller and can interact with other control planes in the same or different strata to perform control functions.
- Management plane. The management plane logically centralizes the management functions for each stratum, including the management of the control and resource planes. Reference [O22] describes the functions of the management plane in an SDN environment. This plane is also part of the SDN Controller and can interact with the corresponding management planes residing in SDN controllers of the same or different strata.
- Resource plane. The resource plane comprises the resources for either the transport or the service functions. In some cases, the service resources can be connected to the transport ones (e.g., being the terminating points of a transport function); in other cases, it can be decoupled from the transport resources (e.g., one database keeping a register for the end user). Both the forwarding and operational planes proposed in [O22] would be part of the resource plane in this architecture of CLAS.

3.1.3 Required features foreseen in CLAS

Since the CLAS architecture implies the interaction of different layers with different purposes and responsibilities, a number of features are required to be supported:

- Abstraction: the mapping of physical resources into the corresponding abstracted resources.
- Service-Parameter Translation: the translation of service parameters (e.g., in the form of SLAs) to transport parameters (or capabilities) according to different policies.

- Monitoring: mechanisms (e.g., event notifications) available in order to dynamically update the status of (abstracted) resources while taking into account, for example, the traffic load.
- Resource Computation: functions able to decide which resources will be used for a given service request. As an example, functions like the Path Computation Element (PCE) could be used to compute / select / decide a certain path.
- Orchestration: the ability to combine diverse resources (e.g., IT and network resources) in an optimal way.
- Accounting: record of resource usage.
- Security: secure communication among components, preventing, for example, Denial of Service (DoS) attacks.

3.1.4 Communication between SDN Controllers

The SDN controllers residing respectively in the Service and Transport Strata need to establish tight coordination. Mechanisms for transferring relevant information for each stratum should be defined.

From the service perspective, the Service SDN Controller needs to easily access transport resources through well-defined APIs to retrieve the capabilities offered by the Transport Stratum. There could be different ways of obtaining such transport-aware information, i.e., by discovering or publishing/subscription mechanisms. In the former case, the Service SDN Controller could be able to handle complete information about the transport capabilities (including resources) offered by the Transport Stratum. In the latter case, the Transport Stratum reveals the available capabilities, for example, through a catalog, reducing the amount of detail of the underlying network.

On the other hand, the Transport Stratum must properly capture the Service requirements. These can include SLA requirements with specific metrics (such as delay), the level of protection to be provided, maximum/minimum capacity, applicable resource constraints, etc.

The communication between controllers must also be secure, e.g., by preventing denial of service or any other kind of threat (similarly, communications with the network nodes must be secure).

3.1.5 Deployment scenarios

Different situations can be found depending on the characteristics of the networks involved in a given deployment.

3.1.5.1 Full SDN Environments.

This case considers that the networks involved in the provision and delivery of a given service have full SDN capabilities.

- Multiple Service Strata Associated with a Single Transport Stratum. A single Transport Stratum can provide transfer functions to more than one Service Stratum. The Transport Stratum offers a standard interface(s) to each of the Service Strata. The

Service Strata are the clients of the Transport Stratum. Some of the capabilities offered by the Transport Stratum can be isolation of the transport resources (slicing), independent routing, etc.

- Single Service Stratum Associated with Multiple Transport Strata. A single Service Stratum can make use of different Transport Strata for the provision of a certain service. The Service Stratum invokes standard interfaces to each of the Transport Strata, and orchestrates the provided transfer capabilities for building the end-to-end transport needs.

3.1.5.2 Hybrid Environments.

This case considers scenarios where one of the strata is totally or partly legacy.

- SDN Service Stratum Associated with a Legacy Transport Stratum. An SDN service Stratum can interact with a legacy Transport Stratum through an interworking function that is able to adapt SDN-based control and management service-related commands to legacy transport-related protocols, as expected by the legacy Transport Stratum. The SDN Controller in the Service Stratum is not aware of the legacy nature of the underlying Transport Stratum.
- Legacy Service Stratum Associated with an SDN Transport Stratum. A legacy Service Stratum can work with an SDN-enabled Transport Stratum through the mediation of an interworking function capable of interpreting commands from the legacy service functions and translating them into SDN protocols for operation with the SDN-enabled Transport Stratum.

3.1.5.3 Multi-domain scenarios in the Transport Stratum

The Transport Stratum can be composed of transport resources that are part of different administrative, topological, or technological domains. The Service Stratum can interact with a single entity in the Transport Stratum in case some abstraction capabilities are provided in the transport part to emulate a single stratum.

Those abstraction capabilities constitute a service itself offered by the Transport Stratum to the services making use of this stratum. This service is focused on the provision of transport capabilities, which is different from the final communication service using such capabilities. In this particular case, this recursion allows multi-domain scenarios at the transport level.

Multi-domain situations can happen in both single-operator and multi-operator scenarios. In single-operator scenarios, a multi-domain or end-to-end abstraction component can provide a homogeneous abstract view of the underlying heterogeneous transport capabilities for all the domains.

Multi-operator scenarios at the Transport Stratum should support the establishment of end-to-end paths in a programmatic manner across the involved networks. For example, this could be accomplished by each of the administrative domains exchanging their traffic-engineered information [O23].

3.1.6 Applicability of CLAS in NFV

The CLAS proposition has been considered also as one of the potential SDN architectures in ETSI NFV [O24]. Figure 3-2 shows the mapping of CLAS into the ETSI NFV architecture. There, the SDN controller with focus on service configuration is considered to be a Tenant SDN Controller, while the one with focus on transport is referred to as Infrastructure SDN controller.

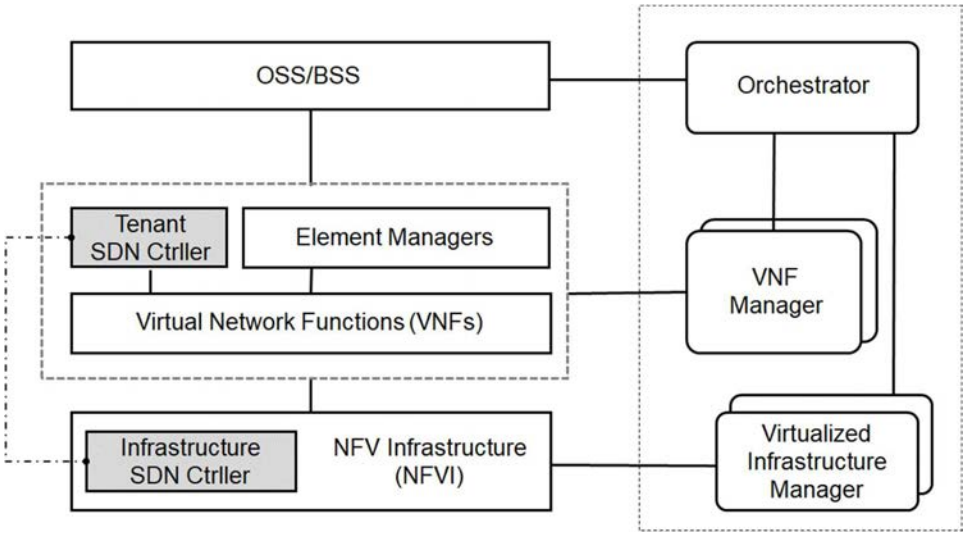

Figure 3-2. Applicability of CLAS in ETSI NFV environments [O24]

3.1.7 Vertical customer's programmability control of network functions and connectivity in a slice-as-a-service schema

On-demand Slice-as-a-Service is foreseen as the solution to be offered to vertical customers in the near future for satisfying their communication needs in the systematic digitalization of industrial segments. While mechanisms for provisioning those slices are progressively being introduced in operational networks, it is not yet resolved how the vertical customers will be able to control not only the functions composing the final service but also the underlying connectivity. In that manner, the vertical can have full control of the allocated resources as if the slice was actually a separate network. Here an architectural approach leveraging on CLAS is proposed, considering some scenarios of applicability.

Once the slice is instantiated, the vertical customer is provided with access to both Function and Network control. The Function control permits to manage and configure the service functions composing the end-to-end vertical service in the different data centers or NFVI Points of Presence (PoPs) where those functions are deployed. This Function control is linked to the service stratum, then associated to the particular service logic of the vertical. The service functions are deployed as dedicated virtual machines or containers, for instance, on commodity servers within the data center. The Function control has a correspondence with the Tenant SDN Controller of the architecture in Section 3.1.6. Such correspondence could be as loose as a simple interaction through a well-defined API (as provided by the specific

provider of the function to be controlled), or can be as tight as a direct implementation of the Tenant SDN Controller itself, e.g. through a virtualized instantiation of such controller for the allocated functions.

The Function controller, through the interaction between the Tenant SDN Controller and the Infrastructure SDN Controller in each data center, can indirectly interact with the networking infrastructure of each DC, that is, the router acting as data center gateway (connecting to the WAN) and the leaf and spine switches that typically serve for interconnecting the compute nodes hosting the functions. This interaction can happen for instance to accompany scaling events at the service function side (e.g., [O25]), thus ensuring that the connectivity associated to the capacity of the network function does not become a bottleneck or, on the contrary, that some connectivity resources can be released, if needed.

The Network controller will assist the vertical customer on the control of the transport connectivity on the WAN. The Network controller has a correspondence with the *WIM agent* described in [A11], and again such correspondence could be loose or tight depending on if it is made through APIs (through an agreed interface, e.g. the ONF Transport API [O26]) or as an actual instantiation of the *WIM agent*.

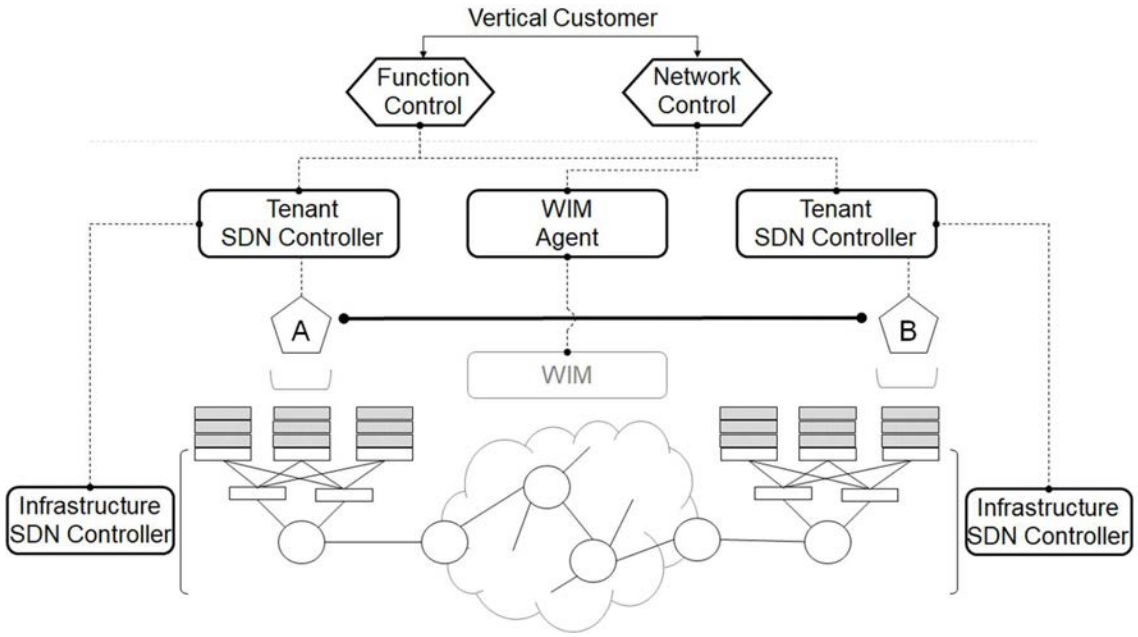

Figure 3-3. Architecture enabling vertical customer control of service functions and their connectivity

As a result, the vertical customer will perceive the slice not only as a dedicated logical network dedicated to it, but also as a *fully controllable* logical network, following the operational model of external slice managed by the vertical customer defined in Section 2.3.

3.1.7.1 Scenarios of applicability

The following scenarios of applicability illustrate how the proposed architecture can assist on the operation of vertical slices.

3.1.7.1.1 Service chain across multiple NFVI PoPs

The ETSI NFV architecture supports multi-Point of Presence (PoP) configurations, where a PoP is defined as the physical location where network functions are instantiated, corresponding to a data center. When expanding more than one single PoP, multi-site connectivity should be performed to connect the functions local to each PoP, as described in [O27]. For a vertical customer running the service end-to-end, it can be necessary to interact with the WAN for ensuring the proper behavior of the service functions interconnection, forming a service chain. Such service chains can be mapped to a dedicated slice. In this manner, if the vertical customer requiring the slicing needs to perform such control actions on both the service and the connectivity, the operator can provide access to both service functions and network resources assisting the vertical in the management and control of the inter-site network slice connectivity.

The architectural schema in Figure 3-3 represents this scenario.

3.1.7.1.2 Non-Public Network (NPN) integrated with operator's network

The digital transformation of productive environments by vertical industries is enabling the emergency of non-public networks (NPN) connected to public operator's networks. Different interconnection models are foreseen ranging from standalone NPNs up to NPNs that show different degrees of integration with public networks [A12]. The latter are the ones of interest here, since such model implies also some level of interaction among the vertical service and the operator infrastructure.

This scenario assumes that the vertical customer leverages on the operator infrastructure for deploying part of its end-to-end service. This requires from the operator side the enablement of the programmable control of functions and connectivity. Taking as example 5G-based services, and assuming that a User Plane Function (UPF), in charge of forwarding and processing the user plane traffic, is deployed in the internals of the NPN, this would imply that the vertical is able to manage and control not only the service functions forming the vertical service (e.g., data bases, content repositories, or even 5G control plane functions) but also the connectivity among the NPN UPF and the rest of elements (UPF and others) deployed at the operator's side. Figure 3-4 graphically describes a generalization of this case.

3.1.8 Summary of the contribution

This section has described an original architecture proposing cooperative SDN layers, named CLAS (released as an IETF RFC [A13]). This same architecture has been reflected in different environments, including ETSI NFV (contributed to [O24]). Furthermore, it has been described how this architecture can permit vertical customers to control the allocated slice, allowing function and network control (under submission as journal paper [A7]).

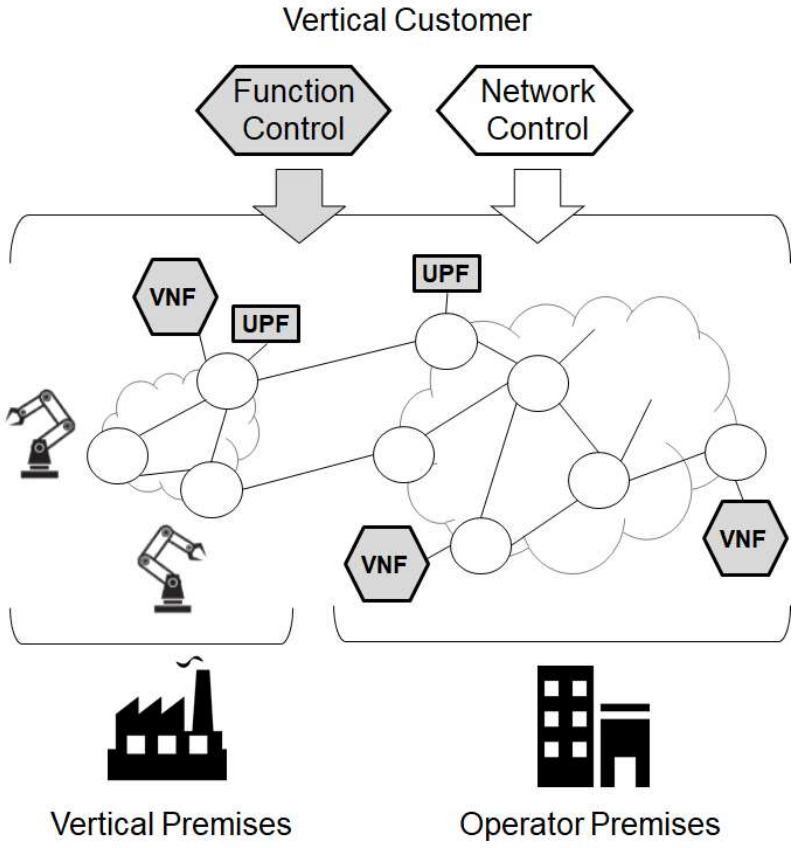

Figure 3-4. Vertical customer control of service functions and their connectivity in NPN scenario

3.2 Interconnection of multi-provider infrastructures for multi-domain service and resource orchestration

This section addresses a number of network interconnection scenarios involving multiple administrative domains.

3.2.1 Complementary components for service orchestration in multi-domain interconnected networks

Multi-domain orchestration allows virtualized network functions to be instantiated in computing facilities available in different administrative domains as reported in [O28] and further elaborated in [O29]. That is the case for service providers offering their NFVI-PoPs to host third party service functions or even offering VNFs to be consumed by others. Thus, the service can be decoupled from the hosting infrastructure.

Multi-domain orchestration implies some level of information exposition between providers, as well as the availability of interfaces or APIs for control and management operations.

Two distinct levels of orchestration could be considered. On one hand, Resource Orchestration (RO), where one operator makes available resources to another operator for instantiating services in the form of interconnected VNFs. On the other hand, Network Service Orchestration (NSO), in which one operator directly offers the instantiated VNFs for composing an end-to-end service together with its own functions. The NSO manages the lifecycle of network services, while the RO provides an overall view of the resources present in the administrative domain to which it provides access.

Assuming that the orchestration domains are based on the ETSI MANO framework, a suitable architecture for accomplishing a multi-domain scenario is the architecture defined by the 5G-Exchange (5GEx) project [O30]. Such architecture allows resources such as networking, connectivity, computing and storage in one operator's domain to be traded among federated operators using this exchange, and thus enabling service provisioning on a global basis. In this vision, 5GEx is the enabler, which facilitates operators to buy, sell, and integrate infrastructure services. It provides the ability to automatically trade resources, verify requested services, and account for billing and charging aligned to resource consumption. An insight on specific use cases and the business dimension of 5GEx can be found in [O31].

A set of APIs and interface protocols implement the exchange from the control plane perspective. Also, from the data plane, it is not necessary for a static and direct connection of physical appliances. Intermediate, transit networks participants of the end-to-end service (and its control procedures) can provide the connecting paths between both parties.

One of the key features inherited by NFV is the separation of services from resources. The NFV architectural model describes a network service (as well as its component VNFs) as a packaging of virtual and physical resources plus an application utilizing them to implement network functions usually executed by a network appliance. The description of the VNFs and services is typically an offline function, eventually populated via a catalog of available services. At operation time, it is possible to select a network service and invoke its provisioning. The key element for accomplishing it is the Orchestrator that decides which infrastructure node is the more convenient for the service deployment counting with the necessary logical resources needed by the service. For that, dynamic information is interchanged among operators with the invoked service description, which contains detailed information about the requirements of the different VNFs (not only in terms of resources, but also geographic affinity, redundancy, etc.). Moreover, information from the underlying interconnected infrastructure(s) is advertised. This resource allocation is dynamically controlled and adapted, based on the monitoring of service performance, SLA enforcement, availability or security triggered reallocations, etc.

Figure 3-5 highlights the scope of 5GEx system by presenting a logical interworking architecture, showing not only functional entities but also the different APIs and operational interfaces between them. This same architecture has been considered as a solution for multi-

domain programmability through SDN [O32] or for the VNF composition across multiple administrative domains [O33].

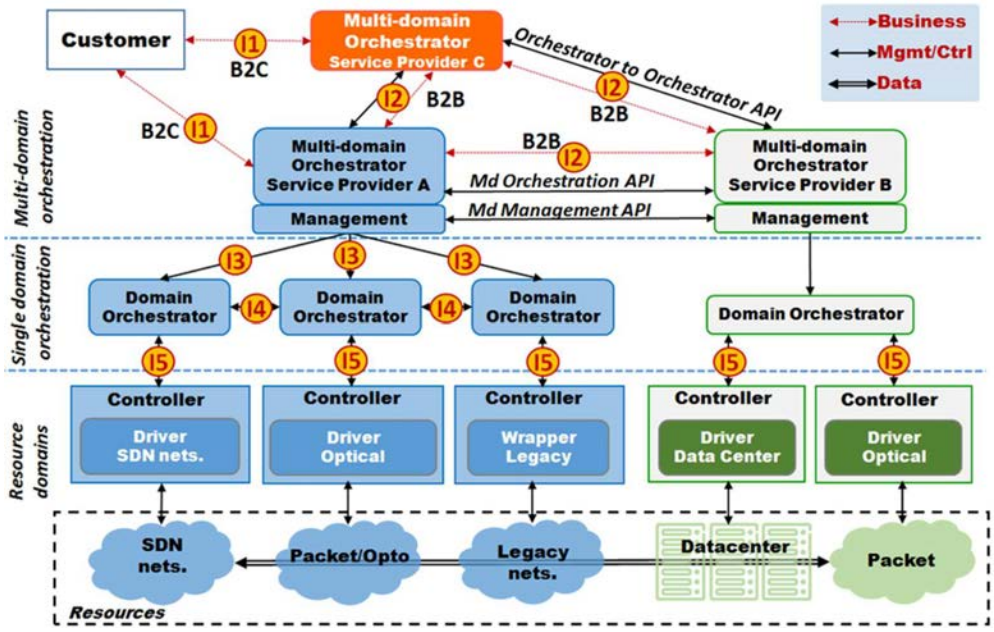

Figure 3-5. 5GEx architecture

The core of 5GEx system is composed of *(i)* the Multi-domain Orchestrator (MdO), *(ii)* various domain orchestrators and *(iii)* collaboration with domain orchestrators and controllers that are in charge of enforcing the requested services on the underlying network, compute, and storage components.

Co-operation between providers takes place at the higher level through the inter-operator orchestration API (I2) that exchanges information, functions and control. This interface also serves for the Business-to-Business relation between operators in complement to the Business-to-Customer API (I1), through which customers request service deployment. The MdO maps service requests into its own resource domains and/or dispatches them to other operators through interface (I2). This interaction is performed at MdO level: each operator MdO can expose to other operators' MdOs an abstract view of its resource domains and available service functions.

The MdO enforces the decision through interface (I3) as exposed by Domain Orchestrators, each one orchestrating and managing resource domains through the northbound interfaces (I5) exposed by technology-specific controllers. Interface (I4) facilitates interaction among different Domain Controllers. 5GEx scope is focused on interfaces I1, I2 and I3. Full 5GEx architecture details can be found in [O34].

Using such an interworking architecture for multi-domain orchestration makes possible to deploy use cases such as the virtualized multi-domain roaming case, described in Section 4.1,

hard to tackle otherwise due to the interaction requirements of multiple heterogeneous actors and technologies.

Despite the 5GEx architecture enables the orchestration across multiple administrative domains, the solely multi-provider orchestration, as defined in [O29], is not enough in most cases since basically it focuses on VNF instantiation and lifecycle management. On one hand, additional functionality for this kind of environments is needed, as capabilities for performing negotiation, charging, etc., in an automated fashion. On the other hand, for some services there is a need of having some other functional modules handling aspects of the service itself which are out of the logic of the multi-domain orchestration.

This is the case of the virtualized roaming solution validated in Section 4.1. Essentially, the goal of that use case is to leverage on virtualized EPC components to deploy mobile packet core entities from the home operator of users in roaming within the premises of the visited operator. In this case, the overall roaming service, besides the instantiation and deployment of a virtualized PGW (vPGW) instance as a VNF requires additional configuration such as the configuration of DNS entries for proper re-direction from the SGW of the visited network to the vPGW of the home network (deployed at the visited network's premises). That additional functionality is out of scope of the multi-domain orchestration tasks but must be accomplished in order to properly deliver the service, with the context and the logic related to them laying outside the MdO.

With that purpose, it is possible to leverage on the management functions of network slices to support communication services as defined by 3GPP [O18]. Despite the concept of network slicing has emerged with the development of 5G, its applicability is agnostic of the technology being contained in the slice. Those functions are (as already described in Section 2.4.2.2.5) CSMF, NSMF and NSSMF. Both NSMF and NSSMF can be considered aggregated as part of the same Network Service realized by means of a slice. This management functions have been mapped [O19] to the ETSI MANO framework. The resulting mapping locates this functionality as part of the broader OSS/BSS part of the framework, as shown in Figure 3-6.

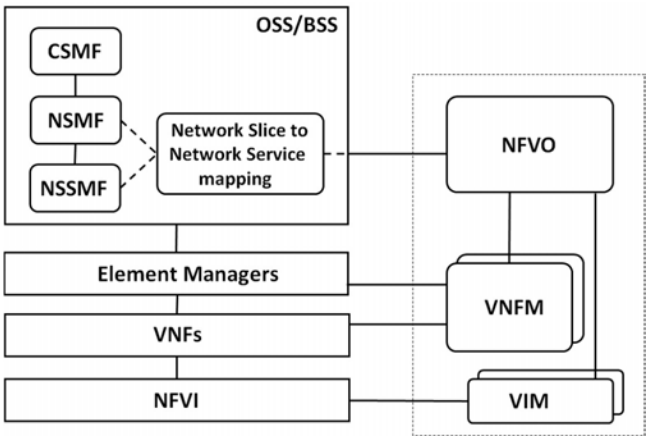

Figure 3-6. Mapping of the 3GPP network slicing concept to the ETSI MANO framework

According to the management functions described above, the CSMF could play the functionality of managing the required service logic for the virtualized roaming service to complement the multi-domain orchestration (e.g., handling the DNS re-configuration in the visited network as triggered by the home network operator for resolving the IP address which corresponds to the vPGW instantiated in the visited network). This kind of action cannot be performed from the logic of the multi-domain orchestration since the service logic could not be fully incorporated at that level. In addition to that, the configuration and management of the services and resources offered by the visited network, once allocated and deployed, can be part of the NSMF (and NSSMF). Figure 3-7 shows the complementary functionality from an architectural point of view.

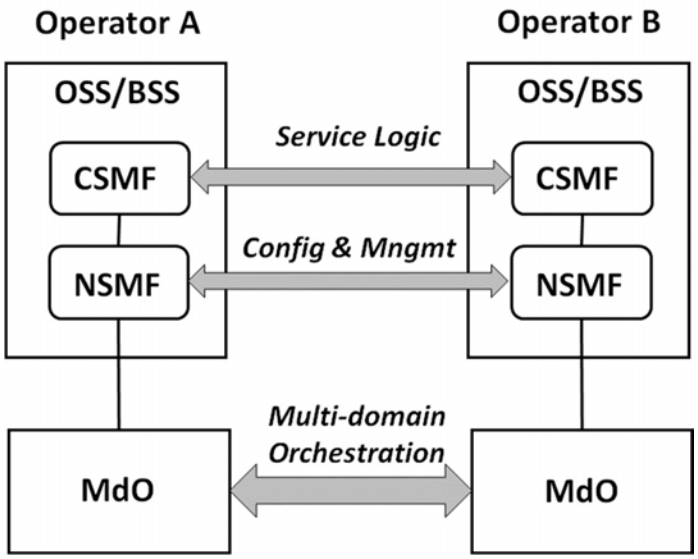

Figure 3-7. Complementary functionality for virtualized roaming service creation

These architectural additions are necessary to enable more complex service orchestration, where mere VNF orchestration cannot been sufficient, since some service aspects can lay outside that the orchestration process. This becomes exemplified in the virtualized roaming solution further developed in Section 4.1.

3.2.2 Integration of multi-domain MEC environments

The Multi-access Edge Computing (MEC) framework is originally defined as an environment managed and administered by a single network operator, which controls a number of edge computing sites defining an area of coverage.

The MEC reference architecture is described in [O35], and graphically represented in Figure 3-8. It is composed on functional components and the reference points between them. It also includes a number of multi-access edge services that complement the overall solution.

As seen in Figure 3-8, the MEC framework differentiates among multi-access edge system and multi-access edge hosts levels. The Multi-access Edge System (MES) consists of a

number of multi-access edge hosts and the multi-access edge management entities necessary to execute multi-access edge applications within an operator network.

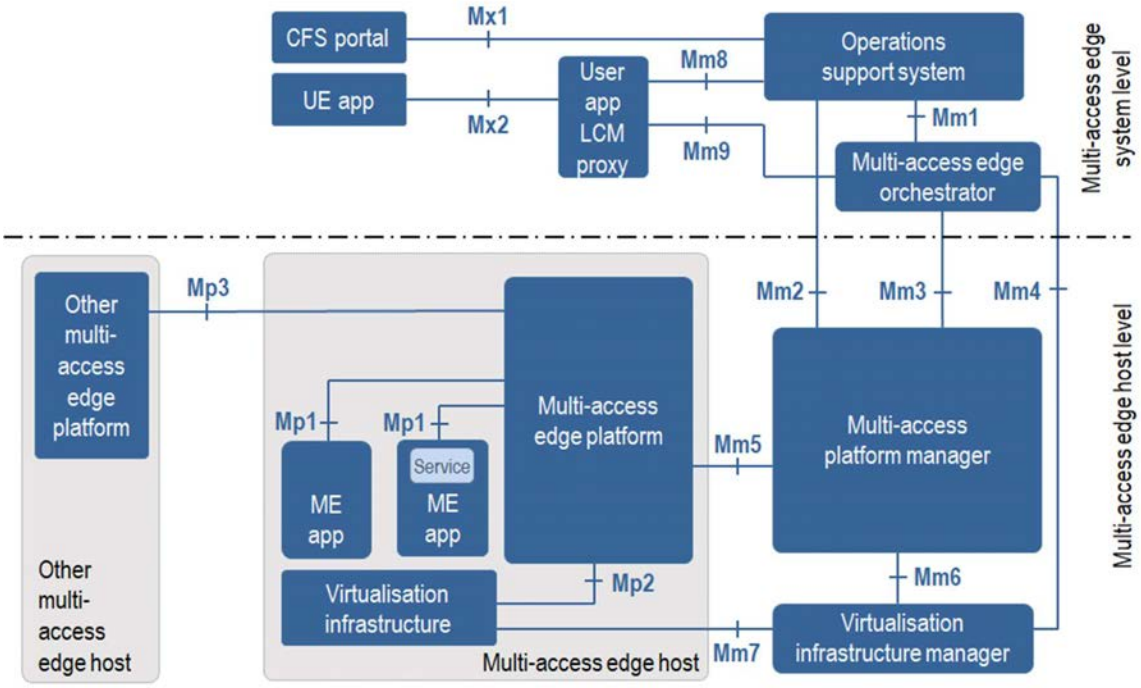

Figure 3-8. MEC reference architecture

The Multi-access Edge Host (MEH) is an entity that contains a Multi-access Edge Platform (MEP) and a virtualization infrastructure.

The MEP provides a functional environment where applications can discover, advertise, consume and offer multi-access edge services. The MEP controls the data-plane in the virtualization infrastructure following an SDN approach, configures the DNS proxy/server in the MEH based on DNS records obtained from the multi-access edge platform manager, and provides access to persistent storage and time of day information.

The virtualization infrastructure is, generally speaking, a Network Function Virtualization Infrastructure (NFVI) as the one described in [O36], which provides compute, storage, and network resources, for running multi-access edge applications on top of it. The virtualization infrastructure includes a data plane that executes the traffic rules received by the MEP, routing the traffic among applications, services, DNS server/proxy, and both local networks and external networks.

Then, the multi-access edge applications, running as virtual machines, are instantiated on the virtualization infrastructure of the multi-access edge host (forming an NFVI point of presence, or NFVI-PoP) based on configuration or requests validated by the multi-access edge management. They can either consume or provide multi-access edge services present in the MEH. These applications could be even relocated to another multi-access edge host, if

supported by the system and the application. Typically, they have associated rules and requirements (e.g., traffic redirection, DNS re-configuration, maximum latency, etc.), that are enforced by the multi-access edge system level management.

The Multi-access Edge Platform Manager (MEPM) acts as the element management of the MEP, performing the management of the application rules, including service authorization, traffic rules, DNS configuration, conflict resolution, etc., and also performing the lifecycle management of the multi-access edge applications, including the notifications towards the orchestrator of any application related lifecycle event. Finally, through the interaction with the VIM, it receives virtual resource fault reports and performance measurements coming from events in the virtualization infrastructure.

Already at system level, the Multi-access Edge Orchestrator (MEO) maintains an overall view of the system and multi-access edge hosts, including available resources, available services and topology. It selects the appropriate host for each application, satisfying its rules and requirements, then triggering the application instantiation, relocation and termination. The MEO is also in charge of on-boarding the application packages.

Finally, the MEC architecture is completed by operation support systems (OSS). These OSSs can receive requests from external entities, either from the user application lifecycle management proxy (UA-LCM proxy), or the customer facing service (CFS) portal, for multi-access edge application instantiation, termination or relocation, determining if such requests can be granted, and in that case, forwarding granted requests to the orchestrator. These OSSs also allow the network operator to trigger management and control actions, including the configuration of policies for the execution of the applications.

Apart from these functional blocks and components, the MEC architecture defined a number of reference points among them. These have been summarized in Table 3-1, presenting the components involved as well as the main scope of each reference point. Here it is analyzed the impact of the multi-domain approach on those reference points for the different scenarios evaluated.

The following sub-sections consider distinct alternatives for multi-domain integration by identifying options for establishing possible administrative domain boundaries with respect to the MEC reference framework in Figure 3-8.

The motivations for going multi-domain can be diverse: savings at the time of deploying full MEC solution; limitation in the access to certain geographical locations; tailored services for specific customer that could require an ad-hoc deployment of MEC capabilities; etc. The following subsections present different alternatives followed by a business rationale for them.

In the accompanying figures the primary domain is labeled as Domain A, while the secondary domain is labeled as Domain B. The different administrative domains are highlighted in different colors in order to easily distinguish the components from each domain in the constitution of the resulting MEC system. Additionally, the reference points requiring multi-domain support are labeled with the prefix “MD-” for clarity.

Table 3-1. Summary of the MEC reference points and their scope.

	Reference point	Components involved	Scope
Multi-access Edge Management	Mm1	OSS-MEO	Triggering the instantiation and the termination of multi-access edge applications in the multi-access edge system
	Mm2	OSS-MEPM	Multi-access edge platform configuration, fault and performance management
	Mm3	MEO-MEPM	Management of the application lifecycle, application rules and requirements and keeping track of the available multi-access edge services
	Mm4	MEO-VIM	Management of virtual resources of the multi-access edge host, including keep track of available resource capacity, and to manage application images
	Mm5	MEPM-MEP	Platform configuration, configuration of application rules and requirements, application lifecycle support procedures, management of application relocation, etc
	Mm6	MEPM-VIM	Management of virtualized resources e.g. to realize the application lifecycle management
	Mm7	VIM-NFVI	Management of the virtualization infrastructure.
	Mm8	OSS-UE LCM proxy	Handling of UE applications requests to run application in the multi-access edge system
	Mm9	MEO-UE LCM proxy	Management of multi-access edge applications requested by UE application
External entities	Mx1	OSS-CFS portal	Requests to the multi-access edge system by the third parties to run applications within the multi-access edge system
	Mx2	UE app-UE LCM proxy	Requests to the multi-access edge system by a UE application to run an application in the multi-access edge system, or to move an application in or out of the multi-access edge system
Multi-access Edge Platform	Mp1	MEP-ME app	Provides service registration, service discovery and communications support for services. It also provides other functionality such as application availability, session state relocation support procedures, traffic rules and DNS rules activation, access to persistent storage and time of day information, etc. This reference point can be used for consuming and providing service specific functionality
	Mp2	MEP-NFVI	Instructs the data plane on how to route traffic among applications, networks, services, etc.
	Mp3	MEP-other MEP	Control communication between multi-access edge platforms

In all of the alternatives presented, it is assumed that the primary domain always retains all the commercial interaction with the MEC customer. This is applicable to the case where the MEC customer wants to deploy an application in the multi-domain MEC systems, but also in the case that the MEC customer wants to make use of an application provided by the secondary domain. In the latter, the primary domain acts as mediator for such interaction.

3.2.2.1 Integration at infrastructure level

A first integration approach would be to consider the usage of infrastructure from a different provider, nominally an infrastructure provider. This situation can be assimilated to an Infrastructure-as-a-Service (IaaS) offering in the cloud computing business. The business motivation for this kind of integration could be that of a MEC provider requiring increasing its footprint in a given geographical area with restrictions for deploying new infrastructure, then leveraging on some available infrastructure, for instance provided by a municipality. Alternatively, it could be the case of an initial and fast deployment of MEC services in a certain location while the own infrastructure is being built for that same area.

Since MEC makes use of NFVI environments for hosting the applications and other virtualized functions, this scenario leads as well to an integration of NFVI environments, probably requiring the interconnection of the overall NFVI substrates used by the MEC provider. Figure 3-9 represents the administrative boundary among providers in this model.

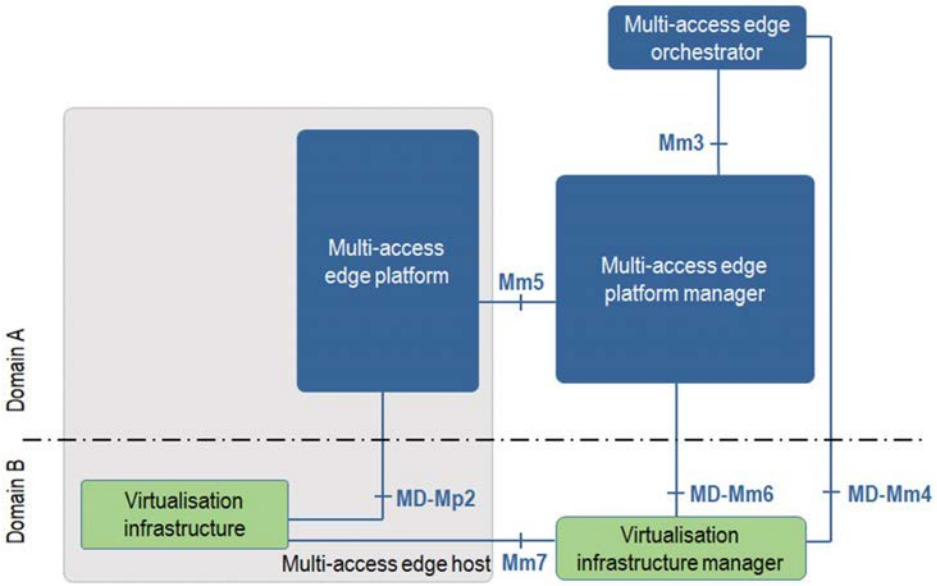

Figure 3-9. Integration at infrastructure level

In this alternative, the MEP from Domain A controls and programs the virtualization infrastructure through the MD-Mp2 reference point following SDN principles, then it is important that the resources allocated from Domain B to Domain A remain isolated from some other resources in Domain B in order to avoid any kind of conflicting configuration action. This could be performed e.g., by providing a specific resource slice to Domain A.

The MD-Mm4 and MD-Mm6 interfaces depend on the VIM used by the infrastructure provider. It can be assumed the usage of some open solution for the VIM, such as e.g. OpenStack.

One component that could be considered apart is the VIM itself. The VIM could be provided or not by the infrastructure provider, that is, Domain B. Alternatively, the MEC provider in Domain A could leverage on the concept of VIM-on-demand [O37] for instantiating a VIM on top of the virtualization infrastructure fully under control of the MEC provider. This could facilitate the integration, since the VIM-on-demand could be prepared in advance with the necessary capabilities for making the integration smooth. This would simplify (or even remove) the requirements to be supported by the MD-Mm4 and MD-Mm6 interfaces, since could appear as being part of the same domain of the MEC provider.

3.2.2.2 Integration at platform level

A different approach could be the integration with a domain that implements the MEP and possibly some specific applications. This approach can be perceived as a Platform-as-a-Service (PaaS) offering, also in analogy with cloud computing world.

The business rationale for this integration model could be the one of a primary provider, Domain A, willing to leverage on the applications of a secondary provider, Domain B, which could retain the rights for the integral exploitation of such applications, including the value added of the functionalities provided by the MEP itself, thus leading to the PaaS concept.

The integration at PaaS level could present two different sub-scenarios: *(i)* integration with the platform provider with infrastructure owned by the primary MEC provider, Domain A; and, *(ii)* integration with the platform provider, Domain B, including its infrastructure. Both scenarios are shown in Figure 3-10 and Figure 3-11 respectively.

The first situation, when the primary MEC provider provides also the infrastructure, implies that the platform provider instantiates in advance the MEP function on top of the primary MEC provider infrastructure. This could be done in the form of a VNF e.g. by leveraging on the integration model of MEC and NFV as defined in [O38].

In this case it can be assumed that the virtualization infrastructure of Domain A is fully controlled by the MEP of Domain B through the MD-Mp2 interface, as result of the indications from the MEPM of Domain A via the MD-Mm5 interface. The MEP from Domain B could interact with other MEHs from either Domain A or Domain B by means of the MD-Mp3 interface.

In the second situation, when the platform provider includes the supporting infrastructure, the platform provider could be a remote provider. In these circumstances, the MD-Mm5 interface behaves as before, however it is required an integration with the VIM, which as mentioned before could be done through open interfaces in case the VIM is an open source solution such as OpenStack. Additionally, as in the IaaS case, the main MEC provider, Domain A, could leverage on the concept of VIM-on-demand for facilitating the integration and control of the resources granted by the platform provider to it.

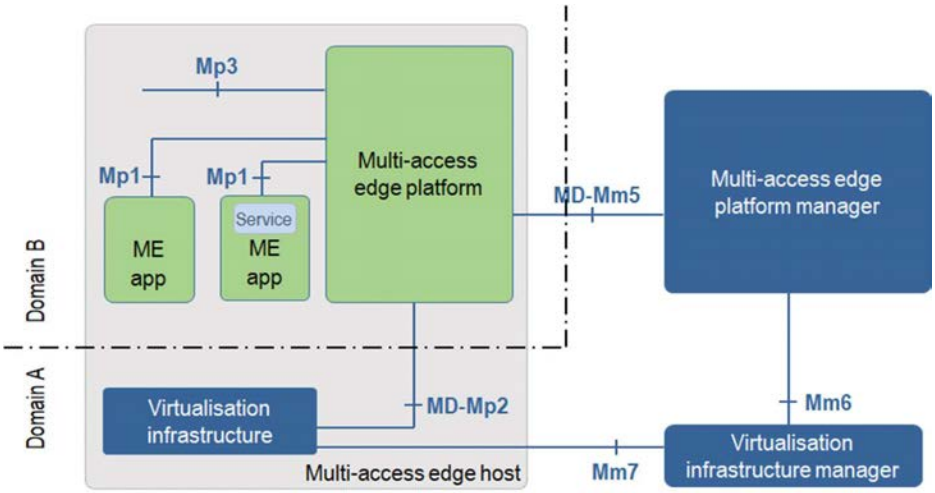

Figure 3-10. Integration at platform level with infrastructure owned by Domain A

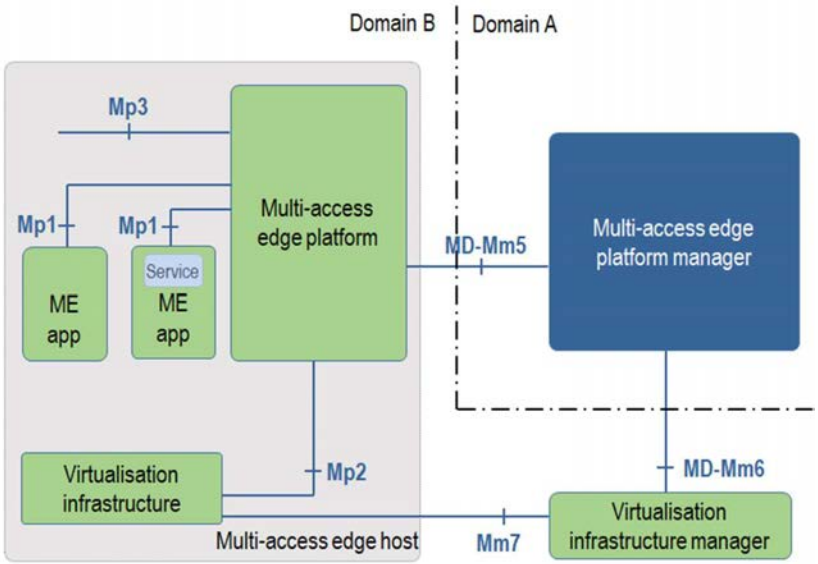

Figure 3-11. Integration at platform level with infrastructure owned by Domain B

3.2.2.3 Integration at MEC service level

In this case, the primary domain, Domain A, implements only the MEO function, interconnecting to the MEPM and the VIM of the secondary domain for the orchestration of the applications as provided or enabled by Domain B. Figure 3-12 graphically depicts this case. Since the secondary domain provides all the capabilities for management of the lifecycle of the applications, this approach can be seen as an outsourcing of all of that functionality from provider in Domain A to provider in Domain B. Then provider in Domain A basically focuses on the commercial relation with the MEC customer (and the end users) and in the decisions about instantiating and running applications in the system.

The business motivation for this integration model could be the one of a main provider acting as aggregator of MEC systems either to increase coverage or to complement its own offer with additional capabilities or applications. The secondary provider retains all the logic for handling the lifecycle of the applications, with the main provider triggering instructing what to do in each moment.

The interaction among providers is done through the management interfaces MD-Mm2, MD-Mm3 and MD-Mm4, then having management interaction from Doman A with the platform and the infrastructure of Domain B.

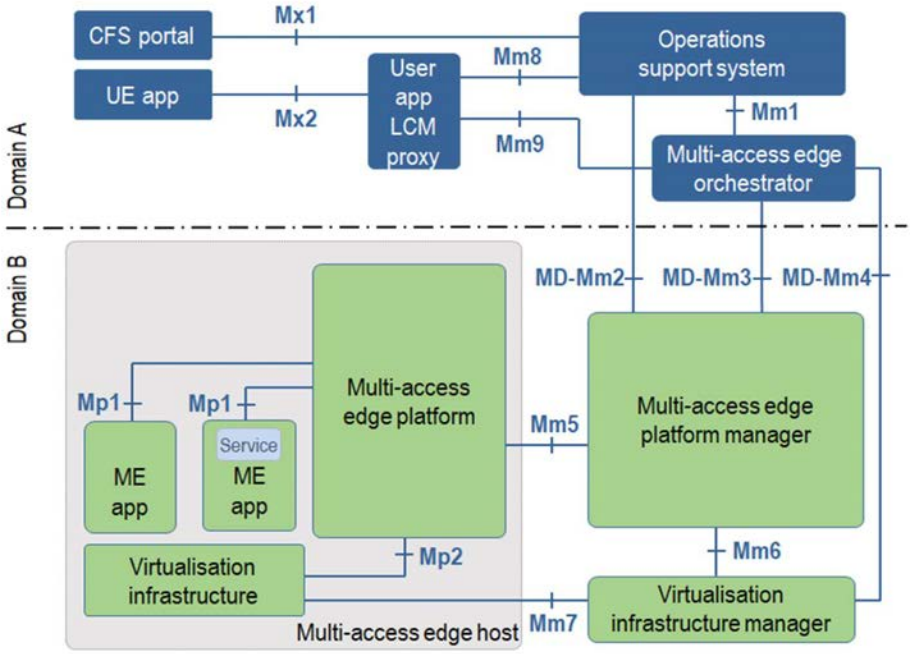

Figure 3-12. Integration at service level

3.2.2.4 Interconnection of MEC systems

The last scenario of integration is the pure interconnection of MEC systems. Here it is considered that such interconnection is done at MEO level, as presented in Figure 3-13, by the definition of a new external interface, named MD-Mx3, in consistency with Mx1 and Mx2 interfaces as already defined in the MEC reference architecture.

The business rationale for this option is the alliance of full MEC providers, which federate for offering a more complete commercial offer to their respective MEC customers. Each of the providers in the federation have its own portfolio and customer base, but they can leverage in the federation in order to constitute a more compelling commercial offer in terms of coverage, services, etc. In the more extreme case of interconnection, it could be even possible for a MEC provider to implement only the MEO, that is, without own resources nor platform. In this situation, such provider would play a role of broker of MEC systems from

some other MEC providers that could participate in a kind of exchange or federation of MEC systems.

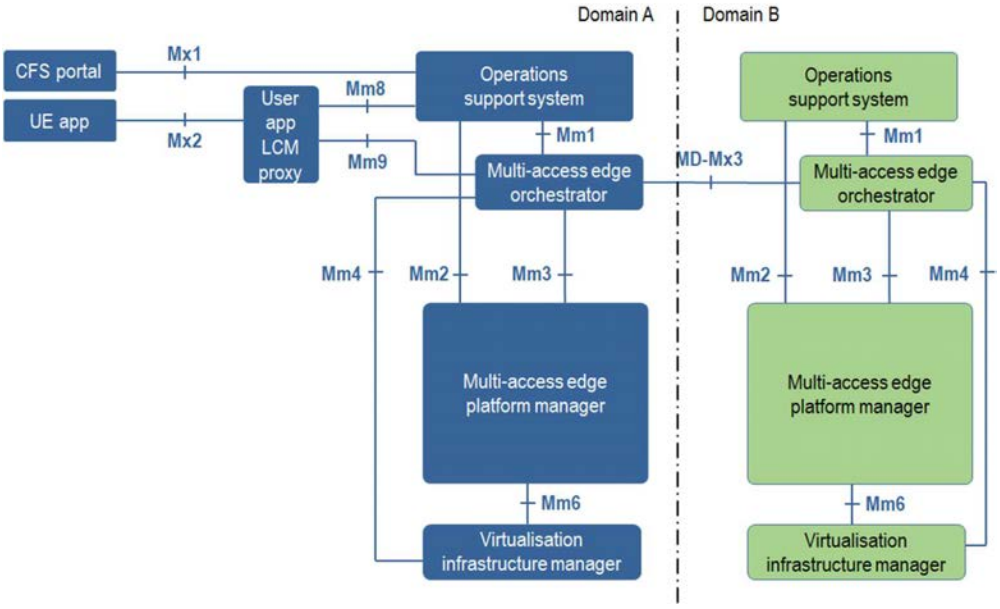

Figure 3-13. Interconnection of MEC systems

3.2.2.5 Summary of alternatives

A number of alternatives have been analyzed depending on where the administrative domain boundary is located in a multi-provider MEC scenario. This will influence the reference points and the MEC components that have to be scoped for multi-domain, potentially by the inclusion of some multi-domain adaptor able to handle the extra functionality needed for multi-domain integration.

Table 3-2 summarizes the findings for each scenario, including the interfaces impacted in each case.

As can be seen, each of the scenarios has different implications on where should reside tthe awareness of the multi-domain interaction, at both MEC component and reference points. Different strategies can be considered and their impacts should be evaluated. A primary indication of the implications at both business and technical levels is also included in Table 3-2. Only in the last case of MEC systems interconnection, there is naturally an impact on the external interfaces declared in the MEC architecture, since the other domain is connected at management system level.

Currently, ETSI MEC is working on a study for inter-MEC systems coordination, as future MEC 035 report. Its latest available version at the time of writing, in [O39], considers basically coordination at system level. Thus, the proposed IaaS, PaaS and SaaS approaches here complement that work.

Table 3-2. Summary of multi-domain integration alternatives

Scenario	Existing interfaces going multi-donain	New interfaces for supporting multi-domain	Comments	Implications
Integration at infrastructure level	Mm4, Mm6, Mp2	--	Resources allocated by Domain B to Domain A have to be isolated (e.g., by means of a slice) to avoid conflicts in the control of them. VIM could be instantiated on-demand by Domain A.	Business – IaaS model for Domain B; SLAs tight to resource capabilities (compute, networking). Technical – Domain B to provide monitoring information of resources; abstraction data models for resources; multi-domain awareness extended to management and platform reference points.
Integration at platform level with infrastructure owned by Domain A	Mm5, Mp2	--	MEP from Domain B can be instantiated as VM on Domain A. MEP from Domain B can interact with other MEHs either from Domain A or B.	Business – PaaS model for Domain B; SLAs related to platform KPIs (e.g., provisioning delay). Technical – Domain B to provide monitoring information of the platform; abstraction data models for MEC platform; multi-domain awareness extended to management and platform reference points.
Integration at platform level without infrastructure owned by Domain B	Mm5, Mm6	--	MEPM from Domain A can interact with the MEP from Domain B remotely. VIM could be instantiated on-demand by Domain A.	Business – PaaS model for Domain B; SLAs extended for including platform and resource related KPIs. Technical – Domain B to provide monitoring information of the platform and resources; abstraction data models for resources and platform; multi-domain awareness retained only on management reference points.
Integration at service level	Mm2, Mm3, Mm4	--	Domain A acts as an integrator of MEC services from other providers e.g., Domain B.	Business – New business model for Domain B by offering MEC host level outsourcing to Domain A; SLAs including platform and resource related KPIs. Technical – Domain B to provide monitoring information of the platform manager, the MEC platform itself and the resources; abstraction data models for resources, platform and platform manager; multi-domain awareness retained only on management reference points.
Interconnection of MEC systems	--	Mx3	The providers from an alliance or federation completing their particular commercial offers when necessary. A new interface is required for this scenario.	Business – Extension to MEC of peering and/or federation business model; SLAs including overall MEC related KPIs. Technical – Domain B to provide MEC monitoring information; abstraction data models for overall MEC system; multi-domain awareness in a new external interface for MEC interconnection.

3.2.3 Multi-domain slicing

The idea of leasing virtualized networking and computing environments is gaining momentum. Thus, Infrastructure Providers (InP) can play the role of facilitators for service providers in order to lower the Total Cost of Ownership (TCO), simplify the network architecture, and streamline the operation and their associated costs.

This can be significantly the case for access and aggregation networks. Uncertainty in the number of end users, their distribution and mobility patterns and heterogeneous service requirements (from data intensive residential-like service to flow-intensive machine-to-machine connections) make unpredictable and dynamic the demand of connectivity and network services.

Specifically, for the aggregation stages, close to the radio access (what is typically known as a conjunction of fronthaul and backhaul areas, or crosshaul) seems quite appealing to introduce flexibility to dynamically adapt the deployed resources to the concrete demand. The demand of dynamic resource allocation involves networking but also computing facilities, in order to flexibly deploy services and host content at the edge, thus saving core network capacity and decreasing service latency.

Crosshaul areas are those integrating fronthaul and backhaul and capable of providing the capillarity to reach the final end users (i.e., those consuming the offerings of the vertical industries). Finding a way of trading crosshaul capabilities becomes relevant to facilitate the deployment of the vertical services in a smooth and normalized manner, without particularization per vertical client. Therefore, open environments are desired for such a one-shop ecosystem.

There is yet a gap to reach the goal of hosting crosshaul in a multi-domain federated infrastructure: a market place where networking and computing facilities are traded. An extension of the traditional concept of telco exchange is needed, covering new needs and capabilities, such as offering resource slices for deployment of the services requested by third party service providers.

Two key characteristics define this new generation of transport networks, which integrate the fronthaul and backhaul segments. On the one hand, the concept of multi-tenancy, where a network is shared among different customers or tenants, being an integral part of the design of future networks, designed from scratch to be shareable, reducing their CapEx and OpEx costs and increasing their usage efficiency. On the other hand, networks become tailor-made for the needs of each tenant. The customization of network services and characteristics allow different network views to coexist on top of a common infrastructure. The advanced concept of slicing is a target on such scenario of multi-domain infrastructure sharing.

In this section, the architectures of the EU H2020 projects 5G-Crosshaul and 5G-Exchange are considered to exemplify the way in which this trading of slices could happen among providers. This section analyzes the way of integrating both architectures presenting the basis for interworking of them, providing hints for integration and identifying architectural gaps.

3.2.3.1 5G-Crosshaul control mechanisms

5G-Crosshaul aims to integrate the fronthaul and backhaul segments providing capillarity for distributed 5G radio access systems. A detailed description is provided in [O40].

5G-Crosshaul is based on three main building blocks: *(i)* a control infrastructure using a unified, abstract network model for control plane integration (Crosshaul Control Infrastructure, XCI); *(ii)* a unified data plane encompassing innovative high-capacity transmission technologies and novel latency-deterministic switch architectures (Crosshaul Forwarding Element, XFE); and *(iii)* a set of computing capabilities distributed across the network (Crosshaul Processing Units, XPUs). Figure 3-14 shows the architectural framework of 5G-Crosshaul.

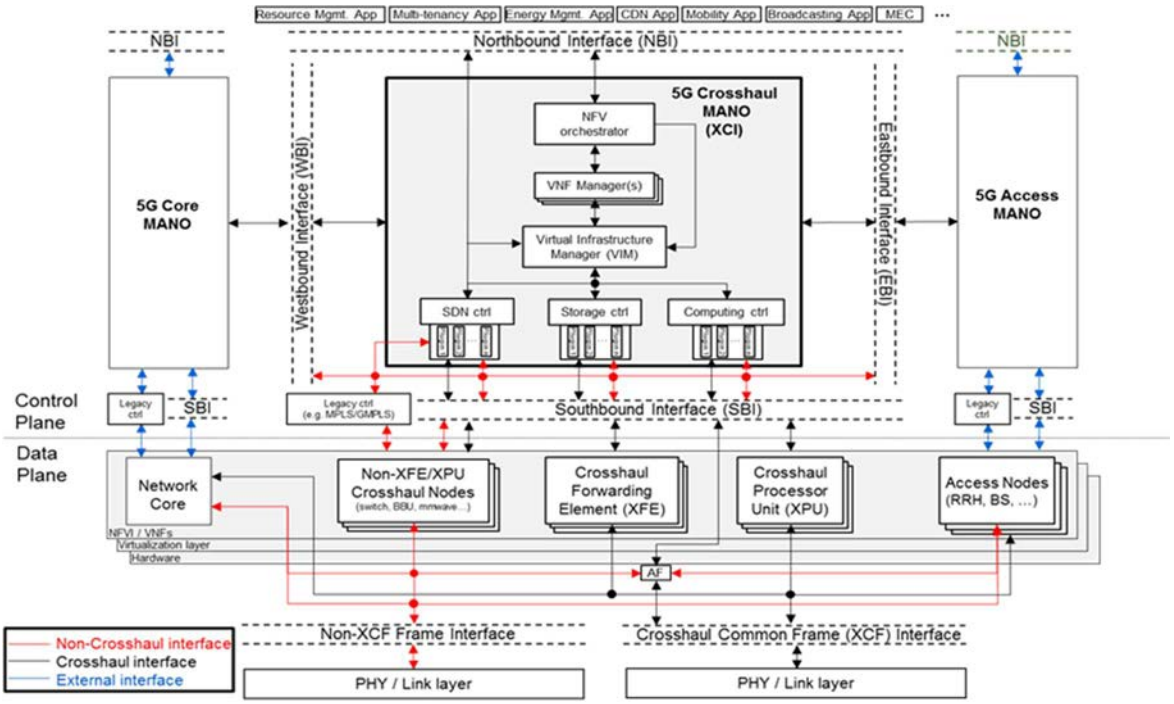

Figure 3-14. 5G-Crosshaul reference architectural framework

The XCI is compliant with the ETSI NFV architecture [O9] with regards to management and orchestration. Additionally, it includes a set of controllers for managing networking, storage and computing resources. Furthermore, a number of applications can be located on top of the XCI. Of special interest for multi-domain environments is the Multi-Tenancy Application (MTA), conceived to support per-tenant infrastructure management in multi-tenancy scenarios. The following sub-sections provide more details in two key aspects of the 5G-Crosshaul architecture: the XCI as the control framework, and the support of multi-tenancy on the same crosshaul infrastructure.

3.2.3.1.1 5G-Crosshaul control framework: the XCI

The XCI is the brain controlling the overall operation of the 5G-Crosshaul. The XCI part dealing with NFV comprises three main functional blocks, namely the NFV orchestrator (NFVO), the VNF Manager (VNFM) and the Virtual Infrastructure Manager (VIM).

In addition to these modules, which are in charge of managing the different VNFs running on top of the 5G-Crosshaul, the XCI includes a set of specialized controllers. The SDN Controller is in charge of controlling the underlying network elements following the conventional SDN paradigm. 5G-Crosshaul aims at extending current SDN support of multiple technologies used in transport networks (e.g., micro-wave links or Ethernet-based

forwarding elements) in order to have a common SDN controlled network substrate which can be reconfigured based on the needs of the network tenants. In addition to that, a Cloud Controller is proposed for handling Storage and Computing resources (e.g., OpenStack). Note that these controllers are functional blocks with one or multiple actual controllers (hierarchical or peer-to-peer structure) that centralize some or all of the control functionality of one or multiple network domains.

XCI components are based on REST APIs by design. Through those APIs, XCI exposes capabilities for different services, namely:

- Network Topology and Inventory, providing information regarding the network, including physical topology, as well as inventory related to network node and port capabilities;
- IT Infrastructure and Inventory, with similar scope as before but focused on the IT part;
- Provisioning and Flow actions, facilitating the request, the installation, and removal of forwarding rules in the network nodes;
- Statistics, for the collection of monitoring information of both network-related and IT-related statistics. This can be complemented by another service, Analytics for Monitoring, in charge of offering to the consumer elaborated information obtained from the processing of the network and computing statistics gathered before;
- NFV Orchestration, VNF Management, and Virtual Infrastructure Management, with similar scope as the defined in ETSI NFV architecture framework, with some augmentation in the latter case by adding planning capabilities (resulting in a Virtual Infrastructure Manager and Planner, VIMaP);
- Local Management Service, for managing the status and properties of the 5G-Crosshaul elements namely XFEs and XPUs.

It is then required to consider this architectural structure for the integration with other administrative domains.

3.2.3.1.2 Multi-tenancy Support

Support of multi-tenancy has a strong impact on the XCI components. The SDN controller must support the provisioning of isolated virtual network infrastructures with a given set of capabilities. Traffic isolation can be achieved through the creation of tagged network connections, configuring the flows at the forwarding elements making use of the multi-tenant features in the 5G-Crosshaul data plane, based e.g., in traffic encapsulation headers with tenant-specific tags to guarantee the proper isolation. The SDN controller needs as well to support the creation and operation of virtual networks assigned to specific tenants, which could be specified e.g., following intent-based network models.

At the VIM level, multi-tenancy is handled through the modelling of the tenant concept, where each tenant has its own view of the VIM capacity, policies to regulate the access to the resources (e.g., a quota of dedicated resources) and, optionally, custom resource flavors and VM images. Requests for new virtual infrastructures must be authenticated and authorized, and they are evaluated based on the resources still available in the tenant's quota. Finally, the access to the instantiated virtual infrastructures is strictly limited to the tenant owning the specific instance.

A similar approach, based on per-tenant profiles and policies, needs to be adopted at the NFV Orchestration level, extending the virtual resources concept to VNF and Network Services entities. Each tenant must have the view and the control on its own VNFs and NSs only; they must be maintained fully isolated from other entities belonging to different tenants, in order to guarantee their security and their desired KPI level independently on the load of other VNFs pertaining to other tenants. New service requests must be granted depending on the tenant's profile, in combination with the tenant-related policies.

3.2.3.2 5G-Exchange as enabler of multi-domain slicing and network sharing

5G-Exchange (5GEx) has been already introduced in Section 3.2.1. The control architecture of 5GEx proposes an ecosystem for the trading of resources (with the slice as extreme case) in a multi-domain environment, as described in [O30]. The initial architecture framework of 5GEx, shown before in Figure 3-5, identifies the main functional components and the interworking interfaces involved in multi-domain orchestration, where different providers participate, each of them representing a distinct administrative domain.

The key 5GEx component is the Multi-domain Orchestrator (MdO), shown at the top of the figure. The MdO handles the orchestration of resources and services from different providers, coordinating resource and/or service orchestration at multi-domain or multi-operator level, orchestrating resources and/or services using Domain Orchestrators belonging to multiple administrative domains.

3.2.3.2.1 5GEx interfaces and APIs

There are three main interworking interfaces and APIs identified in the 5GEx architecture framework. The MdO exposes service specification APIs (Customer-to-Business, C2B) that allow business customers to specify their requirements for a service on interface I1. The MdO interacts with other MdOs via interface I2 (Business-to-Business, B2B) to request and orchestrate resources and services across administrative domains. Finally, the MdO interacts with Domain Orchestrators via interface I3 APIs to orchestrate resources and services within the same administrative domains.

From the perspective of functional capabilities of such interfaces, the functional split considered on each of them is related to service management (-S functionality), VNF lifecycle management (-F), catalogues (-C), resource topology (-RT), resource control (-RC) and monitoring (-Mon). The full identification and specification of these interfaces in terms of protocols and/or APIs is defined in [O34].

3.2.3.2.2 5GEx functional architecture

The 5GEx framework reference architecture has been further developed, defining the different components and interfaces into the functional model shown in Figure 3-15. This architecture extends the ETSI MANO NFV management and orchestration framework [O9], in order to implement Network Service and Resource orchestration across multiple administrative domains, which may belong to different infrastructure operators or service providers.

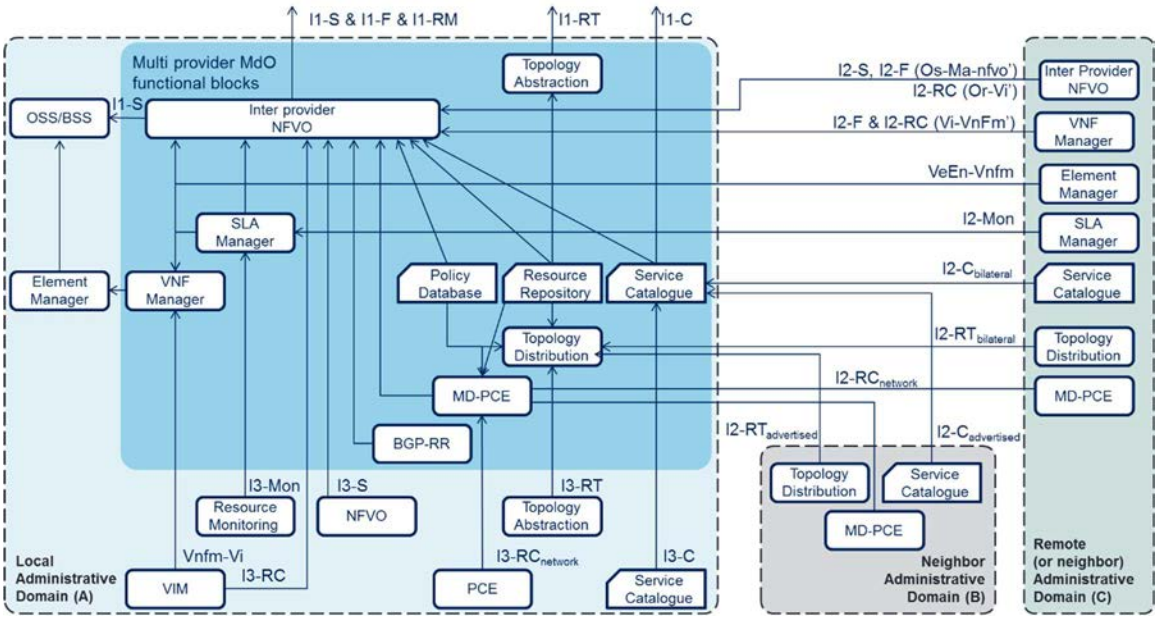

Figure 3-15. Functional model of multi domain orchestration

Figure 3-15 highlights three different administrative domains (A, B and C) involved in the Multi-Domain service/resource orchestration process. All the providers in 5GEx are considered to contain the same components and modules (the Operator-Operator relationships are symmetrical in 5GEx), although in Figure 3-15 the complete view is only shown for the provider on the left for illustration purposes, just showing exemplary consumer-provider roles with arrows from consumer to provider functional blocks. In the figure, Operator Domain A (left-hand) consumes virtualization services of Operator B (transit domain, in the middle) and Operator C (right-hand).

The main functional blocks in 5GEx are:

- The Inter-Provider NFVO is the NFVO that implements multi-provider service decomposition, responsible of performing the end-to-end network service orchestration. The NSO and RO capabilities are contained here;
- The Topology Abstraction module performs topology abstraction elaborating the information stored in the Resource Repository and Topology Distribution modules;
- The Topology Distribution module exchanges topology information with its peer MdOs;
- The Resource Repository that keeps an abstracted view of the resources at the disposal of each one of the domains reachable by the MdO;
- The SLA Manager is responsible for reporting on the performance of its own partial service graph (its piece of the multi-domain service);
- The Policy Database which contains policy information;
- The Resource Monitoring module dynamically instantiates monitoring probes on the resources of each technological domain involved in the implementation of a given service instance;
- The Service Catalogue in charge of exposing available services to customers and to other MdO from other providers.

3.2.3.2.3 5GEx interaction between components and supportive interfaces

For multi provider Network Service orchestration, the Multi-domain Orchestrator (MdO) offers Network Services by exposing an OSS/BSS-NFVO interface to other Multi-domain Orchestrators from other providers. For multi-provider resource orchestration, the MdO presents a VIM-like view and exposes an extended NFVO-VIM interface to other Multi-domain Orchestrators. The Multi Provider MdO exposes I1-S as the northbound interface through which an MdO customer (e.g., a vertical industry) sends the initial request for services. It handles command and control functions to instantiate network services. Such functions can include the request for the instantiation and interconnection of Network Functions (NFs). Interface I2-S is meant to perform similar operations between MdOs of different administrative domains.

Interfaces I2-RT and I3-RT are used to keep an updated global view of the underlying infrastructure topology exposed by domain orchestrators. In addition to that, resource orchestration is complemented with interfaces I2-RC and I3-RC to reflect resource control. The service catalogue exposes available services to customers on interface I1-C and to other MdO service operators on interface I2-C. Finally, I2-Mon and I3-Mon are used for resource monitoring.

3.2.3.3 Integration of 5G-Crosshaul and 5G-Exchange architectures

After the review of both 5G-Crosshaul and 5GEx architectures, it becomes clear that functional adaptation is feasible for allowing the trading of 5G-Crosshaul slices through 5G Exchange, since the common starting point is the ETSI MANO NFV management and orchestration framework. However, there are yet some gaps that would require from certain extension in 5G-Crosshaul for full compliance with a 5GEx ecosystem. This section summarizes both aspects as follows.

3.2.3.3.1 Integration in 5GEx of existing functional blocks from 5G-Crosshaul architecture

Some functional components can be identified as common and existing in both architectures, then foreseeing a straightforward functional integration in this respect. The following components are considered common in both architectures under functionality viewpoint.

1) Statistics and monitoring of crosshaul resources. The 5G-Crosshaul architecture supports the collection of both compute and network statistics, as well as analytic reports derived from the monitoring information. All of this could be reported as part of the 5GEx I2-Mon interface, providing operational information to other administrative domains requesting crosshaul services.
2) Topology and Inventory. The topology information is critical in a multi-domain environment for an efficient and effective placement of functions and connectivity reservation. 5G-Crosshaul supports both network and compute topology and inventory reporting, thus enabling the dissemination of such information outside the crosshaul domain borders. The population of this topology and inventory information can run on top of 5GEx I2-RT interface, for feeding the Resource and Topology functional blocks of the MdOs of the other provider domains in the Exchange.
3) Provisioning and Control of resources. The XCI in 5G-Crosshaul facilitates the control of the networking resources for adapt the underlying forwarding elements to

the needs of the flows to be transported in the crosshaul area. This capability can be easily integrated in 5GEx by mapping it to the I2-RC interface.

4) VNF management and orchestration. 5G-Crosshaul permits to accomplish the full management of the VNF lifecycle via the XCI. The APIs offered by 5G-Crosshaul for this can be homologated to the I2-F interface in 5GEx.

With the integration in a multi-domain environment, the 5G-Crosshaul XCI and the applications on top of it become the 5GEx multi-domain orchestrator. Thanks to the recursiveness properties of XCI, another option could be to dedicate a specific XCI instance to multi-domain aspects interacting as a client with a XCI instance below focused on the Crosshaul domain.

Interestingly, the VIMaP functional block in 5G-Crosshaul provides additional capabilities for planning as an extension to the usual VIM functionality. These planning capabilities can be quite useful on assisting the decisions for placement and connectivity in certain services, as the VNFaaS proposition in 5GEx. In this sense, the I2-F interface from 5GEx could be augmented to support the interaction with the VIMaP module in 5G-Crosshaul in this direction.

3.2.3.3.2 Proposition of additional functional blocks in 5G-Crosshaul for full compliance with 5GEx architecture

There are instead some other functions not fully available in 5G-Crosshaul. The more notorious capabilities are the ones related to business support. Here there is a brief summarization of the findings:

1) Business support. Specially, the population of the services supported in 5G-Crosshaul in terms of catalogue of services is not defined. This capability is necessary for advertising the capabilities of each crosshaul environment in an area in terms of networking and computing resources, as well as some added-value services that could complement the offer.
   In order to complement the 5G-Crosshaul architecture, the proposal here is to define a new functional module on top of XCI in charge of disseminating to other domains the Crosshaul capabilities supported in such domain. This new block, the 5G-Crosshaul Service Catalogue would be placed at the same level as the other applications defined in 5G-Crosshaul (e.g., Resource Management, Energy Management, etc.). In addition, this block is required to support 5GEx I2-C interface for integration on 5GEx ecosystem.
2) Service specification and request. In a multi-domain environment such as 5GEx it is necessary to have a common understanding on the services offered by each of the participants in the Exchange. To do that, the same semantics and abstractions have to be handled by the different administrative domains in order to ensure consistency. Such abstractions at technical level imply the utilization of common information and data models for the resources to be configured and used. In the case of integrating 5G-Crosshaul in a 5GEx environment, the former has to support the request of services through 5GEx I2-S interface.

3.2.3.3.3 Interface adaptation

The previous sub-sections have explored the integration feasibility of 5G-Crosshaul and 5G-Exchange. In general terms, the integration can be achieved through minor adaptations. Some new modules are required for enabling the trading of crosshaul capabilities in a normalized way.

It has been also analyzed the mapping to 5GEx interfaces. While feasible, some adaptations could be also required to this respect. As mentioned before 5G-Crosshaul is based on APIs to retrieve the information in an asynchronous manner. However, 5GEx interfaces are not defined in such format. For example, the I2-RT interface leverages on BGP-LS for the dissemination of the topology information across domains. This means that in some cases, even with an easy conceptual integration between 5G-Crosshaul and 5GEx, some minor functional block could be required for interface adaptation. In the case of I2-RT, for instance, a block in 5G-Crosshaul implementing a BGP-LS speaker facing the multi-domain environment for synchronous exchange of information would be required, at the same time retrieving the internal crosshaul information asynchronously

3.2.4 Summary of the contribution

This sub-section has explored different aspects of the relevant topic of multi-domain service and resource orchestration. First, the basis for multi-domain service orchestration have been established, by describing an architecture enabling it, such as the one of 5GEx, but also identifying gaps necessary for full service activation, which are not part of the multi-domain orchestration itself (included as part of a journal paper [A14]). Second, the specific case of multi-provider multi-access edge computing integration has been analyzed, identifying different architectural alternatives and their implications (published as journal paper in [A15]). Finally, an exemplification of multi-provider integration for crosshaul segments (i.e., integrated fronthaul and backhaul networks) have been described leveraging on 5G-Crosshaul and 5GEx architectures (published in two conferences papers [A16][A17]).

3.3 Determination of appropriate service edge

The cloud computing paradigm has provided a new model for service delivery where Data Centers (DCs) hosting a pool of Information Technology (IT) resources, are able to attend multiple service demands by means of a dynamic assignment of capabilities, such as CPU or storage capacity, either as physical or virtual resources (in the latter, by means of abstraction mechanisms). The virtualization technology in the cloud allows the flexible management of those IT resources, distributing them per service as needed (following an Infrastructure-as-a-Service, IaaS, approach) either among distinct servers into a single data center, or spreading them across several inter-connected data centers, even across multiple administrative domains.

The computing resources are then allocated on-demand depending on the customer (or tenant) requests. This elasticity on resource consumption allows and encourages efficient resource utilization and an agile adaptation to the business and service needs in every moment.

Originally, those DCs were conceived as large, centralized facilities concentrating a significant number of computing resources. However, new service needs (e.g., requiring low latency or benefiting from the proximity to the end-user) are influencing this design by

reconsidering the need to deploy more and more computing capabilities towards the network edge. The emerging approach is to deploy multi-purpose hardware resources at the edge of a telco operator's network to dynamically deploy the application logic close to the end-user device. Such a trend, while initially can be seen as natural, imposes larger investments as well as the introduction of adaptation mechanisms in the control operations of the network in order to offer flexibility for agile connectivity of workloads being variable in time, origin, etc. Therefore, it is essential to understand what the optimal edge is for each of the services to be supported by the network, thus avoiding any overinvestment and over-dimensioning of the network for both computing execution environments and transport capacity connecting them. This is, however, not easy because of the intrinsic uncertainty of the services that will be deployed on those systems, and where (location) and when (time) they will be deployed, especially in the advent of 5G.

Moreover, even assuming a certain degree of distribution of the IT resources across the network, it is not clear what can be the criteria and procedures for identifying the most convenient location for each service at each time, pursuing efficiency in the sense of not starving precious resources for services that could be accommodated in other less important or critical infrastructures. All of this, for sure, depends on the kind of service to be deployed, since whatever can be essential for a service does not necessarily correspond with a key constraint for another one.

3.3.1 Physical edge vs. service edge

The evolutionary roadmap of existing telco networks offer multiple levels of processing and storage facilities (local, edge, remote, and federated cloud). The criteria to decide where to deploy the service (i.e., the NFs defining the service) must be defined by considering a combination of several factors, among which it can be mentioned the service performance parameters, the minimization of energy consumption, the network and cloud load balancing, etc. Such decisions should be transparent to the user.

In [O41] and [O42], 3GPP has defined three types of service categories, namely enhanced Mobile Broadband (eMBB), ultra-Reliable Low Latency (uRLLC), and massive Machine Type Communications (mMTC), as well as the corresponding traffic requirements for two of them, the eMBB and uRLLC slices. Traffic requirements for mMTC are defined by NGMN in [O43]. Finally, 5G PPP has provided several 5G use cases in [O44], identifying a set of basic characteristics for them.

Table 3-3 presents some examples of the characterization of different forthcoming 5G service scenarios, in terms of latency or traffic needs, for example. It is then clear that different constraints can be relaxed or, on the contrary, considered as mandatory at the time of identifying from which point in the network is most appropriate to carry out the service delivery.

For instance, for the motion control in the discrete automation case, the end-to-end delay is limited to 1 ms, requiring at the same time a high traffic density of up to 1 Tbps/km^2. This suggests the need for delivering the service very close to the physical edge of the network. On the contrary, in the remote control for process automation, both the latency and the traffic density can be relaxed up to 50 ms and 10 Gbps/km^2 respectively, which in principle can be accommodated in DCs more deeply in the network, despite it could also be hosted close to the physical edge of the network, as before. This reflects the fact that the physical edge of

the network does not necessarily correspond with the suitable edges for each of the services to be deployed.

Table 3-3. Service characterization per type [O41]

Scenario	End-to-end latency	Jitter	Traffic density
Discrete automation – motion control	1 ms	1 µs	1 Tbps/km^2
Discrete automation	10 ms	100 µs	1 Tbps/km^2
Process automation – remote control	50 ms	20 ms	100 Gbps/km^2
Process automation – monitoring	50 ms	20 ms	10 Gbps/km^2
Electricity distribution – medium voltage	25 ms	25 ms	10 Gbps/km^2
Electricity distribution – high voltage	5 ms	1 ms	100 Gbps/km^2
Intelligent transport systems/ infrastructure backhaul	10 ms	20 ms	10 Gbps/km^2
Tactile interaction	0,5 ms	TBC	[Low]
Remote control	[5 ms]	TBC	[Low]

Here we assume that the 5G services being deployed basically consists of five technical dimensions:

- Bandwidth (*B*), characterized by indicators like data rate, accumulated data volume, etc.;
- Delay (*T*), articulated around parameters like latency, jitter, etc.;
- Computation (*C*), determined by aspects like the processing imposed by the number of sessions to be maintained, processing needs for the service, etc.;
- Storage (*S*), influenced by memory size, the volume of data to be stored, etc.;
- Durability (*V*), defined by the ephemeral duration or permanent behavior of the service to be deployed.

All these dimensions can be taken into account at the time of deciding where and when to deploy a service, or what resources and functions allocate for creating a supportive slice for that given service. Other parameters can also influence, like geographic or regulatory limitations that can condition the number of selectable edges, but these are not addressed here for the sake of brevity. Then, for each service to be deployed, it is required to identify how it maps against the referred technical dimensions, in such a way that compliant DCs can be discriminated and selected as suitable candidates for deployment.

All these dimensions, in short referred as *BTCSV*, can be taken into account at the time of deciding where and when to deploy a service, or what resources and functions allocate for creating a supportive slice for that given service. Other parameters can also influence, like geographic or regulatory limitations that can condition the number of selectable edges, but this is left for further work.

Then, for each service to be deployed it is required to identify how it maps against the referred technical dimensions, in such a way that compliant DCs can be discriminated and selected as suitable candidates for deployment.

Let us define *I* as the set of cloud infrastructures available at the time of deploying a service. These infrastructures are composed of many IT resources *R* for computation and storage in a certain location *L*, and are accessible through the network via a number of links of a given

capacity *A*. Then *I* can be defined as *I ={R, A, L}*. The resources *R* can be checked against the service need in terms of computation *C* and storage *S*; the capacity *A* can be checked against the needed bandwidth, *B*; and finally, the location *L* can be checked against the delay *T*, e.g. by means of monitoring data obtained through active probing between the access point of interest and the targeted cloud sites. Furthermore, the temporal availability of the resources could be restricted for instance due to a pre-scheduled future use because of e.g. a calendaring schema, *Cal*.

When a service is to be deployed, it can be modelled in terms of parameters *B*, *T*, *C* and *S*, and also characterized by the expectation on the durability of the service, *V*. Then, in order to discriminate what of those *I* infrastructures are able to properly host the new service requested, *I**, smart algorithms can be defined. Once identified what are the environments suitable for the deployment of a given service, then this information can be consumed by the orchestration system ensuring that the service KPIs can be satisfied.

Table 3-4 proposes a simplistic algorithm for illustration purposes only. However, discriminating what facilities can properly host the newly requested service and choosing a proper matching of services and infrastructures is known to be a computationally hard task known as the Assignment Problem [O45]. In [A18] it is proposed a solution based on classic network flow-based combinatorial optimization that leverages the particularities of the problem to control its computational complexity. Furthermore, in [A19] a new efficient algorithm, named Vectorial Successive Shortest Path (VSSP) is proposed considering a multi-dimensional assignment problem.

Table 3-4. Simplistic edge discrimination algorithm

INPUT *I*(*R,A,L*), *service (B, T, C, S, V)*
OUTPUT I^*

```
I* ← ∅
for each i ∈ I do
    check R ← C, S
    if FEASIBLE then
        check A ← B
        if FEASIBLE then
            check L ← T
     if V < Cal then
     I* ← i
rank I*
if I*=∅ then return UNFEASIBLE
return I*
```

The next concentrates on the architectural aspects of the service edge selection.

3.3.2 Using ALTO as network entity for exposing integrated network and compute capabilities

This section describes an architecture based on the ALTO concept [O46], where the identification of the most appropriate edge execution environments, as proposed above, could be offered as an ALTO service. The idea is to relay on ALTO for retrieving the recommended edges for a given 5G service. ALTO is identified as valuable mechanism for this due to fact that allows to expose capabilities of the underlying transport network by

providing both a network- and a cost-map between points (e.g., IP addresses) in a given topology (similar entities like the Network Exposure Function –NEF– in the 5G Core are overlay functions without information of the underlying transport network capabilities).

A given network function or application typically shows certain requirements in terms of processing capabilities (i.e., CPU), as well as volatile memory (i.e., RAM) and storage capacity.

Cloud computing providers, such as Amazon Web Services or Microsoft Azure, typically structure their offerings of computing capabilities by bundling CPU, RAM and storage units as quotas, instances or flavors that can be consumed in an ephemeral or temporal fashion, during the actual lifetime of the required function or application.

This same approach is being proposed for characterizing bundles of resources on the so-called Network Function Virtualization Infrastructure (NFVI) Points of Presence (PoPs) being deployed by the telco operators (such as [O10]). Specifically, the Common Network Function Virtualisation Infrastructure Telecom Taskforce (CNTT), jointly hosted by GSMA and the Linux Foundation, is intending to harmonize the definition of flavors for abstracting capabilities of the underlying NFVI facilitating a more efficient utilization of the infrastructure and simplifying the integration and certification of functions (see [O47] or [O48]).

Focusing on the CNTT ongoing work, the flavors or instances are described according to a number of characteristics:

- Type of instance (*T*): the types of instances are characterized as B (Basic), or N (Network Intensive).
- Interface Option (*I*): it refers to the interface bandwidth.
- Compute flavor (*F*): it refers to a certain combination of virtual CPU, RAM, disk, and bandwidth for the management interface.
- Optional storage extension (*S*): to request additional storage capacity.
- Optional hardware acceleration characteristics (*A*): to request specific acceleration capabilities for improving the performance of the infrastructure.

ALTO can assist in the selection of convenient flavors or instances of the computing substrate by taking into consideration both, availability of compute resources and network cost metrics. The function or application to be deployed on top of a given computing flavor is interconnected outside the computing environment where it is deployed, also requiring to guarantee some transport network requirements, such as bandwidth, latency, etc.

The objective then is to leverage on ALTO provide information about the more convenient execution environments to deploy virtualized network functions or applications, allowing the operator to get a coordinated service edge and transport network recommendation.

3.3.2.1 Compute information in ALTO

CNTT proposes the existence of infrastructure profiles through a catalogue collecting the instances available to be consumed. Such kind of catalogue could be communicated to ALTO or even incorporated to it.

ALTO server queries are required to support the encoding of the *{T, I, F, S, A}* characteristics in order to retrieve proper maps from ALTO. Additionally, filtered queries for particular

characteristics of a flavor could also be supported. Additionally, it is required to associate the location of the available instances together with network topological information to allow ALTO construct the overall map.

At this stage three potential solutions could be considered:

- To leverage on (and possibly extend) topology information [A4] for disseminating topology information together with notion of function location (that would require to be adapted to the existence of available compute capabilities). A recent effort in this direction can be found in [A5].
- To extend BGP-LS [O49], already considered as mechanism for feeding topology information in ALTO, to advertise computing capabilities as well.
- To combine information from the infrastructure profiles catalogue with topological information by leveraging on the IP prefixes allocated to the gateway providing connectivity to the NFVI PoP.

The viability of these options is identified as future work (as indicated before, some of them are being contributed to IETF).

3.3.2.2 ALTO architecture for determining serve edge

The following logical architecture defines the usage of ALTO for determining service edges.

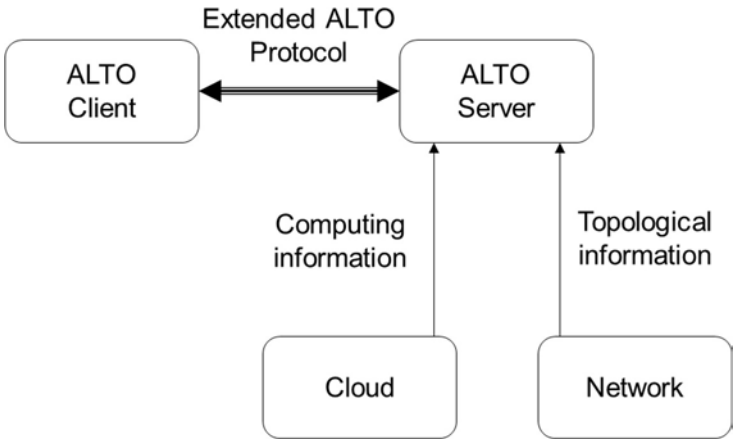

Figure 3-16. Service Edge Information Exchange with ALTO

ALTO, as the entity defined in IETF in charge of providing network and cost maps to applications, is the most convenient element to assist on this selection.

3.3.2.3 Definition of flavors in ALTO property map

The ALTO unified property extension [O50] generalizes the concept of endpoint properties to domains of other entities through property maps. In the context of the CNTT domain, an ALTO property map could be used to expose *{T, I, F, S, A}* information of potential candidate flavors, i.e., potential NFVI PoPs where an application or service can be deployed.

Table 3-5 below shows an illustrative example of an ALTO property map with property values grouped by flavor name.

Table 3-5. Example of ALTO property map for compute information

Flavor name	Type of Instance (T)	Interface Option (I)	Compute Flavor (F) {CPU, RAM, disk and bandwidth}	Storage Extension (S) [Optional]	Hardware Acceleration (A) [Optional]
Small-1	Basic	{1, 2, 3, 4, 5, 6, 7, 8, 9 Gbps}	{1, 512 MB, 1 GB, 1 Gbps}	…	…
Small-2	Network Intensive	{1, 2, 3, 4, 5, 6, 7, 8, 9 Gbps}	{1, 512 MB, 1 GB, 1 Gbps}	…	…
…	…	…	…	…	…
Medium-1	Network Intensive	{25, 50, 75, 100, 125, 150 Gbps}	{2, 4 GB, 40 GB, 1 Gbps}	…	…
…	…	…	…	…	…
Large-1	Compute Intensive	{50, 100, 150, 200, 250, 300 Gbps}	{4, 8 GB, 80 GB, 1 Gbps}	…	…
Large-2	Compute Intensive	{100, 200, 300, 400, 500 Gbps}	{8, 16 GB, 160 GB, 1 Gbps}	…	…
…	…	…	…	…	…

3.3.3 Service edge view based on ALTO

The proposed architecture is illustrated in Figure 3-17. It assumes that an Orchestrator, which is in charge of deploying 5G services that require to instantiate capabilities at the edge, interacts with an ALTO server in order to retrieve an indication of the computing environments that could satisfy the final service requirements. For doing so, the ALTO server interacts with a number of Edge Managers responsible for managing the compute and storage infrastructure of each of those execution environments spread across the network.

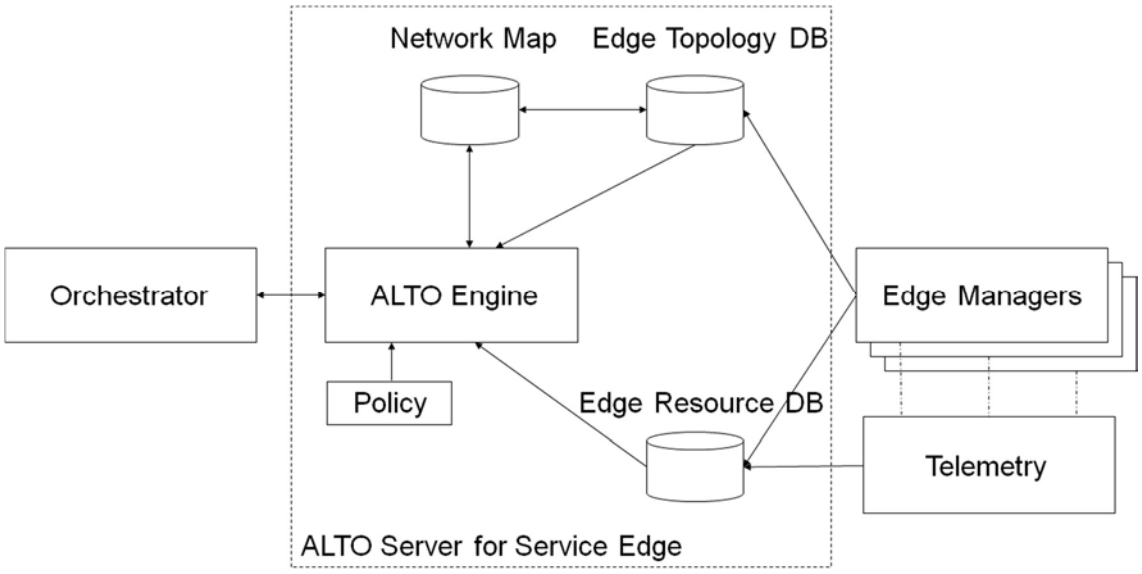

Figure 3-17. Integration with the orchestration system

These Edge Managers are responsible mainly of providing information about the resource availability in each edge node under its control and, at the same time, reporting sufficient information for helping to identify the topological location of the edge node itself.

The resource-based information can be used to feed an Edge Resource database, assisting in the identification of resources available in each node. In order to allow real-time status collection of the resources in each edge node, the system can rely on telemetry information [O51] provided by a monitoring system attached to the underlying infrastructure, interacting with the ALTO server either directly or through the particular Edge Managers. On the other hand, the topology information could be merged or integrated with the ALTO network maps, in order to create an overall topological view of the network and computing capabilities internal to the service provider.

This topological view can be relevant to decide the placement of service components, for example, in the form of Virtual Network Functions (VNFs) or applications that could show some restrictions (e.g., latency). The ALTO server can provide a ranking of convenient edges for the specific service as requested by the Orchestrator. Such ranking is originated based on intelligent discrimination algorithms (e.g., [A18] or [A19]), potentially producing optimal decisions when algorithmically feasible. Notably, the application of some policies can be also foreseen, e.g., through specific modules, in order to assist the edge discrimination, for instance, policies for data sovereignty. Additionally, the profit computed during the assignation process can be used as a priority by the Orchestrator.

Thus, this augmented ALTO server could be integrated into orchestration frameworks for edge computing. For instance, considering the ETSI Multi-access Edge Computing architecture [O35], the Orchestrator in Figure 3-17 could be the Multi-access Edge Orchestrator at MEC system level, while the Edge Managers represented in the figure could be the per-host Virtual Infrastructure Managers in MEC.

It is important to note that ALTO can basically assist in the identification of the best execution environments for certain services, but it will not participate in the resource allocation, which will be the responsibility of the Edge Managers once the Orchestrator, based on ALTO, selects some edge node.

Then, after resource allocation is performed for a given service, whenever any parameter is unmet, the virtual system is adapted as soon as possible to ensure the continuity of the service.

3.3.4 Summary of the contribution

This sub-section has provided contributions in relation to architectures capable of assist on the determination of the proper edge for a given service for a more efficient usage of the scarce resources at the edge (generating a conference paper [A18] and an extension as journal paper [A19]). Such architecture leverages on ALTO as the entity that can expose combined network and compute information by integrating a common network and cost map (being proposed as ALTO extension in IETF [A20]).

3.4 Service blocking in a multi-domain service provision

In a multi-domain providers' federation scenario the local domain interacts with several different overflow domains to provide an efficient orchestration of network slices and to

overcome the possible mismatch between 5G verticals' service expectations and the entry provider owned resources.

In the approach here proposed, each overflow domain participating in a federation is expected to advertise information about resources and latency towards the other participants in the federation. Thanks to populating such information, the local domain can tale an informed decision on which VNFs can be allocated to any of the overflow domains, if needed. However, in case the local domain does not have a priori information about guaranteed latency and resource availability, it can try to allocate a VNF to an overflow domain whose data centers do not fully comply with the VNF latency constrains or do not have enough available resources to deploy the VNF.

In cases where there is no exchange of information among administrative domains, the local domain can only rely on its own data centers and assume that a best effort approach is taken by the overflow domains at the time of orchestrating a multi-domain network slice. That is, the VNFs that cannot be directly allocated in the local domain are passed over some overflow domain with the expectation of being deployed in such other domain. Here, the benefits of interchanging information among domains are analyzed in terms of impacts on service blocking. Eq. (3.1) describes the accounting of available resources at the time *t* when no information exchange is in place, which basically corresponds to the available resources in the local domain.

$$R_t = \sum_{n=1}^{N} \left(C_{DC_n} - U_{t-1} \right), \;\; U_{t-1} < C_{DC_n} \tag{3.1}$$

where

- $R_t \in \mathbb{R}^+$ represents the available resources when allocating VNFs at the time $t \in \mathbb{R}^+$ without exchange communication support
- $C_{DC_n} \in \mathbb{R}^+$ corresponds to the maximum capacity of the local domain for $n \in \mathbb{N}$ of $N \in \mathbb{N}$ data centers that complies with the latency constraints of the service
- $U_{t-1} \in \mathbb{R}^+$ are the utilized resources at the time $(t-1) \in \mathbb{N}$

In this situation, the overflow domains may deploy or not the allocated VNFs depending on their available resources and the latency that they can provide, but that fact is unknown at the time of the request sent from the local domain because of the lack of information. This case is represented by $U_{t-1} \geq C_{DC_n}$, i.e., there are no more resources available in the local domain.

When introducing the exchange of information in terms of guaranteed latency and resource availability that the overflow domains can guarantee, the local domain can perform informed decisions when passing over VNFs to the other domains. The total available resources increase in this case, as reflected in Eq. (3.2), because of the knowledge of the available resources of the overflow domains, which can be considered by the local domain on its decision process.

$$R_t = \begin{cases} \sum_{k=1}^{K} \sum_{n=1}^{N_k} C_{DC_n}^k - R_{t-1}^k, & U_{t-1}^k < C_{DC_n}^k \\ 0, & U_{t-1}^k \geq C_{DC_n}^k \end{cases} \tag{3.2}$$

where K overflow domains containing N_k datacenters that complies with the latency constraints of the service are taken into consideration.

When there are no more resources available in the local or overflow domains, i.e., $U_{t-1}^{k} \geq C_{DC_n}^{k}$, then the system cannot serve the request then blocking service provision.

As stated before in Section 3.3.2.1, it can be considered that the advertisement of the compute capabilities in the different domains is performed by a solution like in [A5].

In order to set a simulation scenario for the case analysis, a system structure is first proposed. This system structure represents the interaction between the multiple stakeholders that interact with each other to provide services to tenants (i.e., vertical customers). It is composed by three modules that oversee different functions, as indicated in Figure 3-18.

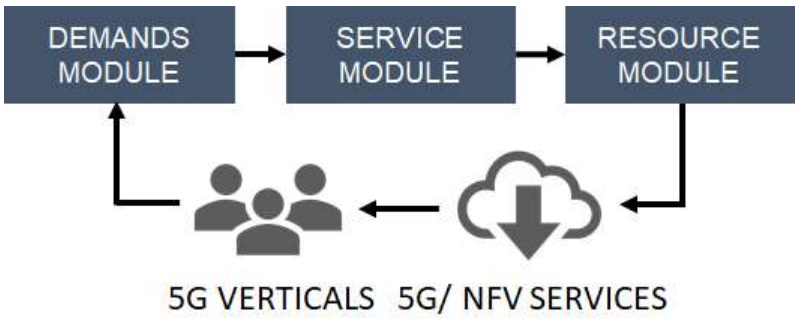

Figure 3-18. System modules

The *Demands Module* generates the service demands (i.e., the vertical network slices) and takes care of the service characterization. The *Service Module* performs the partitioning of the network-graph for a given service by assigning VNFs to an overflow domain complying with the latency constrains, when no resources are left in the local domain. In addition, this module calculates the resources that are required from each domain to make the network slice deployment possible. Finally, the *Resource Module* takes care of allocating the required resources to the data centers of each domain.

3.4.1 Scenario of analysis

One possible scenario has been defined for illustration purposes. The case analysis scenario gathers three types of inputs: services or network slices (as requested by 5G verticals), data centers (for resources enabling the deployment of VNFs) and administrative domains (i.e., multiple providers).

Services are defined by describing the VNF resources in terms of computing capacity (CPU), memory (RAM), and storage (HDD), as well as bandwidth and latency needs. In order to characterize the system behaviour, three different kind of services are considered, as indicated in Table 3-6.

Data centers are defined according to their resource capacity in terms of CPUs, RAM, storage (HDD), bandwidth capacity and their guaranteed latency. Three data centers sizes are considered: small, medium and large. The three of them are defined as a function of the Service 1 parameters defined in Table 3-6.

Table 3-6. Network slice service parametrization

	CPUs	RAM [GB]	Disk [GB]	BW [Gbps]	Latency [ms]
Service 1	C	R	D	B	L
Service 2	5C	5R	5D	5B	5L
Service 3	10C	10R	10D	10B	10L

Concerning the number of resources allocated to each of the data center types, Table 3-7 describes the characterization of each of them.

Table 3-7. Data center characterization

	CPUs	RAM [GB]	Disk [GB]	BW [Gbps]	Latency [ms]
Small	25C	25R	25D	25B	L
Medium	50C	50R	50D	50B	2L
Large	300C	300R	300D	300B	5L

Regarding the number of domains in the federation, three of them are considered: the local domain plus two overflow domains. When deploying a service in an overflow domain, a latency increment of L ms is added to account the physical latency between domains, which has some implications in the consideration of data centres for the allocation of the VNFs, attending to the characterization in Table 3-7.

- *Small datacentre*. When federating a service in a small datacentre, the minimum latency that can be ensured is $L + L = 2L$. Therefore, Service 2 ($2L \leq 5L$) and Service 3 ($2L \leq 10L$) can be allocated to these data centers.
- *Medium datacentre*: When federating a service in a medium datacentre, the minimum latency that can be ensured is $2L + L = 3L$. Therefore, Service 2 ($3L \leq 5L$) and Service 3 ($3L \leq 10L$) can be allocated to one of these data centers, as well.
- *Large datacentre*. When federating a service in a large datacentre, the minimum latency that can be ensured is $5L + L = 6L$. Therefore, only Service 3 ($6L \leq 10L$) can be allocated to a large datacentre.

Accordingly, Service 1 can only be deployed in the local domain and this domain must have at least one small datacentre. Service 2 can be deployed in the local domain or in any overflow domain containing a small or medium datacentre. Lastly, Service 3 can be deployed to any datacentre in any domain.

Finally, Table 3-8 presents the distribution of data centers considered in the analysis across the three proposed domains. The reduced number of data centers in Table 3-8 has been selected to force a service-blocking situation for both the local domain and the overflow domains, in order to compare the effectiveness of the information exchange. This service-blocking occurs when a data center is fully occupied and cannot deploy any further VNF, then forcing the service (i.e., the network slice) to be blocked if no further resources are available in the selected overflow domain.

Table 3-8. Number of data centers per domain

	Small	Medium	Large
Local domain	3	0	0
Overflow domain 1	0	1	0
Overflow domain 2	0	0	1

For all the domains, the allocation of VNFs to data centers depends on three rules:

- *First rule*: On the assumption that the local domain complies with the latency and resource availability constrains, the VNFs are allocated to the local domain. The overflow domains remain unused.
- *Second rule*: When the first rule is not fulfilled, VNFs will be assigned to one of the overflow domains. Here we observe two possibilities depending on the use or not of the exchange of information proposed in this analysis.
    - *Without information exchange*: In this case the local domain is not aware of the state of the overflow domains and does not know in advance whether the VNFs that are intended to be federated could be deployed. One overflow domain will be randomly selected.
    - *With information exchange*: In this case the local domain knows beforehand whether the VNFs that are intended to be federated could be deployed and exactly where can be deployed in the case it is possible.
- *Third rule*: In any case, either in local or overflow domain, the resources of a given data center are totally consumed before attempting to deploy VNFs in another datacentre of the same kind.

After applying the three rules, if it is not possible to deploy all the VNFs of a network slice, the service is reported as blocked. As described only technical constraints have been taken into consideration on the deployment decision (resource availability, performance, etc). Other constraints such as the deployment cost could be included, but are not in the scope of this analysis.

3.4.2 5G Vertical service demand and lifetime definition

The definition of the 5G vertical service demand and duration is done in accordance with the values in Table 3-9. There, the arrival rate and the lifetime for each of the services defined in Table 3-9 can be found.

Table 3-9. Service arrival rate and duration

	Arrival rate	Minimum lifetime [H]	Maximum lifetime [H]	Mean lifetime [H]
Service 1	10	96	240	168
Service 2	25	18	30	24
Service 3	50	0.5	4.5	2.5

Network slice request arrivals are characterized by a Poisson distribution. This probability distribution gives the probability of occurrence of events $P(t;\lambda)$ in a fixed interval of time T (one year in this analysis) provided that these events occur with a known constant rate λ and with independency of the time t since the last event occurred, as indicated by Eq. (3.3):

$$P(t;\lambda) = \frac{e^{-\lambda} \cdot \lambda^t}{t!} \tag{3.3}$$

*w*here

- $\lambda \in \mathbb{R}^+$ is the frequency of service requests based on the type of service as proposed in Table 3-9.
- $t \in \mathbb{R}^+$is the time of arrival.

Service lifetime is characterized by a truncated Gaussian distribution [O52] to ensure that the lifetime values are distributed between the minimum and maximum lifetime values indicated in Table 3-9 for each of the services. It is characterized in Eq. (3.4) [O53] with mean μ and standard deviation σ^2 that lies within the interval (a, b), with $-\infty < a \leq b < \infty$.

$$P(x;\mu,\sigma,a,b) = \begin{cases} \dfrac{\phi\left(\frac{x-\mu}{\sigma}\right)}{\sigma\left(\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)\right)} & a \leq x \leq b \\ 0 \quad Otherwise \end{cases} \tag{3.4}$$

where

- $\mu \in \mathbb{R}^+$ is the mean which corresponds to the mean slice lifetime from Table 3-9.
- $\sigma^2 \in \mathbb{R}^+$ is the standard deviation which is set to 1 in order to ensure a small deviation from the mean.
- $\phi(\xi) = \frac{1}{\sqrt{2\pi}} e^{\left(-\frac{1}{2}\xi^2\right)}$ is the probability density function of the Gaussian distribution.
- $\Phi(\xi) = \frac{1}{2}\left(1 + \text{erf}\left(\frac{\xi}{\sqrt{2}}\right)\right)$ is the probability density function of the cumulative Gaussian distribution.

3.4.3 Analysis

The scenario setting described in the previous sections has served as a baseline for running a simulation of the federation behavior. This scenario tries to be as generic as possible, including for such a purpose three services that represent different load (in terms of needed resources) as well as lifetime and arrival rate. The simulation has been developed in MATLAB, by considering 10,900 events (i.e., service or network slice requests) generated with the assumptions of a one-year timeframe. Of those events, 1,605 belong to Service 1, 3,595 to Service 2 and 5,700 to Service 3 (the number of events was chosen for generating service blocking situations).

VNFs are considered as atomic units. When a VNF of a service cannot be deployed, the entire service is not deployed, and it is accounted as a blocked service. This situation can occur due to two possible situations:

1) Latency constrains: Depending on the latency that each of the datacenter types can guarantee, and the incremental latency of *L* ms due to the fact of deploying the service in an overflow domain, the final latency could be greater than the minimum latency required for a given service. In this case, the service cannot be deployed.
2) Occupation constrains: When the previous constrain is satisfied, the occupation constrain must be analyzed. In order to deploy a VNF in a particular data center, it has to have enough resources for the deployment. Otherwise, the service is also blocked.

Figure 3-19 presents the cumulative blocked services for each of the service types in the cases of exchanging and not exchanging both latency and resource availability information among domains in the federation, during one year in the conditions previously described. Table 3-10 summarizes the absolute number of blocked services.

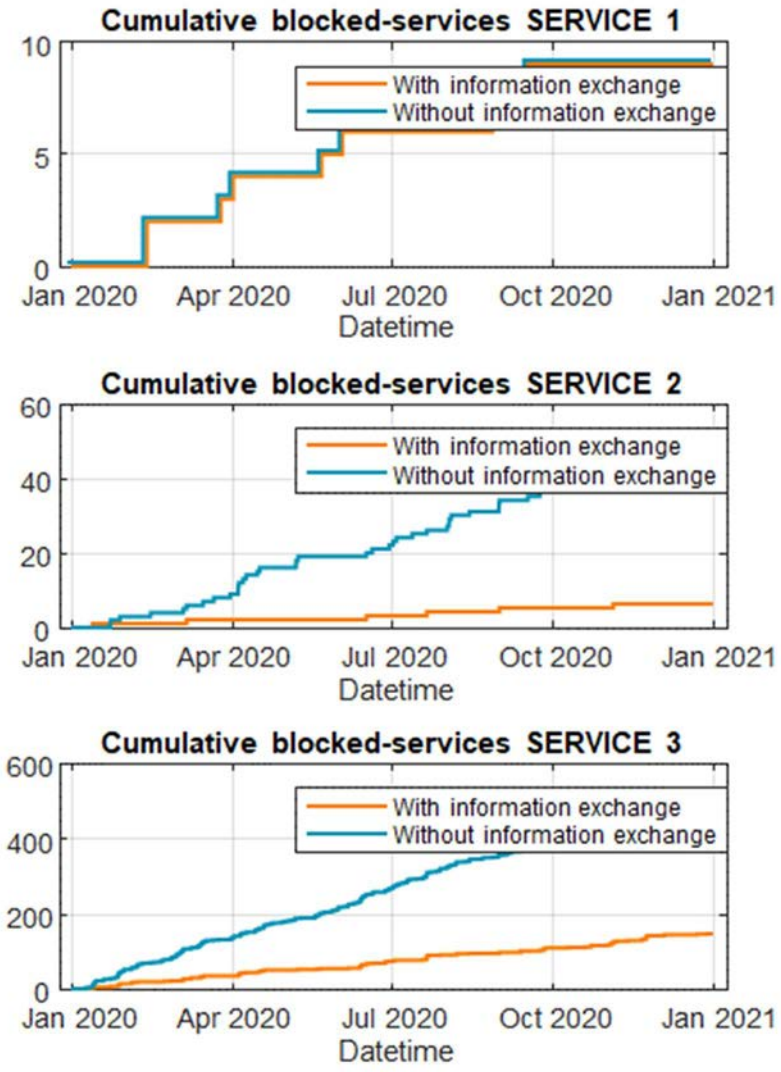

Figure 3-19. Cumulative blocked-services along simulated time per type of service

Because of the latency constrains defined for each kind of service, Service 1 essentially cannot be federated and, in consequence, the results are the same in both cases. So the impacts of enabling information exchange among domains can be only perceived on the results obtained for Service 2 and 3. In this respect, the results obtained in Table 3-10 lead to the conclusion that the exchange of information in terms of guaranteed latency and resource availability highly decrease the ratio of blocked services.

The reduction in the number of blocked services has direct impact on the incomes of the entry point domain as single contractor of the vertical customer, but also for the federation itself, since non-blocked services, making use of resources from overflow domains, also imply the generation of incomes for those other domains, incomes that otherwise are lost. The reduction in the number of blocked services depend on the characteristics of service, reaching in some cases an improvement of 86.66% in the best case.

Table 3-10. Accumulated blocked services

	With information exchange	**Without information exchange**	**Improvement in percentage**
Service 1	9	9	0%
Service 2	6	45	86.66%
Service 3	246	509	51.66%

3.4.4 Protocol support for the information exchange

The interconnection of different networks, each of them representing a different administrative domain (i.e., distinct Autonomous Systems managed by different providers), is commonly performed by leveraging on the Board Gateway Protocol (BGP) protocol [O54]. BGP has a huge capability to scale and has undoubtedly contributed to the success of the Internet by facilitating a standardized manner of interconnecting networks. However, the network slice concept intensively stresses the network capabilities because of the dynamicity introduced in the deployment and lifetime of services. This can be expected also to occur in the multi-domain environment.

Board Gateway Protocol Link-State (BGP-LS) [O49] is the mechanism by which Interior Gateway Protocol (IGP) link-state and traffic-engineering information is collected from the local domain and is shared with other domains using the BGP protocol. This is achieved by using a new Network Layer Reachability Information (NLRI) encoding format. NLRIs are used to advertise link, node, and prefix information in the form of parameters and attributes. It is defined as a set of Type/Length/Value (TLV) fields.

Then a possible way of disseminating information per domain is to define also as TLVs fields related to guaranteed latency supported and resource availability, advertising that between the various telecommunication providers participating in a federation. This same approach has been taken in [O55] to populate traffic engineering performance metrics such as link bandwidth or delay.

Table 3-11 presents the proposed values for the parameters to be exchanged between domains in a federation as the one described here. The proposed parameters can be encoded as TLV fields and sent within each Link-State NLRI updates. As a result, and thanks to the

BGP-LS protocol, all the telecommunication providers participating in the federation can receive up-to-date information about the latency and availability that the datacentres of other providers in the federation can provide. This information is crucial to decide where to deploy the VNFs outside the local domain, reducing blocked services due to not fulfilling the latency contains or not having enough vacancy in the data centers of an overflow provider, as illustrated before through simulations. Table 3-11 shows potential values associated to the parameters of interest with reference to actual equivalent codifications on different RFCs. With approaches like this, it could be possible to build topologies including datacentre capabilities as proposed in [A5]. Integration of these approaches is left for further study.

Table 3-11. Proposed values for BGP TLV fields

	Type	Length	Comment	Reference value
Guaranteed Latency	int	24 bits	Max measured link delay value (in ms) over a configurable interval	Max Unidirectional Link Delay [O55][O56]
CPUs availability	int	32 bits	Number of available virtual CPUs	vmCurCpuNumber [O57]
RAM availability	int	32 bits	Memory size	vmCurMem [O57]
HDD Disk availability	int	32 bits	Virtual storage size	vmStorageAllocatedSize [O57]

3.4.5 Summary of the contribution

This sub-section has served to highlight the relevance of disseminating resource and performance information among providers in a federation. The dissemination of that information allow informed decisions, here exemplified by the improvements on the reduction of service blocking, even for simplistic allocation schemes (published as conference paper in [A21]).

The dissemination of information among providers in a federation can assist architectural solutions as the ones proposed in Section 3.2. Due to the multi-domain nature of the problem, it is key to establish standardized mechanisms for such dissemination of information, in order to avoid integration costs and delays.

3.5 Efficiency gains due to Network Function sharing in CDN-as-a-Servicee slicing scenarios

Video traffic is becoming nowadays the killer application for service providers' networks, and it will be probably the dominant component of the overall traffic in the future. The raise of multiple offers from a variety of video platforms, either directly owned by Internet Service Providers (ISPs) or offered by Over-The-Top (OTT) content providers, such as Netflix, Amazon Prime, HBO, etc., is effectively changing the network demand landscape. Being this fact already true for streaming content, it will be even increased when considering in the near future other flavors of multimedia delivery, such as gaming or virtual reality.

According to analysis from the telco industry [O58], video traffic in 2020 represents the 66% of all mobile data traffic, with the perspective of increasing up to 77% in 2026. The same

occurs for fixed networks, where e.g. TVs generate the largest component of the traffic per device in Western Europe [O59], with residential users moving from offline to online activity, in parallel with the improvement in video formats, especially the Ultra-High-Definition (UHD) or 4K. Some estimations consider that the number of installed flat-panel TV supporting UHD will raise from 33% in 2018 to 66% by 2023 [O60], then favoring the demand of higher resolution on-line contents.

Such a huge amount of content is served leveraging on overlay Content Delivery Networks (CDNs), which allow servicing contents in a scalable manner. A number of distributed delivery points or caches store the content, delivering copies of it locally, alleviating the demand in terms of capacity needed in the transport network, since that content would be required to obtain remotely, otherwise. Those caches can be found at the border of the ISP networks or even internal to them. The latter is the current trend, where multiple caches from different content providers (i.e., from the own ISP but also from third parties) are deployed internally to the network, delivering the content in proximity producing the reduction of bandwidth at higher layers in the network topology, but also the perceived latency, leading to which is named as sub-millisecond Internet [O61].

The advent of network virtualization has brought the attention on the possibility of considering the CDNs and the caching of content as a relevant use case. For instance, it can permit the dynamic deployment and flexible instantiation of caches in the network on top of virtualized infrastructures. Thus, both Network Function Virtualization (NFV) and Multi-access Edge Computing (MEC) paradigms have look at the CDNs in their specifications [O33][O62]. Furthermore, traditional CDN providers have also moved into the virtualization arena by providing virtualized solutions of their CDN solutions like Akamai [O63] or Amazon [O64].

Thinking on the OTT content providers, the aforementioned trend of increase on the video consumption implies that subscribers from different ISPs in a given geographical area practically consume the same kind of content from a reduced number of content providers, if not actually the same, independently of the ISPs they are subscribed to.

On the other hand, it is being common the fact of sharing infrastructures among service providers [O65][O66][O67]. This is due to the need of reducing investments for improving margins. Such a sharing imply to host services of competitor ISPs on top of a single and common infrastructure, and/or hosting directly Virtual Network Operators (VNOs) not having any infrastructure at all, or very limited.

The sharing scenario, as mentioned, can be implemented by one ISP sharing its infrastructure to others, or by neutral operators opening their infrastructures to third party ISPs. Such kind of operators are commonly known as Infrastructure Providers (InPs).

From an operational perspective, the effective way of implementing the sharing of infrastructures is expected to be through the adoption of network slicing [O68][O69], allowing the recreation of a virtually dedicated network for each of the ISPs or VNOs in the area, while using a common physical infrastructure (from now on ISPs is used in a general way for simplicity). Then, a network slice per ISP can contain all the services offered for their respective subscribers, including the necessary caches for providing video contents. This help to reduce the traffic and associated costs in the transport network.

The network functions in the allocated slice, like the caches referred before, are deployed in the form of Virtual Network Functions (VNFs). In principle, that VNFs can be instantiated separately, dedicated per ISP. However, when looking at video cache function itself, thinking on the fact that most of the OTTs are usually common to all the ISPs supported by a certain InP, and also considering that the content offering itself is also the same, it can be questioned the necessity of having multiple instances of the same function instead of sharing a single one. Reasons for this are the potential benefits and efficiencies that can be achieved in terms of consumed compute and storage resources, as well as others like energy efficiency.

Fortunately, NFV specifications consider the possibility of sharing VNFs among virtualized services [O70], then open the door to optimize the deployment and usage of network functions, that in this particular case implies leveraging on the same virtualized cache instance for all the ISPs.

The virtualization paradigm permits to deploy CDN caches in a virtualized fashion, as VNFs. Considering the fact that popular contents are concentrated on the hands of few content providers, ISPs deploying network slices on top of the facilities of the same InP could potentially share the VNF instances implementing such caches. Thus, it is possible to conjugate both the efficiency due to the deployment of CDNs and the efficiencies due to the sharing of caches.

Figure 3-20 illustrates the scenario under evaluation. It is assumed that different ISPs make use of a common infrastructure provided by an InP. Note that one of those ISPs could play also the role of InP in a given geographical area.

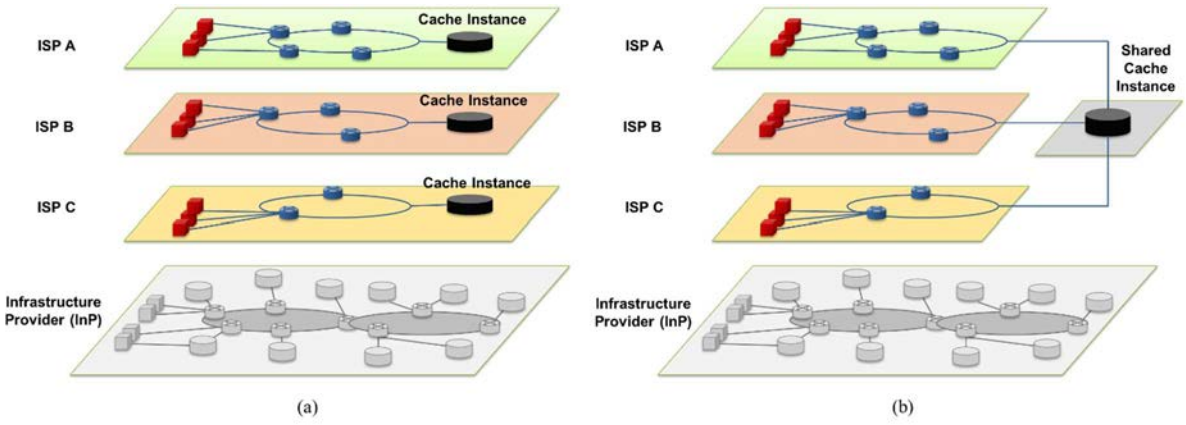

Figure 3-20. Scenarios of comparison: (a) dedicated cache instances per ISP vs (b) shared cache instance

It is also assumed that those ISPs in such area served a common set of contents, provided e.g. by an OTT that maintains separated commercial agreements with all those ISPs. Thus, the ISPs essentially have a similar content offer for all the base of end-users in the area, achieving service differentiation by other means (e.g., price competition, bundle of offers, etc.)

The OTT initially performs the content distribution by means of virtual caches (vCaches), in the form of VNFs, deployed on top of the InP infrastructure. Each vCache stores the contents demanded by the customer base of each of the ISPs in that particular area. When

considering a separate instance of virtual cache per ISP, this means that each vCache will be dedicated for a given ISP.

However, since the catalogue of contents offered by the OTT are the same for all the set of ISPs, in principle it is possible to deliver all the demanded contents from a common virtual cache instance just managing the distinct subscriptions differentiating the delivery of the content for each of the ISPs, e.g. by using different Ethernet VLANs in the connection to each network slice per ISP. That is, a certain content "Content-A", instead of being replicated through different virtual cache instances (one per ISP) is contained in a single virtual cache instance but being delivered to the corresponding ISPs demanding such content.

This is relevant since the efficiency of caching increases with the number of end users being served, since there will be higher coincidence in the kind of content demanded by the users. This is evident for the more popular contents, but also happens for the less popular ones. Thus, sharing the vCache has the effect of concentrating the end-user demands, which should reduce the overall number of contents being cached in an area.

3.5.1 Simulation framework

In order to understand to what extent it is possible to achieve efficiency when following the approach of cache sharing, simulations have been performed based on the number of contents consumed for a base of subscribers of distinct ISPs in different scenarios. This permits to understand on which conditions this can be desirable.

It is considered that the network is divided in areas as the ones represented in Figure 3-20. The objective of the simulation is to understand what the peak number of contents demanded in the network is, and how the peak impacts on the shared vs non-shared vCache scenario.

3.5.1.1 Content popularity model

The preference of content visualization by end users in IPTV, Video-on-Demand (VoD) and cache systems in general is commonly modeled by a power-law distribution known as Zipf function [O71][O72][O73]. The Zipf function states that the occurrence of a certain event (here, the tuning of a multicast channel for IPTV or the selection of a certain unicast content in a VoD system) is determined by:

$$k\frac{1}{x^{\alpha}} \tag{3.5}$$

where k is a constant value, x the rank or popularity of the event (i.e., a given content) in the distribution, and α the factor which characterizes the skewness of the distribution. Then, the frequency or probability that predicts the eligibility of an event is provided by:

$$\frac{\frac{1}{x^{\alpha}}}{\sum_{n=1}^{N}\left(\frac{1}{n^{\alpha}}\right)} \tag{3.6}$$

where N is the total number of ranked elements. As α increases, the popularity of the first ranked events increases, while the distribution tail concentrates less occurrences. Here we will consider 0,6 and 0,9 as reference values for α in line with the observations in [O73].

3.5.1.2 Number of contents

The amount of available on-demand contents to be consumed either live or in an on-demand fashion has continuously increased along the time. Even such increase applies to both types of content, the order of scale differs. In the case of live content, usual values nowadays could stay around few hundreds of contents. For the VoD case, the quantity considered can be even higher than ten thousand.

Several factors have contributed to this. On one hand, the proliferation of OTT video platforms has augmented significantly the number of contents in their respective catalogues, generating a very broad multimedia offer. Secondly, it usually occurs that the same multimedia content is coded differently (e.g., Smooth Streaming, DASH, HLS, etc.) adapting it to multiple receiver platforms and players, then creating differentiated copies of the same content which are consumed also differently depending on the acceptance of a given player.

There are also differences between live and VoD contents regarding their lifetime. Live content usually is stored in the cache during few hours, as much, since further than that time the content can be considered no longer to be live. This time in the cache allows users accessing late to the content, but yet interested in a recent event, to be served. Once that time is exceeded, the profile in the consumption of that content can be considered as passing to the category of on-demand.

The on-demand content can usually stay in the caches for longer, typically up to the time that the capacity of the cache is exhausted, then requiring to make storage space available for newer content being demanded.

For the analysis here it is generically consider the on-demand content as the subject of interest for the end-users. In addition to that, a common and unique coding of the content is assumed for simplification. Here, a typical content is considered to be coded for an average bit rate of 5 Mbps and an average duration of 80 minutes, requiring then ~ 2,8 GB of storage in the vCache.

3.5.1.3 ISPs and end-users

In open, competitive markets, the base of end-users is divided among different competing ISPs addressing such market. The distribution is usually unbalanced, with some ISPs capturing higher share of users than others. Reasons for differentiation can be multiple, such as overall service offering pricing, etc.

For the analysis we will consider the presence of four ISPs with different market shares, as follows:

- Service Provider A: 40%
- Service Provider B: 30%
- Service Provider C: 20%
- Service Provider D: 10%

This differentiation represents a market where a dominant ISP is taken the majority of the share, with two major competitor ISPs as followers plus one challenger entering the market. The specific market share can differ in real scenarios but the assumption here permits to compare at different granular levels between the ISPs as a function of their relative share in

a market. As reference, during Q1 2020 the share in Spain of the four main operators in the country was 38,3%, 25,2%, 20,6% and 10,3% for fixed broadband, and 29,6%, 24,4%, 22,4% and 13,9% for mobile access, respectively.

3.5.1.4 Number of vCaches in the network

The number of locations where to deploy vCaches is another axis of dimensioning. In the context of this analysis, it can be understood as the number of Points of Presence (PoPs), Central Offices (COs) or edge nodes (thinking on a more distributed deployment with high capillarity) where a vCache could be instantiated. It is assumed that on each of those locations there is sufficient compute infrastructure to host the vCache in terms of processing capacity, storage, etc.

The number of locations depends on the size of the country to be served, the distribution of the population (i.e., density) and the availability of physical infrastructures (sites, transmission, etc.). The selected locations do not necessarily need to be at the same hierarchical level in a layered network topology, that is, some of them could be considered at the edge of the network while others being more centralized. The criteria followed in this analysis considers simply the total number of users being served by each vCache, where such total number is divided proportionally among the ISPs as described before.

Here it is analyzed the behavior of distributing the users in a number of locations with vCache ranging from 100 to 1000. The former can represent the number of PoPs concentrating main distribution areas (at region/province level) in a mid-size country, while the latter can represent the number of central offices in such a country. This is compatible, for instance, with the number of aggregation (the former) and pre-aggregation (the latter) sites considered in other reference networks like in [O74]. The results however can be extrapolated to any other number of vCaches present in the network.

3.5.1.5 Modelling the assignment of end-users to desired content, ISPs and vCaches

For each of the simulated preference of end users, the simulation firstly identifies the content desired by the user, then associates such end-user to one of the ISPs, and finally assigns him/her to one of the vCaches. With that, once the modelling is finished, it is possible to quantify how many different contents are stored per vCache. That quantification essentially considers that any solicited content is stored in the vCache with the expectation of being served from the cache for a second or higher request. In this way, the model reflects the situation that would result from a peak demand instant in the aforementioned conditions. This procedure generates two views: *(i)* the view of the solicited unique contents that would be generated in case each of the ISPs maintain separated vCaches, and *(ii)* the view of the solicited unique contents if the ISPs share the vCache. From the comparison of both views, conclusions about the efficiency of the shared cache approach can be obtained.

For the selection of the content, a content is identified according to the Zipf distribution as defined in Section 3.5.1.1. Then, the end user is associated to an ISP according to the market share percentages specified in Section 3.5.1.3. Finally, that user is assigned to a vCache node. The simulation, as first approximation to the problem, assumes a uniform distribution of end users among the vCaches. This implies that, roughly, a similar number of end users is considered per vCache.

3.5.1.6 Description and parametrization of the simulation

The simulations have been performed with MATLAB, running on a server counting on two 2.20GHz vCPUs, 16GB RAM and 200GB storage capacity.

Table 3-12 summarizes the parameters of the simulation. Each simulated scenario run on average for 2 hours and a half, generating 5MB of data. Twenty scenarios were run, with minor deviations among runs, confirming the validity of the obtained results. The results here presented are based on the average values of one of that runs.

Caching contents produces a clearer advantage in the reduction of traffic in the networks, from the peering and transit points up to the locations of the vCaches. However, such efficiency comes at the cost of deploying different levels of computing infrastructure in the network.

Table 3-12. Parameters of the simulation

Parameter	Values	Description
Skew factor (α)	0,6 and 0,9	Power-law factor of the Zipf distribution of content selection
Number of contents	[500, 5000]	Number of selectable contents in the scenario under evaluation
Number of end users	[10000, 1000000]	Population of users simultaneously demanding content
Number of ISPs	4	The ISPs have the following market share over the base of users, respectively: 40%, 30%, 20% and 10%
Number of locations	[100, 1000]	Locations where vCaches are deployed

In order to measure that trade-off, the following ratio R is defined as

$$R = \frac{\#\ stored\ objects}{\#\ end\ users\ requesting\ contents} \tag{3.7}$$

The lower R, less storage is needed for serving the same amount of end users at a given instant. Figure 3-21 shows the results of the simulation, showing the percentage of the total number of cached contents with respect to the number of end users. In general terms, intuitively, the better R is achieved when the larger the number of end users and the lower the number of locations.

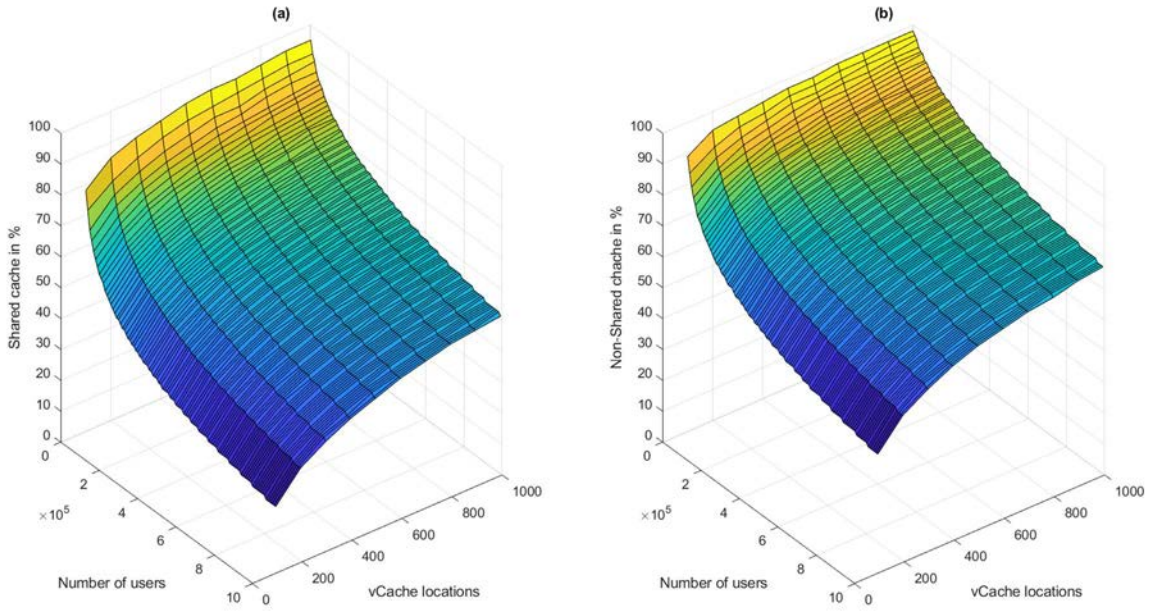

Figure 3-21. Percentage of the total number of cached contents in shared (a) vs non shared (b) approaches with respect to the number of end users (10000 to 1000000), for α = 0,9 and 3000 selectable contents as a function of the number of vCache locations (100 to 1000)

Essentially, concentrating the contents in few locations reduce the need of having multiple replicas of the same content, mainly the popular ones. The counterpart of this approach is that having fewer locations for vCaches imposes certain degree of centralization, then implying more usage of networking resources for distributing the traffic from the vCache locations towards the end users.

It can be also observed that following the shared approach clearly improves the ratio of stored content in the network, that is, for the same number of locations and end users requesting contents, less storage is needed. This is also intuitive in the sense that concentrating the demands from the different ISPs, the number of replicas for the less popular content becomes also reduced.

The trade-off between networking and compute/storage resources for the overall network design with the shared and non-shared vCache approaches is left for further study. The analysis now focuses on understanding the particular efficiencies achieved when sharing vCache VNFs among all the ISPs versus maintaining separated VNFs per ISP in the network.

3.5.2 Efficiency analysis of vCache VNF sharing

The following sub-sections analyze the impact of the distinct factors in the simulation from the perspective of shared VNFs performing the caching of contents. In this respect, it can be considered the achieved efficiency *E* as the ratio between the average stored contents per vCache in both the shared and the non-shared network scenarios. Figure 3-22 provides an overview of *E* when varying the number of contents offered, the user preferences on that contents, the number of locations where the vCaches are instantiated, and the number of end users simultaneously requesting contents.

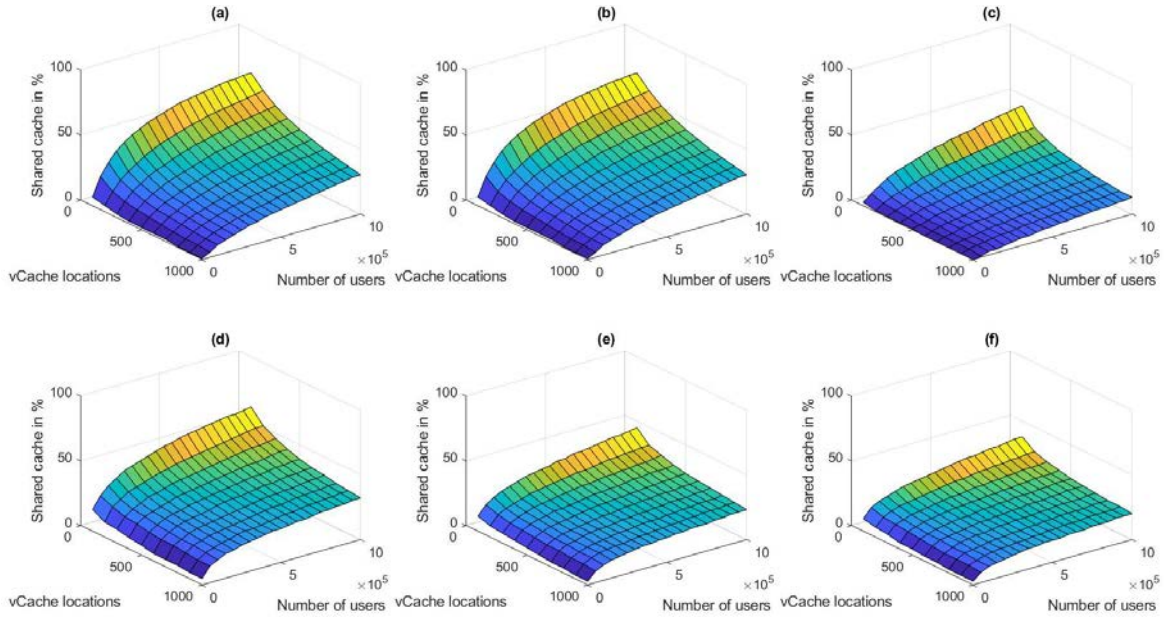

Figure 3-22. Efficiency E (in percentage) in terms of the average contents cached per location considering 100 to 1000 vCaches, and 10000 to 1000000 users uniformly distributed among the vCaches. Graphs show the results when α = 0,6 for 1000 (a), 3000 (b) and 5000 (c) contents, and similarly when α = 0,9 for 1000 (d), 3000 (e) and 5000 (f) contents

3.5.2.1 Impact of the content offer

In general terms, as the content offer increases, i.e. as more contents are available for the end users, the efficiency gain of the shared cache approach diminishes. This is due to the fact that the more contents are selectable, the higher the dispersion of the chosen contents is. Thus, the coincidence of election of contents among ISPs is also reduced, which implies that in shared vCache more individual contents need to be stored.

3.5.2.2 Impact of user preferences

The end user preferences can be more or less disperse. The higher the dispersion, the larger the number of individual content selected. In Figure 3-22 the effect of the dispersion can be observed by comparing the charts with different skew factor. When the selection of content is more concentrated (i.e., higher value of α, or less variance of the distribution), the efficiency becomes higher for the shared vCache case.

As an example, Figure 3-23 represents the efficiency achieved by the shared vCache for 500000 users distributed across 500 locations when 2000 contents are available with both α = 0,6 and α = 0.9. The absolute difference in terms of average cached contents are 177,8 and 166 respectively. Despite the number of contents decreases in absolute value, the relative efficiency increases with greater values of α. That is, the shared approach is better when the preference in the selection of the content is less dispersed.

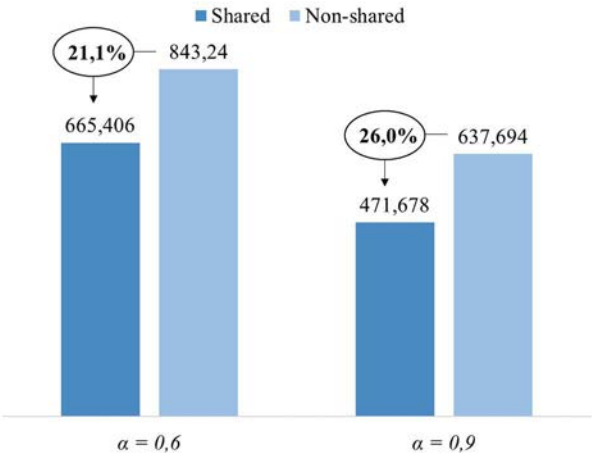

Figure 3-23. Comparison of average number of contents in the shared and non-shared vCache scenarios for 500000 users uniformly distributed in 500 locations with 2000 selectable contents

3.5.2.3 Impact of the number of locations

As long as the number of vCache locations increases, the efficiency of the shared approach decreases. Two effects can be considered here. On the one hand, distributing end users among more vCaches implies the replication of the more popular contents in all the caches. Furthermore, the gain on the less popular contents obtained when concentrating them in less locations is diluted, also provoking the need of storing more objects across the network for those less demanded contents. Figure 3-24 shows graphically that trend.

As observed, when the preferences of the users are more dispersed (α = 0,6) the gain is severely reduced especially for high number of locations. While the percentage of efficiency is similar for low number of locations (e.g., 32,9% when α = 0,9 and 32,6% when α = 0,6 for 100 vCaches), it becomes lessened for high number of them (e.g., 17,6% when α = 0,9 and 9,0% when α = 0,6 for 1000 nodes).

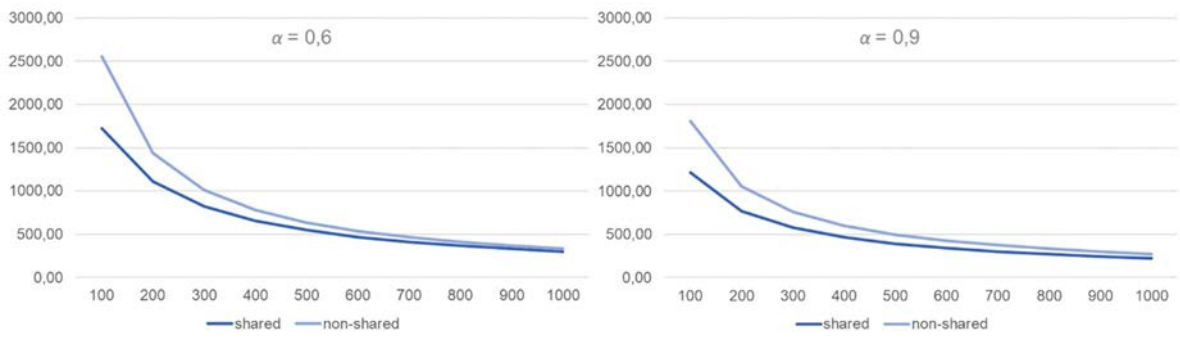

Figure 3-24. Variation on the number of total cached contents when increasing of vCache locations for 3000 selectable contents for 350000 simultaneous end users

3.5.2.4 Impact of the number of end users

As the number of users increases, the number of requested contents also increases. Due to the different popularity of the contents, the number of contents do not grow at the same pace, since certain contents are already cached. Figure 3-25 shows this fact by illustrating the evolution in the number of cached contents for both shared and non-shared situations.

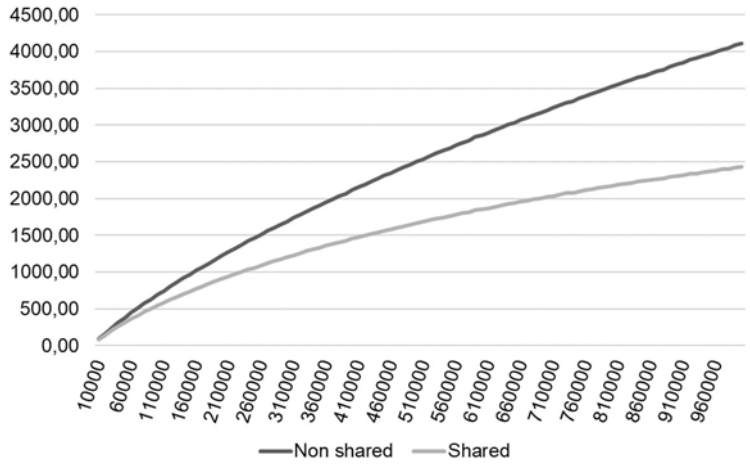

Figure 3-25. Number of total cached contents as the number of end users grows from 10000 to 1000000 for shared and non-shared approaches for α = 0.9, 4000 selectable contents and 100 locations

Interestingly, as the number of end users grows, the efficiency of the shared approach also increase. Clearly, the more popular contents in the shared approach are cached only once instead of four times (one per each of the ISPs), then contributing to the reduction of cached contents. However, as the number of end users increases, also less popular contents are subject of coincidence in the end users' requests, thus contributing to the overall efficiency of the shared scheme.

3.5.3 Economic assessment

In order to perform an assessment of the economics of the shared approach the following cost function per individual vCache location is considered.

$$Cost_{vCache} = \gamma + \rho \times \varepsilon \quad (3.8)$$

where γ represents the costs associated to the instantiation of the vCache in terms of processing and volatile memory, ρ indicates the number of stored contents in the cache, and ε is the storage cost per content. For simplicity, all the vCaches are assumed to require similar CPU and RAM capacity, as well as all the contents are considered to have the same storage needs. For comparison, the total cost should consider all the vCaches instantiated for satisfying the end users' demand.

For the non-shared approach, the total cost is the sum of the individual costs per ISP, i.e., four different VNFs, each of them dimensioned to the specific needs of a particular ISP. In the case of the shared vCache, the cost is the one of the shared VNF aggregating all the contents. How the split of costs is performed for each of the ISP in the latter case is out of scope of this analysis, but it can be assumed that such split could be based on the actual number of content requests by the end users of each ISP, so proportional to the market share.

Figure 3-26 presents one sample case of the average values of content stored per ISP, as well as the resulting stored contents when sharing the vCache.

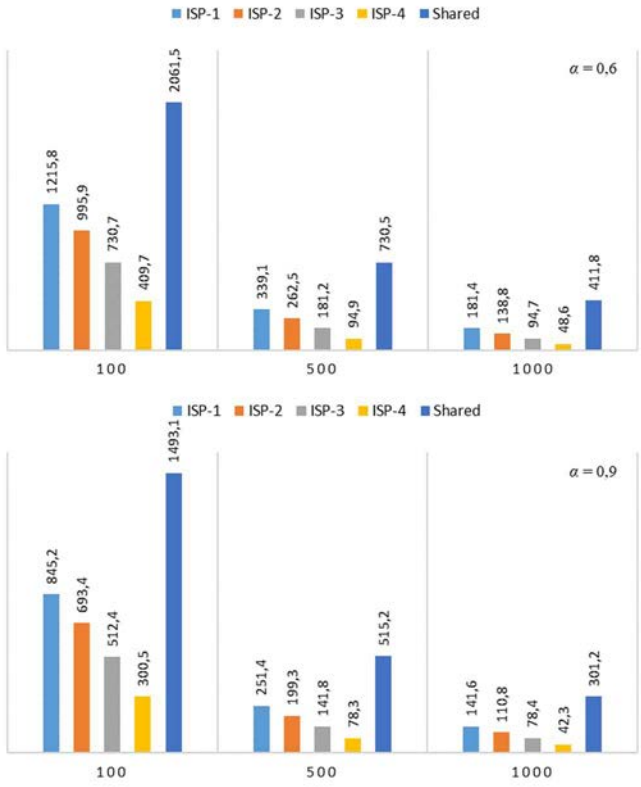

Figure 3-26. Total average cached contents per individual ISP and when sharing the vCache for 100, 500 and 1000 locations, considering 500000 end users and 3000 available contents

To calculate the cost a sensitivity analysis is applied based on the relation among the unitary cost of the CPU and RAM per vCache, Υ, and the unitary cost of storage per content, ε. Three scenarios are evaluated:

- Scenario 1: $\varepsilon = 0.01 \cdot \Upsilon$
- Scenario 2: $\varepsilon = 0.1 \cdot \Upsilon$
- Scenario 3: $\varepsilon = 0.5 \cdot \Upsilon$

With that in mind, and assigning a unitary cost of 1 Cost Units [CU] to Υ, Table 3-13 presents a sample economic assessment considering 500000 end users accessing over 3000 selectable contents, and for 100, 500 and 1000 vCache locations. The assessment is calculated for the two skew factors in the simulation, 0.6 and 0.9 respectively.

The costs for the non-shared case is calculated as the sum of the individual costs per ISP, while in the case of the shared vCache, the cost is calculated considering a single vCache per node for all the ISPs. In order to estimate the costs per ISP the total cost of the shared case is divided proportionally to the assumed market share per ISP.

The table presents the overall [CUs] per solution as well as the percentage of savings for the shared vs non-shared cases. The calculation takes the absolute values for all the vCaches locations.

Table 3-13. Economic assessment (500000 users, 3000 available contents, 100/500/100 nodes)

		100 nodes			500 nodes			1000 nodes		
		0,01 x Υ	0,1 x Υ	0,5 x Υ	0,01 x Υ	0,1 x Υ	0,5 x Υ	0,01 x Υ	0,1 x Υ	0,5 x Υ
α = 0,6	***Non-Shared***	**3752,1**	**33921,2**	**168006,0**	**6388,6**	**45886,1**	**221430,5**	**8634,8**	**50347,6**	**235738,0**
	ISP-1	*1315,8*	*12258,0*	*60890,0*	*2195,6*	*17456,3*	*85281,5*	*2814,3*	*19143,4*	*91717,0*
	ISP-2	*1095,9*	*10059,3*	*49896,5*	*1812,4*	*13623,7*	*66118,5*	*2388,2*	*14881,6*	*70408,0*
	ISP-3	*830,7*	*7406,8*	*36634,0*	*1406,2*	*9561,8*	*45809,0*	*1946,7*	*10466,7*	*48333,5*
	ISP-4	*509,7*	*4197,1*	*20585,5*	*974,4*	*5244,3*	*24221,5*	*1485,6*	*5855,9*	*25279,5*
	Shared	**2161,5**	**20714,6**	**103173,0**	**4152,7**	**37027,4**	**183137,0**	**5118,2**	**42182,1**	**206910,5**
	ISP-1 (40%)	*864,6*	*8285,8*	*41269,2*	*1661,1*	*14811,0*	*73254,8*	*2047,3*	*16872,8*	*82764,2*
	ISP-2 (30%)	*648,4*	*6214,4*	*30951,9*	*1245,8*	*11108,2*	*54941,1*	*1535,5*	*12654,6*	*62073,2*
	ISP-3 (20%)	*432,3*	*4142,9*	*20634,6*	*830,5*	*7405,5*	*36627,4*	*1023,6*	*8436,4*	*41382,1*
	ISP-4 (10%)	*216,1*	*2071,5*	*10317,3*	*415,3*	*3702,7*	*18313,7*	*511,8*	*4218,2*	*20691,1*
	Sh* vs *N-Sh	**42,4%**	**38,9%**	**38,6%**	**35,0%**	**19,3%**	**17,3%**	**40,7%**	**16,2%**	**12,2%**
	ISP-1	*34,3%*	*32,4%*	*32,2%*	*24,3%*	*15,2%*	*14,1%*	*27,3%*	*11,9%*	*9,8%*
	ISP-2	*40,8%*	*38,2%*	*38,0%*	*31,3%*	*18,5%*	*16,9%*	*35,7%*	*15,0%*	*11,8%*
	ISP-3	*48,0%*	*44,1%*	*43,7%*	*40,9%*	*22,6%*	*20,0%*	*47,4%*	*19,4%*	*14,4%*
	ISP-4	*57,6%*	*50,6%*	*49,9%*	*57,4%*	*29,4%*	*24,4%*	*65,5%*	*28,0%*	*18,2%*
α = 0,9	***Non-Shared***	**2751,3**	**23913,4**	**117967,0**	**5354,0**	**35539,9**	**169699,5**	**7731,0**	**41309,5**	**190547,5**
	ISP-1	*945,2*	*8551,5*	*42357,5*	*1757,0*	*13069,8*	*63349,0*	*2416,3*	*15163,2*	*71816,0*
	ISP-2	*793,4*	*7033,5*	*34767,5*	*1496,5*	*10464,8*	*50324,0*	*2108,1*	*12081,0*	*56405,0*
	ISP-3	*612,4*	*5223,7*	*25718,5*	*1208,9*	*7588,8*	*35944,0*	*1783,7*	*8837,1*	*40185,5*
	ISP-4	*142,3*	*3104,7*	*15123,5*	*891,7*	*4416,5*	*20082,5*	*1422,8*	*5228,2*	*22141,0*
	Shared	**1593,1**	**15030,9**	**74754,5**	**3076,0**	**26259,6**	**129298,0**	**4011,7**	**31116,7**	**151583,5**
	ISP-1 (40%)	*637,2*	*6012,4*	*29901,8*	*1230,4*	*10503,8*	*51719,2*	*1604,7*	*12446,7*	*60633,4*
	ISP-2 (30%)	*477,9*	*4509,3*	*22426,4*	*922,8*	*7877,9*	*38789,4*	*1203,5*	*9335,0*	*45475,1*
	ISP-3 (20%)	*318,6*	*3006,2*	*14950,9*	*615,2*	*5251,9*	*25859,6*	*802,3*	*6223,3*	*30316,7*
	ISP-4 (10%)	*159,3*	*1503,1*	*7475,5*	*307,6*	*2626,0*	*12929,8*	*401,2*	*3111,7*	*15158,4*
	Sh* vs *N-Sh	**42,1%**	**37,1%**	**36,6%**	**42,5%**	**26,1%**	**23,8%**	**48,1%**	**24,7%**	**20,4%**
	ISP-1	*32,6%*	*29,7%*	*29,4%*	*30,0%*	*19,6%*	*18,4%*	*33,6%*	*17,9%*	*15,6%*
	ISP-2	*39,8%*	*35,9%*	*35,5%*	*38,3%*	*24,7%*	*22,9%*	*42,9%*	*22,7%*	*19,4%*
	ISP-3	*48,0%*	*42,5%*	*41,9%*	*49,1%*	*30,8%*	*28,1%*	*55,0%*	*29,6%*	*24,6%*
	ISP-4	*-12,0%*	*51,6%*	*50,6%*	*65,5%*	*40,5%*	*35,6%*	*71,8%*	*40,5%*	*31,5%*

As can be observed from the analysis, there are important savings when the shared approach is followed. The savings are higher for lower values of ε (i.e. storage costs) mainly due to the fact that, when sharing, the overhead processing costs of the vCache are reduced from 4 times (one per ISP) just to 1 (the shared vCache). The processing costs impact more severely to the ISPs with lower market share, since that processing costs can be considered as fixed independently of the number of contents to be delivered.

Obviously, for a given number of locations, the total costs increase as ε grows, since the contribution of storage to the total cost increases, as well. The costs also increase as the number of locations increases, essentially because both the processing costs of the vCache (one per location) and the number of contents stored in the network increase, as well.

The impact of the skew factor does not vary too much the relative savings in percentage, with better behavior as α increases, but not significant. In absolute terms, the higher the concentration of user preferences (i.e., higher α), the lower the overall cost.

When looking at the impact on the ISPs, the most benefitted ISPs of following the shared approach are the ones with lower market share, showing very important savings. As the weight of the number of contents in the final cost increases, the savings are reduced. This is because the fixed costs of the vCache processing are diluted when summed up with the storage costs. This is more obvious on the ISPs with higher market share, since they contribute more to the total number of contents requested in the network.

3.5.4 Summary of the contribution

As part of the softwarization process, the sharing of VNFs emerges as an opportunity for cost-efficient deployments by sharing functions as long as infrastructure. This sub-section has contributed with the analysis of the efficiency gains due to sharing vCaches in CDN-as-a-Service offering, as exemplary case of VNF sharing (the analysis shown has been accepted as conference paper [A22]).

3.6 Summary and outlook

This chapter addresses the Objective 2 of this Thesis: *to define novel architectural models that could support such transition towards software-driven telecom networks, addressing novel situations enabled by these new technological paradigms.*

This chapter has overviewed a number of contributions at architectural level, which have produced the following outcomes:

- Architecture describing cooperation in the programmability of service and transport concerns, which has been released in IETF as RFC [A13], and has been contributed to ETSI NFV EVE specifications [O24]. Furthermore, an architecture to allow vertical customer's programmability of slices is currently under submission as journal paper [A7].
- Interconnection of multi-provider infrastructures for service orchestration, considering multi-domain NFV-enabled carrier networks, which has been published as part of a journal paper [A14].
- Options of federation of MEC environments, which has been published as a journal paper [A15].
- Multi-domain slicing integrating fronthaul/backhaul aggregation networks, which has been published as conference papers [A16] and [A17].
- Service edge determination, which has been published in part as conference paper [A18] and contribution to IETF [A20], with further analysis in [A19] .
- Analysis of service blocking probabilities in multi-domain service provisioning, considering compute and network capabilities in different administrative domains, which has been published as conference paper [A21].
- Analysis of efficiency gains when sharing virtualized delivery points in a CDN-as-a-Service slicing scenario, which has been accepted as conference paper [A22].

4 ADVANCES AT SERVICE LEVEL

This chapter provides a number of insights on contributions to advances at service level related to the work in this Thesis.

Specifically, the following aspects have been main subjects of research:

- Analysis of the applicability of SDN and NFV techniques for a virtualization-based roaming solution.
- Analysis of new multicast distribution solutions through the introduction of IGMP/MLD proxy with multiple upstream interfaces

The following sub-sections provide further details on each of these lines of work.

4.1 Virtualized multi-domain roaming solution

Traditionally, technological innovations are introduced in the network for satisfying new service demands, for enhancing existing service, or for delivering the same services in a more cost-efficient manner. However, it may happen that factors external to the pure technological domains, e.g., new regulations, foster the development and the adoption of technological innovations to cope with new market landscapes.

An example of a service influenced by both technical and non-technical aspects is the case of the roaming service. When roaming, a mobile end user from Operator A in Country X (i.e., the home network) is allowed to use the infrastructure of Operator B in Country Y (i.e., the visited network) for accessing mobile services (e.g., for voice, messaging and data). Specifically focusing on data services, a roaming end user typically accesses the home operator specific data services, creating a situation where data traffic must be routed from/to the home network up to the visited network where the end user is connected to (procedure which is commonly known as *home routing*). This fact does not only generate a large amount of interconnection traffic, but also a poorer user experience given by the larger latency in delivering traffic.

In addition to that, recent regulation changes have been introduced by the European Union (EU) with respect to roaming in the European single digital market. This new regulation, known as *Roam Like at Home* (RLAH), has been effective since June 15th, 2017. According to RLAH, roaming is charged at domestic prices, then benefiting the end users that previously were usually billed with expensive roaming tariffs.

This positive change in the end user side has not been accompanied yet by a transformation in the way the roaming services are provisioned by network operators, which maintain the already established interconnection architecture and cost structure. On the other hand, the advent of RLAH promoted new habits among the end users resulting in an ever-increasing demand of data intensive services, even while they roam. Such situation is severely challenging the current mode of operation of roaming services as delivered nowadays by mobile operators.

Figure 4-1 [O75] presents the evolution of the average data consumption per subscriber in roaming in the European Economic Area (EEA)[1]. Looking at the reported data, the total

[1] The EEA includes the countries forming the EU plus the countries of the European Free Trade Association (EFTA), all of them being part of the EU's single market.

increase on the average Gigabyte (GB) consumed per end user after the activation of the RLAH regulation is of 397% for the period Q2 2017 - Q3 2019. Comparing Year-on-Year (YoY) growth, for avoiding seasonal impacts, it can be found that the YoY increase for Q1 2017-2018 is 280%, while the YoY increase for Q1 2018-2019 represents an additional increment of 54%. It is therefore clear that the new regulation has had a significant impact on the overall demand of roaming services, particularly on the data consumption while the end user is abroad.

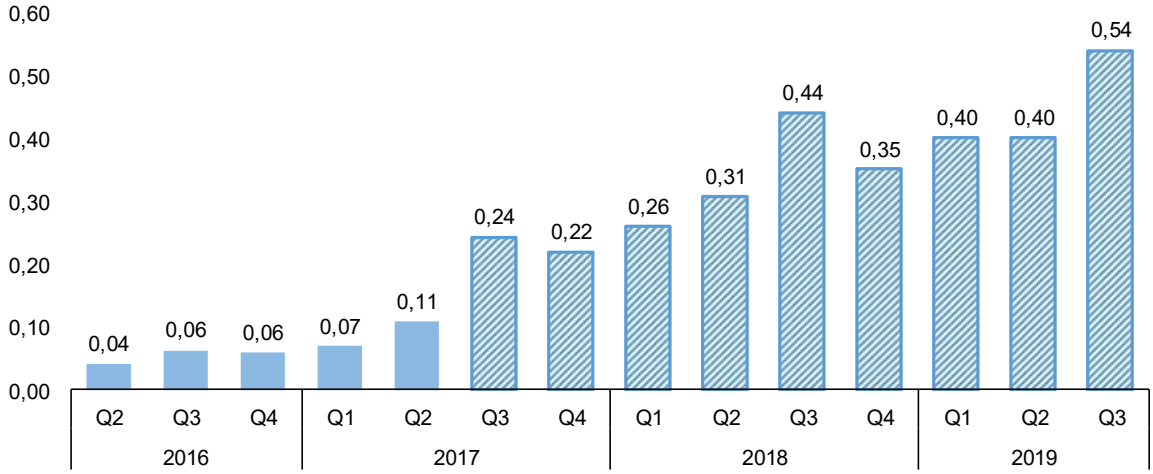

Figure 4-1. Monthly average consumption in GB per roaming subscriber [O75] (solid columns show consumption in quarters previous to the applicability of RLAH)

Looking at the impressive growth percentages triggered by RLAH, it is hence necessary to evolve the way roaming services are provided, transitioning towards a future mode of operation that could easily scale in line with the growth perceived in the users' demand.

To this respect, SDN and NFV bring a new possible strategy to evolve existing interconnection scenarios for roaming services towards more dynamic and better performance scenarios. In this sense, a potential solution to be considered is the deployment of virtualized mobile packet core entities from the home operator into the visited operator premises.

4.1.1 Roaming in existing mobile networks

The roaming service allows an end user to use the infrastructure of the visited network for accessing mobile services. This analysis focuses on the access to data services, i.e. Internet and associated content-related services as subscribed by the end user to the home network operator. It is also important to note that the analysis here is centered on 4G/LTE networks. Nevertheless, 5G networks adopt the same approach as LTE for roaming and it is not envisage any substantial change in the way roaming services are delivered.

4.1.1.1 Roaming architecture

In the LTE architecture, the Evolved Packet Core (EPC) [O76] is in charge of providing IP connectivity and session continuity to the mobile terminal or User Equipment (UE) as it moves around. Figure 4-2 represents the basic entities of the EPC involved in the roaming procedure [O77]. Once the UE of an end user moving to another country attempts to attach to a visited network, the Mobility Management Entity (MME) of the visited network

identifies the newly connected device and tries to register it into the system. During this process, the MME is able to detect through signaling procedures that the UE belongs to a foreign network. In case that the two operators have a valid roaming agreement, the MME of the visited network retrieves the information of the service subscriptions associated to that UE from the Home Subscriber Server (HSS) of the home network. With such information, the UE becomes registered in the home HSS as located in the visited network and the roaming subscriber can start using the home network services from the visited network infrastructure.

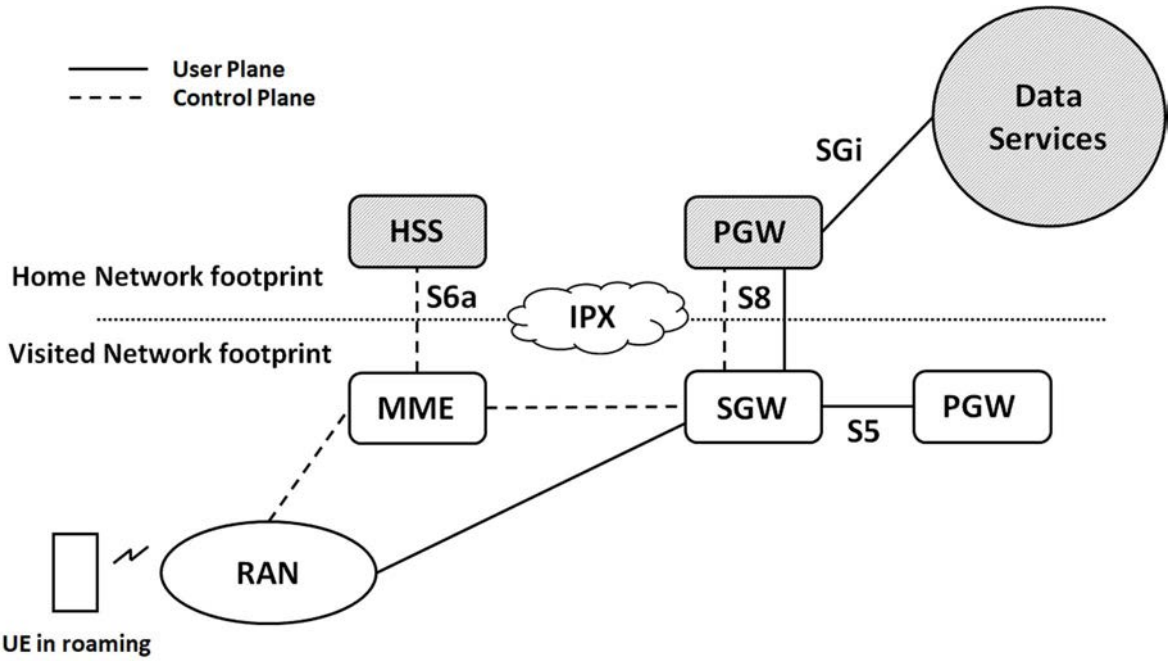

Figure 4-2. LTE roaming architecture (shadowed boxes represent home network elements)

The S5/S8 interface identifies the logical connection between the Serving Gateway (SGW) and the Packet Data Network Gateway (PGW) in the EPC, with the S5 referring to the case when this connection is done to a PGW of the local network while S8 relates to the logical connection to a PGW of a visited network as needed for enabling the roaming services. Usually, when an end user is in roaming and thus attached to a visited network, the connectivity to external networks (e.g., the Internet) via the SGi interface is gained through its home network (home routing). The motivations for that are basically that the local breakout option presents incompatibility issues, such as e.g. different billing systems, and lawful interception obligations.

The logical interconnection represented by the S8 interface is typically arranged by leveraging on an Internetwork Packet Exchange (IPX) provider which enables the interconnection of network operators for the interchanging of IP services with committed QoS. Figure 4-3 depicts a simplified architecture of the IPX interconnection model [O78].

In case network operators do not have a direct interconnection, the access to data services (e.g., Internet) by a roaming end user implies the delivery of the data from the home network to the visited network via the IPX infrastructure. In this way, the home operator incurs in costs due to the usage of the visited network and the IPX infrastructure, which is usually proportional to the volume of traffic transited among providers. Thus, the sustainability of

the service becomes compromised as the data traffic demanded by the roaming users increases.

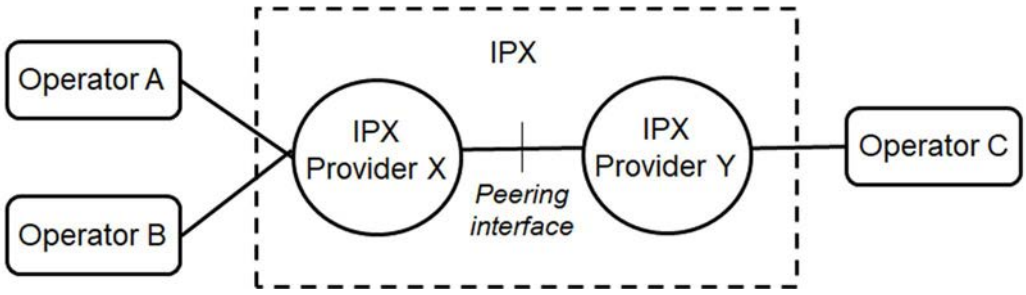

Figure 4-3. Simplified IPX model

4.1.2 Virtualized solution for the support of roaming users

The feasibility of virtualizing EPC components, including functions for complementary services deployed on the SGi interface, offers new possibilities for reconsidering the existing roaming architecture. The concept behind was already described in [A23]. The main idea consists then in the instantiation by the home operator of a PGW (and other complementary components when needed, such as e.g. content delivery points) in the form of a virtual network function (VNF) within the premises of the visited operator. Figure 4-4 depicts the prospected solution.

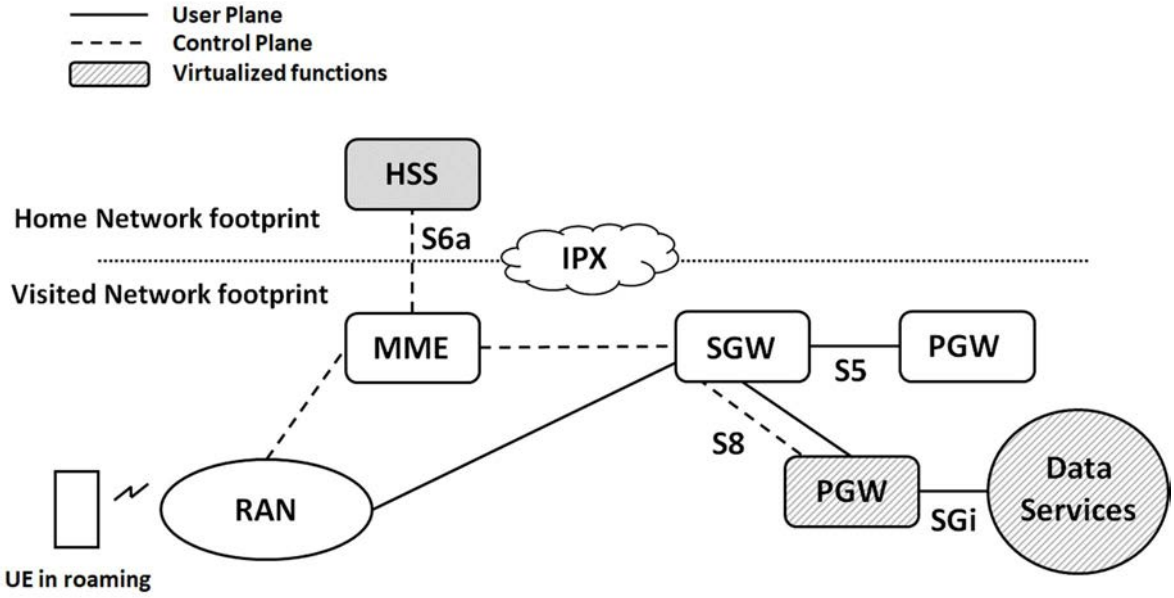

Figure 4-4. Virtualized-based LTE roaming architecture proposal (shadowed boxes represent home network elements, being virtualized the stripped ones)

This can be done by leveraging on the NFVI facilities made available by visited operators to home operators. However, it is yet necessary to define mechanisms that could allow the orchestration of VNFs through administrative domain boundaries, as well as additional configuration actions (e.g., in the DNS of the home network) for a full enablement of the roaming service.

4.1.3 Virtualization-based roaming solution

4.1.3.1 Proposed solution

Three phases can be distinguished in the provision and execution of the virtualized roaming service between Operator A (the home network) and Operator B (the visited network), namely the service preparation, service creation and service activation phases. Figure 4-5 shows a generic flowchart of these three phases, summarizing the main result for each phase, which are described next.

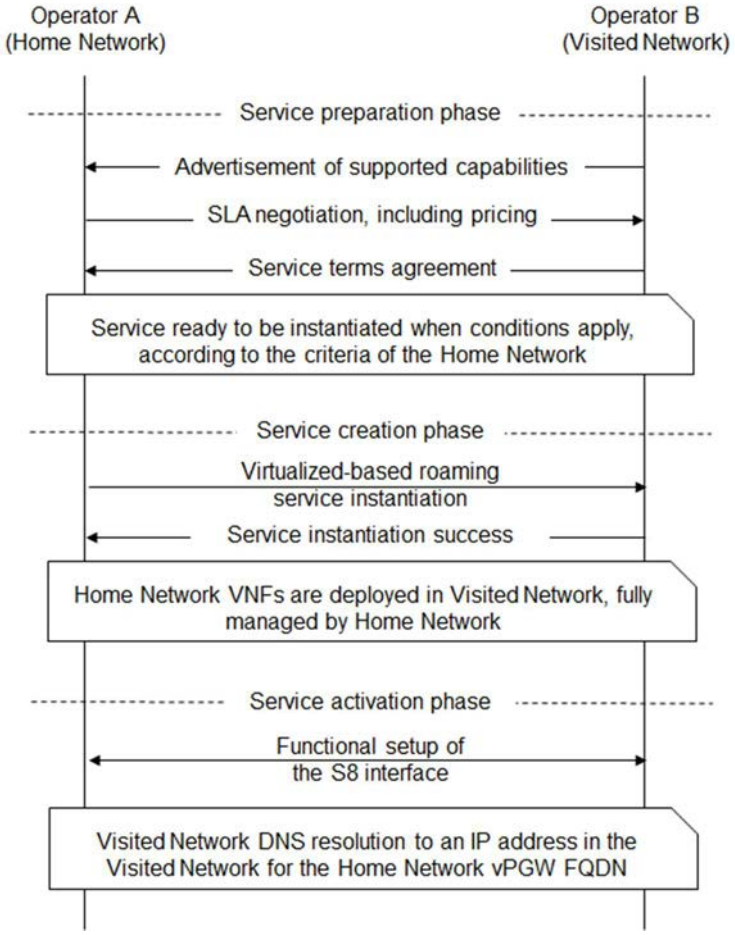

Figure 4-5. Service preparation, creation and activation phases for the virtualized-based roaming service

4.1.3.1.1 Phase I – Advertisement of capabilities between operators and service preparation

Before any kind of service interaction, the operators have to interchange information about their respective supported capabilities. Specifically, for the virtualized roaming service, potential visited operators have to advertise aspects such as resource availability, geographical location, product offerings (e.g., NFVI or VNF as a service), orchestration features (e.g., versioning, supported interfaces), etc. For this case the 5GEx architecture is assumed, as described in Section 3.2.1 for the multi-domain orchestration.

In these scenarios of multi-provider orchestration, this interchange can be published in the form of a service catalog per administrative domain, accessed by potential home operators.

This first phase would also include a number of business-related actions, such as service pricing and SLA negotiation, which are out of scope of this analysis.

4.1.3.1.2 Phase II – Virtualized roaming service orchestration

Figure 4-6 graphically describes a detailed view of the workflow for the creation of the service between Operator A (the home network) and Operator B (the visited network).

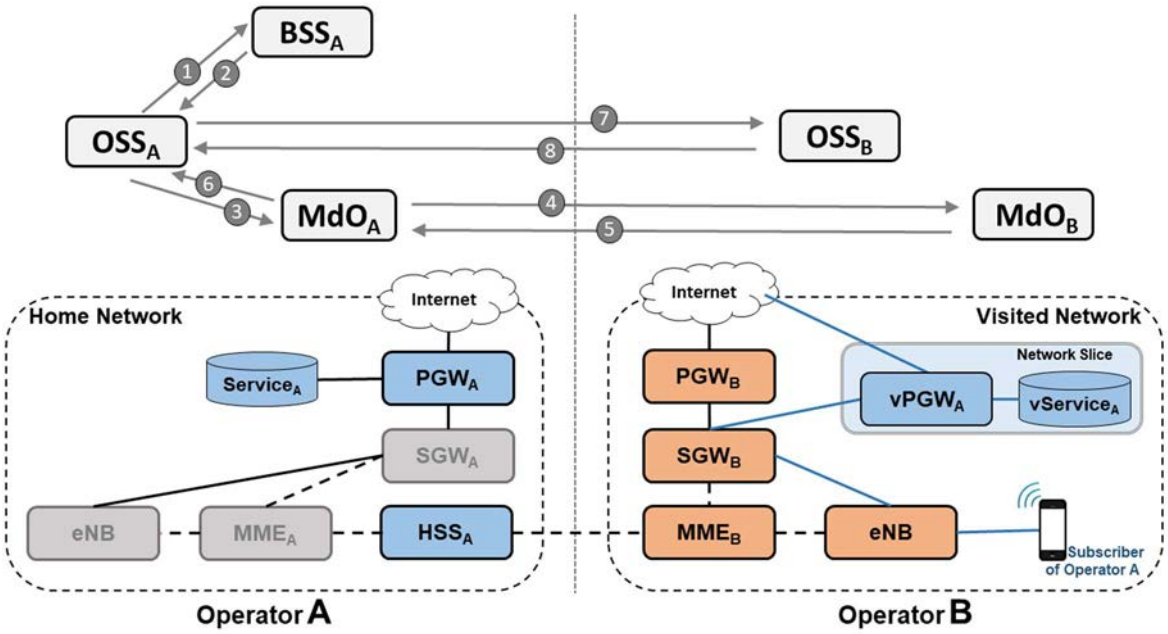

Figure 4-6. Workflow for virtualized roaming service creation

Once the home and visited network operators agree on the terms of provision of the virtualized roaming service, the orchestration is triggered under a number of conditions. In this case, the trigger for orchestrating a vPGW in the visited network is considered to be the number of end users from Operator A roaming in Operator B network.

Step 1. The home network continuously monitor the number of roaming users attached to the visited network. In the event of crossing a given threshold on the number of roaming users, the OSS of Operator A will notify that case to its corresponding BSS.

Step 2. The BSS then instructs the OSS to initiate the provision of the virtualized roaming service.

Step 3. The OSS requests the MdO to deploy the vPGW (and associated virtual functions, noted as vService$_A$) in the visited network, including the necessary connectivity among operators. At this stage Operator A has intelligently decided when and where to deploy the VNFs of the roaming service with the information previously shared by Operator B.

Step 4. The multi-domain orchestration is initiated. The MdO of the home network request to its counterpart, the MdO of the visited network to deploy the VNFs of the service on its premises.

Step 5. Once deployed, the MdO of the visited network indicates success of operation to the MdO of the home network, including information (i.e., identifiers) of the interfaces to access the deployed VNFs for configuration and operation.

Step 6. The MdO of the home network forwards that information to its OSS in such a way that Operator A can configure and monitor the deployed VNFs by OSS, as if they were part of domestic assets.
Step 7. The final configuration of the VNFs involves interaction among OSSs of both operators since the access to the network functions is indirect (i.e., OSS of the visited network mediates in that access).
Step 8. The success of the configuration is confirmed by the OSS of the visited network.

As result, the VNFs of the roaming service have been instantiated and are up and running in the visited network.

4.1.3.1.3 Phase III – Virtualized roaming service activation

As mentioned before, pure orchestration is not sufficient for activating the virtualized roaming service. Part of the logic of the service activation exceeds the purpose of the multi-domain orchestration.

The additional action to accomplish is basically to configure the PGW selection process and the functional setup of the S8 interface between the SGW of the visited operator and the vPGW.

The procedure follows the guidelines of [O77]. When a roaming UE from Operator A sends an attachment request using the visited network, the MME of Operator B identifies and registers such UE into the visited network based on the information acquired from the home network HSS. During this process, the home network HSS provides to the visited network MME the identifier of the home network PGW that the roaming UE must connect to. Typically, that identifier is coded as a Fully Qualified Domain Name (FQDN) [O79], requiring this FQDN to be translated into an IP address by a DNS server within the visited network.

In this case of virtualized roaming service, the home network HSS provides to the visited network MME an FQDN for the vPGW. The mapping of that FQDN and the actual IP address of the vPGW (and IP address of the visited network) should be configured in the DNS server as part of the service creation. With such mapping, once the PGW procedure concludes, the roaming UE is able to connect to the vPGW, completing the virtualized roaming service activation.

4.1.3.2 Experiment setup

For the experiment, a real-life prototype is built and two different operators are considered. Figure 4-7 shows the geographical location and the hardware and software setup of each of these operators, namely Operator A and Operator B. Specifically, Operator A envisages an NFVI-PoP physically located in Berlin, Germany, while Operator B envisages an NFVI-PoP physically located in Madrid, Spain. Each NFVI-PoP is composed of two Intel-based servers running CentOS 7 as Linux-based operating system while OpenStack 11 is employed as Infrastructure-as-a-Service (IaaS) for handling the underlying hypervisor – Kernel-based Virtual Machine (KVM).

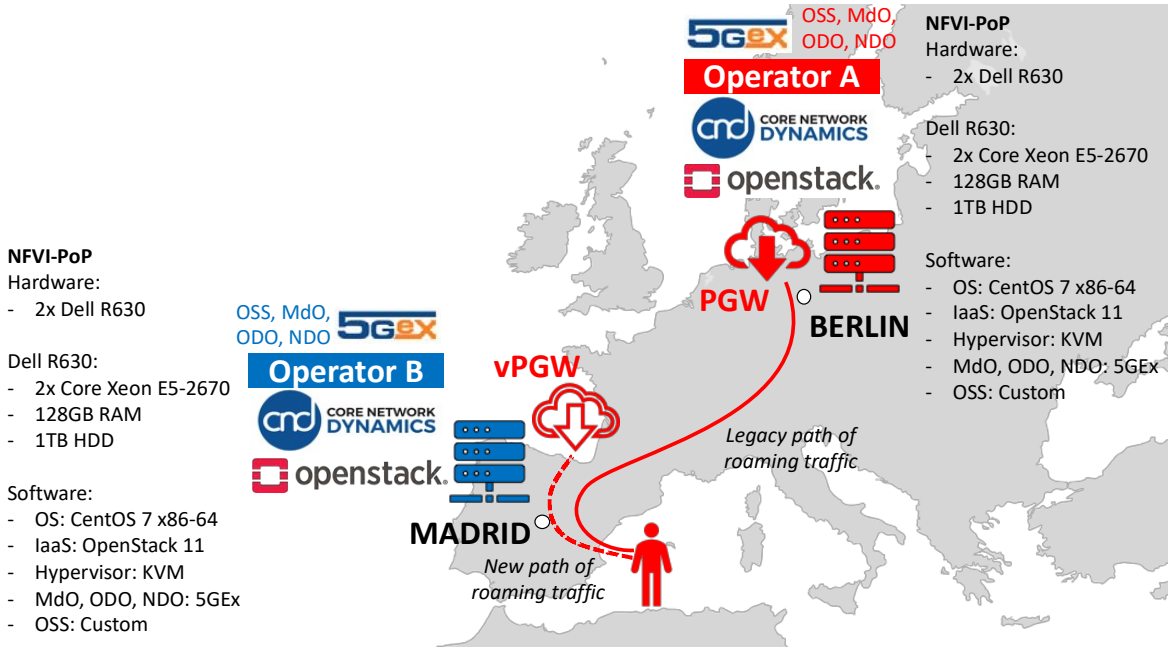

Figure 4-7. Experimental roaming setup consisting of two test-sites: Madrid and Berlin

Moreover, each NFVI-PoP is equipped with the 5GEx implementation of the MdO in charge of the multi-provider orchestration. In addition to this, two domain orchestrators are considered, namely the OpenStack Domain Orchestrator (ODO) and the Network Domain Orchestrator (NDO). The 5GEx implementation of the ODO is in charge of interacting with the virtualized infrastructure for the deployment of the virtual machines supporting the VNFs to be deployed by the home operator, while the 5GEx implementation of the NDO is in charge of resolving the connectivity needed for the new S8 interface.

Each operator also has some OSS/BSS functions that assists on the provision and activation of the service by monitoring the number of roaming users to trigger the deployment of the virtual machines, as described before. The OSS/BSS implementation used in this experiment consists of a mockup of the functionalities required to support the roaming service. Particularly, in this experiment Operator A plays the role of home network while Operator B plays the role of visited network.

Finally, each operator has its own running EPC in advance with a conventional active roaming interconnection established between them emulating an IPX environment. The EPC solution in the experiment is based on OpenEPC (from Core Network Dynamics), which permits the deployment of EPC entities as virtual machines, including virtualized UEs and eNBs. As a result, each NFVI-PoP runs an end-to-end virtualized LTE network (from the RAN to the Core) with a fully compliant 3GPP signaling. It is worth noting that in this virtualized version of OpenEPC, the only layer being emulated is the physical layer, all the others strictly follow the specifications by 3GPP.

In this manner, both the nominal EPC and the VNFs needed for the virtualized roaming services can be deployed and operated in a similar way. The deployment of the EPC on the NFVI-PoPs is done by using an OpenStack Heat Orchestration Template (HOT). Additionally, since the virtualized solution is complementary to the conventional roaming

infrastructure, backward compatibility is ensured, and incremental deployment of the virtual solution can be prospected.

It is assumed that through the advertisement phase, the operators have agreed on the possibility of activating the virtualized roaming service, ensuring that there are enough resources available for such service. Those resources refer not only to networking and computing resources (CPU/Storage) but also to EPC specific resources like the number of supported roaming users.

4.1.3.3 Experiment execution

Figure 4-8 represents the execution operations of the experimental setup described before. The starting point considers that the advertisement phase has been already accomplished through the interaction of the MdOs from Operator A (located in Berlin) and B (located in Madrid), and the service, that was part of the Operator B catalog, has been agreed.

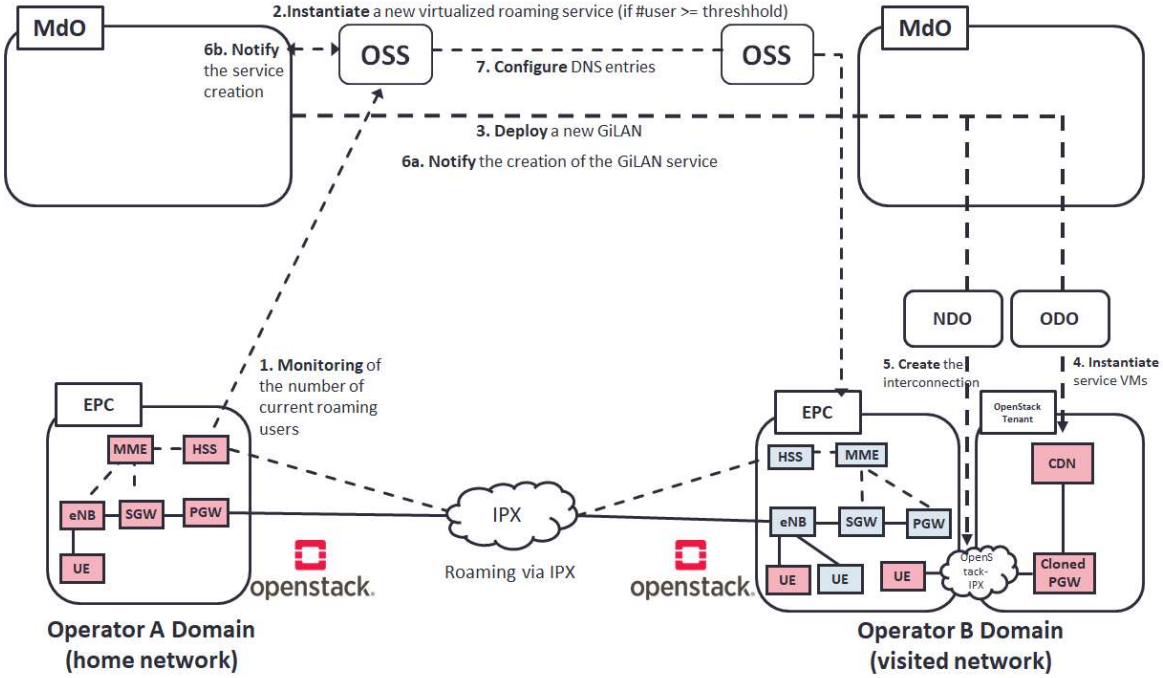

Figure 4-8. Experiment setup of the virtualized roaming service

The detailed workflow is as follows:

Step 1. The monitoring of the number of roaming users in the target visited network is done by checking the HSS of the home network EPC. This can be performed as a management action in the EPC being an internal process of the home network (e.g., performed through OSS).

Step 2. Once the threshold fixed for triggering the virtualized roaming service is exceeded, the OSS/BSS in the home network request its MdO to invoke the instantiation of the VNFs in the visited network.

Step 3. The home network MdO launches the instantiation of the service by interacting with the visited network MdO according with what has been previously agreed.

Step 4. The MdO of the visited network instantiates the creation and deployment of the VNFs of the virtualized roaming service via its ODO. The ODO spawns the

required service virtual machines as associated to a particular tenant of Operator B, being the tenant the Operator A. The VNFs to be deployed consists of a vPGW (cloned from the PGW as facilitated by OpenEPC) and whatever other service that the home network could require to locate in proximity to the roaming users (for instance a CDN end point caching specific contents distributed by Operator A).

Step 5. The MdO of the visited network orchestrates the action of the NDO to creates the required service interconnections, i.e., the connection of the vPGW to the external data services (i.e., the Internet), the setup of the new S8 interface by connecting the SGW with the vPGW, and a tunnel (via the IPX) between the tenant space and the home network EPC to allow the synchronization of non-user related data, e.g. used for the CDN.

Step 6. Once the service has been provisioned, the MdO of the visited network confirms the MdO of the home network the success of the operation. As part of such notification, the MdO of the visited network sends the information about the configuration interface for the recently created VNFs.

Step 7. After the proper configuration of the VNFs, Operator A activates the virtualized roaming service by requesting the configuration of the proper DNS entry in the SGW of Operator B. As consequence of that, the DNS will start forwarding the requests for new roaming users to the vPGW, in the visited network, instead of forwarding them to the PGW in the home network.

From this point on, the new roaming users attaching to Operator B network will use the virtualized roaming service.

4.1.3.4 Experimental results

4.1.3.4.1 Service creation and activation

Figure 4-9 shows the deployment and termination time for the virtualized roaming service, providing details on the time incurred per each component as well as the total time for the service after running 100 experiments. The granularity of the polling request for collecting time information is 1 second, with the figure reporting the average 95th percentile of the different contributions.

For the service deployment phase, the MdO contribution in Figure 4-9 reports the time spent from the point in which the OSS in the home network instructs its MdO to launch the service up to the point in which the MdO in the visited network informs back that instantiation has started from the ODO. The OpenStack (OS) contribution considers the time since the OSS starts polling the OpenStack domain until it reports that the vPGW virtual machine has been created. Finally, the VM contribution contains the time it takes to the vPGW virtual machine to perform the boot up. At the end of this time, the service is up.

For the service termination phase, the MdO contribution reports the time between the request of termination up to the instant in which the visited network MdO notifies back that the termination has been requested to the ODO. Finally, the OS contribution considers the time from the point in which the OSS starts polling OpenStack asking if the vPGW VM and the instant in which all the related networks have been removed.

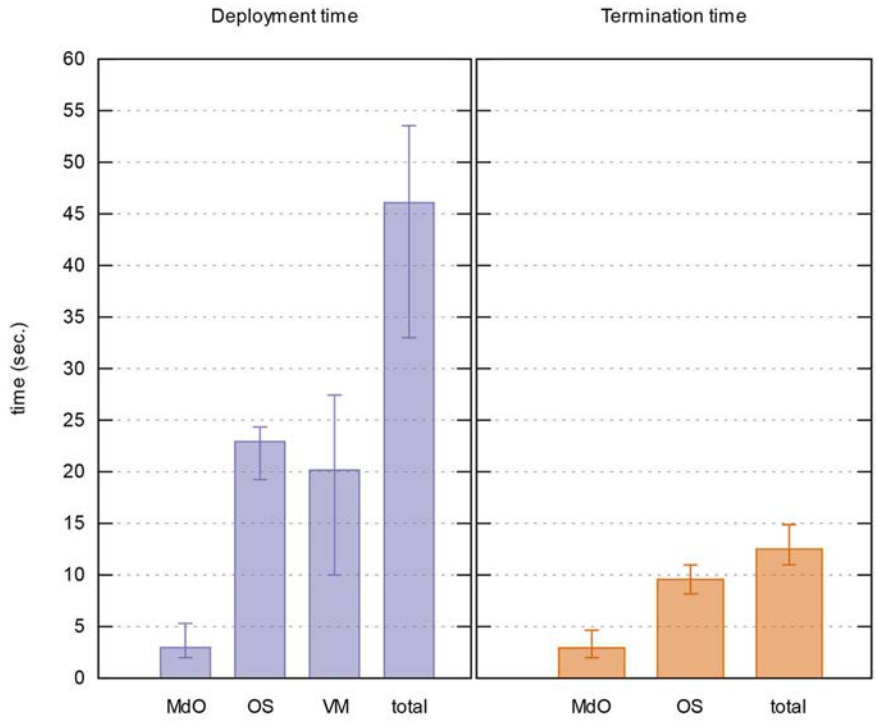

Figure 4-9. Virtualized roaming service deployment and termination time (in total and per involved components) [A14]

According to the results, the virtualized roaming service can be deployed in an automated manner on average in 46 seconds and terminated on average in 12.5 seconds. These figures shown that the service can be created and terminated easily for reacting certain service conditions in a dynamic manner. OpenStack operation is the most time-consuming contribution to the overall process. This is due to the time involved in the creation of the vPGW that implies the disk to be copied, and the time spent in the creation of the interconnection of this same vPGW with the rest of the EPC inside OpenStack.

4.1.3.4.2 Improvement on latency

A relevant effect of deploying a vPGW in the visited network is the possibility of a faster provision of the data services demanded by the roaming user, either a simple access to the Internet or more complex services such the access to specific subscribed content from the home network (if this content is co-located with the vPGW).

During the experiments, the home network components were deployed in Berlin, Germany, while the visited network ones were deployed in Madrid, Spain. The connectivity among home and visited network was established via a basic VPN. Both EPCs at Berlin and Madrid were connected to Internet.

Under these conditions the experiments were run collecting ping traces for an Internet access to google.com from a roaming UE.

The average ping latency in the traditional scenario is ~67 ms. However, by deploying the new proposed solution, the average ping latency drops to ~6 ms [A14]. That is, an improvement on latency of an order of magnitude was observed when moving to the

virtualized solution. The real gain however will depend on the conditions for the conventional roaming, since inherent distance and connectivity conditions between home and visited networks affect in the observed latency, as reported in [O80]. Anyway, for the virtualized solution it can be assumed a latency similar to the one observed in whatever domestic network and similar to the experienced by the local users. Table 4-1 presents a comparison with the related work. The trend shown is the same, reflecting a penalization in terms of delay for roaming users because of the home routing access to contents, including Internet.

Table 4-1. Result comparison to related work.

Study (Date of the experiment)	**Latency penalty for roaming users [ms]**	**Conditions**	**Main takeaways**	**Traffic type**
Mandalari et al. [O81] (not detailed)	~ 60	Real measurements over sixteen European mobile operators from 6 different countries (thus, five different roaming destinations).	Home routing is generalized among mobile operators. The delay penalty varies as a function of the location of the home country. There is variety on the visited network when using 3G or 4G.	HTTP
Michelinakis et al. [O82] (2017)	~ 20	Real measurements involving two Nordic operators (Swedish roamer in Norway).	Experiment on accessing CDNs and Cloud service providers. Due to the home routing, TCP connection time for roamers increases when compared to local users because of the country distance inflation. Roaming users do not benefit from existing local peering or cache agreements.	TCP
Speedtest [O83] (2018), [O84] (2019)	~ 76 ~ 71	Real measurement collected across 39 (28) different European countries by end of 2018 (2019). The measurements were performed through based on Speedtest data from Android devices on 4G LTE cellular connections during Q3-Q4 period (with at least 30 samples).	The provided value reflects the average increment on latency suffered by residents of European countries while roaming in Europe. Roaming agreements vary widely per operator, as well as the roaming destinations per country. Thus, uniform values for a resident across different countries cannot be guaranteed.	ICMP
This work (2019)	~ 67 (when no local instantiation of vPGW) ~ 6 (with local instantiation of vPGW)	Real measurements through a fully operational prototype consisting of an experimental OSS mock-up together with a virtualization-based solution, between Spain and Germany.	The local instantiation of a vPGW from the home network removes the impact of distance inflation observed in conventional home routing solutions.	ICMP

Thus, the ability of deploying a home network environment close to the roamers in a virtualized manner can be an important improvement in the quality of experience of the roaming user that can be translated to a commercial advantage, as well.

4.1.4 Techno-economic insight

The previous sections have shown the technical feasibility of the virtualized roaming solution. The economic viability of it depends on a number of variables to quantify, as the number of roaming users (and from that, the traffic growth), the seasonal effect (or how the roaming users distribute along the time), the geographical footprint (concentration of roaming users per country and within a country), the content offerings and subscriptions of the home network, the evolution of regulation in terms of pricing, the cost evolution of the technology, etc. This makes the calculation complex and particular to each scenario defined.

In order to get a primary insight on economic viability, here a very simplistic model is assumed, just focusing on the impact due to the growth of traffic incentivized by the new regulation. The monetary flows to consider are as follows:

- In the existing mobile data roaming scenario, the home network operator pays a regulated wholesale fee to the visited operator. Wholesale payments by home operator should allow the visited operator to recover both mobile data service origination costs plus transit costs for the home routing. Apart from that, the home operator has to cover the cost of the physical PGW.
- In the proposed virtualization-based roaming solution the home network pays a fee to the visited operator for hosting the virtual PGW. Apart from that, the home operator has to cover the cost of the virtual PGW.

This leads to an analysis that basically compares the incremental cost of the conventional solution, leveraging on higher capacity in both IPX and PGW, versus the incremental cost in the virtualized one, considering the instantiation of the vPGW fitted to the capacity required in the visited network plus the associated costs for hosting it.

An industry analyst forecast [O85] states the declining on the average spend per roaming user in Western Europe in a 60% in the 2016-2023 time frame (moving from $125 to $50). This basically implies that any costs in CAPEX and OPEX derived by the traffic increase as motivated from the RLAH regulation will not be covered by a corresponding increase on revenues. Any increase on incomes will come only from an increment on the number of roaming users. With the new regulation the data consumption by roamers in EU will increase, while the average spend per active roamer will decline in the short term. The same forecast states that the rise in revenues will flatten over the next 2-3 years, starting to decline towards the end of 2021.

Figure 4-10 presents the evolution of the data roaming traffic during the last quarters in the EEA. This evolution on traffic evidences the need of an increase on IPX interconnection capacity for traffic transit in the conventional roaming scenario. Assuming a bandwidth based IPX charging model [O86], such increase implies higher OPEX, even if the price per unit of traffic counts some annual erosion. As reference, the EU has established a wholesale roaming data cap of 4,50 € per GB (plus VAT) in 2019. The price caps act as benchmark prices in wholesale roaming negotiations and any discount on the wholesale roaming market is made from these reference prices. It is expected a decrease of the cap in the next years, according to [O87].

The increment on traffic also implies the need of investing on PGW capacity. The dimensioning of the EPC entities depends on the number of users to be supported (dimension related to the signaling capabilities of the specific entity) and on the traffic that it can deliver (depending on the throughput supported). The PGW is a data intensive entity which provides connectivity to the external data services, thus usually being more limited by traffic than by the number of users (especially in situations when the average traffic per user increases).

Figure 4-10. Data roaming traffic (in millions of GB) in EEA [O75]

Then, an increment in the overall roaming traffic will also force CAPEX investments on PGW. To this respect, the adoption of virtualized solutions for PGW (and EPC in general) is expected to reduce CAPEX and OPEX. For instance, in [O88] a TCO analysis show a significant cost reduction of 69% for the virtualized option. Similar findings are reported in [O89], with savings due to the virtualized solution over the physical one between 49,2% (for a NFVI as a service approach) and 34,2% (for a VNF as a service case). All these calculations, however, are dependent on the particular scenario of analysis. It is worthy to note that the savings reported incorporate the cost of the virtualized infrastructure that in the case of the virtualized roaming solution will be offered by the visited network (translating it into an income).

For the sake of simplicity in the analysis, the following assumptions are considered:

- For both the conventional and the virtualized case, the costs of the existing infrastructure is considered as equivalent. This applies for instance to the cost of the rest of EPC entities, the cost of connection to Internet in either home or visited network (similar in a single market as the one under analysis), etc.
- It is assumed that the SDN and NFV capabilities are already in place at both the home and the visited network and are not only dedicated to the roaming case. This is the general trend, as described before, with previous analysis supporting its viability (e.g., [O90]). In other words, the roaming case would be an incremental case to apply by leveraging on the SDN and NFV capabilities of the operators involved.
- In line with the studies of savings for virtualized EPC, the savings here considered embed the cost in the usage of the virtualized infrastructure, in this case offered by the visited network.

- By moving towards virtualization, the seasonality of the roaming traffic can be better managed adapting the vPGW capacity to the demand in a dynamic way. This implies the possibility of activating vPGW only when needed. This is not the case for PGW and IPX capacity which has to be properly planned and deployed in advance attending the expected peak demand within a given time frame. For instance, in Figure 4-10 there is a clear trend of seasonality showing huge traffic increments on the third quarter each year (comprising Q3 the traditional annual vacation period in Europe). In a usual planning exercise, capacity for the physical PGW and the IPX interconnection should take into account the potential peak reached along the year, without possibility of reducing the capacity in less demanding periods. This dynamicity, however, is not consider here. Instead, comparison for the same capacity is performed in a static manner (dynamicity implies an extra level of cost savings).
- Also leveraging on virtualization (in a similar way as in [O91]), some of the service subscribed by the end user could be provided locally, at the visited network, while roaming, for instance by deploying virtual CDN endpoints with home network contents, actually facilitating a service "like at home", that otherwise is not possible (e.g., distribution rights). This contribution would lower the traffic increase in the virtualized solution versus the conventional one. However, this is not accounted as depends on the specific offerings of each home network operator.

Thus, the analysis can be reduced to the impact of the traffic growth on the PGW platform and the wholesale costs. According to that, it is possible to assume an incremental cost due to the traffic increase for the conventional roaming scenario as

$$\alpha \times T_u + \beta \times T_u \tag{4.1}$$

being T_u the incremental traffic unit (in GB), and both α the PGW and β the wholesale cost (including service originating and transit IPX costs) per traffic unit respectively.

Similarly, in the virtualized roaming case, the incremental cost can be stated as

$$\alpha \times (1 - \gamma) \times T_u \tag{4.2}$$

representing ϒ the percentage of cost savings of the virtualized option of the PGW (including hosting costs) versus the physical one.

Then, the cost ratio between the virtualized and the conventional solutions due to the traffic growth can be established as

$$\frac{\alpha \times (1 - \gamma) \times T_u}{\alpha \times T_u + \beta \times T_u} = \frac{1 - \gamma}{1 + \frac{\beta}{\alpha}} \tag{4.3}$$

Table 4-2 summarizes the parameters used for cost comparison and their values. Figure 4-11 shows the potential savings that could be achieved through the deployment of the virtualized roaming solution.

What can be observed is that the cost level of the virtualized roaming solution with respect the conventional one due to the increase of traffic is dominated by the savings that the vPGW could bring with respect to the physical PGW, since α >> β. From this it could be inferred that equivalent gains could be similarly obtained by absorbing the new demand in the conventional case simply deploying vPGW capabilities in the home network. However, an

additional advantage of the virtualized roaming solution is the minimization of the home network cash-out since no (annually growing) OPEX is payed to third party IPX providers. It is worthy to note that such an OPEX would be typically fixed and uniform along each year since the capacity is dimensioned for the annual expected peak, independently of the actual demand.

Table 4-2. Parameters considered for cost comparison

Parameter	Symbol	Values
PGW cost per traffic unit (€/GB)	α	[100 – 35000]
Wholesale cost per traffic unit (€/GB)	β	4,50
Percentage of cost savings of virtual versus physical PGW	ϒ	20% / 40% / 60%

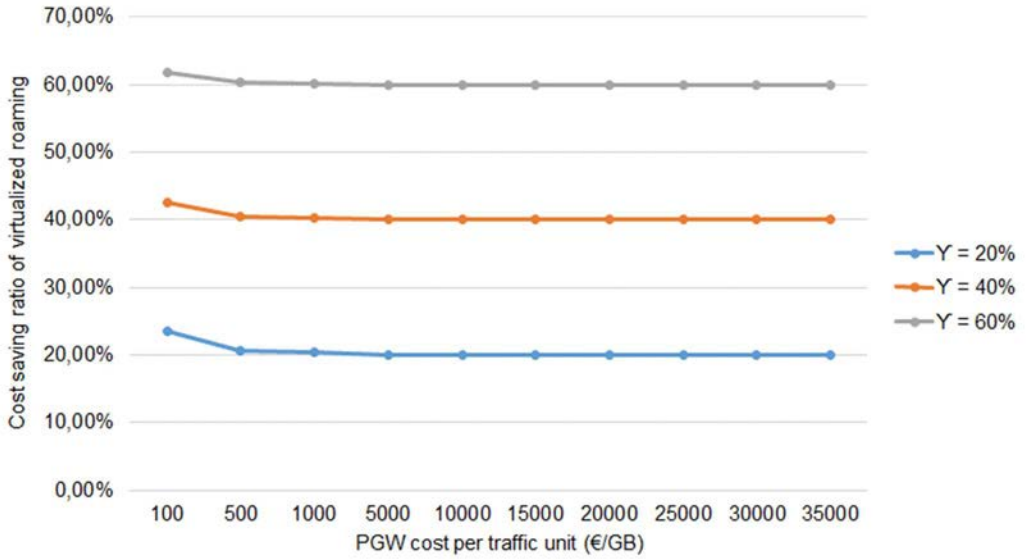

Figure 4-11. Savings of the virtualized roaming solution

4.1.5 Summary of the contribution

This contribution has provided technical insights on a prototype implementation of a virtualized roaming solution, involving multi-domain orchestration. Latency values are provided assessing the improvement achieved by locally deploying visited home vPGWs. Economic insights are also provided, complementing the analysis. This has been published as a conference [A23] and as a journal paper [A14].

4.2 Functionality of IGMP / MLD proxy with multiple upstream interfaces

The Internet Group Management Protocol (IGMP) [O92][O93]for IPv4 and the Multicast Listener Discovery Protocol (MLD) [O93][O94] for IPv6 are the standard protocols for hosts to initiate the process of joining or leaving multicast sessions.

A proxy device performing IGMP/MLD-based forwarding, known as IGMP/MLD proxy [O95], maintains multicast membership information signaled by means of IGMP/MLD protocols on the downstream interfaces, and sends such IGMP/MLD membership report messages via the upstream interface to the upstream multicast routers only when the

aggregated membership information for a channel changes (e.g., when an end user behind the proxy becomes the first user in soliciting a channel, or the last user requesting to leave a given content). The proxy device forwards appropriate multicast packets received on its upstream interface to each downstream interface based on the subscriptions of each downstream interface.

Existing specifications of multicast subscription proxies [O95], for either IPv4 based on IGMP, or IPv6 based on MLD, do allow solely a single upstream interface for requesting the multicast group from the network, with one or more downstream interfaces. An IGMP/MLD proxy device hence performs the router portion of the IGMP or MLD protocol on its downstream interfaces, and the host portion of IGMP/MLD on its upstream interface. The proxy device must not perform the router portion of IGMP/MLD on its upstream interface. This implies that the proxy can only join at one time a unique branch of a multicast distribution tree in the network, and, in consequence, the content is provided solely by a single provider at once.

This fact, even simplifying the real implementation of the proxy, since it merely acts as a relay point for hosts to initiate the action of joining or leaving multicast sessions, produces limitations for certain content distribution scenarios, as described later. For making feasible all of those scenarios, the proxy device shall be able to receive multicast packets simultaneously from different upstream interfaces and forwarding them afterwards to the downstream interfaces according to the subscription requests received.

For all these cases, the existence of multiple upstream interfaces would be needed for simplifying the way in which the end customer receives the desired multicast content. Figure 4-12 graphically presents the global scenario.

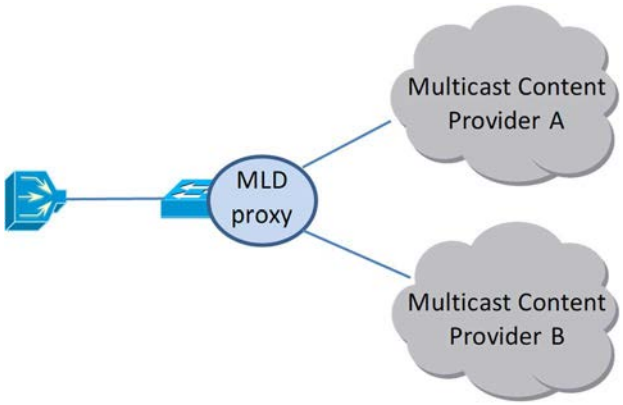

Figure 4-12. MLD proxy with multiple upstream interfaces

Through the different upstream interfaces, the proxy can be simultaneously connected to separate content providers, as in Figure 4-12, or it could be connected to the same multicast content provider but via alternative branches of the same distribution tree. Furthermore, both situations could even co-exist, providing greater flexibility to the system.

4.2.1 Content distribution scenarios favored by multiple upstream interfaces solution

Richer scenarios can be assumed in modern networks, either making possible the co-existence of multiple multicast content providers on top of the same infrastructure, enabling a better and more flexible engineering of the network, or simply facilitating a simpler way of operating the network, which otherwise would require costly mechanisms of solving the

same problem. The scenarios can be divided on those enabling new service offerings, and those facilitating a more efficient network operation.

4.2.1.1 Scenarios enabling new business offerings

In this case, the support of multiple upstream interfaces in the IGMP/MLD proxy essentially permits the subscription to independent contents which otherwise would not be possible except if those contents become integrated into the same platform and delivered by a unique content provider. With this mode of operation, such integration at application level is not needed and the contents are naturally delivered at transport level. Content from independent sources (e.g., from different content providers) is expected to be continuously received through the available upstream interfaces, with the proxy simply collecting the traffic of the channels which become subscribed.

4.2.1.1.1 Enabling multicast wholesale offers for residential services

This scenario refers to the possibility of commercializing complementary multicast services to those offered by the owner of the network where the customer is connected to. This happens when multiple providers offer video services distributed via multicast. Both providers could offer distinct multicast groups. However, more than one subscription to multicast channels of different providers could take place simultaneously.

This case considers the co-existence of two or more different providers. Taking the simplistic example of only two providers, the first provider would represent the one operating the network where the end user is connected (e.g., an incumbent operator), while the other one would provide complementary content services to the end user. Both of them can include in their respective service offerings multicast content, in such a way that the end customer can decide to subscribe both content offerings independently. This is beneficial to the customer since it is giving much more flexibility to the user at the time of consuming video service, independently of which provider distribute the content.

Despite this is exemplified with just two providers, the case could be scaled up to several of them, maybe specialized per kind of content.

Obviously, the complementary providers, as different administrative domains, require to reach agreements with the incumbent provider for using the infrastructure (i.e., the telecom network) to distribute the content in the same way as the incumbent does.

4.2.1.1.2 Neutral network operators

Neutral infrastructure network operators are emerging, being deployed in geographical areas where other operators lack infrastructure. Such a neutral operator can enable the multicast services simultaneously for those other operators making use of its infrastructure through a single proxy supporting multiple upstream interfaces, instead of deploying one proxy function per operator.

When those contents are differentiated by operator, the neutral one needs to manage all those different contents simultaneously, distributing them to distinct customer base, according to the residential service subscriptions.

4.2.1.2 Scenarios facilitating a more efficient network operation

In this other case, the fact of having multiple upstream interfaces assist on a better usage of the available transport infrastructure and permits a smooth and simple operation of such infrastructure. The collection of content through the available upstream interfaces is opportunistic depending on the network conditions or in particular situations, allowing for reconfiguring or reverting the situation in case of any issue in the network. The proxy serves as a mechanism to choose and select the proper content at the right time, as long as the contents can be distributed through different trees in the network, being those trees reachable via the distinct upstream interfaces in the proxy.

4.2.1.2.1 Increased resiliency via fast switching among upstream interfaces

Current solution for multicast delivery relay on the routing capabilities of some protocols to guarantee some levels of resiliency for the distribution of multicast content. This is the case, for instance, of PIM [O95], which helps to rebuild the distribution tree in the event of network failures (like link or node failure).

A simpler scheme could be achieved by implementing the support of differentiated upstream interfaces on IGMP/MLD proxies, providing path diversity through the connection to distinct leaves of a given multicast tree.

It is assumed that only one of the upstream interfaces is active in receiving the multicast content, while the other is up and in standby mode for fast switching. Thus, the proxy can easily switch among the designated upstream interfaces in case of some problems are detected on the original multicast delivery path.

4.2.1.2.2 Load balancing for multicast traffic in the aggregation segment of the telecom provider

Typically, aggregation segments in telecom networks are built in the form of rings, collecting traffic from different aggregation nodes that provide the necessary capillarity to reach the end users. Two nodes act as head-ends of the ring providing connection diversity towards the rest of the network. The capacity of those rings is usually built by using one or several links of either 10 GEth or 100 GEth, depending on the number of nodes in the ring and the traffic collected per node. For multicast distribution, the nodes in the ring can act as IGMP/MLD proxies, including the ones playing the role of head-ends.

A single upstream interface in existing IGMP/MLD proxy functionality typically forces the distribution of all the channels on the same path in the last segment of the network. Furthermore, in the case of having multiple links constituting the ring, the multicast traffic typically uses only one of those links. This is due to the way of working of the hashing algorithms used trying to equally distribute the load among the existing links, which unfortunately do not perform well for the multicast case.

Under situations like the ones described before, the availability of multiple upstream interfaces in the proxies could naturally help to split the demand among the existing links in the aggregation segment, by smartly selecting different channels in each of the upstream interfaces. In this mode, the bandwidth requirements in the metro segment can be alleviated since the load is shared at least in both directions of the aggregation ring.

4.2.1.2.3 Merging of provider networks with different multicast services

The consolidation of service provider's networks is a common fact in mature markets, as a way of improving benefits and economic margins. One of the difficulties faced during the integration of that networks is the integration of the services themselves without impacting in any manner the final user, which can decide to abandon its existing service subscription if service impacts or inconveniences arise. The video or IPTV service enabled by the multicast distribution of content is one of the most sensitive for the end customer, who usually pays a relevant amount of money for accessing such contents.

In some network merging situations, the multicast services provided before in each of the merged networks are maintained for the respective customer base. This is initially conceived to be done in a temporal fashion until the multicast service is redefined in a new single offer, but not necessarily. This could be even not implemented or possible in short term, e.g., because of commercial agreements for each of the previous service offers.

In order to assist such network merging situations, IGMP/MLD proxies with multiple upstream interfaces can help in the transition towards a unique and common video or IPTV offer, simplifying the service provisioning and facilitating service continuity, without affecting the customer perception of the service received.

4.2.1.2.4 Migration of multicast services in an operational network

Any kind of migration in an operational network can produce interruption in the service, which can impact negatively in the perception of the end customer and the reputation of the telecom operator. The severity of the interruption relates to the time the service become unavailable, which depends on the complexity of the service to be migrated. The multicast services are especially sensitive to this fact since the service interruption can affect to the complete customer base since all the customers potentially subscribe to the same set of channels.

This use case considers the situation where a multicast service needs to be migrated. By applying IGMP/MLD proxies with multiple interfaces, the migration can be in principle performed just by switching from one upstream interface to another upstream interface.

In this case the multicast content is initially offered in both upstream interfaces and the proxy dynamically switches from the first to the second upstream interface, according to certain policies, and enabling to shut down the first upstream interface once the migration is completed. The migration can become smooth and without any service interruption at all by delivering migrated channels to new subscriptions while keeping old subscriptions receiving the previous content till the subscription expires.

4.2.2 Solution prospection

The solution targets the reception of multicast sessions or channels through the different upstream interfaces. For that, it is necessary to define a mechanism for selecting the upstream interfaces to be used, as well as the configuration of the potential upstream interfaces to assists on the decision of selection.

Such a mechanism can consider a configuration based on either "channel-based upstream selection" or "subscriber-based upstream selection", or even both of them. Regarding the targeted upstream, the proxy device presents a number of candidate upstream interfaces from

where to select one or more as upstream interfaces. The selection can be performed considering as decision criteria the following parameters: “subscriber address prefix”, “channel/session ID”, and “interface priority value”.

4.2.2.1 Mechanism for selecting upstream interfaces

Two basic mechanisms can be foreseen.

- Channel-based upstream selection. In this option of channel-based upstream selection, an IGMP/MLD proxy device selects one or multiple upstream interfaces from the candidate upstream interfaces in a per channel fashion, based on the configuration of some identifiers for the channel or session.
- Subscriber-based upstream selection. Alternatively, in this options of subscriber-based upstream selection, the IGMP/MLD proxy device selects one or multiple upstream interfaces from the candidate upstream interfaces according to the subscriber requesting the content. Essentially, it bases the decision on the subscriber address prefix.

4.2.2.2 Configuration of the candidate upstream interfaces

The configuration of the following parameters have to be taken into consideration.

4.2.2.2.1 Address prefix configuration

An IGMP/MLD proxy device can be configured with the “subscriber address prefix” and the “channel ID" for each of the candidate upstream interfaces. The channel ID consists of both the “source address prefix” and “multicast group address prefix”.

Once all of this configuration applies, a proxy can decide to select an upstream interface from its candidate upstream interfaces based on the configuration such three-tuple: subscriber, source and multicast group address prefixes (the last two forming what has been previously called channel ID).

An additional tool for prioritizing different rules could apply playing with how specific the prefix provided during configuration is. For instance, the candidate upstream interface having the configuration of an explicit subscriber address prefix is prioritized. On the contrary, if the network operator wants to assign a specific upstream interface for specific subscribers without depending on source and multicast address prefixes, both source and multicast addresses in the address prefix record have to be configured as “null”. Similarly, if the network operator wants to select specific upstream interface without depending on subscriber address prefix, the subscriber address prefix in the address prefix record has to be configured as “null”.

All of this provides great flexibility to adapt the network to specific needs in each moment.

4.2.2.2.2 Channel ID

The configuration of the channel ID consists of the specification of both source and multicast group address prefixes.

In the same way as before, the more explicit the address prefix, the higher priority applies on selecting a given upstream interface. For example, a candidate upstream interface having non-null source and multicast group address configuration is prioritized for the upstream

interface selection. Thus, if a proxy device has two candidate upstream interfaces for the same multicast group address prefix but one of them has a non-null source address configuration, then that candidate upstream interface is selected for the source and multicast address pair.

Furthermore, in the proposed solution, the source address prefix configuration takes priority over the multicast group address prefix configuration.

4.2.2.2.3 Interface priority

The same address prefix may be configured on different candidate upstream interfaces. When the same address prefix is configured on different candidate upstream interfaces, the upstream interface selected for that address prefix can be based on particular interface priority values.

In these conditions, the candidate upstream interface with the highest priority is chosen as the upstream interface.

4.2.2.3 Multiple upstream interface selection for robust data reception

When more than one candidate upstream interface are configured with the same source and multicast addresses for the "channel IDs", and the "interface priority values" are identical, these candidate upstream interfaces jointly act as the upstream interfaces for the channels and receive the packets simultaneously. This multiple upstream interface selection produces duplicate packet reception from redundant paths. It may improve data reception quality or robustness for a given channel, as the same multicast data packets can come from different upstream interfaces at the same time.

However, it is worthy to note that this robust data reception does not guarantee that the packets come from totally disjoint paths. It only ensures that the adjacent upstream routers are different.

4.2.3 SDN control for enabling IGMP/MLD proxies with multiple upstream interfaces

SDN provides a proper framework for the support IGMP/MLD proxy control capabilities. Assuming the existence of a centralized control element, the SDN controller, it is possible to leverage on it to assist on the selection of the proper upstream interface for a given multicast channel. The decision can consider situations like network congestion, shortest multicast path, etc.

Thanks to the flexibility brought by network programming, it is possible to implement sophisticated decision criteria for the delivery of the multicast traffic in the network when compared to traditional multicast service delivery. Conventional multicast delivery is static in essence at the backbone, in the sense that the multicast tree, in case of any failure event, is predetermined by the signaling of the PIM protocol [O95], which essentially provides a static approach from the network. However, towards the access, dynamic behavior is present due to the different customer preferences on specific content, which signaled by explicit indication of subscription to a certain channel through IGMP/MLD snooping [O96], then adding or removing channels to the multicast tree towards the access but not altering the multicast distribution tree at the backbone.

The aforementioned centralized SDN controller can maintain the overall view of the network, including different proxies on it, in order to coordinate actions for the purpose of optimizing the overall multicast delivery. This enables a central point for the planning and operation processes for multicast, usually performed in a manual basis nowadays. SDN control of IGMP/MLD proxies with multiple upstream interfaces is already proposed in [A24].

Two different strategies are to be considered when introducing SDN capabilities for governing the multicast delivery in the network.

- *Alternative 1: Programmability of forwarding elements to perform proxy functionality.* The SDN controller can instruct conventional forwarding elements to implement the proxy function by extending their behavior to support the multiple upstream interface concept, populating forwarding rules to those devices and supporting the proxy logic behavior on the centralized controller.
- *Alternative 2: Configuration of the proxy functionality through standard data models.* In this case, the proxy functionality resides in the device, which is properly configured by the SDN controller. The internal logic of the device acts accordingly, interpreting the controller's instructions that are driven by policies. That instructions are populated through standard data models, e.g. by extending existing YANG models such as [O97].

Summarizing, in the first alternative the logic for using the multiple upstream interfaces is retained by the controller, while in the second case the decision logic is in the device with the policies being set by the controller. Next section described a proof of concept based on *Alternative 1*, evaluating the viability of the proposal for supporting multiple upstream interfaces in IGMP/MLD proxies as a proof of concept, considering *Alternative 2* as future work while accompanying the evolution of multicast standard YANG models in such direction.

4.2.4 Proof of concept of IGMP/MLD proxy with multiple upstream interfaces

4.2.4.1 Proxy implementation

Mupi-proxy (Multiple UPstream Interfaces multicast Proxy) is a proof of concept implementation of the extensions defined in [A24] to support multiple upstream interfaces in IGMP/MLD proxies. It has been implemented on Linux using an SDN application running over Ryu controller. The application controls and Open vSwitch which is in charge of relaying the downstream multicast data flows and the upstream IGMP/MLD control traffic.

The generic working scenario of mupi-proxy is presented in Figure 4-13, where several IP multicast flow providers are connected to upstream interfaces and several subscribers consuming multicast flows are connected to downstream interfaces. In this scenario, mupi-proxy is in charge of relaying the control and data multicast flows among clients and providers. Basically it performs the following actions:

- Relays the IGMP/MLD control requests sent by the clients to the providers following the configured selection policies. For that purpose, mupi-proxy includes a Multicast Upstream Routing Table (MURT) that allows to specify how the upstream interface is selected in terms of the client's IP address, the multicast group and the IP address of the multicast flow source.

- Relays the multicast data traffic sent by the providers to the clients, creating the required switch flow entries that fulfil the policies defined for upstream interface selection.
- Works as a standard Ethernet switch for other unicast traffic.

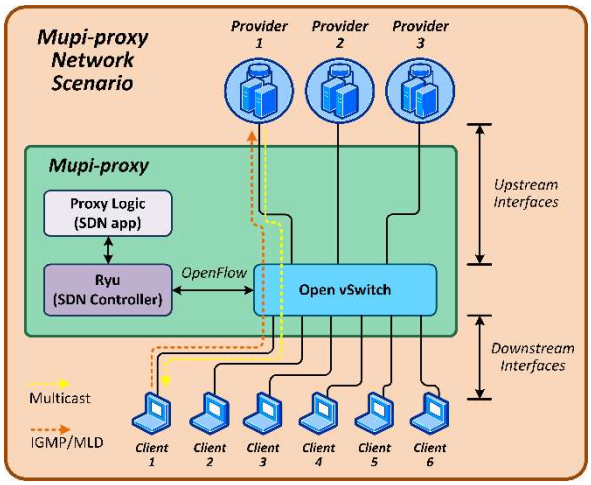

Figure 4-13. Mupi-proxy general scenario

The main task to be carried out when configuring mupi-proxy is the definition of the MURT. Each entry of the MURT is a 5-tuple with the format *(client_ip, mcast_group, mcast_src_ip, upstream_if, priority)*, being:

- *client_ip*: an IP address or prefix used to define the range of client IP addresses the entry applies to.
- *mcast_group*: an IP multicast group or a prefix of multicast groups the entry applies to.
- *mcast_src_ip*: an IP address or prefix used to define the range of source IP addresses the entry applies to.
- *upstream_if*: the identifier of the upstream interface to be used.
- *priority*: the priority of the table entry.

The first three fields *(client_ip, mcast_group and mcast_src_ip)* can be empty. An empty value in these fields is equivalent to default entry 0.0.0.0/0.

When a request to join an IP multicast group is received, the mupi-proxy extracts the client source IP address, the IP multicast group and, if specified, the multicast source IP address. With these three values, it checks the MURT entries to find the ones matching the requested values, selecting the one with the highest priority and relaying the IGMP/MLD request to the upstream interface specified in that entry. In case two or more entries with the same priority are selected, the request is sent to all the upstream interfaces specified in the MURT entries.

Table 4-3. Example of MURT configuration

Num	client_ip	mcast_group	mcast_src_ip	ups_if	pri
1	10.100.0.0/26	224.0.122.5	10.100.0.21	7	30
2		224.0.122.5		8	20
3	0.0.0.0/0	224.0.0.0/4	0.0.0.0/0	9	0

For example, if the MURT is configured with the values shown in Table 4-3, the following queries would be directed to the upstream interfaces specified below:

- Q1: (10.100.0.20, 224.0.122.5, 10.100.0.21). In this case, all entries match. The one selected is entry 1, due to its highest priority. Request is sent to upstream interface 7.
- Q2: (10.100.0.70, 224.0.122.5, 10.100.0.21). Entries 2 and 3 match; entry 2 is selected (interface 8).
- Q3: (10.100.0.70, 224.0.122.6, 10.100.0.21). Only entry 3 match; request sent to interface 9.

4.2.4.2 Exemplary use case: different multicast providers for residential services

This section describes in more detail the general use case where there are subscribers requesting content from different multicast providers, including the functional assessment of the mupi-proxy proposition. Specifically we consider the case of the neutral infrastructure operator that is leveraged by other operators in a certain geographical area where they lack of Points-of-Presence (PoPs).

Such neutral operator will be required to enable the same kind of services that the other operators provide, including multicast contents.

If we assume that the neutral operator operates on cloudified environments (i.e., its PoPs are based on solutions such as [O98] or [O99]), then there is extremely high flexibility at the time of instantiating functions and services by the other operators in a dynamic manner. This also applies to the distribution of the multicast services, where a proxy could be easily instantiated and deployed for enabling the multicast service per operator, creating multiple upstream interfaces towards the different contents. Figure 4-14 illustrates the experimental setup used for the functional assessment.

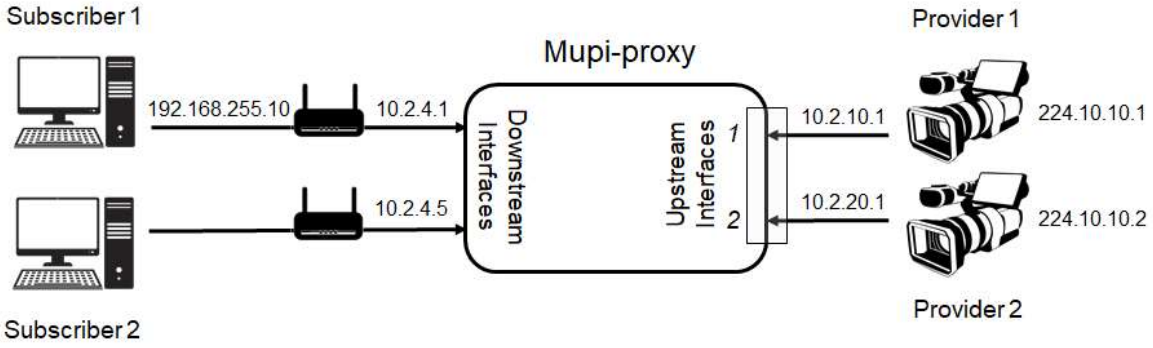

Figure 4-14. Experimental setup for the proof of concept

Two residential networks pertaining to two distinct subscribers are considered. Those subscribers get some routed IP addresses 10.2.4.1 and 10.2.4.5 (they play the role of "public" addresses in a real setup, while for the experiment also private address ranges are used for that purpose). Such subscribers are configured to receive multicast traffic from Content Provider 1 (through the proxy upstream interface 1) and Content Provider 2 (via the proxy upstream interface 2) respectively. The multicast upstream routing table (MURT) at this stage presents the values shown in Figure 4-15.

```
# Multicast upstream routing table config
#
# Format:    Client IP      Multicast     Multicast    Upstream  Priority
#            Addr/Prefix    group         source IP    If Id
#                           Addr/prefix   Addr/Prefix
murt_entry = 10.2.4.1,      ,             ,            1,        10
murt_entry = 10.2.4.5,      ,             ,            2,        10
```

Figure 4-15. Initial MURT configuration

At some point of time, Subscriber 1 acquires the rights of accessing a specific content from Content Provider 2. This can be enabled through the configuration of a policy in the MURT in such a way that Subscriber 1 can simultaneously access contents from both Content Providers through the same mupi-proxy. In the example, Subscriber 1 will request to access the content identified by the multicast group address 224.10.10.2 from Content Provider 2. The MURT is presented in Figure 4-16.

```
# Multicast upstream routing table config
#
# Format:    Client IP      Multicast     Multicast    Upstream  Priority
#            Addr/Prefix    group         source IP    If Id
#                           Addr/prefix   Addr/Prefix
murt_entry = 10.2.4.1,      ,             ,            1,        10
murt_entry = 10.2.4.1,      224.10.10.2,  ,            2,        20
murt_entry = 10.2.4.5,      ,             ,            2,        10
```

Figure 4-16. MURT enabling Subscriber 1 to receive contents from Provider 2

The capture in Figure 4-17 shows a system in the network of Subscriber 1 with internal address 192.168.255.10, first requesting to join the multicast group and port 224.10.10.1:22345, receiving the flow from Content Provider 1 via 10.2.10.1. After that, Subscriber 1 joins the multicast group 224.10.10.2:22345, receiving the flow from Content Provider 2 through 10.2.20.1, due to the new specific MURT entry that has a higher priority than the general entry.

	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	192.168.255.1	224.0.0.1	IGMPv2	46	Membership Query, general
2	2.147154089	192.168.255.10	224.10.10.1	IGMPv2	46	Membership Report group 224.10.10.1
3	4.930789556	192.168.255.10	224.0.0.251	IGMPv2	46	Membership Report group 224.0.0.251
4	6.191502232	10.2.10.1	224.10.10.1	UDP	58	22345 → 22345 Len=16
5	8.771451555	192.168.255.10	224.10.10.1	IGMPv2	46	Membership Report group 224.10.10.1
6	11.189570020	10.2.10.1	224.10.10.1	UDP	58	22345 → 22345 Len=16
7	13.287911911	192.168.255.1	224.0.0.1	IGMPv2	46	Membership Query, general
8	14.494829103	192.168.255.10	224.10.10.1	IGMPv2	46	Membership Report group 224.10.10.1
9	16.189814234	10.2.10.1	224.10.10.1	UDP	58	22345 → 22345 Len=16
10	17.950803076	192.168.255.10	224.10.10.2	IGMPv2	46	Membership Report group 224.10.10.2
11	19.776234638	192.168.255.10	224.10.10.2	IGMPv2	46	Membership Report group 224.10.10.2
12	21.189551312	10.2.10.1	224.10.10.1	UDP	58	22345 → 22345 Len=16
13	21.281670221	10.2.20.1	224.10.10.2	UDP	58	22345 → 22345 Len=16
14	22.846925430	192.168.255.10	224.0.0.251	IGMPv2	46	Membership Report group 224.0.0.251
15	26.189965918	10.2.10.1	224.10.10.1	UDP	58	22345 → 22345 Len=16
16	26.280995544	10.2.20.1	224.10.10.2	UDP	58	22345 → 22345 Len=16
17	26.571514353	192.168.255.1	224.0.0.1	IGMPv2	46	Membership Query, general
18	28.095039822	192.168.255.10	224.10.10.2	IGMPv2	46	Membership Report group 224.10.10.2
19	29.250989596	192.168.255.10	224.10.10.1	IGMPv2	46	Membership Report group 224.10.10.1
20	31.189565406	10.2.10.1	224.10.10.1	UDP	58	22345 → 22345 Len=16
21	31.280726169	10.2.20.1	224.10.10.2	UDP	58	22345 → 22345 Len=16

Figure 4-17. Wireshark capture showing functional behavior of mupi-proxy

4.2.5 Summary of the contribution

This sub-section presented the proposal of a novel IGMP/MLD proxy supporting multiple upstream interfaces supported by SDN control (published as conference paper [A25] and proposed in IETF [A24][A26]). This novel kind of proxy enable a variety of services in operational networks. Proof of concept results validate the proposed functionality (published as conference paper [A27]).

4.3 Summary and outlook

The chapter addresses the Objective 3 of this Thesis: *to assess the feasibility of services leveraging on the aforementioned new paradigms, understanding issues and gaps for fully enabling them.*

This chapter has overviewed a number of contributions at service level, which have produced the following outcomes:

- Analysis of the applicability of SDN and NFV techniques for a virtualization-based roaming solution, which has been published as part of a journal paper [A14] and as a conference paper [A23].
- Analysis of new multicast distribution architectures through the introduction of IGMP/MLD proxy with multiple upstream interfaces, which has been published as conference papers [A25] and [A27], and as contributions to IETF [A24] and [A26].

5 ADVANCES AT TRANSPORT LEVEL

This chapter reports some advances at transport network level related to the contributions of this Thesis.

Specifically, the following aspects at transport level have been main subjects of research:

- Experimentation of SDN-based control for wireless transport networks (WTNs), i.e. microwave radio links.
- Proposal of a mechanism for defining and handling levels of isolation in transport slices.

The following sub-sections provide further details on each of these lines of work.

5.1 Programmability of backhaul transport networks: applicability of SDN to Wireless Transport Networks

The mobile backhaul is the network segment that extends network capillarity collecting and delivering aggregating traffic from the radio network access (i.e., the radio base stations). The mobile backhaul comprises a number of technologies, including IP, Ethernet, optics and wireless transport networks. The latter has been traditionally based on microwave radio links, ultimately leveraging on higher frequency bands such as millimeter wave.

The traditional mode of operation in a network is the deployment of several vendors of the same technological segment in order to promote competition in price and functionality, preventing as well dependencies or lock-ins from particular implementations along the time. Consequence of that, in some circumstances, a given geographical area could be served from nodes from different vendors, and for the Wireless Transport Networks (WTNs) case, even connecting radio links back-to-back. The drawback of such approach is the need of operating different solutions, usually by means of specialized Network Management Systems (NMS) per vendor. Furthermore, due to this particularization, customized tools or solutions that could be necessary for the operation (e.g., radio-link planning tools) either require multiple integrations for each vendor solutions, or even cannot be used for all the vendors at all.

There are a number of incentives in the operator side for the introduction of programmability in this particular technology segment of the WTNs. Some of the motivations are:

- Definition of a common way of controlling and managing Wireless Transport Networks (e.g., microwave links) avoiding the need of handling diverse proprietary solutions for the provision of the same kind of service in the network.
- Road to simplification, since the operational workflows can be homogenized across the different implementations present in the network thanks to a common model of operation, independent of the microwave node manufacturer.
- Enablement of advanced control plane features for rich functionalities, including multilayer coordination (e.g., by combining actions in other technology segments such as the IP network).
- Road to automation, permitting the integration of the control mechanisms of the WTNs with an overarching SDN control approach, allowing for control loop mechanisms to be applied (e.g., for energy saving purposes).

The multi-layer and multi-technology reality of the backhaul networks presents some challenges at the time of providing a common and unified service across this segment.

During service provision any action should be coordinated in time and manner for ensuring service consistency.

The per-vendor specialization in the control of the WTNs produces a lack of agility that prevent deploying advanced services over these infrastructures. There is no way of defining common actions or possibilities for dynamic provision of the WTN substrate for a whole network in a simple way.

SDN brings the opportunity to simplify the operation on one hand, and allow the dynamic provision of advanced services on the other. The key concept to reach those goals is the idea of network programmability, which permits the instantiation of transport network capabilities on-demand and by means of common interfaces, usually in the form of Application Programming Interfaces (APIs). This is especially relevant in the WTN because of the lack of any other existing alternative to do that. The control and management actions in SDN come from a logically centralized controller, keeping the logic to instruct the underlay devices in order to take actions in the data plane.

The static provision of transport capabilities is no longer an efficient way of handling the network; neither the approach of segmenting the network both per-technology and per-vendor to configure each of those segments in an isolated manner. A uniform control and management capability across all those segments is required to simplify the operation of new networks and to enable more efficient resource management such as the reutilization of unused resource from one operator to another, which is not possible with a proprietary approach.

5.1.1 Scenarios of applicability of SDN for WTNs

A number of situations can benefit from the application of SDN to WTNs. Some examples are described next.

5.1.1.1 Selective Traffic Shape and Re-route

Wireless links have a number of special attributes and among them is the fact that the links "breath". This means that the links capacity can change depending on the weather and environment conditions such as rain.

In the traditional WTN, the capacity degradation is handled by statically applying QoS techniques which basically drop packets locally based on their class of service.

SDN provides the opportunity to handle this situation in a more optimal and dynamic way. For instance, if we consider a centralized controller maintaining a global view of the whole network, in case there is a capacity degradation event it is possible to leverage on this to take actions deeper in the network.

The controller can identify the flows that are supposed to go through the degraded link, and apply shaping on them by instructing network elements (the ones of WTN or others in the core). In other words the system can command ahead of time and upfront, the source of this flow to stop, for a moment, transmitting it (or decreasing the bit rate). This way it is possible not only lowering the burden on this degraded link but also avoiding the data of this flow from being discarded too late after having achieved sometimes "a long trip" toward the bottleneck and thus wasting the resources of the intermediate nodes as well. It can even reroute some of the flows to new routes that have available capacity.

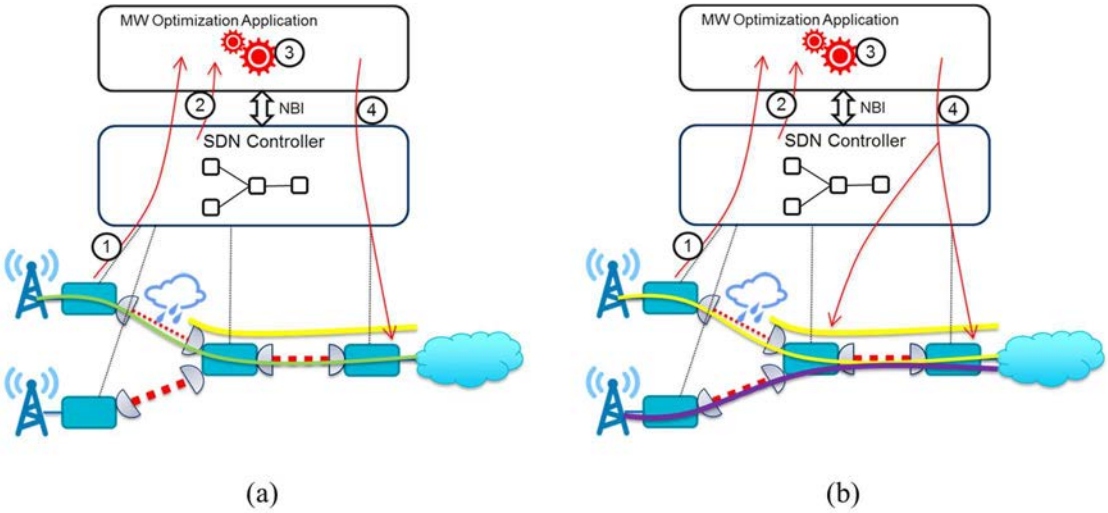

Figure 5-1. Selective traffic shaping (a) and re-routing (b)

According to Figure 5-1, the sequence of actions to perform the resource optimization is the following:

1. Because of a degradation in the transmission conditions of the radio part (e.g., due to rain), an event for change in the capacity link is notified to the optimization application, via the SDN controller.
2. The application generates an updated view of the topology and the flows riding on top of it.
3. With the previous information, the application calculates the optimal point for applying shaping for the selected flows, in such a way that the resulting global resource usage becomes optimal.
4. At the end of this process, the application instructs the SDN controller for populating the necessary rules in the intended network elements for committing the necessary changes.

Figure 5-1 represents two different situations that can benefit from this automatic control. The first scenario, Figure 5-1 (a), assumes an original "yellow" flow with a given capacity which suffers degradation due to weather conditions in some point of time. In that situation shaping is applied to the "yellow" flow to match the available end to end capacity, being transformed in the less bandwidth consuming "green" flow. The second scenario, Figure 5-1 (b), assumes the same affection to the original "yellow" flow, but in this case the action taken is increasing the amount of traffic of another flow sharing the aggregation network, which is the "purple" flow in this example.

5.1.1.2 Dynamic Spectrum Allocation

In regular microwave (MW) point-to-point planning two links that are aggregated at the same point need to use different channels if the angle of separation between the two aggregating antenna at the aggregation node is not large enough (more than 90 degrees). In Figure 5-2 two different base stations (namely *BS1* and *BS2*) are shown being connected via MW links to the aggregation node *AN*. The backhaul link 1 for *BS1* is directed to antenna 1 (*ant.1*) and link 2 for BS2 to antenna 2 (*ant.2*). In this case there are only few degrees of

separation between the antennas, which requires two completely separated channels namely *f1* and *f1'* for DL and UL of link 1, and *f2* and *f2'* for DL and UL of link 2, respectively.

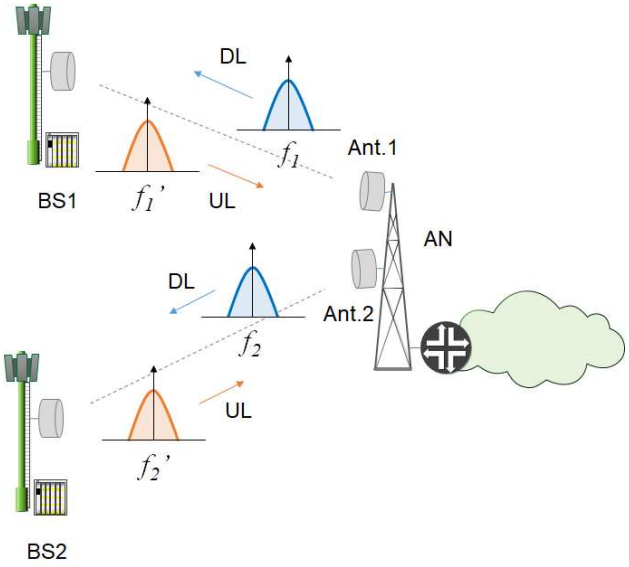

Figure 5-2. Low spectral efficiency case

However, the aforementioned planning leads to low spectral efficiency. Most of the time even small angles in the antenna separation are enough to allow both links to share the same frequency. In Figure 5-3 it is shown a similar scenario as before but with higher antenna separation which enables both link to use the two channels.

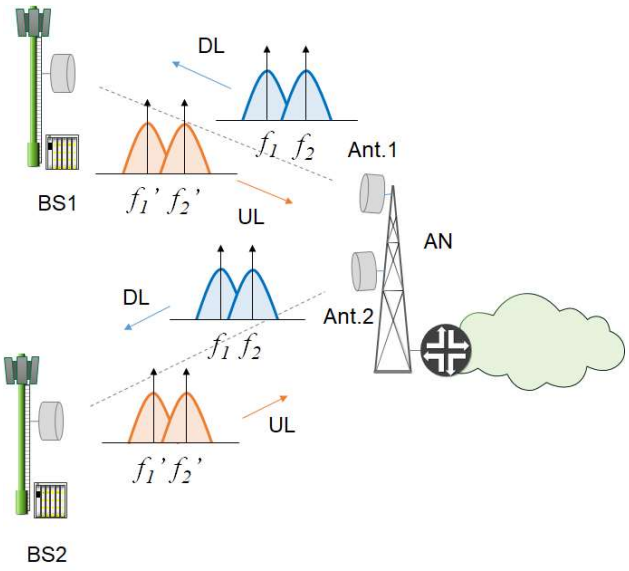

Figure 5-3. Full frequency reuse due to enough antenna separation

A problem can occur when there is an asymmetrical fading situation. For instance, it could happen that link 1 undergoes a fading due to weather condition while at the same time link 2 has excellent conditions. The problem arising is that the interference generated from *BS2* on link 1 is now relatively stronger due to a decreasing received signal from *BS1* at Ant.1.

Network programmability can provide a good solution for these situations. A spectrum allocation system that sits at a central SDN controller can dynamically allocate channels to the links based on interference status with a priority based paradigm. Figure 5-4 schematically represents this situation.

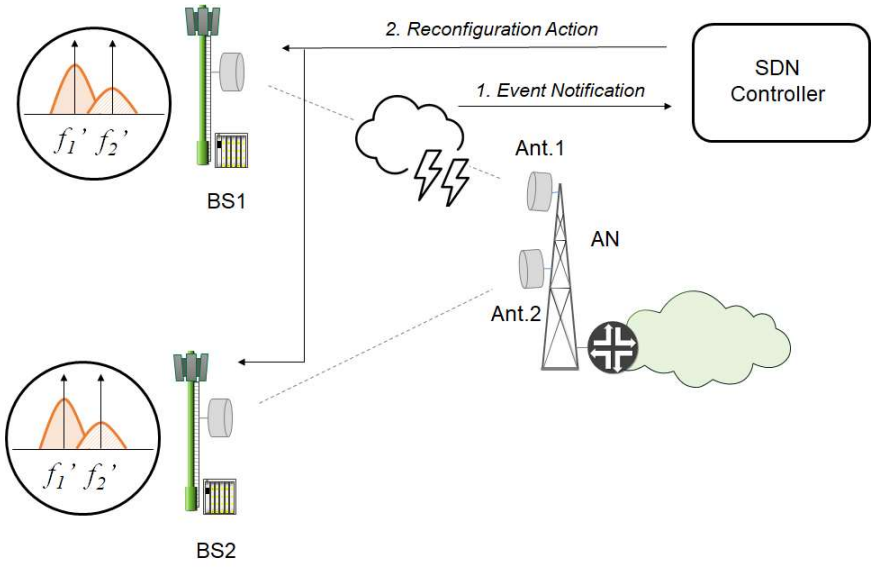

Figure 5-4. Dynamic spectrum allocation

As soon as the degradation due to weather condition is detected by Ant.1 an event is announced at the SDN controller. Then the controller (after running some spectrum allocation algorithms, e.g., implemented by an application on top of the controller) can configure the power transmitted by BS2 on channel *f2'*, for instance reducing it by some factor, allowing the channel *f2'* of the link with BS1 to maintain good conditions.

Such a spectrum allocation system can be aware of the reception signal quality for all the links in a region, and it can detect interferences if they occur. When the spectrum allocation system detects signal interference in a link on its priority channel, the system automatically can request the other links that previously got permission to use this channel as lower priority link to cease using it. This way the link which has the priority over this channel can use it without any degradation. This can be managed assuming that each link gets the right to use a channel marked as high priority, allocating a second channel with a lower priority.

Also, on demand schemas can be supported, providing the system with rich flexible capabilities to serve dynamically traffic needs.

5.1.1.3 Power management

The programmability enabled by SDN can be also used for automating the consumption of power in the radio links, pursuing an energy efficient backhaul operation.

When deploying radio links, they are configured in a way that could provide the maximum bandwidth possible in the worst conditions. Such configuration is maintained along the time, which implies to consume the same power level even in situations of no traffic in the aggregated radio base station (e.g., during the night).

The power management can be performed by adaptively configuring in time WTN equipment to consume just the required power according to the existing climatic conditions (e.g. reducing transmit power in good situation) or workload (e.g. using only the necessary radio resources in N+N configurations) along the time.

5.1.1.4 Multi-layer coordination

Multi-layer coordination implies the coordinated management of multiple technologies, e.g., packet-switched and wireless transport. That coordination can be exploited in a dynamic manner in a software driven control environment.

This multi-layer approach can leverage on the separation between logical or service connectivity from physical or transport connectivity. At logical level the constraints are expressed in terms of capacity, utilization, latency, availability, packet delay variation (PDV), packet loss (PL) and general cost functions (user preferences). In contrast, at physical level the constraints are related to link power consumption, spectrum frequency, modulation, receive signal level (RSL), etc. Additionally, the logical connectivity usually exceeds the scope of the mobile backhaul, extending the connection from access to the network core, and results agnostic about the conditions experienced by the underlying network. An effective coordination between all the network segments traversed can permit a joint optimization of the resources, taking profit of re-configurability and adaptability on the underlying transport capacity.

5.1.2 Programmable control of backhaul networks: microwave nodes

The experimental results are related to a first experience that was performed framed as part of the Wireless Transport Project within the Open Networking Foundation (ONF), as first attempt to define a (unified and standard) control plane for microwave systems. This section reports on some specific experiment from the ones carried out in that occasion.

For doing so, a number of OpenFlow extensions defined by ONF were validated in a multi-vendor interworking, multi-layer control, network-wide coordination, as described next.

5.1.2.1 Experiment description

The experiment described ran during the first proof of concept of applicability of SDN to Wireless Transport Networks. The experiment took place at 5TONIC lab in Madrid.

Radio link capacity can suffer variations impacted, for instance, by bad weather conditions, as illustrated in Figure 5-5. When this happens, the throughput is reduced when compared to normal conditions. Such alteration on the capacity can produce congestion in the communication end-to-end.

By coordinating actions between the routers and the microwave equipment it is possible to adjust the traffic flows offered end-to-end and avoid the saturation of capacity limited links.

OpenFlow provides a common standard interface between a centralized SDN controller and forwarding nodes. Being a vendor-neutral protocol, it is possible to control microwave nodes from different manufacturers. In order to accomplish the expected coordination scenario, traffic policing was considered for reducing the throughput of the flows on a router connecting the microwave node to avoid the saturation of the wireless link capacity.

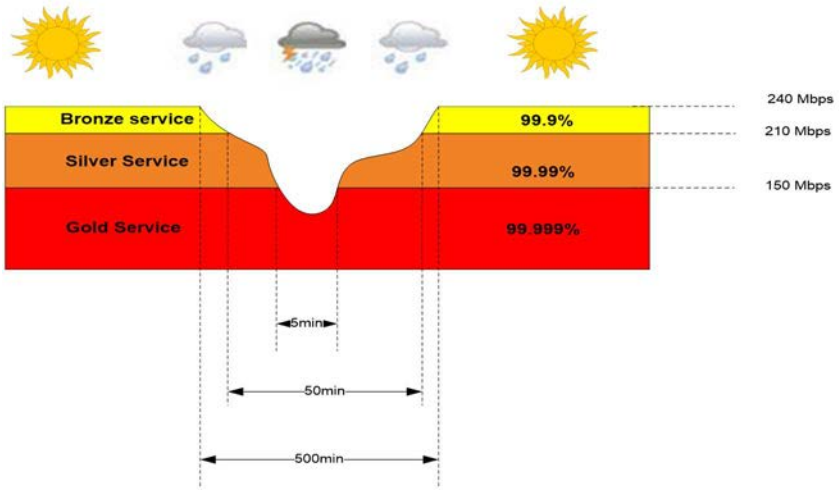

Figure 5-5. Example of capacity availability on the radio link due to weather conditions, reflecting the impact on link capacity and availability

5.1.2.2 Experiment configuration

The test topology was as shown in Figure 5-6, involving different vendors of microwave nodes (Ceragon, Ericsson, Huawei, NEC and SIAE Microelettronica) interacting with a router manufacturer (Coriant, now Infinera).

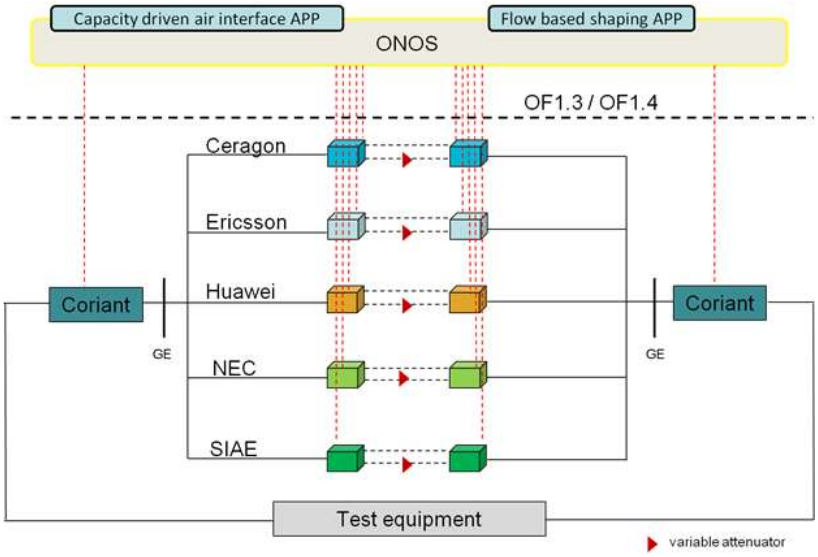

Figure 5-6. Test topology

The SDN controller used was ONOS. Figure 5-7 shows a screenshot of the ONOS controller with access to all the devices of the setup.

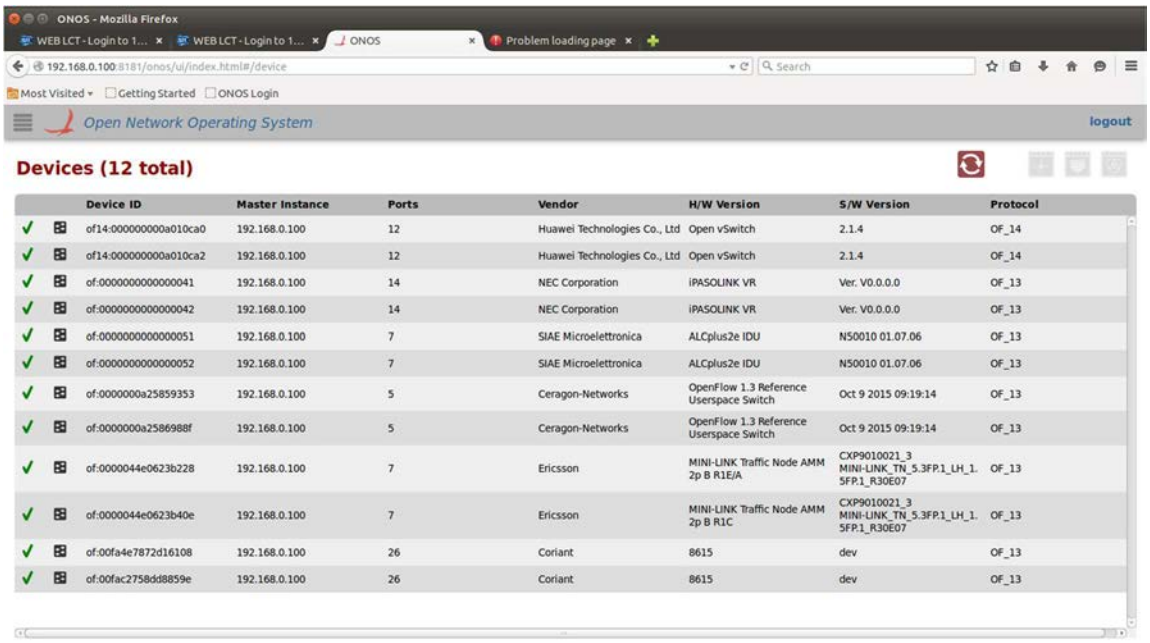

Figure 5-7. ONOS screenshot with the devices under control

Each wireless transport node has at least four OpenFlow ports (one wireless transport layer-1 LAG port, two underlying wireless transport physical ports and one Ethernet port). The SDN controller sends Flow Modification messages to install flow entries in each node. Complementing the experimental setup, a Traffic Generator generates data traffic injected in the routers, which forward it according to the flow entries configured from the controller. The low level setup is as described in Figure 5-8.

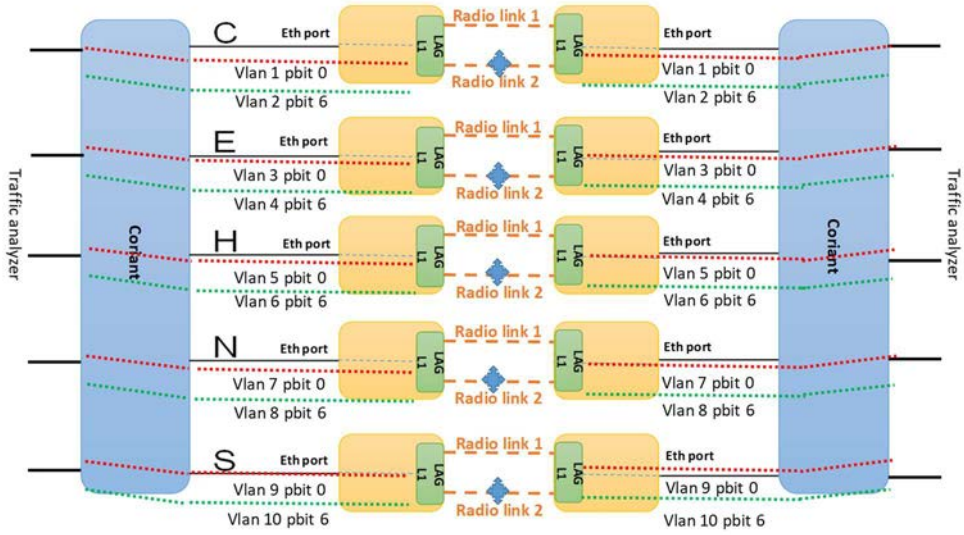

Figure 5-8. Low level setup

5.1.2.3 Experiment results

In order to emulate the impact of bad climate conditions on the link affecting the transmission channel, a radio signal attenuator is used, resulting in a decrease of the radio link capacity. Figure 5-9 shows the effect for one of the links.

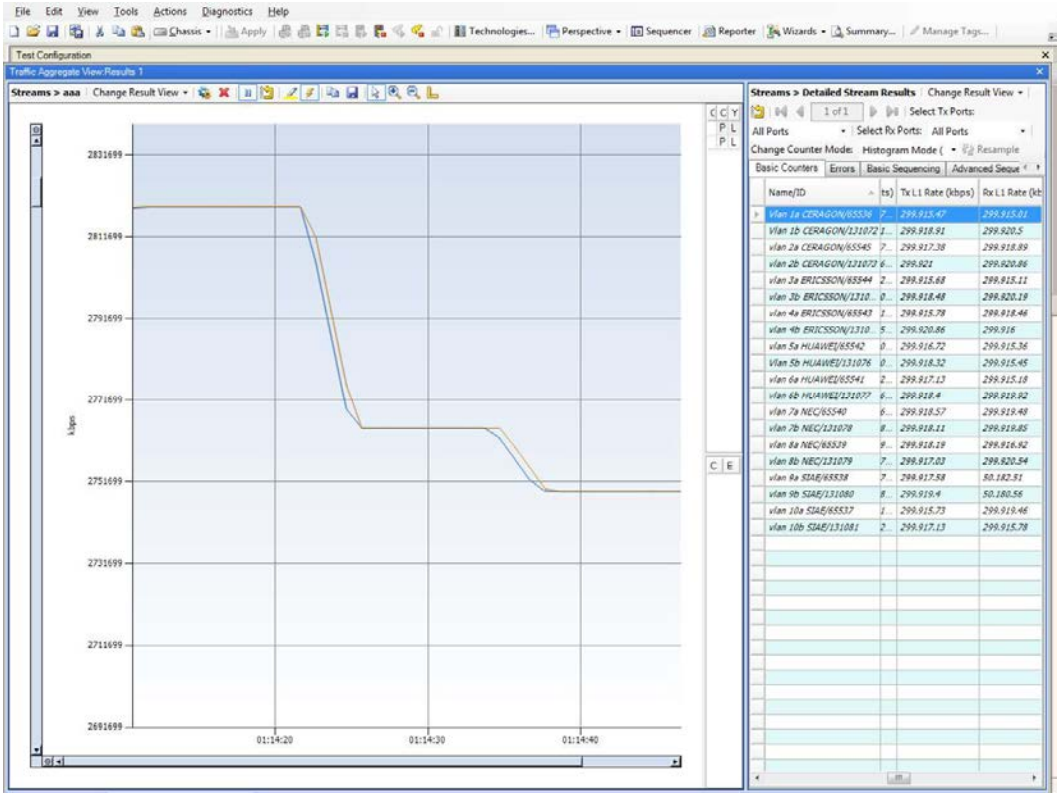

Figure 5-9. Capacity going below threshold after applying attenuation in the radio signal

The controller periodically collects port descriptions from the microwave nodes through new OpenFlow messages (i.e., OFPWTIPPT_TX_CURRENT_CAPACITY) reporting to the controller the transmission capacity of the wireless (radio) port each time.

When the actual capacity of the radio link goes below a pre-established threshold, the ONOS controller sends a Meter Modification message to the routers. The flow rule is updated with the meter ID. A Flow Modification message is further sent to the routers to apply the updated flow rule, policing the traffic flows to avoid the saturation of the wireless links by decreasing the flow throughput. Figure 5-10 presents how the system adapts to the available capacity. As the capacity limitation on the radio link occurs, the microwave node starts dropping traffic up to the point in which the router is properly configured with new rules, then preventing the traffic to be shaped in advance.

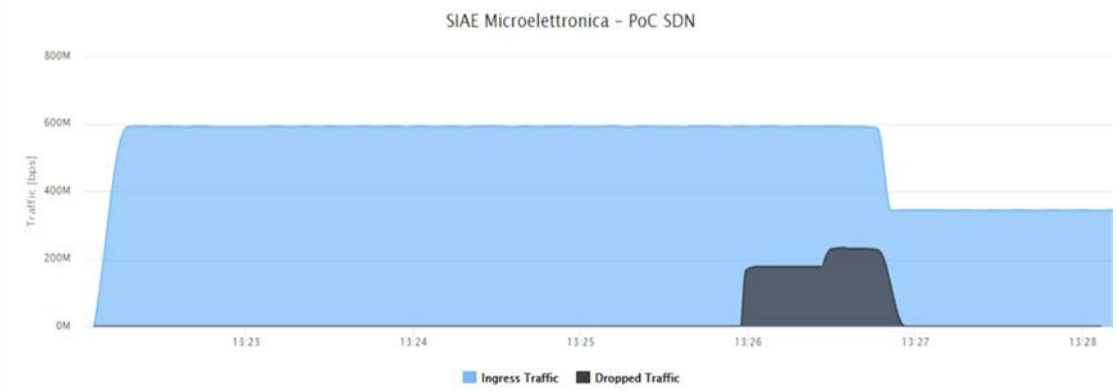

Figure 5-10. Screen shot showing the automated shaping of the traffic flow

Acting on the radio signal attenuator, when the capacity of the radio link returns to the nominal capacity exceeding the threshold, as shown in Figure 5-11, the SDN controller sends a Flow Modification message to the routers for removing the previous flow rule limiting the throughput on them.

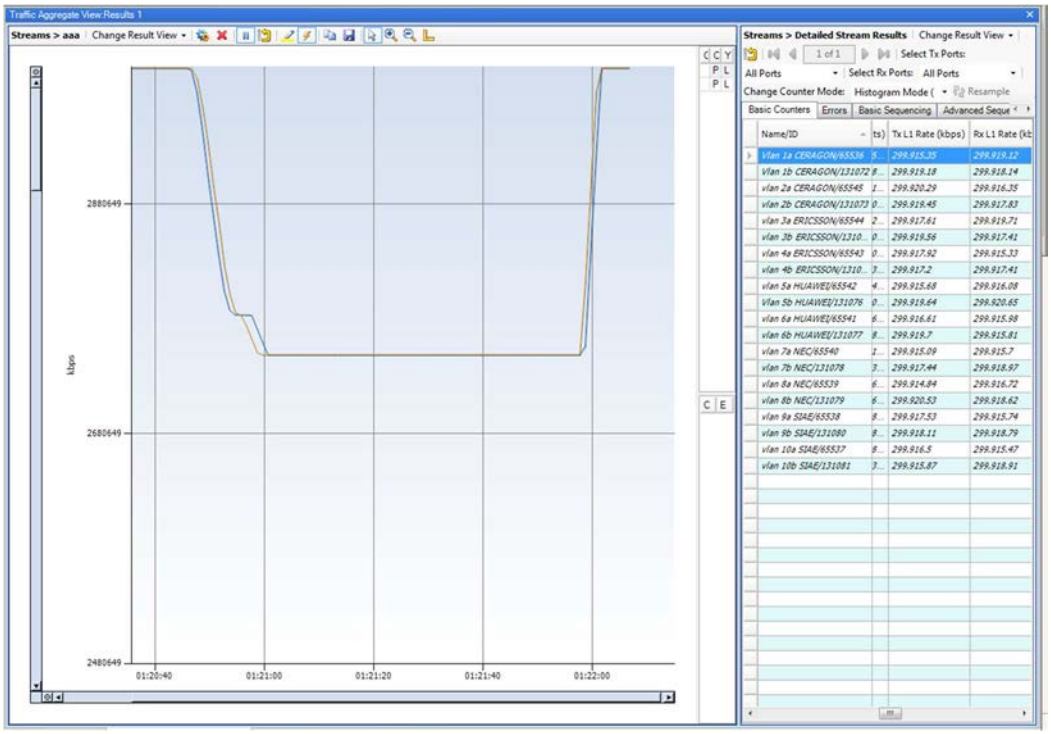

Figure 5-11. Capacity recovery after removing attenuation in the radio signal

5.1.3 Summary of the contribution

This sub-section has presented the applicability of SDN control to Wireless Transport Networks justified by advanced use cases that can benefit network operations (published as a journal paper [A28]). The contribution also reports as result the combined multi-layer interaction between WTN and routing devices for performing traffic shaping, as exemplary case of the advantage of introducing programmability control of WTNs, obtained in the first proof of concept on the topic (generating a number of contributions in ONF [O111][O112][O113] and IETF with an RFC [A30]) .

5.2 Isolation in transport slicing

Slicing at transport network level will leverage on SDN for the establishment, maintenance and release of connectivity capabilities among network functions, all in all composing and end-to-end network slice.

In this direction, from the perspective of a telecom network operator, it is essential to define standard based mechanisms that could allow interoperability, facilitating an easier integration and vendor-agnostic approach in the realization of the transport slices.

New technical challenges accompany the development of the transport slicing concept. For instance, the idea of isolation imposes the need of ensuring the reservation of some resources to specific slices. When ported to transport networks, especially to those providing statistical

multiplexing gains, that reservation cannot be always guaranteed, since there are technological dependencies for achieving a strict dedication of resources.

5.2.1 Carrier SDN architecture enabling transport slicing

Network operations are heavily influenced by the control and management capabilities available. Key aspects of operations such as network provisioning or troubleshooting can then benefit from advanced tools and mechanisms. The paradigm of SDN becomes the central point of network evolution, covering the majority of existing gaps.

This journey has been initiated by a number of major vendors in the world [O100][O101]. The use of the SDN principles is aimed to simplify the network operation and allowing a fast reaction and adaptability to network changes motivated by traffic variability or simply because of service configuration. Besides network element control functions, SDN is being considered also as a mean to provide support for management functions, such as collection of real-time information that could permit the automatic configuration creation and activation in network elements, as triggered by the OSSs.

Vendor-agnostic operation is also fundamental to have a common way of controlling and managing the transport infrastructures. SDN, through the deployment of fully standard South-Bound Interfaces (SBI) towards the network elements performing forwarding, and North-Bound Interfaces (NBI) to OSSs and other management elements, and it is capability to abstract network resources to upper layers, represents an important enabler. It seems convenient to have a clear definition of services and network, so the operator can deploy services across multiple technologies and even administrative domains though a common network services API.

iFUSION is a reference model architecture (proposed by Telefónica [A30]) permitting the separation of concerns for both network and service layers (in line with CLAS architectural view in Section 3.1 and [A13]). The proposed architecture can be shown in terms of components and relationship among them, as depicted in Figure 5-12.

Firstly, the architecture considers differentiated SDN controllers per technology, referred as Domain Controllers that cover respectively the IP, Optical and Microwave (i.e., WTN) transport networks. Each particular Domain controller unifies the device configuration interface for a particular technology leveraging on standard-based device models, thus facilitating a vendor-agnostic network configuration as well as monitoring and resource discovery. Secondly, there is an overarching control element named Software Defined Transport Network (SDTN) Controller which is responsible to orchestrate those Domain controllers providing an E2E transport network vision, in a comprehensive manner. Similar approach is being promoted and supported by industrial initiatives like in [O102].

The key elements of the SDN iFUSION architecture are the following:

- SDN Domain: It is a set of network elements under the supervision of the same SDN Controller. There are several possible levels in the decoupling of control and data planes.
- SDN Domain Controller: This controller is in charge of a set of network elements. It has standard southbound interfaces that depend on the technology, but not in the equipment vendor, to communicate with the network elements. It also has a northbound interface to communicate with the SDN Orchestrator and the OSS.

- Software Defined Transport Network (SDTN) Controller: In case several SDN Domains are needed, the SDN Transport Controller is in charge of providing services through several domains by stitching the different domains/layers/technologies.
- Southbound Interface: It is the interface, based on a standard, between the SDN Domain Controller and the Network Element. Both the communication protocol and the data model used should be standard.
- Northbound Interface: It is the interface, based on a standard, between the SDN Domain Controller and the OSSs and SDN Transport.
- Service SDN controller: An additional SDN layer that takes into account the programmability of services that might be needed. This controller aligns with the CLAS proposition described in Section 3.1.

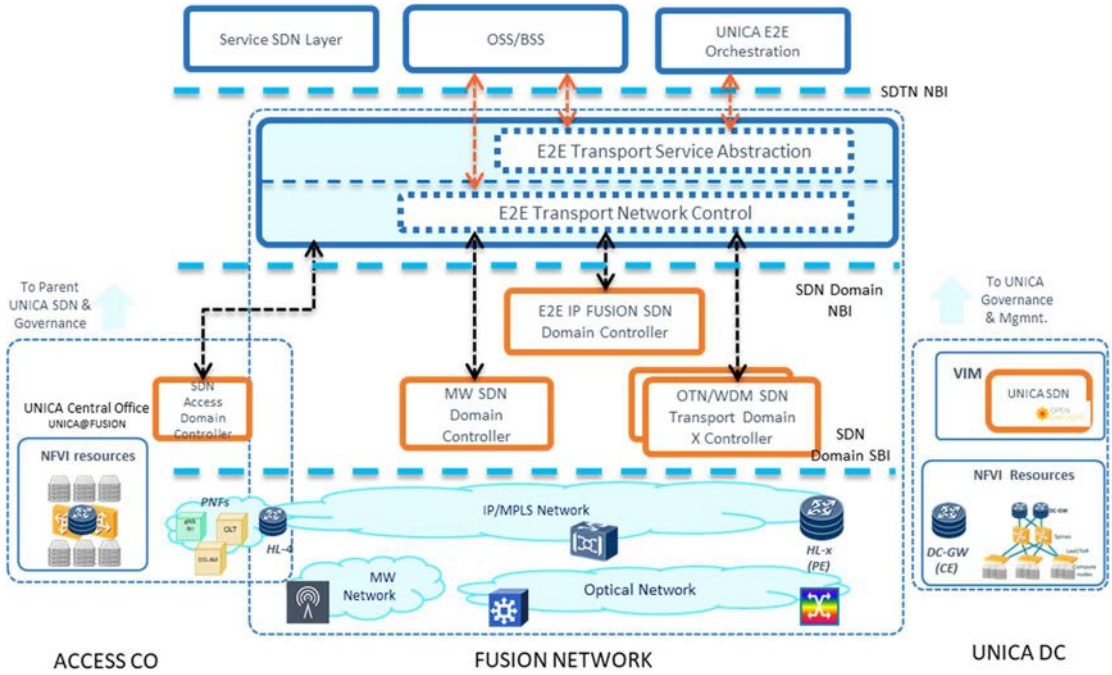

Figure 5-12. Telefonica's iFUSION architecture [A30]

The iFUSION architecture is designed as a hierarchical model where each network segment is controlled by a dedicated SDN Domain controller. The transport network, due to its wide scope and complexity, is divided in three main technology domains: IP, Microwave (MW) for wireless transport, and Optical for transmission.

The Software Defined Transport Network (SDTN) Controller is responsible to orchestrate the respective SDN Domain controllers within the transport segment (IP, Optical and MW) through the Domain Controllers' NBI, providing an end-to-end transport network vision. The SDTN Controller aggregates demands from the management and services layer exposing a unified NBI which should provide resource configuration abstraction and technology agnostic service definition. The SDTN entails two main building blocks: *(i)* end-to-end Transport Network Control and *(ii)* end-to-end Transport Service Abstraction.

The E2E Transport Network Control is the functional component inside the SDTN Controller, responsible of provide E2E network control by coordinating the different technologies through the corresponding SDN Domain controller. SDTN Controller is aimed

to provide per-layer E2E resources visualization (i.e., per-layer topology composition), and stateful control of provisioned network services.

The E2E Transport Service Abstraction component is defined as the element responsible of providing the service interface towards client applications and its translation to network service configurations. It is assumed that all service models exposed through this component are network and technology agnostic. It performs the mapping or translation between business service and network level models.

The SDTN Controller has full visibility of the overall transport segments in the network. It has the ability to expose an abstracted topology view of the entire network to the different consumers of the network capabilities through its North-Bound APIs. There could be multiple consumers of network capabilities, such as the OSS/BSS, different other service orchestrators (or hierarchical SDN controllers), the NFV orchestrator for virtualized services, etc. The level of abstraction could be different according to the specific needs of each of those consumers.

The SDN Domain controllers, on the other hand, are in charge of all the devices in the domain. Each SDN Domain controller unifies the device configuration interface and provides vendor-agnostic network configuration, monitoring and resource discovery. Besides, the Domain Controller exposes high-level network services abstraction to OSS and BSS layers through its North Bound Interface (NBI). Therefore, the abstraction of device specific configuration from network service definition is one of the main features that the SDN controller implements. Moreover, the SDN Domain Controllers entail the function of Path Computation Element to manage and optimize traffic engineering in the domain.

SDN technology is also the base of internal connectivity inside the virtual infrastructure domains. Some mechanisms to coordinate with the Transport SDN will be required. The interaction will be done at a horizontal level between the WIM (Wide area Network Infrastructure Manager) [O9] and the Transport Controller.

The SDTN Controller keeps visibility of all the transport network segments. It exposes an abstracted topology view of the network resources and the available set of network services to different clients through its North-Bound APIs.

One of the main drivers of deploying an SDTN controller is service automation. SDTN enables it, progressively, facilitating that services and network configurations carried out manually today become automated and available through this abstraction layer. The level of abstraction can be different according to the needs of the northbound client (e.g. OSS, service orchestrators/SDN controllers, NFV orchestrator, etc.).

5.2.2 Transport slice controller

Transport network plays a key role in the slice construction, as it is in charge of providing data plane connectivity across the entire slice. In [A31] a slice at transport level is defined as *"logical network topology connecting a number of endpoints using a set of shared or dedicated network resources that are used to satisfy specific Service Level Objectives (SLOs)"*. This connectivity does not only cover the backhaul segment (i.e. access-to-core), but also may include front-/middle-haul segments (i.e. intra-access connectivity, in case of radio access node disaggregation) and segments beyond the core network (i.e. for connectivity towards operator-hosted value-added services or the Internet). This wide variety

of network segments, each conveying traffic using multiple technologies (e.g. IP, optical, microwave) on top of different infrastructure topologies (e.g. hub-spoke, ring, mesh) requires considering a multiplicity of facets when developing slicing on the transport network domain.

However, it is not yet clear how a carrier network operator would be able to satisfy the demand of creating network slices, specifically for the transport network, in a common and standardized way. This point is extremely important since different consumers (i.e., vertical customers but also operator's business units) will demand slice capabilities from the transport network. Having a standardized way of providing slices would simplify and unify all the slice request process coming from different slice consumers as well as facilitate the integration and interaction among distinct vendor implementations. Furthermore, the transport network needs to satisfy such different set of demands ensuring fair and equal treatment of service requests as well as a common understanding of the particular needs from each request.

At this point there is not yet a standard approach in the industry for the definition of a common Transport Slice Controller component able to facilitate the provision of connectivity services as slices. This component will be responsible of handling the slice requests and associated procedures (e.g., transport slice lifecycle).

The transport network is in charge of enabling connectivity between the end users and the service functions composing a given E2E service as requested by a customer. According to the characteristics of the supported service, such connectivity can have distinct properties or characteristics, and pursue different SLOs. The transport network leverages on multiplexing gains to convey traffic from different E2E services. These services are usually segregated by applying different mechanisms, with the ambition of providing the notion of virtually dedicated network for the customer. The mechanisms to be applied vary depending on the network layer where they are performed, typically based on distinct protocol encapsulation mechanisms together with a specific engineering of the paths to traverse. Transport slicing is not so much different in that sense, but nominally incorporates additional relevant capabilities with respect to simple traffic segregation.

First, it inherits from the general concept of network slicing the idea of isolation, which basically means that slices from different customers gracefully coexist without interfering each other in the sense that whatever misbehavior or unforeseen demand from one customer happens, it could not affect the communication service received by any other customer supported by the same physical transport infrastructure.

Second, the customer could require capabilities of control and management of the slice as if it was actually a dedicated network, dealing to the need of implementing mechanisms allowing that, again, without impacting the services (i.e., slices) of other customers in the network.

Third, because of the progressive trends towards a full network softwarization, the dynamicity in the creation of those services could be higher if compared with the existing one, allowing a great flexibility in the provisioning of services with the need of handling from ephemeral to long-lasting services, of very different characteristics in terms of requirements, on top of the common transport substrate.

In order to support the transport slicing capabilities, it seems convenient to define a new component, the Transport Slice Controller (TSC), in charge of control the provision of the transport slices and managing their lifecycle, retaining on that component the awareness of slicing at the transport level. In order to be effective, especially in relation with the expectation of isolation, additional capabilities will be required at both control and data plane levels, in order to preserve the strict allocation of resources to different transport slices.

In line with the definition in [O103] the Transport Slice Controller supports both a Northbound and a Southbound Interface (NBI and SBI respectively). Through the NBI, the different customers proceed to request transport slices adapted to the specific needs of each E2E service. Thus, the Transport Slice Controller acts as single entry point to the transport network for these requests, resolving potential conflicts and/or incompatible requests. On the other side, once the requests are processed, the Transport Slice Controller instructs a number of network controllers to proceed with the proper actions associated to the previous requests. Those requests correspond to both provisioning actions and transport slice lifecycle management. Summarizing, the NBI supports the transport slice description, while the SBI enforce the transport slice realization. Figure 5-13 schematically shows the role of the Transport Slice Controller.

It is worth noting that the NBI can be expected to be technology agnostic, where basically only transport slice characteristics and needs are expressed, without notion of the technology used to realize it. The mapping to specific transport technologies is performed by the Transport Slice Controller, which according to the technology mapping decision interacts later on with specific technology controllers for the overall slice provision. In this manner, the NBI can be common to the different kind of customers (probably not all of them consuming the same capabilities of the NBI [A32]), while there will be multiple incarnations of SBI depending on the specific transport technology.

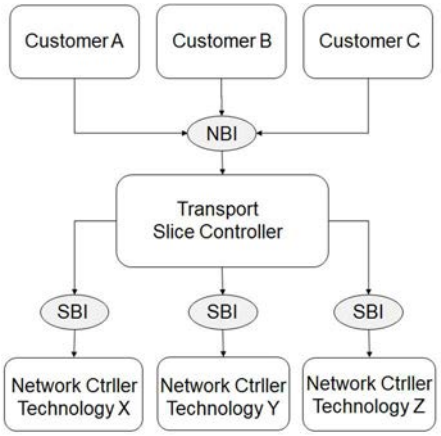

Figure 5-13. Transport Slice Controller concept

In this way it is possible to anticipate [A33] a potential modular structuring of the TSC capabilities as follows:

- A module, defined as *Mapper*, responsible of collecting the customer-facing view of the transport slice for further processing the transport slice request. Thus, the Mapper

module would integrate the customer-facing view on the provider view for triggering slice configuration, control and management actions.

- An additional module, here defined as *Realizer* module, in charge of coordinating different actions on a number of network controllers (below the TSC) for effectively creating the transport slice for the original customer request. This Realizer would manage the workflows for the transport slice provision, as well as for its lifecycle.

Figure 5-14 presents the proposed structure elaborating on top of previous Figure 5-13, including the possible existence of a complementary Intent-based module (e.g., as in [A34]). The figure shows the possibility for the customer (i.e., an upper management system) to interact with the TSC by the intent-based module to directly through the NBI, providing sufficient flexibility.

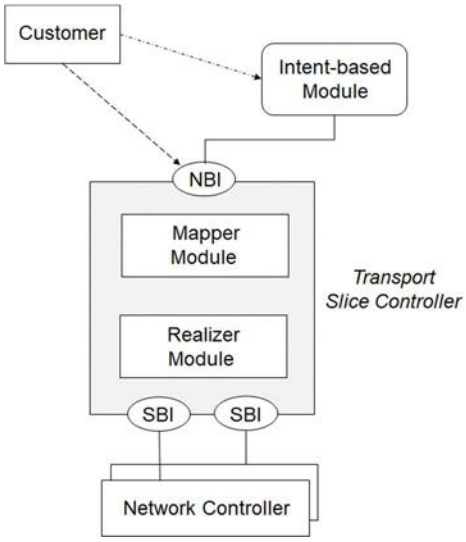

Figure 5-14. Proposed modular structure for TSC

Transport slicing should be integrated in a smooth manner with the network programmable capabilities in carrier networks. One potential example is iFUSION [A30], described in the previous section. Figure 5-15 shows the proposed integration on iFUSION architecture, as well as the positioning of TSC as part of it.

The SDTN Controller has full visibility of the overall transport segments in the network. It has the ability to expose an abstracted topology view of the entire network to the different consumers of the network capabilities through its North-Bound APIs. There could be multiple consumers of network capabilities, such as the OSS/BSS, different other service orchestrators (or hierarchical SDN controllers), the NFV orchestrator for virtualized services, etc. The level of abstraction could be different according to the specific needs of each of those consumers.

In this respect, some of the considerations for providing transport slices are already in place in architectures as iFUSION, however some pieces are yet missing. In essence, it is necessary to add a level of awareness with respect to the slicing concept, as well as to support the logic for the management of the slice lifecycle (i.e., the mapping function). On the contrary, the capability of orchestrating diverse network controllers is already in place (i.e., the realization function), then facilitating the integration.

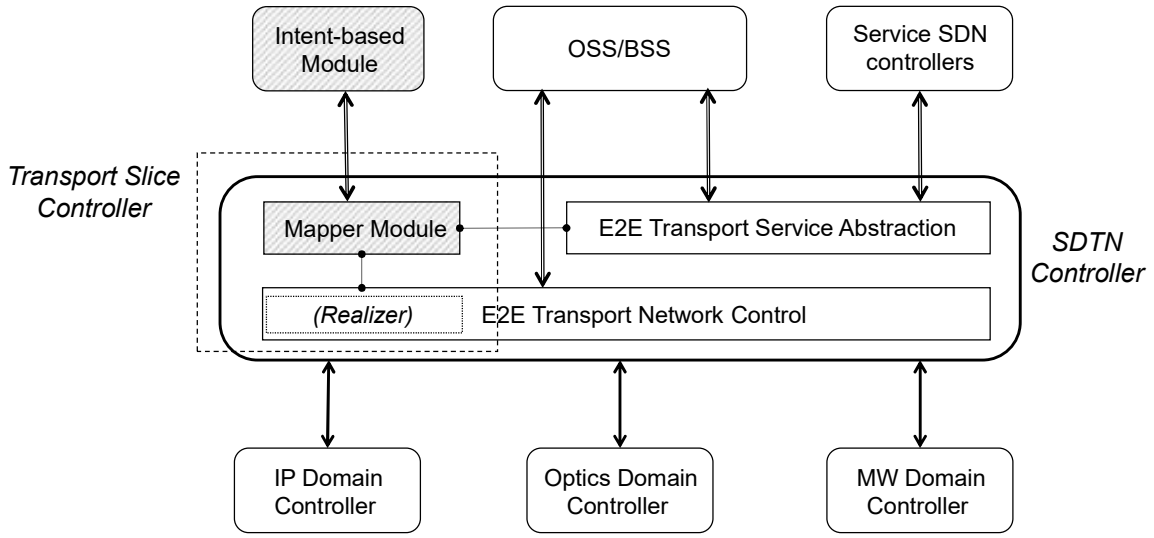

Figure 5-15. Intregration of TSC in iFUSION architecture

Regarding the slicing concept, aspects like the handling of isolation concerns between slices, or the association of the transport slice to E2E slices are missing from the existing definition. Thus, simply a new building block seems to be needed for adding these additional capabilities necessary for support transport slicing on top of the existing architecture, in a smooth and straightforward manner.

This new block is represented in Figure 5-15 as the Mapper Module, within the SDTN. This module offers the TSC NBI, complemented potentially by an Intent-based Module (following for instance the approach in [A34]), in case the transport slice customer produces the requests as intents, instead of directly using the TSC NBI.

The Mapper module maintains the data model for the provider view of the transport slice, as e.g. [A35]. It can consume capabilities offered by both the SDTN E2E Transport Service Abstraction and the Network Control building blocks. The former could provide non-specialized abstract service constructs for building transport slices not needing specific requirements. The latter is used essentially for triggering the configuration of the transport slice in the underlying network substrate, especially in the case of specialized transport slices with very specific requirements that cannot be satisfied by generic service models.

The Mapper maintains the necessary associations between the idea of slice and the realization of it in terms of connectivity constructs forming the transport slice. The Mapper also supports all the necessary procedures for managing the transport slice lifecycle. In this manner, the introduction of slicing capabilities in the transport network results incremental on top of the SDN capabilities being currently deployed.

On the other hand, as reflected in the figure, the E2E Transport Network Control component can assume the role of Realizer within the TSC architecture, leveraging on the functional description of that component provided in Section 5.2.1. Finally, the Intent-Based module, if present, directly translates the slice intent to a set of requirements and procedures towards the Mapper.

5.2.3 Handling of isolation at transport level

As established by 3GPP [O42], a Network Slice Instance (NSI) for a given customer may be logically and/or physically isolated from another NSIs, in a full or partial manner.

In order to homogenize the way in which slices could be specified at the time of being requested, the GSMA is developing a universal and generic blueprint that can be used by any vertical customer to order the deployment of a NSI based on a specific set of service requirements. Such a blueprint is actually a network slice descriptor called Generic network Slice Template (GST) [O105]. The GST contains multiple parameters named *attributes* (up to 36 at the time of writing) that can be used to characterize a network slice. A particular template filled with values assigned to those attributes generates what is called a specific NEtwork Slice Type (NEST). In essence, the NEST is a filled-in version of the GST that allows the operator and the customers to agree on the Service Level Objectives (SLOs) and other characteristics of the slice. The customer uses the NEST to request the provisioning of an NSI able to satisfy a particular set of service requirements.

One of the attributes defined in the GST allows specifying the level of isolation required for a network slice. The specification of this attribute in the NEST means selecting one of the values shown in Table 5-1.

Table 5-1. Parametrization of isolation in GST.

Value	Level of Isolation	Example sub-values
0	No isolation	
1	Physical isolation	0 - Process and threads Isolation 1- Physical memory Isolation 2 - Physical network isolation
2	Logical isolation	0 - Virtual resource isolation 1 - Network Function isolation 2 - Tenant/Service isolation

The GST characterizes the requirement of isolation in base of the physical vs logical distinction. For the physical isolation it is expected a total separation, including hardware, site or even rack separation. In the case of logical separation, it is considered some sharing of physical resources among slices. This characterization, apart from resulting vague, is quite oriented to the overlay compute domain. To take decisions on how to realize the isolation in the TN domain, a more grounded and networking-facing characterization is required.

Next, a number of options are surveyed to assist on the required characterization, identifying slice isolation definitions applicable to the transport part.

5.2.3.1 Segregation vs isolation

Traditionally, network services per customer in transport network have been supported by the concept of Virtual Private Networks (VPN). In a nutshell, VPNs allows operators to setup dedicated tunnels per customer, and even particular service per customer, being these tunnels differentiated through distinct kind of encapsulation schemes. VPNs can have specific inbuilt capabilities (e.g., traffic engineered paths, encryption, etc.) and include different

schemes of connectivity (e.g. virtual leased line, hub and spoke, etc.), either as layer-2 or layer-3. This results in the provisioning of L2VPN and L3VPNs, respectively.

Nowadays, VPNs are typically deployed as a virtual overlay solution on top of packet-switched networks, simultaneously supporting some other variety of services. This solution allows VPNs to provide customer traffic segregation while making an efficient resource usage, taking advantage of the multiplexing gains of the packet-switched networks. However, this statistical multiplexing behavior has a noticeable counterpart: some stringent requirements (e.g., throughput) cannot be permanently guaranteed. This prevents the operator from providing predictive service assurance, with the possibility of certain events provoking that customer expectations are not met, dealing to claims or even penalties. For instance, situations of congestion, such as the ones generated by flash crowd events, failures in parts of the network, or even exceptional situations such as the traffic pattern increase during the generalized lockdown motivated by the COVID-19 pandemic, can eventually impact customer's SLO. Once more, the non-cost-efficient overprovisioning emerges as the main manner of guaranteeing service levels by the network operators in these situations.

While overlay solutions facilitate customer service segregation, they do not allow actual service isolation. Service isolation here represents the way of dedicating resources to customers, guaranteeing that those resources allocated to a given customer are respected and not used for others in situations as the ones described before. It is worth noting that VPNs are created at the edge nodes in the network, where certain level of service separation awareness can be retained. However, core backbone nodes have no such service awareness, thus not being able to contribute to the distinction of services with different needs from different customers. According to the above reasoning, it is therefore needed something more advanced than a VPN to provide a complete service isolation.

Firstly, service isolation can be related to the concept of network partitioning, in the sense that a number of resources are allocated to a specific customer and granted to the customer by the network provider during the slice lifetime. These resources can be physical and/or or virtual ones. On the one hand, the physical infrastructure partitioning (e.g. dedicated fiber access connection) is not flexible at all, and would be limited, commonly, to cases where provider diversity is required by the customer. Physical isolation however could also apply to situations where path disjointness is required, then implying the need of deploying disjoint paths ensuring that those paths traverse physically separated infrastructures internally to the network operator (e.g., site, backbone fiber, etc.). On the other hand, logical partitioning is also possible by allocating different instances of logical network transport constructs, e.g. a router port. The great potential of this approach can be exploited by the general introduction nowadays of softwarization and network programmability capabilities in operational carrier networks [A30].

A second important aspect of isolation is the possibility of enabling certain levels of control and management, by the vertical customer itself, over the allocated resources. This cannot happen when the resources are shared among customers, since the control and management actions of one customer could collide with the similar actions decided by another one. The only manner of avoiding conflicts is the partition and separated allocation of resources per customer. Once partitioned, some level of control could be allowed, e.g. for steering the traffic through different paths. The customer can perceive the slices as controllable and manageable, while the network operator retains the full control and management of the resources, resolving conflicts in a transparent manner for the customers. In that direction

propositions such as WIM-on-demand [A11] could assist in order to facilitate advanced transport network capabilities for customers demanding specialized services.

In summary, the concept of isolation provides a step forward with respect current levels of customer service segregation.

5.2.3.2 Isolation approaches in the transport network

The transport network is rich on technology options and functional capabilities. Not all the technologies offer the same level of isolation possibilities, so there are a number of options to consider at the time of providing isolation. This implies that certain logic is required at the Transport Slice Controller for decision of the more convenient option for each transport slice request.

The allocation of resources for achieving isolation does not imply that the allocated resources are permanently assigned (during slice lifetime) to a specific customer. That is, the concrete resource items can change along the time while the overall amount and type of resources is provided. This flexibility is maintained by the network programmability capabilities present in the network.

This section overviews different isolation approaches assumable for the transport network. These approaches can serve as basis for the definition and exemplification of isolation feasibility indicators later on.

5.2.3.2.1 Control plane isolation

Not all the customers require control capabilities for the allocated slice. This is probably limited to some advanced and specialized customers needing a deeper control of the transport slice resources.

When enabling the possibility of controlling and managing the allocated resources for some customers, control plane isolation permits each of those customers to act on the assigned resources.

Network softwarization can be applied at both network provider and customer levels. In the case of the network provider, it is an enabler for the support of transport slicing, while at customer level it is an advanced feature allowing to perceive the slice as a dedicated fully controllable network (as described in Section 3.1.7). Isolation in this context refers to the ability of separating control and management concerns from different slice customers.

A first approach for separation requires the direct implementation on the network provider control plane elements of different virtual spaces per slice customer, allowing each customer only access to the proper virtual space. A second, more advanced, approach includes the usage by the customer of a dedicated independent control element, interacting directly with the network provider control element. Such control plane element per customer would be limited in functionality since it is restricted to the control and management of the virtual resources assigned. Finally, that control element could be either owned by the customer itself, or facilitated by network provider (as in the aforementioned WIM-on-demand case). The latter simplifies the interoperability among control plane elements, since full compatibility can be managed beforehand, while the former could require some integration effort to make the different control plane elements to interoperate.

5.2.3.2.2 Topological isolation

Usually, topological diversity is used in order to avoid affectation in the primary and backup routes of a service on the event of failure. Thus, at the time of deploying such service, smart decisions can be taken to distinguish the resources allocated for that routes (i.e., links, nodes). This is commonly done by identifying shared risk group at resource level, then enabling to compute routes compliant with the isolation requirement.

Architectural components such as the Path Computation Element (PCE) [O106] can assist on the pertinent identification of separated routes. However, when moving to the network softwarization approach, in order to ensure full isolation some additional consideration should be required in order to ensure that virtual topologies are effectively isolated at topological level.

This same approach of topological isolation described above can be used when dealing with overlay solutions such as the conventional VPNs.

5.2.3.2.3 Resource isolation at device level

Different levels of isolation can also be provided at device level. The partition of the device can be considered at either hardware or software (leveraging on the available operating systems capabilities).

At hardware level, for instance, different ports or boards (within associated resources such as queueing ports) could be allocated for conveying the traffic of distinct slice customers.

On the other hand, at software level, it is possible to instantiate multiple logical devices acting as virtual nodes, leveraging even in several degrees of resource differentiation at hardware level, as before. This is the case, for example, of Juniper's approaches known as logical system, virtual router or node slicing [O107].

It is evident that there is always a dependency of the same hosting device that cannot be avoided, in the sense that, depending on the case, some common parts are shared among customers. For instance, in the allocation of different ports there is dependency of the supportive board, in the allocation of boards there is dependency of the chassis (e.g., switching matrix, or fans), etc. However, these dependencies do not imply that the customers are not isolated, but that customers share the same risk group at node level.

5.2.3.2.4 Resource isolation at data plane level

The isolation at data plane level intrinsically depends on the particular characteristics of each transport technology. Strict allocation of connectivity resources is only available in certain solutions. For instance, it could be possible to allocate for a specific slice a concrete lambda in Dense Wavelength Division Multiplexing (DWDM), or a number of calendar slots in Flex Ethernet [O20].

Other technologies do not allow that kind of strict resource allocation, thus some levels of contention in the usage of shared resources could be expected. This is the case of the conventional packet-switched networks. Certain mechanisms could mitigate the contention impacts, such as advanced QoS mechanisms, either in traditional [O108] or programmable [O109] networks. Depending of the particular needs of the customer, this could be sufficiently acceptable or not.

5.2.3.2.5 Summary of approaches

The approaches described before can be applied independently or being combined. The final decision on which approach is used depends on how explicit is the customer request in that respect. If not specific at all, the network provider, through the Transport Slice Controller, can take the final decision, according to predefined policies or in order to ensure some of the other SLOs requested.

It is worth noting that isolation is orthogonal to performance SLOs, but isolation can assist on meeting those performance SLOs. For example, a SLO of throughput could be enforced through isolation (e.g., as mentioned before, dedicating a lambda of enough capacity for that customer slice) or without isolation (e.g., by grooming that traffic over lambdas with a sufficient capacity to carry all the client signals).

Table 5-2**¡Error! No se encuentra el origen de la referencia.** summarizes the isolation options here described providing a qualitative comparison in terms of potential dependencies from full isolation perspective, scalability limitations, complexity issues, and generic ways for implementation of each of the approaches. Those approaches are not exclusive and could be combined for the provision of a slice with isolation requirements. Finding a simple way of comparing transport slice requests when combining more than one option can be achieved through the definition of indicators as introduced next.

Table 5-2. Summary of isolation approaches.

Isolation Approach	Description	Dependencies	Scalability	Complexity	Implementation
Control Plane	Dedicated control of the transport resources by the customer	(Logically) centralized control element of the network operator	As the control element is software based, it can be supported as long as the computing capability is increased accordingly	The central control element should separate slice contexts among customers	Instantiation of dedicated control plane entities per slice
Topology	Diversity in the routes for a given slice	Node origin where the customer is connected	Limited by the total number of nodes and links in the network	Computation of constrained paths to ensure disjointness	Need of tools for calculating diverse paths to be allocated per slice
Device	Partition of device resources (either hardware or software)	Device common parts	Limited by the hardware resources available in the device	It is required to keep awareness of resource allocation by some control element	Association of device resources per slice
Data Plane	Allocation of connectivity resources	Port(s) where the connectivity resource is associated to	Limited to the capabilities of each technology (e.g., FlexE calendar slots)	Configuration of data plane constructs (e.g. labels or extension headers) for assignment to particular slices	Association of data plane capabilities per slice

5.2.3.3 Isolation feasibility indicators

It could be expected the need of handling multiple transport slice requests along the time, each of them expressing different needs in different aspects, including isolation.

Isolation is impacted by the scarcity of resources not only at the time of provisioning the slice but also during its lifetime, when e.g. some network events could require the reconfiguration and reallocation of resources for the slices. In all those situations it is

important to understand the feasibility of guaranteeing the demanded level of isolation for respecting the original transport slice request, especially when multiple and different resources are involved on the realization of such transport slice.

Transport slices compete on resources of diverse nature, not becoming a simple task to establish comparison among them for assessing isolation feasibility along their individual lifetime. Here it is introduced a modelling methodology of how feasible is to achieve the isolation of a transport slice request in order to easily compare among slices for assisting on decisions like what slice to accommodate first if reconfiguration is needed.

With that purpose, it is proposed an isolation feasibility vector that could take into account different isolation levels or factors based on the particular characteristics of each slice. For simplicity, those factors are related to the isolation approaches described before.

For modelling the vector we follow some of the ideas described in [O110] but with a distinct approach. In that previous work slices for both Radio Access (RAN) and Core Networks (CN) are characterized by different properties or traits that in some cases correspond to parameters that have a numeric, measurable value while in others are simply Boolean variables. The numeric traits can be classified as *rising*, when having a higher value is better, *falling*, the opposite, or *Gaussian*, when the values can be described by a normal-like function. In order to account for the effect of all the parameters in the same manner, those parameters become normalized, finally fitting its value into the range [0, 1]. That work then defines a vector with the obtained values that can be further merged into a single and unique value as a kind of comparable index.

In [O110], the parameters or traits are considered as defining a certain level of isolation. Example of parameters considered in that analysis are the stream cipher key's length of a radio or fiber link, the operating system of the device, the average time between vulnerabilities assessments for a router device in the CN, etc. Here, differently to that work, when referring to some transport resources being allocated to a given slice for providing isolation, it is assumed that isolation *per-se* is guaranteed with such transport resources, so no way of asserting that some slice with more dedicated resources is more isolated that another one which demands less resources. In other words, both are equally isolated at the transport layer. Thus, when applying this conceptualization to the transport network, a different angle is taken. Instead of considering it for defining an isolation level, here the traits are used to obtain a comparable value of how feasible is to keep the requested isolation level along the time. That is, the higher the value, the more feasible is to accommodate the transport slice request with isolation (at either provisioning or reconfiguration time).

Thus, the parameters in this analysis basically form an isolation feasibility vector, that when merged, result into an isolation feasibility index. With that index, it is possible to easily compare among slice requests for taking informed decisions, such as reallocation of resources, by Transport Slice Controllers.

The following sub-sections exemplify how both the vector and the index can be generated. To that end, it is assumed that the same kind of transport slice is requested, just varying in terms of resources needed to be isolated in each case.

5.2.3.3.1 Control plane

A sophisticated customer could need to have control capabilities on the resources allocated for the transport slice (e.g., as discussed in Section 3.1.7). Such requirement can be modeled as a Boolean variable indicating the binary option of having or not having allocated a dedicated control instance.

Defining *C* as the variable for indicating the need for a dedicated control, the values for that variable are defined as *true*, in case of needing a control instance, or *false*, on the contrary.

5.2.3.3.2 Topology

Topology diversity enables different alternative paths in a network. To guarantee isolation the paths should be disjoint, not sharing any common single element such as node or link.

The transport slice request could lead to having topological isolation, which can be realized by reserving some specific paths for those transport slices in a manner that failures of misbehaviors in other slices do not affect those requesting isolation.

Being *P* the total number of alternative and equivalent paths, a transport slice could obtain different isolation guarantees depending on the number of paths $p \in P$ that can be reserved for that slice.

The higher the value of *p*, the more difficult is for a provider to ensure isolation for such a slice during its lifetime, since a larger number of disjoint paths is needed. *P* is a linearly increasing function with the number of *p*. However, from the perspective of isolation feasibility, the topological isolation has to be assimilated to a falling trait in the sense that a transport slice with lower requirements of topology isolation would be more feasible than other with higher requirements in this respect.

In consequence, the following normalization function $f_n(x)$ is proposed

$$f_n(x) = 1 - \left[\frac{r - l}{h - l}\right] \tag{5.1}$$

being *r* the requested value for the transport slice, *l* the lower possible value and *h* the highest one.

It should be noted that the minimum value for topological diversity is to have available at least 2 disjoint paths, thus the range of values for topology would be $[2, p]$.

5.2.3.3.3 Device

In the case of device partition, it can be assumed that there is a process of allocation of ports per transport slice. The partition can apply to both physical downlink (or client) and uplink (or network) ports for a true separation of services among customers, or the allocation of some transport construct such as e.g. a lambda in a DWDM device or a calendar slot in FlexE links. This view is provided from the perspective of a single device, but the same idea can be easily extended to a situation involving a set of devices.

Here it is assumed that the limiting factor is on the client ports. The device has a number of client ports of different types, but the slice is considered to require ports of the same kind and bit rate. Being *D* the total number of downlink or client ports, a transport slice could demand a given number of ports $d \in D$, allocated for that slice.

Again, the higher the value of d, the more difficult is to ensure isolation for such slice during its lifetime. At both provisioning time and slice reconfiguration events, it implies to select devices with such number of ports available. Since this trait follows a falling behavior from the perspective of isolation feasibility, the same normalization function described in Eq. (5.1) is used. Lower values facilitate the feasibility of the transport slice. It should be noted that the minimum value in the case of ports is 1, thus the range of admissible values is [1, d].

5.2.3.3.4 Data plane

There are different technology alternatives at the data plane. For illustration, it is considered a single data plane link of 100 Gbps based on FlexE. In the case of a FlexE link, the full capacity link is divided in 20 different calendar slots of 5 Gbps each. Being S the total number of calendar slots, a transport slice could demand a given number of slots $s \in S$.

If we assume that the total capacity of the transport slice is less than 100 Gbps, a vertical could request some value in the range [1, 20] of calendar slots for its traffic. This trait is also falling in the sense that the less slots are demanded, the more feasible is to keep the isolation of the requested service along the time. Because of that, Eq. (5.1) is also the normalization function for this parameter.

5.2.3.4 Isolation feasibility vector example

With the examples before, it is possible to build an isolation feasibility vector in the form $\{c_i, p_i, d_i, s_i\}$, with $c_i \in C$ the need for dedicated control element instance, $p_i \in P$ the number of disjoint paths, $d_i \in D$ the number of client ports, and $s_i \in S$ the number of calendar slots for the transport slice TS_i.

Table 5-3 summarizes the feasibility vector for two different exemplary slice requests, for comparison. The lower and higher values of the numerical traits in this example are, respectively, P = [2, 4], D = [1, 24] and S = [1, 20].

As can be observed, the obtained isolation feasibility vectors obtained are TS_1 = {*true*, 1, 0.391, 0.947} and TS_2 = {*true*, 0.5, 0.521, 0.894}.

Table 5-3. Example of isolation feasibility vectors.

Transport Slice Request	Trait	Requested value (r)	Normalized value
TS_1	c_1	*true*	--
	p_1	2	1
	d_1	15	0,391
	s_1	2	0,947
TS_2	c_2	*true*	--
	p_2	3	0,5
	d_2	12	0,521
	s_2	3	0,894

Reference [O110] also proposes the merging of the vectors in order for obtaining single values to facilitate comparison. A merged value from the isolation feasibility vector is then referred as isolation feasibility index. Such a merging could consider different weights per

parameter. For the example here, it is just considered the merging of the numerical traits assuming equal importance and contribution of all the parameters to obtain an overall isolation feasibility index. The Boolean parameters can help to group the vectors taking advantage of the binary value of the Boolean variables, later on comparing among the ones in a group through the merged numerical value.

The function $f_m(x_1, x_2, \ldots, x_n)$ used for merging the values of the numerical traits in the vector is as follows.

$$f_m(x_1, x_2, \ldots, \ x_n) = n\left(\sum_{i=1}^{n} \frac{1}{X_i}\right)^{-1} \tag{5.2}$$

Applying Eq. (5.2) to the previous vectors for TS_1 and TS_2, the isolation feasibility index TSI_i are, respectively $TSI_1 = 0{,}65$ and $TSI_2 = 0{,}595$. Since $TSI_1 > TSI_2$, this can be interpreted in the way that TS_1 has higher feasibility, in general, than TS_2. Then, both the vector and the index can be used to discriminate among transport slice requests assisting on decisions for the realization of the slice. If some specific dimension of the vector is critical for the operator, the comparison of the indexes can be complemented by the comparison on a specific dimension, then permitting to qualify the feasibility of a transport slice not only on general terms but also in a particular aspect. For instance, in the example above, despite TS_1 having overall higher feasibility, from the point of view of the number of client ports demanded, TS_2 has a better indicator. If this was a critical aspect, the Transport Slice Controller could take it as valuable input for configuration decisions (e.g., prioritization of TS_2 vs TS_1).

5.2.4 Isolation handling component

Since isolation is technology dependent, the component handling the isolation should reside in the Realizer module in the TSC, described in Section 5.2.2. This is the module with full awareness of the technologies below, then being able to match isolation requirements to the specific techniques available in the transport network.

The technologies can vary across networks, thus the isolation handling component needs to be adapted for considering the specific capabilities of the technologies in place. Then it should be customized per each particular deployment.

As described, the isolation handling component acts not only at provisioning time but also in case of network reconfiguration as could happens on failure events. The isolation index can serve as guidance for prioritizing the reconfiguration of slices.

5.2.5 Summary of the contribution

This sub-section has overviewed control plane approaches for slicing the transport networks, including the definition of a Transport Slice Controller (with a number of contributions in IETF [A31][A32][A33][A34][A35], and a paper under submission [A36]) integrated on an operational SDN framework (described in a conference paper [A30]). As part of it, a relevant characteristic of slicing such as isolation has been analyzed, defining a mechanism that can permit an easier and faster comparison of isolation requirements on transport slices, which can facilitate control actions such as re-configurability of resources in situations of physical network failures (being accepted as conference paper [A37]).

5.3 Summary and outlook

The chapter addresses the Objective 4 of this Thesis: *to develop mechanisms for introducing network softwarization aspects at transport level, in order to fully exploit the expected advantages of this approach.*

This chapter has overviewed a number of contributions at transport level, which have produced the following outcomes:

- Applicability of SDN concepts to wireless transport networks, exemplified on microwave radio equipment, which has serve as basis for a journal paper [A28] and contributions to standardization bodies including several contributions to ONF [O111][O112][O113] and an RFC in IETF [A29]
- Advances on the integration of transport slicing topic on operating networks, generating a number of papers either published [A30], accepted for publication [A37], or under submission [A36], as well as extensive contributions to IETF [A31][A32][A33][A34][A35], being [A31] a working group adopted document at the time of writing.

6 CONCLUSIONS AND FUTURE WORK

Telecom operators are experiencing an important transformation driven by the so called Network Softwarization trend. During the last years, the telecommunication industry has experienced the bloom of network programmability and virtualization which have promoted important changes in traditional telecom operator networks. The last step on this journey is the emerging of network slicing as model for service delivery.

This Thesis has covered a variety of technical aspects related to architecture, service and transport network topics, as well as general analysis of impacts due to SDN, NFV and Network Slicing.

Those mentioned changes impact on the networks in different dimensions and at different levels. Here, a collection of scenarios, technologies and situations have been reported identifying novel steps to follow or assessing technical propositions.

Some of the achieved results or developed ideas have been (or are being moved) to standardization, in a way of transferring them to the industry. Some others stay as research contributions to the state of the art, presenting a base for future work development. All in all they relate to use cases and scenarios applicable to operational networks by improving existing services or enabling new ones.

Regarding the work done, the general analysis of SDN, NFV and Network Slicing has dig into the implications and challenges to be considered by telecom operators for dealing with these new paradigms.

The advances produced at architectural level have explored novel architectures for facilitating interconnection and federation of infrastructures from different providers, the coordinated interaction between transport and network concerns, selection of the appropriate service edge, as well as analysis of service blocking in federated providers and efficiencies due to VNF sharing.

The advances surveyed at service level have reported experiences on the virtualization of roaming solution in multi-domain infrastructures and the validation of the novel proposition of a multicast proxy with multiple upstream interfaces.

Finally, the advances at transport level have presented the programmability of wireless transport networks and the handling of isolation for transport slicing.

As direct, short-term future work, it is yet necessary further research and standardization effort in different fronts. Here the main lines identified are briefly described:

- Architecture level: as main line of work, a more complete characterization of the gains in VNF sharing scenarios is planned, including energy efficiency characterization and network costs as a function of the delivery scenarios.
- Service level: as main line of work, for the multicast proxy with multiple upstream interfaces, the development of a solution for the configuration of the proxy functionality through standard data models is foreseen, then interacting with the proxy in a programmatic manner rather than directly program the proxy behavior.
- Transport level: as main line of work, in the case of slicing, when applied to transport networks, it is needed to close the architectural frameworks for implementation of the Transport Slice Controller concept, as well as the mapping to existing protocols and encapsulation techniques. In case of technical gaps for satisfying slice

requirements, it could be necessary the development of new transport technologies or capabilities.

However it is important to note that Network Softwarization does not terminate here. Novel capabilities for automatic and autonomic behavior are expected to emerge leveraging on the development and integration of Artificial Intelligence (AI) and Machine Learning (ML) techniques. This is what is referred as Smart Networks [O114], envisaged as the very next technological step on the evolution of telecom networks. The key concept behind Smart Networks is that the network should be able to continuously take into consideration the specific context of the end-users and the status of the network itself (i.e., understanding network in broad sense, that is, connectivity circumstances, availability of computing facilities, virtualization mechanisms, etc), all in a zero-trust environment.

In addition to that, the proper evolution of 5G services, known as Beyond 5G or even 6G, are expected to facilitate a Human Centric Internet through combining and exploiting together a variety of technologies, just emerging in the market or currently being under development. In this direction, future service and network capabilities such as the ones described in [O115] will provide a richer experience for end users, imposing stringent requirements to the telecom networks.

In any case, the solutions described in this Thesis serve as the foundation of such future service, paving the way for any future development in the direction of automation. Then, there is a need of defining an evolutionary path from the existing Network Softwarization architectures and methodologies to that ones targeting the beyond 5G future, which can be considered as mid-/long-term area of future work.

In summary, the advances here reported represent basic building blocks of the future fully automated and (to some degree) autonomous telecom operational networks.

REFERENCES

The references listed here below are divided in tow blocks. The first block, with references denoted as [Ax], is formed by the reference articles and contributions from the author which are related to the topics covered in this Thesis. The second block, denoted as [Ox], is formed by any other references included in the Thesis.

Authored publications and contributions related with the Thesis

[A1] L.M. Contreras, D.R. López, "A Network Service Provider Perspective on Network Slicing", *IEEE Softwarization*, January 2018. [Online]: https://sdn.ieee.org/newsletter/january-2018/a-network-service-provider-perspective-on-network-slicing

[A2] L. M. Contreras, V. López, O. González De Dios, A. Tovar, F. Muñoz, A. Azañón, J. P. Fernandez-Palacios, J. Folgueira, "Toward cloud-ready transport networks", in *IEEE Communications Magazine*, Vol. 50, No. 9, pp. 48-55, 2012.

[A3] L. Cominardi, L.M. Contreras, C.J. Bernardos, I. Berberana, "Understanding QoS applicability in 5G transport networks", *IEEE International Symposium on Broadband Multimedia Systems and Broadcasting*, Valencia, June 2018.

[A4] I. Bryskin, X. Liu, Y. Lee, J. Guichard, L.M. Contreras, D. Ceccarelli, J. Tantsura, D. Shytyi, "SF Aware TE Topology YANG Model", draft-ietf-teas-sf-aware-topo-model-07 (work in progress), February 2021.

[A5] Y. Lee, X. Liu, L.M. Contreras, "DC aware TE topology model", draft-llc-teas-dc-aware-topo-model-00 (work in progress), November 2020.

[A6] L.M. Contreras, V. López, R. Vilalta, R. Casellas, R. Muñoz, W. Jiang, H. Schotten, J. Alcaraz-Calero, Q. Wang, B. Sonkoly, L. Toka, "Network management and orchestration", in P. Marsch, Ö. Bulakci, O. Queseth, & M. Boldi (Eds.), *5G System Design: Architectural and Functional Considerations and Long Term Research*, John Wiley & Sons, 2018.

[A7] L.M. Contreras, S. Barguil, R. Vilalta, V. Lopez, "Architecture for integrating vertical customer's programmability control of network functions and connectivity in a slice-as-a-service schema", under submission, 2021.

[A8] L.M. Contreras, P. Doolan, H. Lønsethagen, D.R. López, "Operation, organization and business challenges for network operators in the context of SDN and NFV", *Computer Networks*, Volume 92, pp. 211-217, 2015.

[A9] L. M. Contreras, "Slicing challenges for operators", in *Emerging Automation Techniques for the Future Internet*, M. Boucadair and C. Jacquenet (Eds.), IGI Global, 2019.

[A10] L. M. Contreras, "Slicing challenges for operators", in *Research Anthology on Developing and Optimizing 5G Networks and the Impact on Society*, IGI Global, 2021.

[A11] S. Clayman, F. Tusa, A. Galis, L.M. Contreras, "WIM on-demand – A modular approach for managing network slices", *IEEE Conference on Network Softwarization (NetSoft)*, Ghent, Belgium, June 2020.

[A12] J. Ordonez-Lucena, J. Folgueira, L.M. Contreras, A. Pastor Perales. "The use of 5G Non-Public Networks to support Industry 4.0 scenarios", *IEEE Conference on Standards for Communications and Networking (CSCN)*, Granada, Spain, October 2019.

[A13] L.M. Contreras, C.J. Bernardos, D. Lopez, M. Boucadair, P. Iovanna, "Cooperating Layered Architecture for Software-Defined Networking (CLAS)", RFC 8597, May 2019.

[A14] L.M. Contreras, L. Cominardi, J. Martín-Pérez, C.J. Bernardos, "Applicability of SDN and NFV techniques for a virtualization-based roaming solution", *Journal of Network and Systems Management*, vol. 28, no. 3, pp. 576-604, 2020.

[A15] L.M. Contreras, C.J. Bernardos, "Overview of Architectural Alternatives for the Integration of ETSI MEC Environments from Different Administrative Domains", *Electronics*, September 2020.

[A16] L.M. Contreras, C.J. Bernardos, A. de la Oliva, X. Costa-Pérez, "Sharing of Crosshaul Networks via a Multi-Domain Exchange Environment for 5G Services", *IEEE Conference on Network Softwarization (NetSoft)*, Bologna, 2017.

[A17] L.M. Contreras, C.J. Bernardos, A. de la Oliva, X. Costa-Pérez, R. Guerzoni, "Orchestration of Crosshaul Slices from Separated Administrative Domains", *European Conference on Networks and Communications (EuCNC)*, Athens, 2016.

[A18] L.M. Contreras, J. Baliosian, P. Martinez-Julia, J. Serrat, "Computing at the edge, but what edge?", *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Budapest, Hungary, April 2020.

[A19] J. Baliosian, L.M. Contreras, P. Martinez-Julia, J. Serrat, "An Efficient Algorithm for Fast Service Edge Selection in Cloud-based Telco Networks", under submission, 2021.

[A20] L.M. Contreras, D. Lachos, C.E. Rothenberg, "Use of ALTO for Determining Service Edge", draft-contreras-alto-service-edge-02 (work in progress), November 2020.

[A21] A. Solano, L.M. Contreras, "Information Exchange to Support Multi-Domain Slice Service Provision for 5G/NFV", *2020 IFIP Networking Conference (Networking)*, Paris, France, June 2020.

[A22] L.M. Contreras, A. Solano, F. Cano, J. Folgueira, "Efficiency Gains due to Network Function Sharing in CDN-as-a-Service Slicing Scenarios", accepted in *IEEE International Conference on Network Softwarization (NetSoft)*, 2021.

[A23] J. Kim, L.M. Contreras, P. Greto, H. Magnusson, H. Woesner, D. Fritzsche, L. Cominardi, C.J. Bernardos, "GiLAN Roaming: Roam Like at Home in a Multi-Provider NFV Environment", *IEEE International Symposium on Networks, Computers and Communications (ISNCC)*, pp. 1-4, Rome, Italy, June 2018.

[A24] H. Asaeda, L.M. Contreras, "Multiple Upstream Interface Support for IGMP/MLD Proxy", draft-asaeda-pim-multiif-igmpmldproxy-04 (work in progress), March 2020.

[A25] L.M. Contreras, H. Asaeda, C.J. Bernardos, N. Leymann, "Enabling new multicast distribution architectures through the introduction of IGMP/MLD proxy with multiple upstream interfaces", *IEEE International Symposium on Broadband Multimedia Systems and Broadcasting*, Valencia, June 2018.

[A26] L.M. Contreras, C.J. Bernardos, H. Asaeda, N. Leymann, "Requirements for the extension of the IGMP/MLD proxy functionality to support multiple upstream interfaces", draft-ietf-pim-multiple-upstreams-reqs-08 (work in progress), November 2018.

[A27] D. Fernández, L.M. Contreras, R. Flores Moyano, S. García, "NFV/SDN Based Multiple Upstream Interfaces Multicast Proxy Service", *24th Conference on Innovation in Clouds, Internet and Networks (ICIN)*, 2021.

[A28] D. Bercovich, L.M. Contreras, Y. Haddad, A. Adam, C.J. Bernardos, "Opportunities for Software-Defined Wireless Transport Networks to Deploy a Flexible Mobile Backhaul", *ACM/Springer Mobile Applications and Networks*, Volume 20, Issue 6, pp. 793-801, 2015.

[A29] J. Ahlberg, M. Ye, X. Li, L.M. Contreras, C.J. Bernardos, "A Framework for Management and Control of Microwave and Millimeter Wave Interface Parameters", RFC 8432, October 2018.

[A30] L.M. Contreras, Ó. González, V. López, J.P. Fernández-Palacios, J. Folgueira, "iFUSION: Standards-based SDN Architecture for Carrier Transport Network", *IEEE Conference on Standards for Communications and Networking (CSCN)*, 2019.

[A31] R. Rokui, S. Homma, K. Makhijani, L.M. Contreras, J. Tantsura, "Definition of IETF Network Slices", draft-ietf-teas-ietf-network-slice-definition-01 (work in progress), February 2021.

[A32] L.M. Contreras, S. Homma, J. Ordonez-Lucena, "IETF Network Slice use cases and attributes for Northbound Interface of Controller", draft-contreras-teas-slice-nbi-03 (work in progress), October 2020.

[A33] L.M. Contreras, R. Rokui, J. Tantsura, B. Wu, X. Liu, D. Dhody, S. Belloti, "IETF Network Slice Controller and its associated data models", draft-contreras-teas-slice-controller-models-01 (work in progress), February 2021.

[A34] L.M. Contreras, P. Demestichas, J. Tantsura, "IETF Network Slice intent", draft-contreras-nmrg-transport-slice-intent-04 (work in progress), November 2020.

[A35] X. Liu, J. Tantsura, I. Bryskin, L.M. Contreras, Q. Wu, S. Belloti, R. Rokui, "Transport Network Slice YANG Data Model", draft-liu-teas-transport-network-slice-yang-02 (work in progress), November 2020.

[A36] L.M. Contreras, J. Ordonez-Lucena, S. Barguil, A. Mayoral, Ó. González de Dios "On the Definition of Standard Mechanisms for Transport Network Slicing", under submission, 2021.

[A37] L.M. Contreras, J. Ordonez-Lucena, "On Slice Isolation Options in the Transport Network and Associated Feasibility Indicators", accepted in *IEEE International Conference on Network Softwarization (NetSoft)*, 2021.

Other references

[O1] E. Rojas, R. Doriguzzi-Corin, S. Tamurejo, A. Beato, A. Schwabe, K. Phemius, C. Guerrero, "Are we ready to drive software-defined networks? A comprehensive survey on management tools and techniques", *ACM Computing Surveys*, vol. 51, no. 2, pp. 1-35, February 2018.

[O2] E. Kaljic, A. Maric, P. Begovic, M. Hadzialic, "A Survey on Data Plane Flexibility and Programmability in Software-Defined Networking", *IEEE Access*, vol. 7, pp. 47804–47840, 2019.

[O3] N.M.M.K. Chowdhury, R. Boutaba, "A survey of network virtualization", *Computer Networks*, vol. 54, pp. 862–876, 2010.

[O4] R. Mijumbi, J. Serrat, J.-L. Gorricho, N. Bouten, F. De Turck, R. Boutaba, "Network function virtualization: State-of-the-art and research challenges", *IEEE Communications Surveys & Tutorials*, vol. 18, issue 1, pp. 236 – 262, First Quarter 2016.

[O5] ONF TR-504, "SDN Architecture, Issue 1", November 2014. [Online] https://opennetworking.org/wp-content/uploads/2014/11/TR_SDN-ARCH-1.0-Overview-12012016.04.pdf

[O6] ONF TR-521, "SDN Architecture, Issue 1.1", February 2016. [Online] https://opennetworking.org/wp-content/uploads/2014/10/TR-521_SDN_Architecture_issue_1.1.pdf

[O7] A. Doria, et al., "Forwarding and Control Element Separation (ForCES) Protocol Specification", RFC 5810, March 2010.

[O8] OpenFlow Switch Specification, v. 1.5.1, March 2015. [Online]: https://opennetworking.org/wp-content/uploads/2014/10/openflow-switch-v1.5.1.pdf

[O9] ETSI GS NFV MAN 001, "Networks Functions Virtualization (NFV); Management and Orchestration", V1.1.1, December 2014.

[O10] D. Cooperson, C. Chappell, "Telefónica's UNICA architecture strategy for network virtualisation", Analysys Mason white paper, July 2017. [Online]: https://www.analysysmason.com/globalassets/x_migrated-media/media/analsyys-mason-telefonica-unica-architecture-strategy-network-virtualisation-white-paper-20174.pdf

[O11] J. Ordonez-Lucena, P. Ameigeiras, D. Lopez, J. J. Ramos-Munoz, J. Lorca and J. Folgueira, "Network slicing for 5G with SDN/NFV: Concepts, architectures and challenges", *IEEE Communications Magazine*, vol. 55, no. 5, pp. 80–87, May 2017.

[O12] M. Boucadair, C. Jacquenet, "Software-Defined Networking: A Perspective from within a Service Provider Environment", RFC 7149, March 2014.

[O13] ITU-T Recommendation X.700, "Management framework for Open Systems Interconnection (OSI) for CCITT applications", September 1992.

[O14] R. D. Doverspike, K. K. Ramakrishnan, C. Chase, "Structural Overview of ISP Networks", in C. R. Kalmanek, S. Misra, & Y. R. Yang (Eds.), *Guide to Reliable Internet Services and Applications*, Springer-Verlag, 2010.

[O15] M.-P. Odini, "SDN and NFV Evolution Towards 5G", *IEEE Softwarization*, September 2017. [Online]: https://sdn.ieee.org/newsletter/september-2017/sdn-and-nfv-evolution-towards-5g

[O16] L. Velasco, A. Castro, D. King, O. Gerstel, R. Casellas, V. López, "In-Operation Network Planning". *IEEE Communications Magazine*, Vol. 52, No. 1, pp. 52–60, 2014.

[O17] A. Kaloxylos, et al., "Network Slicing", in P. Marsch, Ö. Bulakci, O. Queseth, M. Boldi (Eds.), *5G System Design: Architectural and Functional Considerations and Long Term Research*, John Wiley & Sons, 2018.

[O18] 3GPP TR 28.801 "Study on management and orchestration of network slicing for next generation network (Release 15)", V15.1.0, 2018.

[O19] ETSI GR NFV-EVE 012, "Report on Network Slicing Support with ETSI NFV Architecture Framework", V3.1.1, 2017.

[O20] OIF, "Flex Ethernet 2.0 Implementation Agreement", IA OIF-FLEXE-02.0, June 2018.

[O21] ITU-T Recommendation Y.2011, "General principles and general reference model for Next Generation Networks", October 2004.

[O22] E. Haleplidis, K. Pentikousis, S. Denazis, J. Hadi Salim, D. Meyer, O. Koufopavlou, "Software-Defined Networking (SDN): Layers and Architecture Terminology", RFC 7426, January 2015.

[O23] A. Farrel, J. Drake, N. Bitar, G. Swallow, D. Ceccarelli, X. Zhang, "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", RFC 7926, July 2016.

[O24] ETSI GS NFV-EVE 005, "Report on SDN Usage in NFV Architectural Framework", v1.1.1, December 2015.

[O25] T.V.K. Buyakar, A.K. Rangisetti, A. Franklin, B. R. Tamma, "Auto Scaling of Data Plane VNFs in 5G Networks", *13th International Conference on Network and Service Management (CNSM)*, Tokyo, Japan, November 2017.

[O26] ONF, "TAPI Reference Implementation Agreement", v2.1.3, June 2020. [Online]: https://wiki.opennetworking.org/display/OTCC/TR-547+TAPI+v2.1.3+Reference+Implementation+Agreement

[O27] ETSI GR NFV-IFA 032, "Interface and Information Model Specification for Multi-Site Connectivity Services", V3.2.1, April 2019.

[O28] ETSI GR NFV-IFA 028, "Report on architecture options to support multiple administrative domains" V3.1.1, January 2018.

[O29] ETSI GS NFV-IFA 013, "Os-Ma-Nfvo reference point - Interface and Information Model Specification", V3.4.1, June 2020.

[O30] C.J. Bernardos, B. Gerö, M. Di Girolamo, A. Kern, B. Martini, I. Vaishnavi, "5GEx: Realizing a Europe wide multi-domain framework for Software Defined Infrastructures", *Transactions on Emerging Telecommunications Technologies*, Vol. 27, No. 9, pp. 1271-1280, September 2016.

[O31] G. Biczók, M. Dramitinos, L. Toka, P.E. Heegaard, H. Lønsethagen, "Manufactured by software: SDN-enabled multi-operator composite services with the 5G Exchange", *IEEE Communications Magazine*, Vol. 55, No. 4, pp. 80-86, April 2017.

[O32] ONF TR-534, "Framework and Architecture for the Application of SDN to Carrier Networks", July 2016. [Online]: https://www.opennetworking.org/wp-content/uploads/2014/10/TR-534_SDN_Carrier_Grade_Framework.pdf

[O33] ETSI GR NFV 001, "Network Functions Virtualisation (NFV); Use Cases", V1.2.1, May 2017.

[O34] 5G-Exchange project, Deliverable 2.2, "5GEx Final System Requirements and Architecture", December 2017. [Online]:

https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5b77e3330&appId=PPGMS

[O35] ETSI GS MEC 003, "Framework and Reference Architecture", v2.1.1, January 2019.

[O36] ETSI GS NFV-INF 001, "Infrastructure Overview", v1.1.1, January 2015.

[O37] S. Clayman, F. Tusa, A. Galis, "Extending Slices into Data Centers: the VIM on-demand model" *9th IEEE International Conference on Network of the Future (NoF 2018)*, Poznań, Poland, November 2018.

[O38] ETSI GR MEC 017, "Deployment of Mobile Edge Computing in an NFV environment", v1.1.1, February 2018.

[O39] Draft ETSI GR MEC 035, "Study on Inter-MEC systems and MEC-Cloud systems coordination", V2.0.14, February 2021.

[O40] X. Costa-Perez, A. Garcia-Saavedra, X. Li, T. Deiss, A. de la Oliva, A. di Giglio, P. Iovanna, A. Mourad, "5G-Crosshaul: An SDN/NFV Integrated Fronthaul/Backhaul Transport Network Architecture", *IEEE Wireless Communications*, Vol. 24, Issue 1, pp. 38-45, February 2017.

[O41] 3GPP TS 22.261, "Service requirements for next generation new services and markets", V15.5.0, 2018.

[O42] 3GPP TS 23.501, "System Architecture for the 5G System", V15.2.0, 2018.

[O43] NGMN, "5G White Paper", 2015. [Online]: https://www.ngmn.org/5g-white-paper/5g-white-paper.html

[O44] 5G-PPP, "5G PPP use cases and performance evaluation models", v1.0, April 2016. [Online]: https://5g-ppp.eu/wp-content/uploads/2014/02/5G-PPP-use-cases-and-performance-evaluation-modeling_v1.0.pdf

[O45] Z. Drezner, H. W. Hamacher, *Facility location: applications and theory*, Springer, 2002.

[O46] R. Alimi, et al., "Application-Layer Traffic Optimization (ALTO) Protocol", RFC7285, September 2014.

[O47] CNTT, "Common NFVI for Telco Reference Model, Release 4.0", September 2020. [Online]: https://cntt-n.github.io/CNTT/doc/ref_model/

[O48] GSMA, "Cloud Infrastructure Reference Model", NG.126, v1.0, November 2020. [Online]: https://www.gsma.com/newsroom/wp-content/uploads//NG.126-v1.0-2.pdf

[O49] H. Gredler, J. Medved, S. Previdi, A. Farrel, S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, March 2016.

[O50] W. Roome, S. Randriamasy, Y. Yang, J. Zhang, K. Gao, "Unified Properties for the ALTO Protocol", draft-ietf-alto-unified-props-new-16 (work in progress), February 2021.

[O51] H. Song, Z. Li, P. Martinez-Julia, L. Ciavaglia, A. Wang, "Network Telemetry Framework", draft-ietf-opsawg-ntf-07 (work in progress), February 2021.

[O52] N.T. Thomopoulos, *Probability Distributions with Truncated, Log and Bivariate Extensions*, Springer, 2018.

[O53] S. Kortum, "Normal Distribution", Universidade de São Paulo, November 2002. [Online]: https://edisciplinas.usp.br/pluginfile.php/2028147/mod_resource/content/0/Normal truncada.pdf

[O54] Y. Rekther, T. Li, S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[O55] L. Ginsberg, S. Previdi, Q. Wu, J. Tantsura, C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, March 2019.

[O56] L. Ginsberg, S. Previdi, S. Giacalone, D. Ward, J. Drake, Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, March 2019.

[O57] H. Asai, M. MacFaden, J. Schoenwaelder, K. Shima, T. Tsou, "Management Information Base for Virtual Machines Controlled by a Hypervisor", RFC 7666, October 2015.

[O58] Ericsson white paper, "Ericsson Mobility Report", November 2020. [Online]: https://www.ericsson.com/4adc87/assets/local/mobility-report/documents/2020/november-2020-ericsson-mobility-report.pdf

[O59] R. Wood, J. Konieczny, "Fixed network data traffic: worldwide trends and forecasts 2019–2025", Analysis Mason, February 2020.

[O60] Cisco white paper, "Cisco Annual Internet Report (2018–2023)", 2020. [Online]: https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.pdf

[O61] M. Trevisan, D. Giordano, I. Drago, M.M. Munafò, M. Mellia, "Five Years at the Edge: Watching Internet From the ISP Network", *IEEE/ACM Transactions on Networking*, Vol. 28, No. 2, pp. 561-574, April 2020.

[O62] ETSI GS MEC 002, "Mobile Edge Computing (MEC); Technical Requirements", V1.1.1, March 2016.

[O63] Akamai white paper, "The Case for a Virtualized CDN (vCDN) for Delivering Operator OTT Video", 2017. [Online]: https://www.akamai.com/us/en/multimedia/documents/white-paper/the-case-for-a-virtualized-cdn-vcdn-for-delivering-operator-ott-video.pdf

[O64] Amazon white paper, "Secure Content Delivery with Amazon CloudFront", November 2016. [Online]: https://d1.awsstatic.com/whitepapers/Security/Secure_content_delivery_with_CloudFront_whitepaper.pdf

[O65] A. Khan, W. Kellerer, K. Kozu, M. Yabusaki, "Network Sharing in the Next Mobile Network: TCO Reduction, Management Flexibility, and Operational Independence", *IEEE Communications Magazine*, Vol. 49, pp. 134–142, 2011.

[O66] D.-E. Meddour, T. Rasheed, Y. Gourhant, "On the role of infrastructure sharing for mobile network operators in emerging markets", *Computer Networks*, Vol. 55, Issue 7, pp. 1576-1591, May 2011.

[O67] I. Szczesniak, P. Cholda, A. R. Pach and B. Wozna-Szczesniak, "Interoperator fixed-mobile network sharing", *International Conference on Optical Network Design and Modeling (ONDM)*, Pisa, Italy, 2015.

[O68] I. Afolabi, T. Taleb, K. Samdanis, H. Flinck, "Network Slicing and Softwarization: A Survey on Principles, Enabling Technologies, and Solutions", *IEEE Communications Surveys & Tutorials*, Vol. 20, No. 3, pp. 2429-2453, Third Quarter, 2018.

[O69] A.A. Barakabitze, A. Ahmad, R. Mijumbi, A. Hines, "5G network slicing using SDN and NFV: A survey of taxonomy, architectures and future challenges", *Computer Networks*, Vol. 167, February 2020.

[O70] ETSI GS NFV-INF 001, "Network Function Virtualizarion (NFV); Infrastructure Overview", V1.1.1, January 2015.

[O71] Z. Avramova, D. De Vleeschauwer, S. Wittevrongel, H. Bruneel, "Dimensioning Multicast-Enabled Networks for IP-Transported TV Channels", *Proc. of the International Teletraffic Congress (ITC20)*, pp. 6-17, June 2007.

[O72] D. T. van Veen, M. K. Weldon, C. C. Bahr, E. E. Harstead, "An analysis of the technical and economic essentials for providing video over fiber-to-the-premises networks", *Bell Labs Technical Journal*, Vol. 10, No. 1, pp. 181-200, 2005.

[O73] B. Berg, et al., "The CacheLib Caching Engine: Design and Experiences at Scale", *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2020.

[O74] B. Naudts, M. Kind, F.-J. Westphal, S. Verbrugge, D. Colle, M. Pickavet, "Techno-economic analysis of software defined networking as architecture for the virtualization of a mobile network", *European Workshop on Software Defined Networking*, 2012

[O75] BEREC BoR (20) 31, "International Roaming BEREC Benchmark Data Report April 2019 – September 2019", March 2020. [Online]: https://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/9031-international-roaming-berec-benchmark-da_0.pdf

[O76] 3GPP TS 23.401, "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access", V14.11.0, December 2019.

[O77] GSMA IR.88 version 18, "LTE and EPC roaming guidelines", 2018. [Online]: https://www.gsma.com/newsroom/wp-content/uploads//IR.88-v18.0.pdf

[O78] GSMA IR.34 version 14.0, "Guidelines for IPX Provider networks", 2018. [Online]: https://www.gsma.com/newsroom/wp-content/uploads//IR.34-v14.0.pdf

[O79] 3GPP TS 23.003, "Numbering, addressing and identification", V15.8.0, 2019.

[O80] A. Asrese, E.A. Walelgne, V. Bajpai, A. Lutu, Ö. Alay, J. Ott, "Measuring Web Quality of Experience in Cellular Networks", *20th International Conference on Passive and Active Measurement (PAM)*. Springer, LNCS 11419, 2019.

[O81] A.M. Mandalari, A. Lutu, A. Custura, A.S. Khatouni, Ö. Alay, M. Bagnulo, V. Bajpai, A. Brunstrom, J. Ott, M. Mellia, G. Fairhurst, "Experience: Implications of Roaming in Europe", *24th Annual International Conference on Mobile Computing and Networking (MobiCom)*, New Dehli, India, 2018.

[O82] F. Michelinakis, H. Doroud, A. Razaghpanah, A. Lutu, N. Vallina-Rodriguez, P. Gill, J. Widmer, "The Cloud that Runs the Mobile Internet: A Measurement Study of Mobile Cloud Services", *IEEE Conference on Computer Communications (INFOCOM)*, Honolulu, 2018.

[O83] I. McKetta (Speedtest), "An expansive analysis of European mobile roaming speeds and behaviors", 2019. [Online]: https://www.speedtest.net/insights/blog/roaming-in-europe-2019

[O84] M.F. Valenzuela Gómez (Speedtest), "How roaming affects mobile speeds in Europe", 2020. [Online]: https://www.speedtest.net/insights/blog/european-roaming-2020

[O85] Juniper Research white paper, "Roam Like At Home' impact explained", 2019. [Online] https://www.juniperresearch.com/document-library/white-papers/roam-like-at-home-impact-explained

[O86] GSMA, "Next-generation Interconnection and Roaming Analysis for Mobile Services", 2017. [Online]: https://www.gsma.com/futurenetworks/wp-content/uploads/2017/03/IPX-Business-Analysis-V1-0-061016-1.pdf

[O87] BoR (19) 168, "BEREC supplementary analysis on wholesale roaming costs. BEREC", 2019. [Online]: https://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/8756-berec-supplementary-analysis-on-wholesal_0.pdf

[O88] T.M. Knoll, "Life-cycle cost modelling for NFV/SDN based mobile networks", *IEEE Conference in Telecommunication, Media and Internet Techno-Economics (CTTE)*, Munich, 2015.

[O89] C. Bouras, P. Ntarzanosy, A. Papazois, "Cost modeling for SDN/NFV based mobile 5G networks", *8th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, Lisbon, 2016.

[O90] B. Naudts, M. Kind, S. Verbrugge, D. Colle, M. Pickavet, "How can a mobile service provider reduce costs with software defined networking?", *International Journal of Network Management*, Vol. 26, No. 1, pp. 56-72, 2016.

[O91] P. Grønsund, K. Mahmood, G. Millstein, A. Noy, G. Solomon, A. Sahai, "A solution for SGi-LAN services virtualization using NFV and SDN", *European Conference on Networks and Communications (EuCNC)*, Paris, 2015.

[O92] B. Cain, S. Deering, I. Kouvelas, B. Fenner, A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[O93] H. Liu, W. Cao, H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790, February 2010.

[O94] R. Vida, L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[O95] B. Fenner, H. He, B. Haberman, H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.

[O96] M. Christensen, K. Kimball, F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, May 2006.

[O97] X. Liu, et al., "A YANG Data Model for the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD)", RFC 8652, November 2019.

[O98] ONF Central Office Re-architected as a Datacenter (CORD) project. [Online]: https://www.opennetworking.org/cord/

[O99] BBF TR-384, "Cloud Central Office Reference Architectural Framework", January 2018.

[O100] AT&T, "Domain 2.0 Vision White Paper", November 2013.

[O101] Verizon, "SDN-NFV Reference Architecture", version 1.0, February 2016.

[O102] Telecom Infra Project, "Open Transport SDN Architecture Whitepaper", November 2020. [Online]: https://cdn.brandfolder.io/D8DI15S7/at/jh6nnbb6bjvn7w7t5jbgm5n/OpenTransportArchitecture-Whitepaper_TIP_Final.pdf

[O103] E. Gray, J. Drake, "Framework for IETF Network Slices", draft-ietf-teas-ietf-network-slice-framework-00 (work in progress), March 2021.

[O104] 3GPP TR 23.799, "Study on Architecture for Next Generation System", V14.0.0, 2016.

[O105] GSMA PRD NG.116 v4.0, "Generic network Slice Template", November 2020. [Online]: https://www.gsma.com/newsroom/wp-content/uploads//NG.116-v4.0-1.pdf

[O106] A. Farrel, J.-P. Vasseur, J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC4655, August 2006.

[O107] Juniper, "Logical Systems User Guide for Routers and Switches", March 2020. [Online]: https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-logical-systems/config-guide-logical-systems.pdf

[O108] J. Soldatos, E. Vayias, G. Kormentzas, "On the Building Blocks of Quality of Service in Heterogeneus IP Networks", *IEEE Communications Surveys & Tutorials*, Vol. 7, No. 1, pp. 70-89, First Quarter 2005.

[O109] M. Karakus, A. Durresi, "Quality of Service (QoS) in Software Defined Networking (SDN): A survey", *Journal of Network and Computer Applications*, Vol. 80, pp. 200-218, 2017.

[O110] Z. Kotulski, T.W. Nowak, M. Sepczuk, M.A. Tunia, "5G networks: Types of isolation and their parameters in RAN and CN slices", *Computer Networks*, Vol. 171, April 2020.

[O111] ONF TR-532 "Microwave Information Model", December 2016, [Online]: https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-532-Microwave-Information-Model-V1.pdf

[O112] ONF White Paper, "Wireless Transport SDN Proof of Concept 2 Detailed Report", June 2016. [Online]:

https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/Wireless_Transport_SDN_PoC_White_Paper.pdf

[O113] ONF White Paper, "Wireless Transport SDN Proof of Concept White Paper", October 2015. [Online]: https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/ONF_Microwave_SDN_PoC_White_Paper%20v1.0.pdf

[O114] Networld2020, "Smart Networks in the context of NGI", 2020. [Online]: https://bscw.5g-ppp.eu/pub/bscw.cgi/d367342/Networld2020%20SRIA%202020%20Final%20Version%202.2%20.pdf

[O115] ITU-T FG NET-2030 Sub-G2, "New Services and Capabilities for Network 2030: Description, Technical Gap and Performance Target Analysis", October 2019. [Online]: https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/Deliverable_NET2030.pdf

ANNEX I – PUBLISHED, ACCEPTED AND SUBMITTED CONTENT IN THE THESIS

According to the anti-plagiarism principles dictated by the Law 14/2011, as well as the Code of Good Practices defined by the Doctoral School of the Universidad Carlos III de Madrid, I detail here the published, accepted for publication or submitted papers and contributions from where I have extracted content used as part of this Thesis, being all of them either authored or co-authored by me. They are listed in chronological order.

- L.M. Contreras, P. Doolan, H. Lønsethagen, D.R. López, "Operation, organization and business challenges for network operators in the context of SDN and NFV", in *Elsevier Computer Networks*, Volume 92, pp. 211-217, 2015. http://doi.org/10.1016/j.comnet.2015.07.016. Included in this Thesis as reference [A8].
  - This article is wholly included, being part of Chapter 2;
  - The role of the author of this Thesis in the included content is related to the analysis of challenges due to SDN and NFV for real operational telco networks;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- D. Bercovich, L.M. Contreras, Y. Haddad, A. Adam, C.J. Bernardos, "Opportunities for Software-Defined Wireless Transport Networks to Deploy a Flexible Mobile Backhaul", *ACM/Springer Mobile Applications and Networks*, Volume 20, Issue 6, pp. 793-801, 2015. https://doi.org/10.1007/s11036-015-0635-y. Included in this Thesis as reference [A28].
  - This article is wholly included, being part of Chapter 5;
  - The role of the author of this Thesis in the included content is related to the contribution on the conceptualization of the solution and the description of the use cases;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, C.J. Bernardos, A. de la Oliva, X. Costa-Pérez, R. Guerzoni, "Orchestration of Crosshaul Slices from Separated Administrative Domains", *EuCNC* 2016. http://doi.org/10.1109/EuCNC.2016.7561036. Included in this Thesis as reference [A17].
  - This article is partially included, being part of Chapter 3;
  - The role of the author of this Thesis in the included content is related to the analysis of the architectural implications of multi-domain slicing, and identification of the gaps necessary to be solved;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, C.J. Bernardos, A. de la Oliva, X. Costa-Pérez, "Sharing of Crosshaul Networks via a Multi-Domain Exchange Environment for 5G Services", *2017 IEEE Conference on Network Softwarization (NetSoft)*, Bologna, 2017, pp. 1-6. http://doi.org/10.1109/NETSOFT.2017.8004233. Included in this Thesis as reference [A16].
  - This article is partially included, being part of Chapter 3;
  - The role of the author of this Thesis in the included content is related to the analysis of the architectural implications of multi-domain slicing, and identification of the gaps necessary to be solved;

  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, D.R. López, "A Network Service Provider Perspective on Network Slicing", *IEEE Softwarization*, January 2018. [Online]: https://sdn.ieee.org/newsletter/january-2018/a-network-service-provider-perspective-on-network-slicing. Included in this Thesis as reference [A1].
  - This article is partially included, being part of Chapter 2;
  - The role of the author of this Thesis in the included content is related to the analysis and implications of network slicing in operational telco networks, and the definition of distinct types of slices from an operational perspective;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- J. Kim, L.M. Contreras, P. Greto, H. Magnusson, H. Woesner, D. Fritzsche, L. Cominardi, C.J. Bernardos, "GiLAN Roaming: Roam Like at Home in a Multi-Provider NFV Environment", *IEEE International Symposium on Networks, Computers and Communications (ISNCC)*, pp. 1-4, Rome, Italy, June 2018. http://doi.org/10.1109/ISNCC.2018.8530973. Included in this Thesis as reference [A23].
  - This article is partially included, being part of Chapter 4;
  - The role of the author of this Thesis in the included content relates to the contribution to the analysis of the virtualized roaming use case and the conceptualization of the solution;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, H. Asaeda, C.J. Bernardos, N. Leymann, "Enabling new multicast distribution architectures through the introduction of IGMP/MLD proxy with multiple upstream interfaces", *IEEE International Symposium on Broadband Multimedia Systems and Broadcasting*, Valencia, June 2018. http://doi.org/10.1109/BMSB.2018.8436830. Included in this Thesis as reference [A25].
  - This article is wholly included, being part of Chapter 4;
  - The role of the author of this Thesis in the included content is related to the analysis of the use cases, requirements and potential solutions;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, C.J. Bernardos, H. Asaeda, N. Leymann, "Requirements for the extension of the IGMP/MLD proxy functionality to support multiple upstream interfaces", draft-ietf-pim-multiple-upstreams-reqs-08 (work in progress), November 2018. https://tools.ietf.org/html/draft-ietf-pim-multiple-upstreams-reqs-08. Included in this Thesis as reference [A26].
  - This standardization contribution is partially included, being part of Chapter 4;
  - The role of the author of this Thesis in the included content is related to the analysis of the use cases, elicitation of requirements and the conceptualization of the solution;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L. M. Contreras, "Slicing challenges for operators", in *Emerging Automation Techniques for the Future Internet*, M. Boucadair and C. Jacquenet (Eds.), IGI Global, 2019.

https://www.igi-global.com/chapter/slicing-challenges-for-operators/214431. Included in this Thesis as reference [A9].
  - This book chapter is wholly included, being part of Chapter 2;
  - The role of the author of this Thesis in the included content relates to the analysis of the implications of network slicing in operational telco networks;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, C.J. Bernardos, D. Lopez, M. Boucadair, P. Iovanna, "Cooperating Layered Architecture for Software-Defined Networking (CLAS)", RFC 8597, May 2019. https://tools.ietf.org/html/rfc8597. Included in this Thesis as reference [A13].
  - This standardization contribution is wholly included, being part of Chapter 3;
  - The role of the author of this Thesis in the included content is related to the conceptualization of the layered SDN architecture proposed and its functional analysis;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, Ó. González, V. López, J.P. Fernández-Palacios, J. Folgueira, "iFUSION: Standards-based SDN Architecture for Carrier Transport Network", *IEEE Conference on Standards for Communications and Networking (CSCN)*, 2019. http://doi.org/10.1109/CSCN.2019.8931386. Included in this Thesis as reference [A30].
  - This article is partially included, being part of Chapter 5;
  - The role of the author of this Thesis in the included content is related to the contribution to the overall definition of the architecture proposed;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- H. Asaeda, L.M. Contreras, "Multiple Upstream Interface Support for IGMP/MLD Proxy", draft-asaeda-pim-multiif-igmpmldproxy-04 (work in progress), March 2020. https://tools.ietf.org/html/draft-asaeda-pim-multiif-igmpmldproxy-04. Included in this Thesis as reference [A24].
  - This standardization contribution is partially included, being part of Chapter 4;
  - The role of the author of this Thesis in the included content relates to the conceptualization of the SDN control of the proxy and its functional analysis;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, J. Baliosian, P. Martinez-Julia, J. Serrat, "Computing at the edge, but what edge?", *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Budapest, Hungary, April 2020. http://doi.org/10.1109/NOMS47738.2020.9110342. Included in this Thesis as reference [A18].
  - This article is partially included, being part of Chapter 3;
  - The role of the author of this Thesis in the included content relates to the conceptualization of the service edge problem and definition of the architecture proposed (leveraging on ALTO);
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, L. Cominardi, J. Martín-Pérez, C.J. Bernardos, "Applicability of SDN and NFV techniques for a virtualization-based roaming solution", *Journal of Network*

and Systems Management, vol. 28, no. 3, pp. 576-604, 2020. https://doi.org/10.1007/s10922-020-09534-z. Included in this Thesis as reference [A14].
  - This article is wholly included, being part of both Chapter 3 and 4;
  - The role of the author of this Thesis in the included content relates to the conceptualization of the problem, definition of the architecture proposed, analysis of the experiment results, and the techno-economic analysis;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.

- A. Solano, L.M. Contreras, "Information Exchange to Support Multi-Domain Slice Service Provision for 5G/NFV", *2020 IFIP Networking Conference (Networking)*, Paris, France, June 2020. https://ieeexplore.ieee.org/document/9142819. Included in this Thesis as reference [A21].
  - This article is wholly included, being part of Chapter 3;
  - The role of the author of this Thesis in the included content relates to the conceptualization of the problem, the definition of the simulations, the analysis of results and the proposition of the solution for information dissemination among administrative domains;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, C.J. Bernardos, "Overview of Architectural Alternatives for the Integration of ETSI MEC Environments from Different Administrative Domains", *Electronics*, vol. 9, issue 9, 2020. https://doi.org/10.3390/electronics9091392. Included in this Thesis as reference [A15].
  - This article is wholly included, being part of Chapter 3;
  - The role of the author of this Thesis in the included content is related to the conceptualization of the problem, definition of the architectural options proposed, and analytical comparison of them;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, D. Lachos, C.E. Rothenberg, "Use of ALTO for Determining Service Edge", draft-contreras-alto-service-edge-02 (work in progress), November 2020. https://tools.ietf.org/html/draft-contreras-alto-service-edge-02. Included in this Thesis as reference [A20].
  - This standardization contribution is partially included, being part of Chapter 3;
  - The role of the author of this Thesis in the included content is related to the conceptualization of the problem, the proposal of ALTO as architectural component for addressing it, and the parameters to be included;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L. M. Contreras, "Slicing challenges for operators", in *Research Anthology on Developing and Optimizing 5G Networks and the Impact on Society*, IGI Global, 2021. https://www.igi-global.com/chapter/slicing-challenges-for-operators/270211. Included in this Thesis as reference [A10] (being a re-print of [A9]).
  - This book chapter is wholly included, being part of Chapter 2;
  - The role of the author of this Thesis in the included content relates to the analysis of the implications of network slicing in operational telco networks;

  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, R. Rokui, J. Tantsura, B. Wu, X. Liu, D. Dhody, S. Belloti, "IETF Network Slice Controller and its associated data models", draft-contreras-teas-slice-controller-models-01 (work in progress), February 2021. https://tools.ietf.org/html/draft-contreras-teas-slice-controller-models-01. Included in this Thesis as reference [A33].
  - This standardization contribution is partially included, being part of Chapter 5;
  - The role of the author of this Thesis in the included content is related to the conceptualization of the structure and the applicability of the data models;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- D. Fernández, L.M. Contreras, R. Flores Moyano, S. García, "NFV/SDN Based Multiple Upstream Interfaces Multicast Proxy Service", *24th Conference on Innovation in Clouds, Internet and Networks (ICIN)*, 2021. https://doi.org/10.1109/ICIN51074.2021.9385529. Included in this Thesis as reference [A27].
  - This article is wholly included, being part of Chapter 4;
  - The role of the author of this Thesis in the included content is related to the conceptualization of the problem, validation of the experiments and the analysis of results;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, J. Ordonez-Lucena, "On Slice Isolation Options in the Transport Network and Associated Feasibility Indicators", accepted in *IEEE International Conference on Network Softwarization (NetSoft)*, 2021. Included in this Thesis as reference [A37].
  - This article is wholly included, being part of Chapter 5;
  - The role of the author of this Thesis in the included content relates to the conceptualization of the problem and definition of the solution;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, A. Solano, F. Cano, J. Folgueira, "Efficiency Gains due to Network Function Sharing in CDN-as-a-Service Slicing Scenarios", accepted in *IEEE International Conference on Network Softwarization (NetSoft)*, 2021. Included in this Thesis as reference [A22].
  - This article is wholly included, being part of Chapter 3;
  - The role of the author of this Thesis in the included content relates to the conceptualization of the problem, definition of the simulations, analysis of the results, and techno-economic analysis;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, J. Ordonez-Lucena, S. Barguil, A. Mayoral, Ó. González de Dios "On the Definition of Standard Mechanisms for Transport Network Slicing", under submission, 2021. Included in this Thesis as reference [A36].
  - This article is partially included, being part of Chapter 5;
  - The role of the author of this Thesis in the included content relates to the conceptualization of the problem and the definition of the architecture;

  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.
- L.M. Contreras, S. Barguil, R. Vilalta, V. Lopez, "Architecture for integrating vertical customer's programmability control of network functions and connectivity in a slice-as-a-service schema", under submission, 2021. Included in this Thesis as reference [A7].
  - This article is partially included, being part of both Chapters 2 and 3;
  - The role of the author of this Thesis in the included content relates to the conceptualization of the problem and the definition of the architecture;
  - The content from this reported source included in the Thesis has not been singled out with any typographic means or references.

ANNEX II - OTHER MERITS RELATED TO RESEARCH ACTIVITIES

This Annex reports additional research activities performed by the author of this Thesis. First, additional (co-)authored publications are listed, being those publications either not related to the topic of this Thesis or not included as a reference within. Second, results from the participation in standardization activities are provided, just reporting those including explicit evidences of contribution or authorship. Third, (co-)authored patents and inventions are mentioned. Fourth, it is described the role taken in several international research and innovation where the author has been involved. Fifth, research visits realized up to date. Sixth, it is reported a number of activities related to editorship of special issues on magazines and journals, as well as organization of scientific events. Finally, a list of diverse talks (keynotes, invited talks, etc.) at different scientific and industrial venues is provided.

Additional publications

The following is a list of additional co-authored publications not included as reference work in this Thesis (publications under submission not related to the Thesis are not listed).

Journals

1. C. Guimarães, M. Groshev, L. Cominardi, A. Zabala, L.M. Contreras, S.T. Talat, C. Zhang, S. Hazra, A. Mourad, A. de la Oliva, "DEEP: A Vertical-Oriented Intelligent and Automated Platform for the Edge and Fog", accepted for publication in *IEEE Communications Magazine*, 2021.
2. I. Sarrigiannis, L.M. Contreras, K. Ramantas, A. Antonopoulos, C. Verikoukis, "Application as a Service Function Chain in a Fog-enabled C-V2X Architecture", *IEEE Network*, vol. 34, no. 5, pp. 120-126, September/October 2020.
3. P. Iovanna, G. Bottari, F. Ponzini, L.M. Contreras, "Latency Driven Transport for 5G", *IEEE/OSA Journal of Optical Communications and Networking (JOCN)*, Vol. 10, Issue 8, pp. 695-702, 2018.
4. J. M. Fabrega, M. Svaluto Moreolo, L. Nadal, F. J. Vílchez, R. Casellas, R. Vilalta, R. Martínez, R. Muñoz, J. P. Fernández-Palacios, L.M. Contreras, "Experimental Validation of a Converged Metro Architecture for Transparent Mobile Front-/Back-Haul Traffic Delivery using SDN-enabled Sliceable Bitrate Variable Transceivers", *IEEE/OSA Journal of Lightwave Technology (JLT)*, Vol. 36, Issue 7, pp. 1429-1434, April 2018.
5. R. Guerzoni, I. Vaishnavi, D. Perez-Caparros, A. Galis. F. Tusa, P. Monti, A. Sgambelluri, G. Biczók, B. Sonkoly, L. Toka, A. Ramos, J. Melián, O. Dugeon, F. Cugini, B. Martini, G. Giuliani, R. Figueiredo, L. M. Contreras, C. J. Bernardos, R. Szabó, "Analysis of End-to-End Multi-Domain Management and Orchestration Frameworks for Software Defined Infrastructures: an Architectural Survey", *Transactions on Emerging Telecommunications Technologies*, Vol. 28, Issue 4, 2017.
6. Ll. Gifre, M. Tornatore, L.M. Contreras, B. Mukherjee, L. Velasco, "ABNO-driven Content Distribution in the Telecom Cloud", *Elsevier Optical Switching and Networks*, vol. 26, pp. 25-38, 2017.
7. F. Morales, M. Ruiz, L. Gifre, L. M. Contreras, V. López, L. Velasco, "Virtual Network Topology Adaptability based on Data Analytics for Traffic Prediction", *IEEE/OSA Journal of Optical Communications and Networking (JOCN)*, Vol. 9, Issue 1, pp. A35-A45, 2017.

8. Asensio, M. Ruiz, L. M. Contreras, L. Velasco, "Dynamic Virtual Network Services to Support 5G Backhauling", *IEEE/OSA Journal of Optical Communications and Networking (JOCN)*, vol. 8, pp. B93-B103, 2016.
9. M. Ruiz, M. Germán, L. M. Contreras, L. Velasco, "Big Data-backed Video Distribution in the Telecom Cloud", *Elsevier Computer Communications*, Volume 84, pp. 1-11, 2016.
10. L.M. Contreras, L. Cominardi, H. Qian, C.J. Bernardos, "Software-Defined Mobility Management: Architecture Proposal and Future Directions", *ACM/Springer Mobile Applications and Networks*, Volume 21, Issue 2, pp. 226-236, 2016.
11. L. Velasco, L. M. Contreras, G. Ferraris, A. Stavdas, F. Cugini, M. Wiegand, J. P Fernández-Palacios, "A Service-Oriented Hybrid Access Network and Cloud Architecture", *IEEE Communications Magazine*, Vol. 53, No. 4, pp. 159-165, 2015.
12. Ll. Gifre, F. Paolucci, O. González de Dios, L. Velasco, L. M. Contreras, F. Cugini, P. Castoldi, V. López "Experimental Assessment of ABNO-driven Multicast Connectivity in Flexgrid Networks", *IEEE/OSA Journal of Lightwave Technology (JLT)*, Vol. 33, Issue 8, pp. 1549-1556, 2015.
13. M.R. Sama, L. M. Contreras, J. Kaippallimalil, I. Akiyoshi, H. Qian , H. Ni, "Software-Defined Control of Virtualized Mobile Packet Core", *IEEE Communications Magazine*, Vol. 53, No. 2, pp. 107-115, 2015.
14. C.J. Bernardos, A. de la Oliva, P. Serrano, A. Banchs, L.M. Contreras, H. Jin, J.C. Zúñiga, "An Architecture for Software Defined Wireless Networking", *IEEE Wireless Communication Magazine*, Vol. 21, No. 3, pp. 52-61, 2014.
15. Belter, J. Rodriguez Martinez, J. I. Aznar, J. Ferrer Riera, L.M. Contreras, M. Antoniak-Lewandowska, M. Biancani, J. Buysse, C. Develder, Y. Demchenko, P. Donadio, D. Simeonidou, R. Nejabati, S. Peng, L. Drzewiecki, E. Escalona, J. A. Garcia-Espin, S. Gheorghiu, M. Ghijsen, J. Gutkowski, G.Landi, G. Carrozzo, D. Parniewicz, P. Robinson, S. Soudan, "The GEYSERS Optical Test-bed: a Platform for the Integration, Validation and Demonstration of Cloud-based Infrastructure Services", *Elsevier Computer Networks*, Vol. 61, pp. 197-215, March 2014.
16. J. Buysse, M. De Leenheer, L. M. Contreras, J. I. Aznar, J. Rodriguez Martinez, G. Landi, C. Develder, "NCP+: an integrated network and IT control plane for cloud computing", *Elsevier Optical Switching and Networking*, Vol. 11, pp. 137-152, 2014.
17. L.M. Contreras, C. J. Bernardos, I. Soto, "RAMS: A protocol extension to PMIPv6 for improving handover performance of multicast traffic", *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, Vol.2, No. 2, pp. 67-82, 2011.
18. C.J. Bernardos, M. Gramaglia, L. M Contreras, M. Calderón, I. Soto, "Network-based localized IP mobility management: Proxy Mobile IPv6 and current trends in standardization", *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, Vol.1, No. 2/3, pp. 16-35, 2010.

Conferences

1. A. Alcalá, S. Barguil, V. López, L.M. Contreras, C. Manso, P. Alemany, R. Casellas, R. Martínez, D. González-Pérez, X. Liu, J.M. Pulido, J.P. Fernandez-Palacios, R. Muñoz, R. Vilalta, "Multi-layer Transport Network Slicing with Hard and Soft Isolation", *OSA Optical Fiber Communication Conference (OFC)*, June 2021.
2. J. Zhang, L.M. Contreras, K. Gao, F. Cano, P. Diez Cano, A. Escribano, Y. R. Yang, "Sextant: Enabling Network-Aware Applications in Carrier Networks", *IFIP/IEEE*

International Symposium on Integrated Network Management, Bordeux, France, May 2021.

3. J.T. Infiesta, C. Guimarães, L.M. Contreras, A. de la Oliva, "GANSO: Automating the Provisioning of Network Slices over SDN Infrastructures", *IEEE Conference on Network Function Virtualization and Software Defined Networks (IEEE NFV-SDN)*, Madrid, Spain, November 2020.
4. K. Gao, L.M. Contreras, S. Randriamasy, "Bi-directional Network and Application Interaction: Application Intents upon Abstracted Network Information", *Workshop on Network Application Integration/CoDesign (NAI 2020), ACM SIGCOMM*, New York, USA, August 2020.
5. D. A. Lachos, Q. Xiang, C.E. Rothenberg, S. Randriamasy, L.M. Contreras, B. Ohlman, "Towards Deep Network & Application Integration: Possibilities, Challenges, and Research Directions", *Workshop on Network Application Integration/CoDesign (NAI 2020), ACM SIGCOMM*, New York, USA, August 2020 .
6. S. Bryant, U. Chunduri, T. Eckert, A. Clemm, L.M. Contreras, P. Diez Cano, "A novel hybrid distributed-routing and SDN solution for Traffic Engineering", *Applied Networking Research Workshop (ANRW)*, Madrid, Spain, July 2020.
7. D. A. Lachos, C.E. Rothenberg, Q. Xiang, Y.R. Yang, B. Ohlman, S. Randriamasy, L.M. Contreras, K. Gao, "Multi-Domain Information Exposure using ALTO: The Good, the Bad and the Solution", *Applied Networking Research Workshop (ANRW)*, Madrid, Spain, July 2020.
8. J. Ordonez-Lucena, C. Tranoris, J. Rodrigues, L.M. Contreras, "Cross-domain Slice Orchestration for Advanced Vertical Trials in a Multi-Vendor 5G Facility", *European Conference on Networks and Communications (EuCNC)*, Dubrovnik, Croatia, 2020.
9. M. Lashgari, C. Natalino, L.M. Contreras, P. Monti, "Cost Benefits of Centralizing Service Processing in 5G Network Infrastructures", *Asia Communications and Photonics Conference (ACP 2019)*, November 2019.
10. J. Garcia-Reinoso, M. Molla, E. Kosmatos, G. Landi, G. Bernini, R. Legouable, L.M. Contreras, M. Lorenzo, "The 5G-EVE Multi-site Experimental Architecture", *IEEE 5G World Forum conference*, October 2019.
11. D.A. Lachos, C.E. Rothenberg, Q. Xiang, Y.R. Yang, B. Ohlman, S. Randriamasy, F. Weaver, L. M. Contreras, "Supporting Multi-domain Use Cases with ALTO", *Applied Networking Research Workshop (ANRW)*, Montreal, Canada, July 2019.
12. L.M. Contreras, L. Luque, G. Landi, G. Bernini, G. Carrozzo, J. García-Reinoso, M. Mollà-Roselló, "Interworking of softwarized infrastructures for enabling 5G multi-site slice orchestration", *IEEE Conference on Network Softwarization (NetSoft)*, June 2019.
13. L. Mo, W. Cheng, L.M. Contreras, "ZTE 5G Transport Solution and Joint Field Trials with Global Operators", *Optical Fiber Conference (OFC)*, March 2019.
14. P. Iovanna, S. Stracca, F. Ubaldi, F. Cavaliere, G. Vall-Llosera, L.M. Contreras, "Network Convergence in 5G Transport", *Optical Fiber Conference (OFC)*, March 2019.
15. F.S. Dantas Silva, M.O.O. Lemos, A. Medeiros, A.V. Neto, R. Pasquini, D. Moura, C. Rothenberg, L. Mamatas, S.L. Correa, K. Vieira Cardoso, C. Marcondes, A. Abelem, M. Nascimento, A. Galis, L.M. Contreras, J. Serrat, P. Papadimitriou, "NECOS Project: Towards Lightweight Slicing of Cloud Federated Infrastructures", *IEEE Conference on Network Softwarization (NetSoft)*, Montreal, June 2018.

16. M.Á. Vázquez-Castro, L.M. Contreras, "Softwarization of Network Coding Functions and Logical Mapping to SDN", *IEEE International Symposium on Networks, Computers and Communications (ISNCC)*, Rome, June 2018.
17. R. Vilalta, R. Casellas, R. Martínez, R. Muñoz Y. Lee, H. Zheng, Y. Lin, V. López, L. M. Contreras, "Fully automated peer service orchestration of cloud and network resources using ACTN and CSO", *22nd International Conference on Optical Network Design and Modeling (ONDM)*, 2018.
18. S. González, A. de la Oliva, L.M. Contreras, C.J. Bernardos, "Towards a resilient OpenFlow channel through MPTCP", *IEEE International Symposium on Broadband Multimedia Systems and Broadcasting*, Valencia, June 2018.
19. L. Cominardi, L.M. Contreras, C.J. Bernardos, I. Berberana, "Understanding QoS applicability in 5G transport networks", *IEEE International Symposium on Broadband Multimedia Systems and Broadcasting*, Valencia, June 2018.
20. P. Trakadas, P. Karkazis, H.-C. Leligou, T. Zahariadis, W. Tavernier, T. Soenen, S. van Rossem, L.M. Contreras, "Scalable Monitoring for Multiple Virtualized Infrastructures for 5G Services", *7th International Conference on Networks (ICN 2018)*, Athens, April 2018.
21. B. Gerö, D. Jocha, R. Szabó, J. Czentye, D. Haja, B. Németh, B. Sonkoly, M. Szalay, L. Toka., C.J. Bernardos, L.M. Contreras, "The Orchestration in 5G Exchange – a Multi-Provider NFV Framework for 5G Services", *IEEE Conference on Network Function Virtualization and Software Defined Networks,* 2017, Berlin, Germany.
22. J.M. Fabrega, M. Svaluto Moreolo, L. Nadal, F.J. Vílchez, R. Casellas, R. Vilalta, R. Martínez, R. Muñoz, J.P. Fernández-Palacios, L.M. Contreras, "Experimental Validation of a Converged Metro Architecture for Transparent Mobile Front-/Back-Haul Traffic Delivery using SDN-enabled Sliceable Bitrate Variable Transceivers", *43rd European Conference and Exhibition on Optical Communications (ECOC),* 2017.
23. T. Deiß, J. Baranda, L. Cominardi, L.M. Contreras, J. Gomes, S. Gonzalez, P. Iovanna, J. Mangues-Bafalluy, N. Molner, J. Núñez-Martínez, A. de Oliva, S. Stracca, "Dataplane Measurements on a Fronthaul and Backhaul Integrated Network", *CLEEN* 2017.
24. J.M. Fabrega, M. Svaluto Moreolo, L. Nadal, F.J. Vílchez, J.P. Fernández-Palacios, L.M. Contreras, "Mobile front-/back-haul delivery in elastic metro/access networks with sliceable transceivers based on OFDM transmission and direct detection", *ICTON* 2017.
25. R. Casellas, R. Vilalta, A. Mayoral, R. Martínez, R. Muñoz, L. M. Contreras, "Control Plane Architectures Enabling Transport Network Adaptive and Autonomic Operation", *ICTON* 2017.
26. V. López, R. Jiménez, O. González de Dios, L.M. Contreras, J.P. Fernández Palacios, "Open Source Netphony Suite: Enabling Multi-layer Network Programmability", *21st International Conference on Optical Network Design and Modeling (ONDM)*, 2017.
27. R. Vilalta, Y. Lee, H. Zheng, Y. Lin, R. Casellas, A. Mayoral, R. Martinez, R. Muñoz, L. M. Contreras, V. Lopez, "Fully Automated Peer Service Orchestration of Cloud and Network Resources Using ACTN and CSO", *Optical Fiber Conference (OFC)*, March 2017.
28. A. Sgambelluri, A. Milani, J. Czentye, J. Melian, W. Y. Poe, F. Tusa, O. Gonzalez de Dios, B. Sonkoly, M. Gharbaoui, F. Paolucci, E. Meini, G. Giuliani, A. Ramos, P. Monti, L. M. Contreras, I. Vaishnavi, C. J. Bernardos, R. Szabó, "A Multi-Operator Network Service Orchestration Prototype: The 5G Exchange", *Optical Fiber Conference (OFC)*, March 2017.

29. F. Paolucci, V. Uceda, A. Sgambelluri, F. Cugini, O. Gonzalez De Dios, V. Lopez, L.M. Contreras, P. Monti, P. Iovanna, F. Ubaldi, T. Pepe, P. Castoldi, "Interoperable Multi-Domain Delay-aware Provisioning using Segment Routing Monitoring and BGP-LS Advertisement", *42nd European Conference and Exhibition on Optical Communications (ECOC)*, Düsseldorf, Germany, September 2016.
30. V. López, L.M. Contreras, Ó. González de Dios, J.P. Fernández-Palacios, "Towards a Transport SDN for Carrier Networks: An Evolutionary Approach", *21st European Conference on Network and Optical Communications (NOC)*, Lisbon, Portugal, June 2016.
31. G. Gardikis, S. Costicoglou, H. Koumaras, Ch. Sakkas, G. Xilouris, F. Arnal, L. M. Contreras, P. Aranda, M. Guta, "NFV Applicability and Use Cases in Satellite Networks", *25th Edition of the European Conference on Networks and Communications (EuCNC)*, Athens, Greece, June 2016.
32. X. Li, G. Landi, J. Núñez-Martínez, R. Casellas, S. González, C.-F. Chiasserini, J. Rivas, D. Siracusa, L. Goratti, D. Jimenez, L. M. Contreras, "Innovations through 5G-Crosshaul Applications", *25th Edition of the European Conference on Networks and Communications (EuCNC)*, Athens, Greece, June 2016.
33. Ll. Gifre, L.M. Contreras, V. López, L. Velasco "Big Data Analytics in Support of Virtual Network Topology Adaptability", *Optical Fiber Conference (OFC)*, Anneheim, March 2016.
34. A. Asensio, M. Ruiz, L.M. Contreras, L. Velasco, G. Junyent, "Dynamic Customer Virtual Network Reconfiguration with QoS and Bandwidth Guarantees", *41st European Conference on Optical Communication (ECOC)*, Valencia, Spain, September 2015.
35. F. Ubaldi, P. Iovanna, F. Giurlanda, S. Noto, A. Priola, L.M. Contreras, V. López, J.P. Fernández-Palacios, "Effective Elasticity for Data Centers Interconnection in Multi-Domain WAN: Information Modelling and Routing", *41st European Conference on Optical Communication (ECOC)*, Valencia, Spain, September 2015.
36. P. Iovanna, F. Ubaldi, T. Pepe, L.M. Contreras, V. López, J.P. Fernández-Palacios, "Main Challenges on WAN due to NFV and SDN: Multi-layer and Multi-domain Network Virtualization and Routing", *Optical Networking Design and Modeling (ONMD)*, May 2015.
37. A. Asensio, L.M. Contreras, M. Ruiz, V. López, L. Velasco, "Scalability of Telecom Cloud Architectures for Live-TV Distribution", *Optical Fiber Conference (OFC)*, March 2015.
38. V. López, O. González de Dios, L.M. Contreras, J. Foster, H. Silva, L. Blair, J. Marsella, T. Szyrkowiec, A. Autenrieth, C. Liou, A. Sasdasivarao, S. Syed, J. Sun, B. Rao, F. Zhang, J.P. Fernández-Palacios, "Demonstration of SDN Orchestration in Optical Multi-Vendor Scenarios", *Optical Fiber Conference (OFC)*, March 2015.
39. Ll. Gifre, F. Paolucci, J. Marhuenda, A. Aguado, L. Velasco, F. Cugini, P. Castoldi, O. Gonzalez de Dios, L.M. Contreras, V. Lopez, "Experimental Assessment of Inter-datacenter Multicast Connectivity for Ethernet services in Flexgrid Networks", *European Conference on Optical Communications (ECOC)*, 2014.
40. T. Szyrkowiec, A. Autenrieth, J-P. Elbers, W. Kellerer, P. Kaczmarek, V. López, L.M. Contreras, O. Gonzalez de Dios, J. P. Fernández-Palacios, R. Muñoz, R. Vilalta, R. Casellas, R. Martínez, A. Mayoral, M. Channegowda, S. Peng, R. Nejabati and D. Simeonidou, "Demonstration of SDN Based Optical Network Virtualization and

Multidomain Service Orchestration", *Third European Workshop on Software Defined Networking (EWSDN)*, Sep 2014.
41. J. González, F. Álvarez, L.M. Contreras, Ó. González, "Interdomain Monitoring and Internetworking Connectivity in Federated Infrastructures based on Software Defined Networking", *2014 European Conference on Networks and Communications (EuCNC)*, Bologna, June, 2014.
42. R. Muñoz, R. Vilalta, R.Casellas, R. Martínez, L.M. Contreras, V. López, J.P. Fernández-Palacios, O. González de Dios, S. Peng, M. Channegowda, R. Nejabati, D. Simeonidou, X. Cao, N. Yoshikane, T. Tsuritani, A. Autenrieth, M. Schlosser, "Network Virtualization, Control Plane and Service Orchestration of the ICT STRAUSS Project", *European Conference on Networks and Communications (EuCNC)*, Bologna, June, 2014.
43. E. Escalona, J.I. Aznar, L.M. Contreras, O. González de Dios, G. Cossu, E. Salvadori, F.M. Facca, "Using SDN for Cloud services provisioning: the XIFI use case", *1st IEEE Workshop on Software Defined Networks for Future Networks and Services (SDN4FNS)*, Trento, Italy, 2013.
44. L.M. Contreras, C.J. Bernardos, "Optimal Distribution of Remotely-Subscribed Multicast Traffic within a Proxy Mobile IPv6 Domain by Using Explicit Multicast", *Proceedings of JITEL*, Granada, Spain, 2013.
45. Y. Demchenko, C. Ngo, C. de Laat, J. A. Garcia-Espin, S. Figuerola, J. Rodriguez, L.M. Contreras, G. Landi, N. Ciulli, "Intercloud Architecture Framework for Heterogeneous Cloud based Infrastructure Services Provisioning On-Demand", *27th International Conference on Advanced Information Networking and Applications Workshops (AINA)*, pp. 777-784, Barcelona, Spain, 2013.
46. A-F. Antonescu, P. Robinson, L.M. Contreras, J. Aznar, S. Soudan, F. Anhalt, J. A. Garcia-Espin, "Towards Cross Stratum SLA Management with the GEYSERS Architecture", *10th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA)*, pp. 527-533, Leganés, Spain, 2012.
47. L.M. Contreras, C. J. Bernardos, I. Soto, "On the efficiency of a dedicated LMA for multicast traffic distribution in PMIPv6 domains", *Fifth ERCIM Workshop on e-Mobility*, pp. 51-62, Vilanova i la Geltrù, Spain, 2011.
48. L.M. Contreras, F. Escrihuela. "VoIP transmission in Fixed Wireless Access networks based on DECT Packet Radio Service", *XIII International Symposium on Services and Local accesS (ISSLS 2000)*, Stockolm, Sweden, June 2000.

Book chapters

1. M. Gramaglia, A. Kaloxylos, P. Spapis, X. Costa, L. M. Contreras, R. Trivisonno, G. Zimmermann, A. de la Oliva, P. Rost, P. Marsch, "E2E Architecture", in *5G System Design – Architectural and Functional Considerations and Long Term Research*, P. Marsch, Ö. Bulakci, O. Queseth, M. Boldi (Eds.), Wiley, 2018.
2. G. Biczók, M. Dramitinos, H. Lønsethagen, L.M. Contreras, G.D. Stamoulis, L. Toka, "Towards multi-operator IPTV services over 5G networks", in *A Comprehensive Guide to IPTV Delivery Networks*, Suliman Mohamed Fati, Saiful Azad, Al-Sakib Khan Pathan (Eds.), Wiley, 2018.
3. A. Asensio, L.M. Contreras, M. Ruiz, L. Velasco, "Dynamic connectivity services in support of future mobile networks", in *Provisioning, Recovery and In-Operation Planning in Elastic Optical Networks*, L. Velasco, M. Ruiz (Eds.), Wiley, October, 2017.

4. L.M. Contreras, V. López, O. González de Dios, F. Jiménez, J. Rodríguez, J.P. Fernández-Palacios, "Migration Path Towards Cloud-Aware Core Networks", in *Communication Infrastructures for Cloud Computing: Design and Applications*, H.T. Mouftah and B. Kantarci (Eds.), IGI Global, September, 2013.

Standardization

The following are references to documents, white papers, reports and specifications officially released by Standards Development Organizations (SDOs) in which I have either co-authored (i.e., taking leadership on the edition of the document) or contributed.

Other released contributions or documents where no explicit reference to the author is given are not listed.

IETF (only RFCs)

X. Xu, B. Decraene, R. Raszuk, L.M. Contreras, L. Jalil, "The Tunnel Encapsulations OSPF Router Information", RFC-to-be 9013, pending of approval for publication, 2021.

L.M. Contreras, C.J. Bernardos, D. Lopez, M. Boucadair, P. Iovanna, "Cooperating Layered Architecture for Software-Defined Networking (CLAS)", RFC 8597, May 2019.

C.J. Bernardos, A. Rahman, J.C. Zuniga, L.M. Contreras, P. Aranda, P. Lynch, "Network Virtualization Research Challenges", RFC 8568, April 2019.

J. Ahlberg, M. Ye, X. Li, L.M. Contreras, C.J. Bernardos, "A Framework for Management and Control of Microwave and Millimeter Wave Interface Parameters", RFC 8432, October 2018.

L.M. Contreras, C.J. Bernardos, I. Soto, "PMIPv6 multicast handover optimization by the Subscription Information Acquisition through the LMA (SIAL)", RFC 7161, March 2014.

J.C. Zuniga, L.M. Contreras, C.J. Bernardos, S. Jeon, Y. Kim, "Multicast Mobility Routing Optimizations for Proxy Mobile IPv6", RFC 7028, September 2013.

ETSI NFV

ETSI GR NFV-IFA 028, "Report on architecture options to support multiple administrative domains", v3.1.1, January 2018, http://www.etsi.org/deliver/etsi_gr/NFV-IFA/001_099/028/03.01.01_60/gr_NFV-IFA028v030101p.pdf

ETSI GS NFV-EVE 005, "Report on SDN Usage in NFV Architectural Framework", v1.1.1, December 2015, http://www.etsi.org/deliver/etsi_gs/NFV-EVE/001_099/005/01.01.01_60/gs_NFV-EVE005v010101p.pdf

ETSI MEC

ETSI MEC white paper 28, "MEC in 5G networks", June 2018, http://www.etsi.org/images/files/ETSIWhitePapers/etsi_wp28_mec_in_5G_FINAL.pdf

ETSI MEC white paper 23, "Cloud RAN and MEC: A Perfect Pairing", February 2018, https://www.etsi.org/images/files/ETSIWhitePapers/etsi_wp23_MEC_and_CRAN_ed1_FINAL.pdf

ETSI IP6

ETSI IP6 white paper 35, "IPv6 Best Practices, Benefits, Transition Challenges and the Way Forward", August 2020, https://www.etsi.org/images/files/ETSIWhitePapers/etsi_WP35_IPv6_Best_Practices_Benefits_Transition_Challenges_and_the_Way_Forward.pdf

ONF

ONF TR-532 "Microwave Information Model", December 2016, https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-532-Microwave-Information-Model-V1.pdf

ONF TR-534 "Framework and Architecture for the Application of SDN to Carrier Networks", July 2016, https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-534_SDN_Carrier_Grade_Framework.pdf

ONF TR-521 "SDN Architecture – Issue 1.1", January 2016, https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-521_SDN_Architecture_issue_1.1.pdf .

ONF White Paper, "Wireless Transport SDN Proof of Concept 2 Detailed Report", June 2016, available at: https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/Wireless_Transport_SDN_PoC_White_Paper.pdf

ONF White Paper, "Wireless Transport SDN Proof of Concept White Paper", October 2015, available at: https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/ONF_Microwave_SDN_PoC_White_Paper%20v1.0.pdf

Telecom Infra Project

TIP Disaggregated Cell Site Gateway Technical Specification, https://telecominfraproject.com/wp-content/uploads/TIP_OOPT_DCSG_Technical_Specification_v1.1_FINALAPPROVED.pdf

O-RAN

O-RAN, "X-Haul Packet switched architecture and solutions", O-RAN.WG9.XPSAAS-v01.00, November, 2020.

Patents and inventions

"Method and system for transmitting IP traffic in a radiocommunication system". [Online]: https://www.google.com/patents/EP1179918A3?cl=en

"Method and system for handling IT information related to cloud computing services". [Online]: https://www.google.com/patents/WO2014060226A1?cl=en

"System, Method and live streaming optimizer server for live content distribution optimization over a content delivery network". [Online]: https://www.google.com/patents/WO2014090794A1?cl=en

"System and Method to trigger cross-layer optimizations in a network". [Online]: https://www.google.com/patents/EP2770431A1?cl=en

"A method to minimize network and information technology resources consumption in converged networks when provisioning services", [Online]: https://www.google.com/patents/WO2013087535A1?cl=en

"Method and Apparatus for cross-stratum path computation in a network". [Online]: https://patents.google.com/patent/EP2887621B1/en

"A network controller and a computer implemented method for automatically define forwarding rules to configure a computer networking device". [Online]: https://patents.google.com/patent/EP3021534B1/en

International research and innovation projects

The following is a list of funded projects where I have actively participated.

ROOT – Rolling Out OSNMA for the secure synchronization of Telecom networks – Grant agreement 101004261 (https://www.gnss-root.eu/)

- Project objective: ROOT aims to assess the benefits introduced by the Galileo authenticated signals (OSNMA) in the specific context of synchronization of 5G telecommunications networks. Specifically, ROOT purpose is to estimate the increased resilience that OSNMA can bring to GNSS-based timing sources. This is obtained through an experimental approach which enables the measurement of the increased level of robustness provided by the use of a mix of emerging technologies.
- Main role in the project: leader of Work Package 1 on the definition of secure 5G network synchronization architecture and requirements.

5G-DIVE - eDge Intelligence for Vertical Experimentation – Grant agreement 859881 (https://5g-dive.eu/)

- Project objective: 5G-DIVE targets end-to-end 5G trials aimed at proving the technical merits and business value proposition of 5G technologies in two vertical pilots, namely Industry 4.0 and Autonomous Drone Scout. These trials put in action an end-to-end 5G design tailored to the requirements of the applications targeted in each vertical pilot, such as digital twinning and drone fleet navigation applications.
- Main role in the project: leader of Task 1.1 for the analysis of the vertical industry use cases including their business, functional, and technical requirements, as well as the techno-economic analysis of the project.

5GROWTH - 5G-enabled Growth in Vertical Industries – Grant agreement 856709 (https://5growth.eu/)

- Project objective: the goal of 5GROWTH is the technical and business validation of 5G technologies from a vertical's point of view, following a field-trial-based approach on vertical sites. With that purpose, the project defines an AI-driven automated and sharable 5G end-to-end solution allowing vertical industries to simultaneously achieve their respective key performance targets.
- Main role in the project: leader of Task 1.3 on business layer modelling and SLAs.

5G-EVE – 5G European Validation platform for Extensive trials – Grant agreement 815074 (https://www.5g-eve.eu/)

- Project objective: 5G EVE is one of the three European 5G validation infrastructure platform for extensive trials with the goal of implementing and testing advanced 5G vertical services in Europe. 5G-EVE interconnects European sites in in Greece, Spain, France, and Italy.
- Main role in the project: leader of Task 6.2 on standardization activities related to the project outcomes.

NECOS - Novel Enablers for Cloud Slicing – Grant agreement 777067 (http://www.h2020-necos.eu/)

- Project objective: A unique service provider sometimes owns the data centers supporting the cloud, but there is a trend to a fully distributed architecture, where several data center providers will have to cooperate through the interconnection of their computation and storage infrastructure. The interconnection network is a fundamental part in the performance of such a distributed cloud environment. The set of resources that will enable a service to be deployed fulfilling its specific requirements over time gives rise to the Lightweight Slice Defined Cloud (LSDC) concept.
- Main role in the project: leader of Work Package 2 on Service Requirements, covering use cases, system requirements and techno-economic analysis.

5G-Transformer - 5G Mobile Transport Platform for Verticals - Grant agreement 761536 (http://5g-transformer.eu/)

- Project objective: Transformation of today's mobile transport network into an SDN/NFV-based Mobile Transport and Computing Platform (MTP), which brings the "Network Slicing" paradigm into mobile transport networks by provisioning and managing MTP slices tailored to the specific needs of vertical industries.
- Main role in the project: leader of Task 1.3 addressing techno-economic implications of 5G-Transformer solution.

5GEx – 5G-Exchange – Grant agreement 671636 (https://cordis.europa.eu/project/id/671636)

- Project objective: The 5GEx project is creating an agile exchange mechanism for contracting, invoking and settling for the wholesale consumption of resources and virtual network services which can be provisioned in less than 90 minutes and rapidly invoked. This will enable network operators, applications providers and other stakeholders in the 5G supply chain to deliver new service value for 5G customers

and at the same creating and enhancing revenue-generating potential for 5G providers, third party verticals and others in the supply chain
- Main role in the Project: leader of WP2 for Architecture and Use Cases, and leader of Task 5.2 for Standardization

5G-Crosshaul - The 5G Integrated fronthaul / backhaul - Grant agreement 671598 (https://cordis.europa.eu/project/id/671598)

- Project objective: 5G-Crosshaul aims at developing an adaptive, sharable, cost-efficient 5G transport network solution integrating the fronthaul and backhaul segments of the network. This transport network will flexibly interconnect distributed 5G radio access and core network functions, hosted on in-network cloud nodes
- Main role in the Project: leader of Task 1.1 on Use Cases and Requirements

ESA CloudSat - Scenarios for integration of satellite components in future networks - Activity Code: 1C.017 (https://artes.esa.int/projects/cloudsat)

- Project objective: The CloudSat study focuses on the applicability of virtualisation and softwarisation technologies to satcom platforms and determining the benefits and the challenges associated with the integration of satellite infrastructures into future cloud Networks.
- Main role in the Project: leader of WP1 devoted to perform an extensive state-of-the-art survey of emerging virtualization technologies

XIFI - eXperimental Infrastructures for the Future Internet – Grant agreement 604590 (https://cordis.europa.eu/project/id/604590)

- Project objective: The goal of XIFI project aimed to support advanced experiments on the FI-PPP core platform, leveraging on advanced test infrastructures and Future Internet services to cope with large trial deployments involving users. This was achieved through a core federation of test infrastructures, and by coordinating efforts with ongoing FI infrastructures and pilots (FIRE, EIT ICT Labs, CIP pilots, Living Labs) assisted by pan-European infrastructures such as GÉANT.
- Main role in the project: leader of Task 3.1 for SDN control of virtual network infrastructures.

GEYSERS - Generalised architEcture for dYnamic infraStructure sERvices – Grant agreement 248657 (https://cordis.europa.eu/project/id/248657/)

- Project objective: GEYSERS proposed an architecture enabling optical network infrastructure providers to compose logical infrastructures including IT resources and rent them out to network operators by means of integrated control and management techniques.
- Main role in the project: co-leader of WP2 defining the architecture of the project and the interfaces for interacting among the different components proposed.

Research visits

Visiting Researcher (June / July 2017) at the Nakao Research Laboratory, the University of Tokyo, Tokyo, Japan, hosted by Prof. Akihiro Nakao.

Guest Editor and organization of scientific events

Editorship

Guest editor of the Special Issue "Machine Learning in WSN and IoT" in *Journal of Sensor and Actuator Networks* (https://www.mdpi.com/journal/jsan/special_issues/ML_WSN_IoT).

Guest editor of the Feature Topic "5G for Verticals: from Theory to Practice and Beyond" in *IEEE Communications Magazine* - Q1 journal under JCR categories for "Engineering, Electrical & Electronic" and "Telecommunications" (https://www.comsoc.org/publications/magazines/ieee-communications-magazine/cfp/5g-verticals-theory-practice-and-beyond).

Guest editor of the Special Issue "Novel Cloud-based Service/Application Platforms and Ecosystems" in *Electronics* - Q2 journal under JCR category for "Engineering, Electrical & Electronic" (https://www.mdpi.com/journal/electronics/special_issues/Cloud_Service_Platforms).

Organization and chairing of scientific events

Industry chair of the ACM SIGCOMM 2021 Workshop on Network-Application Integration (NAI 2021), virtual, August 2021.

Technical Programme Committee co-chair of the 4th Workshop on Advances in Slicing for Softwarized Infrastructures (S4SI), co-located with the IEEE Conference on Network Softwarization (NetSoft 2021), Tokyo, Japan, 2021.

Local Workshop Organisation co-chair of the 3rd Mobility Support in Slice-based Network Control for Heterogeneous Environments (MOBISLICE), co-located with the 6th IEEE Conference on Network Function Virtualization and Software Defined Networks (IEEE NFV-SDN), Madrid, Spain, 2020.

Co-organizer of the 3rd Workshop on Advances in Slicing for Softwarized Infrastructures (S4SI), co-located with the IEEE Conference on Network Softwarization (NetSoft 2020), Ghent, Belgium, 2020.

Co-organizer of the 2nd Workshop on Advances in Slicing for Softwarized Infrastructures (S4SI), co-located with the IEEE Conference on Network Softwarization (NetSoft 2019), Paris, France, 2019.

Co-organizer of the International Workshop on Cross-Stratum Optimization for Cloud Computing and Distributed Networked Applications, co-located with the 10th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA 2012), Leganés, Spain, 2012.

Lead sponsor and coordinator of the successful first-ever SDN Proof of Concept for Wireless Transport Networks carried out in October 2015 in the framework of the Open Transport WG of ONF, involving the companies Coriant, Ceragon, Huawei, NEC, SIAE and Ericsson, leveraging on ONOS as controller (https://www.opennetworking.org/news-and-events/press-releases/2572-open-networking-foundation-completes-industry-s-first-wireless-transport-sdn-proof-of-concept)

Talks

1. L.M. Contreras, "Transformation of Transport Networks through Softwarization", keynote at the *IEEE International Conference on Network Softwarization (NetSoft)*, Tokyo, Japan (going virtual), June 2021.
2. L.M. Contreras, "Towards B5G – an operator's viewpoint", keynote at the *24th Conference on Innovation in Clouds, Internet and Networks (ICIN '21)*, Paris, France (going virtual), March 2021.
3. L.M. Contreras, "The role of service providers' network for supporting edge computing in 5G and beyond", presentation at the workshop "Post 5G edge cloud computing: why, what and when?", co-located with the *46th European Conference on Optical Communication (ECOC 2020)*, Brussels, Belgium, (going virtual), December 2020.
4. L.M. Contreras, "Transport slicing – ongoing work at IETF with a personal view", keynote at the "Second International Workshop on Network Slicing – Network Slicing 2020", co-located with the *IFIP Networking 2020 Conference (NETWORKING 2020)*, Paris, France (going virtual).
5. L.M. Contreras, "Evolutionary trends in operators' networks for beyond 5G", presentation at the workshop "The role of computing in the post 5G-era: Architectures and enabling technologies", co-located with the *24th International Conference on Optical Network Design and Modelling (ONDM 2020)*, Castelldefels, Spain (going virtual).
6. L.M. Contreras, "Towards a standardized transport slicing architecture in operator networks", keynote at the "5th IEEE International Workshop on Orchestration for Software Defined Infrastructures (O4SDI)", co-located with the *2020 IEEE/IFIP Network Operations and Management Symposium (NOMS 2020)*, Budapest, Hungary, April 2020 (going virtual).
7. L.M. Contreras, "Network 2030 – Implications of the new technologies for an operator", Sixth ITU Workshop on Network 2030, Lisbon, Portugal, January 2020.
8. L.M. Contreras, "Networking the cloud, cloudifying the network", invited special session at the *IEEE International Conference on Cloud Networking (CloudNet)*, Coimbra, Portugal, November 2019.
9. L.M. Contreras, "5G + Cloudification = Slicing", NECOS Industrial Workshop, Campinas, Brazil, October 2019.
10. L.M. Contreras, R. Lafetá, "Lightweight Cloud Slicing: The NECOS Project", 5G Core Summit, Madrid, Spain, September 2019.
11. L.M. Contreras, "Federating MEC and Telco Cloud environments for multi-domain slice provision", in the workshop "Multi-provider, multi-vendor, multi-player orchestration: from distributed cloud to edge and fog environments in 5G", at the *European Conference on Networks and Communications (EuCNC) 2018*, Ljubljana, Slovenia, June 2018.
12. L.M. Contreras, "Ongoing transformations in transport networks: how making all them work", keynote at the *IEEE/IFIP Network Operations and Management Symposium*, Taipei, Taiwan, April 2018.
13. L.M. Contreras, "Telefonica path towards 5G – exploration of new capabilities for transport networks", 39th meeting of the Wireless Word Research Forum, Castelldefels, Spain, October 2017.
14. L.M. Contreras, "Slicing across multiple administrative domains", 39th meeting of the Wireless Word Research Forum , Castelldefels, Spain, October 2017.

15. L.M. Contreras, "Programmability, Virtualization, Automation and Slicing as Foundation of Next 5G Telco Networks", keynote at the *3rd Open International Workshop on Elastic Networks and 5G* organized by the Spanish academic interest group ElasticNetworks, Castelldefels, Spain, October 2017.
16. L.M. Contreras, "The ingredients of the new networks – SDN, NFV and Slicing in the evolution towards 5G", at the Network System Research Institute, National Institute of Information and Communications Technology (NICT), Tokyo, Japan, July 2017.
17. L.M. Contreras, "The ingredients of the new networks – SDN, NFV and Slicing in the evolution towards 5G", at the Nakao Research Laboratory, Daiwa Ubiquitous Computing Research Building, University of Tokyo, Japan, June 2017.
18. L.M. Contreras, "Microwave POC overview and demo", Workshop on OpenDayLight and NFV/SDN Orchestration at the 5TONIC Laboratory, Leganés, Spain, October 2016.
19. L.M. Contreras, "Connecting Multiple SDN/NFV Administrative Domains", 4th Annual Network Virtualization & SDN Europe, Madrid, Spain, June, 2016.
20. L.M. Contreras, "Control Plane for High Capacity Networks", 5th International Workshop on Trends in Optical Technologies, Campinas-São Paulo, Brazil, May, 2016.
21. L.M. Contreras, "Operation, organization and business challenges for network operators in the context of SDN and NFV", FUSECO Forum, Berlin, November 2015.
22. L.M. Contreras, "5G backhauling", FUSECO Forum, Berlin, November 2015
23. P. Congdon, L.M. Contreras, A. de la Oliva, S. Manning, R. Marks, J.C. Zuniga, "Wireless SDN in Access and Backhaul", IEEE 802 Plenary, Dallas, November, 2013.
24. L.M. Contreras, "Hiperlan 2", within the course "21st Century Telecommunications in the Home and Office: WLAN and WPAN networks", XVI Laredo Summer Courses, Universidad de Cantabria, Laredo, Spain, 24th -28th of July 2000.

ACRONYMS

5G	Fifth generation communication system
5GEx	5G-Exchange
ALTO	Application Layer Traffic Optimization
API	Application Programming Interfaces
B2B	Business-to-Business
B2B2C	Business-to-Business-to-Customer
B2C	Business-to-Customer
BSS	Business Support Systems
CapEx	Capital Expenditure
CDN	Content Delivery Network
CFS	Customer Facing Service
CLAS	Cooperative Layered Architecture for SDN
CN	Core Network
CNTT	Common Network Function Virtualization Infrastructure Telecom Taskforce
COTS	Commercial Off-The-Shelf
CSMF	Communications Service Management Function
DC	Data Center
DNS	Domain Name Server
DoS	Denial of Service
DWDM	Dense Wavelength Division Multiplexing
E2E	End-to-End
EEA	European Economic Area
eMBB	Enhanced Mobile Broadband
EoL	End of Life
EoS	End of Support
EPC	Evolved Packet Core
EU	European Union
FCAPS	Fault, Configuration, Accounting, Performance and Security
FlexE	Flexible Ethernet
ForCES	Forwarding and Control Element Separation
FQDN	Fully Qualified Domain Name
GST	Generic network Slice Template
HOT	Heat Orchestration Template
HSS	Home Subscriber Service
IaaS	Infrastructure-as-a-Service
ID	Identifier
IGMP	Internet Group Management Protocol
InP	Infrastructure Provider
IPTV	IP Television
IPX	Internetwork Packet Exchange
ISP	Internet Service Provider
IT	Information Technology
KPI	Key Performance Indicator
LAG	Link Aggregation Group
LTE	Long-Term Evolution
MD	Muti-Domain
MdO	Multi-Domain Orchestrator
MEC	Multi-access Edge Computing
MEH	Multi-access Edge Host
MEO	Multi-access Edge Orchestrator
MEP	Multi-access Edge Platform
MEPM	MEP Manager
MES	Multi-access Edge System
MLD	Multicast Listener Discovery
MM	Mobility Management
MME	MM Entity
mMTC	Massive Machine-Type Communications
MTA	Multi-Tenancy Application
MURT	Multicast Upstream Routing Table
MVNO	Mobile Virtual Network Operator
MW	Microwave
NBI	North-Bound Interface

NDO	Network Domain Orchestrator
NE	Network Equipment / Network Element
NEST	Network Slice Type
NFV	Network Function Virtualization
NFVI	NFV Infrastructure
NFVO	NFV Orchestrator
NGN	Next Generation Network
NMS	Network Management System
NPN	Non-Public Network
NSMF	Network Slice Management Function
NSSMF	Network Slice Subnet Management Function
NSO	Network Service Orchestrator
ODO	OpenStack Domain Orchestrator
OpEx	Operational Expenditure
OS	Open Source
OSS	Operation Support Systems
OTT	Over-The-Top
PaaS	Platform-as-a-Servicee
PCE	Path Computation Element
PDV	Packet Delay Variation
PGW	Packet Data Network Gateway
PIM	Protocol Independent Multicast
PL	Packet Loss
PMO	Present Mode of Operation
PoP	Point of Presence
QoS	Quality of Service
RAN	Radio Access Network
RLAH	Roam Like At Home
RO	Resource Orchestrator
RSL	Received Signal Level
SaaS	System-as-a-Service
SBI	South-Bound Interface
SDN	Software Defined Networking
SDO	Standards Development Organizations
SDTN	Software Defined Transport Network
SGW	Serving Gateway
SLA	Service Level Agreement
SLO	Service Level Objective
SlaaS	Slice-as-a-Service
TCO	Total Cost of Ownership
TE	Traffic Engineering
TSC	Transport Slice Controller
UA-LCM	User Application – Lifecycle Manageement
UE	User Equipment
UHD	Ultra-High Definition
UPF	User Plane Function
uRLLC	Ultra-Reliable Low Latency Communications
vCache	Virtual Cache
VIM	Virtualized Infrastructure Manager
VIMaP	VIM and Planner
VM	Virtual Machine
VNF	Virtual Network Function
VNFM	VNF Manager
VNO	Virtual Network Operator
VoD	Video on Demand
VoIP	Voice over IP
vPGW	Virtual PGW
VPN	Virtual Private Network
WAN	Wide Area Network
WIM	WAN Infrastructure Manager
WTN	Wireless Transport Network
XCI	Crosshaul Control Infrastructure
XFE	Crosshaul Forwarding Element
XPU	Crosshaul Processing Unit
YANG	Yet Another Next Generation (network data model)