

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Mission Critical Communications Support with 5G and Network Slicing

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Mission Critical Communications Support with 5G and Network Slicing / Borsatti D.; Grasselli C.; Contoli C.; Micciullo L.; Spinacci L.; Settembre M.; Cerroni W.; Callegati F.. - In: IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT. - ISSN 1932-4537. - ELETTRONICO. - 20:1(2023), pp. 9899747.595-9899747.607. [10.1109/TNSM.2022.3208657]

Availability:

This version is available at: <https://hdl.handle.net/11585/904768> since: 2022-11-21

Published:

DOI: <http://doi.org/10.1109/TNSM.2022.3208657>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

D. Borsatti *et al.*, "Mission Critical Communications Support With 5G and Network Slicing," in *IEEE Transactions on Network and Service Management*, vol. 20, no. 1, pp. 595-607, March 2023.

The final published version is available online at:

<https://doi.org/10.1109/TNSM.2022.3208657>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Mission Critical Communications Support with 5G and Network Slicing

Davide Borsatti*, Chiara Grasselli[†], Chiara Contoli[‡], Luigia Micciullo[§], Luca Spinacci[¶],
Marina Settembre^{||}, Walter Cerroni* and Franco Callegati[†]

*Department of Electrical, Electronic and Information Engineering - DEI - University of Bologna (Italy)

[†]Department of Computer Science and Engineering - DISI - University of Bologna (Italy)

[‡]Department of Pure and Applied Sciences - DiSPeA - University of Urbino “Carlo Bo” (Italy)

[§]Cyber Security Division, Broadband solutions Engineering, Leonardo S.p.A. (Italy)

[¶]Liguria Digitale S.p.A. (Italy)

^{||}Fondazione Ugo Bordononi (Italy)

Abstract—Mission Critical (MC) communications take a pivotal role to achieve effective Public Protection and Disaster Relief (PPDR) actions. Even though 3GPP standards define MC applications and services in an architectural framework compatible with current 5G mobile networks, real-life experiments and applications of these concepts are still at the very beginning. In this paper, we present an architectural study and related experimental activity on network slicing for MC communications. We implemented these services in a fully virtualized environment, and deployed and tested them in a multi-domain network slicing scenario compliant with the ETSI NFV-MANO specifications. Our work aligns with the 5G approach separating control and data planes. The level of automation in service deployment and the slice isolation features are demonstrated, showing the benefits in terms of application performance, management flexibility, scalability, and quality of service differentiation capabilities.

Index Terms—Mission Critical Communications, MCX, 5G, SDN, NFV, Network Slicing.

I. INTRODUCTION

Public Protection and Disaster Relief (PPDR) forces, such as the police, first response medical teams, and firefighters, rely on dedicated communication networks, which are usually called Mission Critical (MC) networks and play a crucial role in the success of the missions. As discussed in more detail in the following section, mission critical networks have been based, so far, mainly on dedicated infrastructures. However, the current technological trends suggest that they may progressively integrate with the public mobile networks. For this reason, mission critical communications and services were declared a key priority by 3GPP [1], which standardized the support of such services over LTE and 5G.

A key concept of 5G is the capability to serve, effectively and efficiently, vertical applications thanks to the network slicing paradigm [2], which enables a single physical infrastructure to host different logical networks with diverse requirements. Enhanced Mobile Broadband (eMBB), Massive IoT (MIoT), Ultra-Reliable Low Latency Communications (URLLC), and Vehicle to Everything (V2X) are the four

slice types that, to date, have been standardized by 3GPP [3]. Complete isolation between different slice instances needs to be guaranteed. In other words, two slices with distinct characteristics should not influence each other. The network slicing idea relies on the widespread adoption of virtualization technologies, with Virtual Network Functions (VNFs) hosted in cloud-like facilities, as briefly summarized in Figure 1, which shows that diverse VNFs (green boxes) can be used to implement different slices. The slices are isolated and can span multiple data centers, taking advantage of Commercial Off The Shelf (COTS) hardware. With network slicing, telecom infrastructures will undergo a major revolution in design and implementation techniques, and the telco ecosystem will increasingly rely on cloud-based infrastructures (often called telco-cloud). The relevance of this trend motivated initiatives such as the Cloud Infrastructure Telco Task Force (CNTT), which designed guidelines to drive this innovation. Furthermore, ETSI-driven Industry Specification Groups (ISGs), like Network Function Virtualization (NFV) and Multi-access Edge Computing (MEC), are defining frameworks to manage and orchestrate the deployment of VNFs and edge applications capable of accessing network information via standard APIs.

The PPDR communications ecosystem may also greatly benefit from this architectural evolution. A cloud-based network infrastructure promises to be more flexible, with the possibility to scale as needed. In addition, the operator could activate it everywhere COTS hardware is available with enough computing resources. Therefore, these infrastructures would become inherently more agile in case of disruptive events, which is a very relevant and strategic characteristic for mission critical networks. This work aims at demonstrating the feasibility and the effectiveness of running a PPDR communication service as a network slice, satisfying well-defined performance and functional requirements. The network slice must include all the components of the mobile network as well as the 3GPP-compliant MC communications application elements. In this work, the network slice deployment is fully virtualized, based on VNFs, and its lifecycle is managed according to the Network Function Virtualization Management and Orchestration paradigm (NFV-MANO), as defined in the ETSI NFV-MANO specifications [4]. Moreover, following an

This work was performed while Marina Settembre and Luca Spinacci were with Leonardo S.p.A. and Chiara Contoli and Chiara Grasselli were with CIRI-ICT University of Bologna.

Manuscript received February 28, 2022.




Fig. 1. The network slicing concept.

essential design principle of 5G, the control and user planes are fully separated (Control User Plane Separation or CUPS). In other words, all network signaling is logically separated from the user data that will be transported in the network for service implementation.

The contributions of this work are the design of a multi data center network slice architecture, the demonstration of full automation of the network slice deployment by means of the Open Source MANO platform, and the validation of the effectiveness of the implementation in terms of performance of the PPDR communication services. Everything was deployed in a test-bed that will be described later in the manuscript.

The paper is organized as follows: in Section II we provide background information on the considered scenario and discuss some related work on network slicing. In Section III we illustrate the approach chosen for network slice design and lifecycle management. In Section IV we describe our implementation choices and the resulting network slice blueprint and implementation approach. In Section V we report the results obtained from the experimental test-bed and in Section VI we draw some conclusions.

II. BACKGROUND AND RELATED WORK

A. Networks for PPDR forces: state of the art

Communications for PPDR forces represent a key strategic asset. This is well understood by the European Union, which funded the Broadmap project¹ in the framework of the Horizon 2020 Research and Innovation program. The goal was to “collect and validate the PPDR (Public Protection and Disaster Relief) organisations’ existing requirements with the aim to establish a core set of specifications, and roadmap for procurement, to achieve future evolution of EU broadband applications and interoperable radio communication solutions.” State of the art says that PPDR communications rely on dedicated networks and radio access spectrum, with legacy technologies like TERrestrial Trunked RAdio (TETRA), Tetrapol, Digital Mobile Radio (DMR), and Project 25 (P25); all of them are narrowband technologies with reduced capabilities for broadband applications. The organization schemes for the ownership and operation vary from country to country, going from the government-owned, government-operated to the contractor-owned, contractor-operated solution. The intermediate approach of government-owned, contractor-operated

is also adopted. These differences are mainly related to the diverse legal and economic frameworks available in the various countries.

Despite being efficient for narrowband voice communication, the aforementioned technologies provide limited packet data throughput. Moreover, in current PPDR narrowband networks, small messages (e.g., GPS location, short data services, data queries, status) go through the control channel that can be overloaded, causing problems with voice communication. For this reason, in many cases the PPDR forces complement their narrowband services with a mobile broadband offering, typically seeking collaboration from Mobile Network Operators (MNOs). But for many reasons, PPDR networks should implement priorities and pre-emption, which is not possible in public networks since the EU net neutrality regulation states that users should be treated equally [5]. Under these circumstances, the apparently most effective solution for the PPDR organizations would be to implement dedicated broadband networks. However, this is considered prohibitively expensive [6]. Therefore, one outcome of the Broadmap project was that both technical and economic factors suggest going for the next generation of PPDR networks based on 5G technologies. Network slicing promises to offer a solution to the performance and network isolation problem, leveraging a new platform for service delivery defined by 3GPP.

The 3GPP entered the application domain by standardizing Mission Critical services [1]. It started with Push-to-Talk (MCPTT) in Release 13 [7], and then added further MC services and enhancements, such as MCData and MCVideo, in Release 14 [8]. In general, these services can be referred to as Mission Critical Everything (MCX). In Release 15, MC services were enhanced with interconnections between native MC systems and legacy ones, such as TETRA for voice and short data services [9]. With Release 16 [10] and beyond, the 3GPP is addressing the MC services in the context of 5G. The 3GPP MC Specifications define the mission critical requirements in the Technical Specification (TS) 22 series [11], the functional architecture and procedures in the TS 23 series [12], and the MC protocols in the TS 24 series [13]. The 3GPP provides a network architecture view for MC services, describing the required functional entities, the interfaces between them, as well as possible implementation scenarios. To date, the first deployments of MCPTT have already begun, so the 3GPP approach to MC communications will be tested in real life on the infrastructures of MNOs.

The results of the Broadmap project were the basis of the Broadway project², which aims to “Procure Innovation activity to develop and demonstrate TRL8 technologies that will enable a pan-European interoperable broadband mobile system for PPDR, validated by sustainable testing facilities”. Overall, Broadmap and Broadway projects have a twofold objective: on one side, to understand the expectation for future PPDR networks in order to outline a possible roadmap to be implemented by EU member states to reach such goals; on the other side, to understand how to achieve interoperability and

¹<http://www.broadmap.eu>

²<https://www.broadway-info.eu>

integration of services between PPDR networks of different member states.

The general framework aims at leaving complete freedom to EU member states to adopt new solutions at their own pace, but also envisages an overarching system that will provide connectivity and integration at the EU level. When, at some point, all national networks will be aligned with the 3GPP standard, this infrastructure could offer a federated authentication and authorization platform that will allow a public officer of one nation to roam into another nation's network. In the meantime, it should implement all the missing building blocks, acting in practice as a pan-European PPDR network that brings together the countries with a level of integration that depends on the level of deployment of the national networks.

The *PPDR4Europe* consortium was led by Leonardo S.p.A. and grouped operators, manufacturers, and research centers to participate in the Broadway phase 1 and 2. Besides solutions incrementally improving legacy technologies, the consortium aimed at demonstrating the deployment of 3GPP MC services in a fully virtualized environment, together with network slicing support for QoS management and network isolation. This was called the *innovation ecosystem* because still a subject of research and not yet market-ready. In this manuscript, we report the lessons learned and the solutions proposed to achieve such a demonstration.

As discussed later, the architecture presented in this work (Section IV) is in line with the 3GPP approach and envisages a solution in which the functional entities can be located at will, across different data centers.

B. 5G Network Slicing

Network slicing is a concept that has been discussed by several standardization bodies and industry associations, such as the ITU-T [14], the NGMN Alliance [15], the GSMA [2], and the 3GPP [3]. These documents describe the paradigm of network slicing from various points of view, mostly related to the specific focus of the organization that drafted them. In general, the underlying concept is the same: 5G is expected to serve effectively and efficiently vertical applications to promote new business solutions and open new markets to support the massive investments needed for its full deployment [16]. To meet the heterogeneous service requirements of new vertical applications, different network slices with diverse characteristics and performance will be instantiated on the same physical infrastructure [2], [15]. Network slicing will enable operators to serve customers with suitably designed network slices, which may be composed and orchestrated to maximize efficiency and effectiveness [3].

In the last few years, network slicing has been a hot research topic with several published survey papers. In [17] Kaloxylos provides an overview of the network slicing concept as envisioned by 3GPP and discusses opportunities and challenges to its applicability. In [18], Foukas *et al.* provide a framework to collect and compare the existing work on network slicing. Despite the large amount of work reviewed, they outline significant open challenges, such as service composition strategies when using fine-grained virtual network

functions. Furthermore, they express the need for approaches to implement end-to-end slice orchestration that can guarantee specific performance and functionalities.

In [19] Afolabi *et al.* provide a survey about principal concepts and enabling technologies that contribute to the end-to-end network slicing and how network slicing impacts the evolution of 5G networks, focusing in particular on slice orchestration and management. Here, they refer to the network slicing definition proposed by the NGMN Alliance. The idea is to have a common physical infrastructure on top of which multiple self-contained logical networks are built, thus enabling flexibility and integration. In [20] Barakabitze *et al.* present a survey about projects and industrial initiatives that push for 5G network slicing adoption, accelerated by Software Defined Networking (SDN) and Network Function Virtualization (NFV) technologies. They also discuss the management and orchestration of network slices. In [21] Matencio-Escobar *et al.* provide their definition of network slicing, which complements traditional ones found in the literature by extending the concept to the network traffic over the data path on a given network slice. More recently, in [22] Khan *et al.* present an extensive review of the literature on network slicing to design a taxonomy of the network slicing concept, with particular reference to the field of application. Interestingly, mission critical communications do not appear in the list of explored application domains they include in their survey. That is not surprising since, as explained in section II-A, even though the use of 5G and network slicing for this specific application domain is taken for granted, at least at the EU level, the related examples of applications are still limited. To the best of our knowledge, the work described in this paper is the first extensive report on the topic.

In technical terms the various aspects of the problem of implementing a network slicing strategy were explored. To cite just some of the work in the literature, in [23] Abe *et al.* investigate the virtualization of a mobile core network by considering user/control plane separation to lower the impact of large numbers of Machine-to-Machine/IoT terminals on the mobile core network itself. In [24] Schiller *et al.* propose a slice-based 5G architecture together with an NFV-based network store, whose goal is to provide on-the-fly resource reservation, deployment, and slice management that matches end-users demand. In [25] Taleb *et al.* propose PERMIT, a framework able to customize a mobile network by instantiating slices per application, users (or groups), and devices. Several papers also focus on the design of algorithms to optimally select the resources to allocate to the network slice. The concept is similar to virtual network embedding, which is studied, for instance, in [26], and turns into a resource selection problem with specific constraints, as explained, for example, in [27] and [28]. Again, none of these papers specifically consider the application scenario of this manuscript. At the same time, specific resource allocation problems are not the purpose of our work. As we will explain in the following, our goal was to demonstrate the feasibility of a service implementation as a distributed network slice, capable of satisfying some specific performance requirements. We assume that the network operator (or operators) that will manage the network slice setup

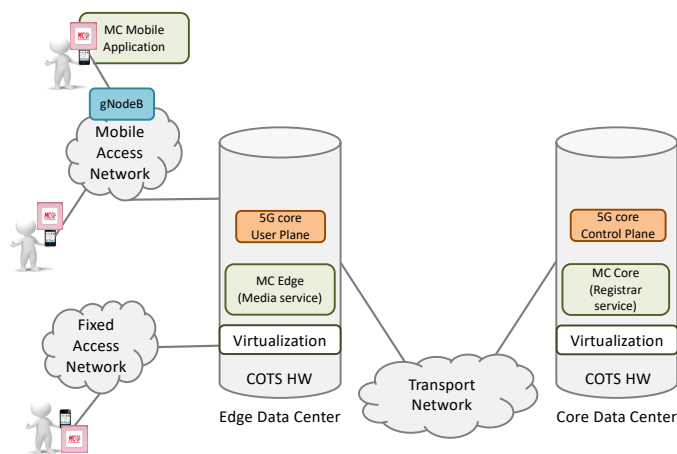


Fig. 2. The high-level architecture of the network slice for the MC services with the related main components.

will apply the optimization and resource allocation strategies of choice. What is important for the goal of the work here reported is that the network slice may be set up and that its design allows full support of the required functionality.

III. NETWORK SLICING FOR MC APPLICATIONS

The network slice considered in this work is split into four logical sections:

- 1) access network, either mobile or fixed;
- 2) edge network components located in an Edge Data Center (DC), virtualizing the user plane part of the mobile core network and the MC proxies for the exchange of the media data flows (voice, video, etc.);
- 3) core network components located in a Core Data Center, virtualizing the control plane part of both the mobile core network and of the MC system;
- 4) interconnection network between the data centers, which could be either a public network or a private geographical interconnection.

Figure 2 shows a high-level view of the network slice. The implementation poses some interesting challenges, most notably:

- multi data center and possibly multi-domain orchestration;
- traffic management for QoS guarantee in the transport network;
- cross data center applications and traffic management.

These challenges require a network slice design properly tailored to support them, as we will discuss later. Before doing that, it is important to recall some basic concepts that will be used extensively in the remainder of this manuscript.

A. Actors and Roles

Network slicing is an approach that involves three main types of actors.

- Infrastructure Providers or infrastructure owners: they own the infrastructure and provide all the infrastructural management actions. A single network slice may span

multiple infrastructure domains. Therefore, its deployment and lifecycle management require interactions with each Infrastructure Provider involved.

- Network Slice Provider: the provider of the communication service implemented with the network slice.
- Network Slice Customers: the users of the communication service.

According to their respective roles, these actors must have different rights, with Infrastructure Providers and Network Slice Provider having specific management roles to keep the infrastructure and the service up and running.

In our case, an Infrastructure Provider can be identified as a mobile network operator, a mobile virtual network operator, or a public body operating the infrastructure for the PPDR forces. At the same time, the Network Slice Provider is the entity that directly manages the mission critical communication services. Depending on the organizational model chosen, this could be a public body serving all the various PPDR forces of the country or a specific body inside a PPDR force (e.g., police, firefighters, etc.). Consequently, the Network Slice Customers are the PPDR forces that will use the service for communication. Moreover, in the framework of the architecture considered in the Broadmap and Broadway projects, the Network Slice Provider could be a body at the European level, which provides an overarching service to integrate the MC communications from different countries in the case of operations spanning across borders.

The very brief discussion above makes clear that the organizations acting as Infrastructure Providers and Network Slice Providers might be different from case to case, either being very closely bound to each other or just linked by a conventional commercial agreement. Therefore, to adapt to these diverse organizational models the slice architecture must be very flexible, allowing a seamless co-existence of these actors, while providing to all of them the required functionalities. For example, management is obviously an important issue for both infrastructure and service providers, since no service can be properly set up or guaranteed in real production environments without management capabilities. This is well understood and explained in the NFV-MANO architecture, where the management components are very clearly outlined. Specifically, we assumed that:

- 1) the Infrastructure Provider must have management access to the whole infrastructure, including all the VNFs, to be able to interact with the various components active in the cluster whenever some high-level general configurations or recovery actions are needed;
- 2) the Network Slice Provider must have management access to its own infrastructure and VNFs to implement all the management actions related to the production phase of the service, including modification of the VNF configurations, performance monitoring, etc.

B. Network Slice Description

A correct interaction between all the actors mentioned above has to be guaranteed. The GSMA standard specifies how to describe the characteristics of each network slice in a

TABLE I
EXAMPLE OF CHARACTERIZING NEST PARAMETERS FOR THE MC COMMUNICATIONS NETWORK SLICE

ATTRIBUTE	VALUE
Coverage	Local (Outdoor)
Guaranteed Downlink Throughput per Network Slice	391600 (391.6Mbps, band 3, channel 20MHz(100RB), 256QAM, 4x4MIMO)
Mission Critical Support	1: mission critical
+ Mission Critical Capability Support	1: Inter-user prioritization, 2: Pre-emption, 3: Local control
+ Mission Critical Service Support	1: MCPTT, 2: MCDData, 3: MCVideo




Fig. 3. The network slice blueprint for a single data center.

standardized way, starting with the Generic Slice Template (GST) [29].

The GST can be used to describe a network slice type. It is a dictionary containing common slice attributes, such as supported throughput/functionality and provided Application Programming Interfaces. Once the GST is filled with values based on specific vertical use cases, it gives birth to the Network Slice Type (NEST), which can be used by vendors, vertical industry customers, and network operators to reach their objectives. Table I shows an example of a NEST for the typical network slice considered in this work. Once the NEST is available, it has to be translated into a description that allows the real-life implementation of the network slice. Neither the GST nor the NEST specifies the steps required to achieve this. The collection of all the technical details that are necessary to implement a particular network slice is usually referred to as the slice *blueprint*. This description depends on the technological approach taken by each Infrastructure Provider and it is not standardized.

In the following section, we will describe the general slice blueprint that we designed to match the high-level implementation goals and the NEST specifications.

C. Network Slice Blueprint

At first, we designed and tested a network slice blueprint that was quite general, to be sure it was suitable to meet all the requirements mentioned above in terms of architecture, role splitting, and performance characteristics. We started by considering a single data center and designed the blueprint plotted in Figure 3. In the figure, the horizontal or vertical bars represent virtual networks defined in the data center, whereas the computer icons represent VNFs. This blueprint aims at satisfying the following characteristics:

- separate management networks for both the Infrastructure Provider and the Network Slice Provider;
- isolation and protection of the VNFs providing the required functionalities, avoiding the direct exposure of their network interfaces to external networks;
- maximum flexibility of interconnection between the VNFs composing the service.

We introduced two separate management networks since the Infrastructure Provider must be able to talk to all its customers (tenants) at once, whereas the Network Slice Provider, acting as a tenant of the Infrastructure Provider, must be able to talk to its dedicated infrastructure only, isolated from those of other tenants. Therefore, two different management networks were implemented:

- the *infrastructure management network*, set up at system start-up, devoted to the Infrastructure Provider, and shared among all tenants;
- the *tenant management network*, set up as part of the network slice, seen only by the Network Slice Provider running the slice.

This general architecture can be composed to create network slices spanning across multiple data centers, according to the schematic presented in Figure 4. These data centers might belong to the same provider or different ones. Regardless of that, this should be transparent from the Network Slice Provider's point of view, given the existence of the interfaces required for these interactions. The basic idea of this design is the following: production VNFs run inside the data center, connected to two different management networks, the former devoted to the Infrastructure Provider and the latter to the Network Slice Provider. Moreover, a slice-specific management console connected to the management network is provided to the Network Slice Provider, thanks to which it can manage the slice components directly from the data center where they are deployed.

The VNFs of a slice section, like the one depicted in Figure 3, are not directly connected to the data center networks providing access to the outside world, but there are gateways in between. This choice is motivated by two main reasons:

- security: the network gateway provides the required traffic isolation and acts like a firewall protecting the production section of the slice;
- functionality: the network gateways can work as endpoints of tunnels (in this example VXLAN tunnels) providing an overlay network between the involved data centers, thus allowing seamless slice management over different sections, even belonging to distinct providers.

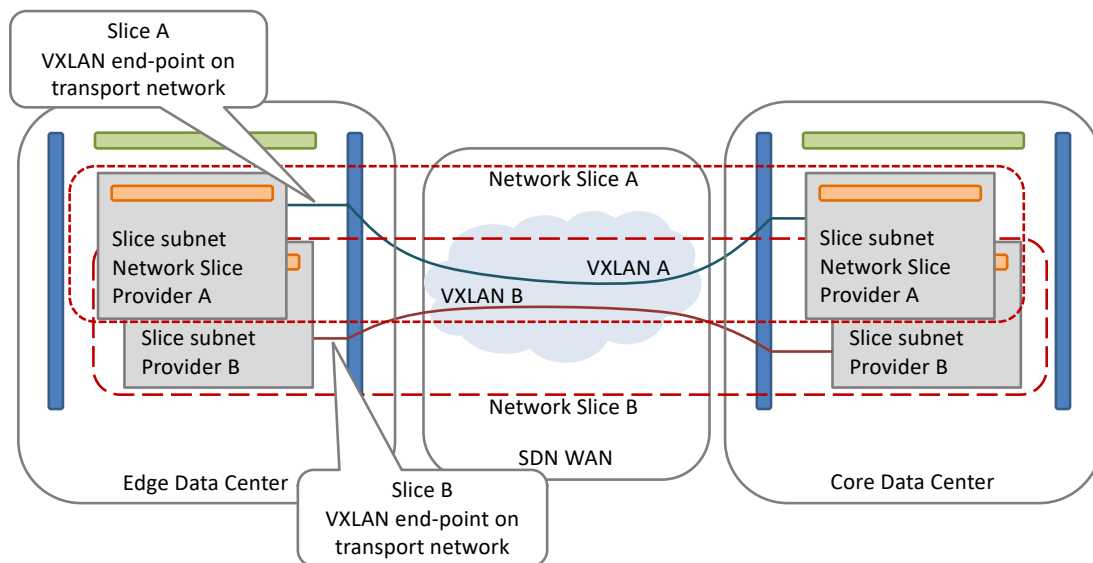


Fig. 4. Full network slice spanning two data centers. In this figure an example with two slices deployed in parallel is shown to provide a better understanding of the different management infrastructures for Infrastructure Provider and Network Slice Provider.

D. Network Slice Delivery and Lifecycle Management

The management and orchestration of the network slice lifecycle is a crucial issue to allow effective usage of this technical approach. 3GPP and ETSI describe the lifecycle of a network slice in their documents [30] and [31], well identifying all the steps required to provide performance requirements for network slice design, as well as to instantiate, run and terminate it. The steps implementing the complete network slice lifecycle management are depicted in Figure 4.3.1.1 in [30]. The figure outlines that there are two main phases; a *preparation phase* containing the description of the network slice blueprint and the preparation of its run-time environment, and a *lifecycle management phase* where the network slice instance is created, run, and eventually terminated. The former phase is indeed a matter of the Infrastructure Providers, which will prepare all the necessary components based on the NEST provided by the Network Slice Provider and the chosen blueprint. In our case, it refers to the networks in the cloud platform that must be shared between slices and must exist before the single network slice instance is started. In particular, these are:

- the management network of the Infrastructure Provider, which will be connected to the parts of the network slice that the provider has to control to handle some emergency event (either collaborating with or overriding the management actions from the Network Slice Provider);
- the inter-DC interconnection network;
- the physical interconnection to the access networks, either mobile or fixed.

Furthermore, the NEST and the slice blueprint are translated to a set of Network Slice Templates and/or Network Service Descriptors, which are then onboarded in the orchestrator platform. These descriptive files represent the list of VNFs and their interconnections for each slice segment (e.g., for each Infrastructure Provider domain), adopting a language

understandable by the NFV-MANO system. On the other hand, the latter phase involves the Network Slice Provider that can start, run, modify, monitor, and stop the network slice at will, using the interfaces provided by the Infrastructure Providers, or through the native interfaces of the applications deployed in it.

The test-bed described in this work follows this paradigm, as described later. But before its description, it is relevant to introduce an approximation we adopted for the test-bed realization. To simplify the deployment process, we imagined having a single Infrastructure Provider offering two data centers, one at the edge and the other at the core. However, the same considerations made both in the blueprint description and slice lifecycle management hold.

IV. IMPLEMENTATION APPROACH AND COMPONENTS

In this section we describe the approach employed to build the slice and the system supporting it, following the general description given in the previous parts of the paper. To this end, we will also introduce the software tools we chose to build the proposed system.

To support the performance requirements of the service, the network slice is split into access and core parts, the former hosted in the Edge DC and the latter in the Core DC. The actual implementation of the access and core parts of the network slice are plotted respectively in Figures 5 and 6. For the sake of readability, the connections of the various VNFs with the Infrastructure Provider and Network Slice Provider management networks (green bar at the bottom and orange bar at the top, respectively) are omitted, but they follow the general blueprint template in Figure 3. The basic idea is to keep everything user-related as close as possible to the user itself. This should minimize the load in the network core and maximize the performance (e.g., by reducing the latency) for the end-users. It is also in line with the Control User Plane Separation principle.




Fig. 5. Architecture of the access section of the network slice in the Edge DC.

The VNFs of the access section in Figure 5 are the User Plane Function (the packet forwarder for the 5G data plane) and the MCX edge component acting as a media server, forwarding media streams from and to users. The core section in Figure 6 is simpler since there is no “transit” traffic and a single internal interconnection network is enough. The specific VNFs are all the components of the 5G control plane and the MCX control element, which acts as a registrar server for the MC applications, managing user registrations and their communication profiles.

It is worth noting that the slice architectures presented here are just a graphical sketch. In practice, following the NFV-MANO architecture [32], each of the VNFs is actually deployed as a pair of virtual machines: the former for production and the latter for management, with an additional network in between to connect them. Firstly, we introduce the tools for managing the infrastructure supporting the architecture proposed (i.e., the tools of the Infrastructure Provider). Afterward, we present the software components running inside the VNFs of the slice. Following the directives proposed by the CNTT group, we chose OpenStack [33] as the cloud management platform for the two data centers. Specifically, we deployed the Stein release with Kolla Ansible in both data centers. Then, to orchestrate the virtual functions of the slice over these virtualization infrastructures, we opted for Open Source MANO (OSM) [34]. OSM is an open-source project backed by ETSI, realizing a standard-compliant NFV-MANO platform. For this work, we used OSM Release 10 with descriptors following the ETSI SOL006 specifications [35].

Finally, to emulate the behavior of the transport network between the two data centers, we introduced emulated delays in the outgoing interfaces connected to this network. Furthermore, we added a single SDN-enabled Open vSwitch [36] switch controlled by the ONOS [37] controller acting as the WAN Infrastructure Manager (WIM).

Following the slice lifecycle presented in Section III-D, after the network environment preparation phase, we designed and onboarded the descriptors and configuration files required by the NFV-MANO platform. The implementation of the network slice considered in this work is rather complex; for this reason, the deployment was split into three steps to make the configuration and debugging process more controllable.




Fig. 6. Architecture of the control plane section of the network slice in the Core DC.

Three NFV-MANO descriptors have been designed for the two sections of the slice. Descriptors can be reused and some are common to the various slice sections. Therefore, the network slice blueprint is represented by the complete set of these descriptors and related configuration files (called packages).

Specifically, the descriptors provide to the NFV-MANO platform all the information about:

- 1) the VNF packages to be run in the slice;
- 2) the interconnections between them (Virtual Links in NFV-MANO terminology), described in the Network Service Descriptors (NSDs) and Virtual Link Descriptors (VLDs);
- 3) the Network Slice Template (NST) as a combination of Network Service Descriptors;
- 4) the details of the Virtualized Infrastructure Managers where the network slice has to be instantiated;
- 5) the VNF Forwarding Graph Descriptor (VNFFGD), specifying the traffic path from one VNF to another, which has to be implemented in the network slice.

Recalling the blueprint description (Section III-C), the Infrastructure Provider management network is already part of the cloud environment, even before the deployment of the slice. Therefore, the first step of the slice deployment phase is the creation of the Network Slice Provider management elements. This initialization deploys the Network Slice Provider management network and the management VNF, which is connected to the data center external networks and the Network Slice Provider management network. Furthermore, the Network Slice Provider management VNF can automatically create an overlay network (e.g., VXLAN) on top of the inter-DC network, allowing a seamless interconnection between components deployed in different data centers. In the second step, the NFV-MANO triggers the deployment of the 5G Core Network elements based on the Open5GS [38] software package. It is an open-source implementation of a hybrid 4G/5G Core Network, compliant with 3GPP Release 16. Finally, the orchestrator instantiates the control and data plane elements of the MCX application provided by Leonardo S.p.A.. More details on this will be given in the remainder of this section.

A. Mission Critical Components

In the experiment reported in this work, the MC services were deployed with a product provided by Leonardo S.p.A. The Leonardo Mission Critical Services is part of the Leonardo Communications Service Platform product family [39]. It extends the portfolio of standard solutions for PPDR communications, ranging from DMR to TETRA technologies, with next-generation broadband capabilities. It is a complete Mission Critical solution compliant with 3GPP MCX standards. It offers Push-to-Talk communication, enhanced with voice, video, multimedia chat, and a set of APIs for third-party application development. It can be deployed over both commercial and private mobile networks and can provide users with advanced functionalities as:

- instantaneous group and private high-quality voice communications;
- mobile broadband multimedia applications (real-time video streaming, multimedia messaging, file/video/photo transfer, database access);
- location-based services;
- emergency, man-down/immobility and Land Mobile Radio standard interaction via InterWorking Function for augmentation of traditional systems [40].

The full solution to provide MCX services is made of the following components:

- An Android client designed for on-field operations with a complete set of functionalities providing all the MC service implementations as per the 3GPP standard, namely MCData, MCVoice, and MCVideo. It can be installed in off-the-shelf smartphones, as well as on ad-hoc terminals, with a fully customizable Human Machine Interface. It can be customized to provide differentiated User Experience, ranging from a traditional push-to-talk radio to a multimedia client similar to a conventional smartphone.
- A web-based dispatcher, providing control, monitoring, and management of the operations of the teams.
- A dedicated interface for the management and monitoring of the platform KPIs.
- A Session Initiation Protocol (SIP) Core server for user registration, location and authorization, as well as for call signaling management as per the 3GPP standard.

The SIP core is a cloud-native platform designed to be deployed either as a virtual machine or as a containerized application. It also supports a full separation of user and control planes, according to the aforementioned 5G Control User Plane separation principle. In particular, the registration server used for signaling can be decoupled from the media servers, which will manage and deliver the media streams. Moreover, the internal SIP Core component may be easily plugged “in” and “out” at run-time with just a few mouse clicks by using the MCX dashboard. External IP Multimedia Subsystem (IMS) core servers are supported for large-scale deployments.

In this experiment, we took advantage of the control and user plane separation offered by the mission critical components, by deploying them in the core and edge data centers, respectively. In detail, we refer to “MCX Core” in Fig. 7 for

the web-based dispatcher and the SIP components in charge of registration and signaling, while to “MCX Edge” in Fig. 6 for the external media server used to exchange users’ voice or video messages. This choice allows us to keep the media servers as close as possible to the final users, thus guaranteeing optimal performance in line with the 5G edge computing concepts.

V. EXPERIMENTAL EVALUATION AND NUMERICAL RESULTS

All the experiments reported here were run in a private data center with two separate OpenStack clusters, one for the Edge DC and the other for the Core DC. Each one of them is composed of two physical servers, equipped as follows: 64 GB of RAM; 40 CPUs; 1.2 TB of disk; 1 Gbit/sec interfaces; Ubuntu 18 LTS as OS. We emulated the 5G Radio Access Network (RAN) elements with UERANSIM [41], both the gNodeB base station and the 5G User Equipments (UEs). This scenario is a limited laboratory infrastructure, thus the presented results must not be considered absolute performance values. It is rather obvious that a more powerful setup will likely lead to better performance overall. Nonetheless, the goal of our work was not to assess the absolute optimal performance achievable with this approach. We intended to show its feasibility and provide insight into the relative performance issues when comparing the various phases of the complete slice deployment to identify possible critical bottlenecks.

A. Network Slice Instantiation

At first, we report results about the time needed to create the virtual infrastructure presented in this paper. In particular, we measured the time required to instantiate all the network slice components. As explained in the previous section, the process is split into three phases:

- 1) initialization of the Network Slice Provider management infrastructure in the data centers, with the dedicated management network and consoles;
- 2) deployment of the 5G mobile core network and gateways, with all the required components split among the edge and the core data centers;
- 3) provisioning of the service to the Network Slice Customers, running the MCX edge and core components respectively in the edge and core data centers.

For each phase, ETSI MANO SOL006 [35] compliant descriptors were implemented and onboarded in the OSM platform for the deployment in each OpenStack cluster.

As reported in Table II, in terms of virtual components, in the first phase, we instantiate 1 VNF on a single virtual machine and 1 virtual network per data center for the Network Slice Provider management infrastructure. During the second phase, we activate in the Edge DC 3 VNFs running on 6 virtual machines (one for the management and one for the service functionalities) and 5 virtual networks for the 5G mobile core and the gateways components at the edge. Instead, in the Core DC, we instantiate 5 VNFs on 10 virtual machines and 6 virtual networks for the other slice components at the core.

TABLE II
NUMBER OF VIRTUAL COMPONENTS OF THE NETWORK SLICE
INSTANTIATED IN EACH PHASE.

	VNFs	Virtual Machines	Virtual Networks
Init. Core Data Center	1	1	1
Init. Edge Data Center	1	1	1
5G core net. Core Data Center	5	10	6
5G core net. Edge Data Center	3	6	5
MCX Core Data Center	1	1	0
MCX Edge Data Center	1	1	0

TABLE III
TIME REQUIRED BY THE ORCHESTRATION SYSTEM TO INSTANTIATE THE
VARIOUS COMPONENTS OF THE NETWORK SLICE. THE VALUES ARE
AVERAGED OVER 10 DIFFERENT EXPERIMENTS. THE 95% CONFIDENCE
INTERVAL IS REPORTED ALONG WITH THE ESTIMATED AVERAGE.

	Average	Min (95%)	Max (95%)
Init. Core Data Center	49.9 s	45.72 s	54.08 s
Init. Edge Data Center	63.6 s	59.95 s	67.25 s
5G core net. Core Data Center	404.8 s	394.36 s	415.24 s
5G core net. Edge Data Center	273.3 s	267.41 s	279.20 s
MCX Core Data Center	75.5 s	67.23 s	83.77 s
MCX Edge Data Center	64.9 s	58.95 s	70.85 s

In the third phase, we run 1 VNF on a single virtual machine per data center for the MCX services.

We performed ten instantiation experiments, measuring the time required to complete the various deployment phases. Table III reports the average time needed to instantiate the network slice components at every step and the related 95% confidence level. The second phase is more complex than the others since it involves a higher number of virtual components. So, as expected, it takes more time for deployment. The other two phases are of similar complexity and involve a smaller number of components than the second phase. Therefore their instantiation requires a shorter time.

The variability of the measured values is due to the fact that we run the experiments on physical servers in a realistic cloud environment. Even though there are no other active network slices, the servers still run the basic management tasks required by the cloud management platform. These tasks share the CPU with all the others and introduce some random delay in the execution of the slice instantiation. Intuitively this sort of “white noise” in the measurements should affect more the short tasks, while it should be less evident in longer ones. That is what happens in practice. It is possible to see in Table III that the confidence interval is between 15% and 20% for the phases that require less time, while it is around 5% for the second phase, which takes more time.

In Figure 7 we also graphically report the entire time needed to instantiate the edge and core network slice sections in each of the ten experiments and their averages. The figure shows that the complete network slice instance can be deployed in a few minutes. There is some variability, as discussed above, but the reported averages provide a rather clear indication of




Fig. 7. Total time required to instantiate the full network slice. The plot shows 10 measurements taken from 10 different experiments on the same infrastructure. The horizontal dashed lines represent the average values.

the values we are facing. According to the scenarios explained in Section II-A, the PPDR communication infrastructure supported by the considered network slice should run for a reasonably long time, from a few days in the case of an ad-hoc deployment because of a specific emergency to months or even years for a stable deployment at the national or European level. Therefore, an initialization time of the whole slice within a few minutes is reasonable. This is also true in case of failure or infrastructure disruption. If the images of the various virtual network functions are available, the entire network slice can be restarted in a reasonably small amount of time at a different location.

B. MCX Service Delivery

The Mission Critical service delivery scenario and the paths of the various traffic flows are shown in Figure 8. From the SIP Uniform Resource Identifier (URI) point of view, the domain is called `test` and two UEs are registered as `user1@test` and `user2@test`. As discussed, the test-bed guarantees separation between the control and data plane. That is true at the 5G level, as the standard implies, but also at the MC service level since the core MCX server is dedicated to handling signaling traffic, such as SIP registration and call set-up messages, while the edge MCX server acts only as a media server. Therefore the signaling for service and call set-up follows a different path than the data flows carrying the communication media streams.

Coming to the experiments, we tested at first the correct functional splitting of roles of the two MCX servers according to the planned split of workloads. Figure 9 shows the flow of an MCVIDEO call from the point of view of the caller (`user1@test 10.250.123.101`) and of the callee (`user2@test 10.250.123.102`). The two packet sequences shown in the figure were obtained by capturing the packet traffic with Wireshark. The core MCX is located at `10.250.2.249`, while the edge MCX is located at `10.250.2.35`. The figure shows that the split of roles is correctly implemented. As planned, the SIP traffic required to




Fig. 8. Overall scenario and an example of the data flows showing the separation between control and data plane of the 5G network as well as of the MCX infrastructure.

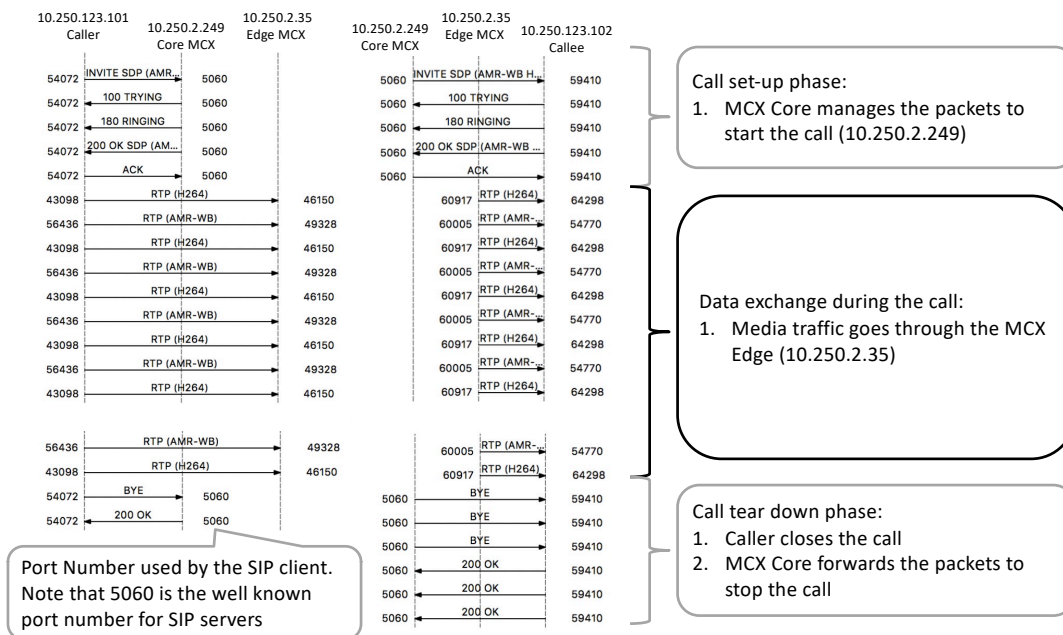


Fig. 9. Packet flows of an MCXVIDEO call, capturing the traffic on the caller (10.250.123.101) and on the callee (10.250.123.102).

set up and close the multimedia call between the two users goes through the core MCX. All SIP signaling messages such as INVITE, TRYING, and RINGING flow to and from the core MCX server (10.250.2.249). Instead, the Real-time Transport Protocol (RTP) media traffic flows to and from the edge MCX server (10.250.2.35). Then to prove the effectiveness of the Control User Plane Separation approach, we exploited a built-in feature of the MC mobile app. This feature provides a series of evaluation tools for measuring network latency and capacity, as shown in Figure 10 and Figure 11.

To emulate a greater latency when connecting to the core infrastructure, we forced a delay of $T = 200$ ms on the inter-DC connection by setting up the Linux traffic control on the interface towards the transport network of the Core DC network gateway. We asked the app to register on both the MCX core and the MCX edge. Obviously, the MCX core is

the only one that allows the registration of a SIP user since it is the only one running the control functions. When we ask the MC app to register on the MCX edge, acting only as a media server, the registration is not successful, but the app still allows the execution of the performance test, even though in a limited way. As a consequence, the two screenshots are different. For the scope of this research, the relevant fields to compare are: 2. CONNECT TCP and 3. HTTP PING. The former reports the time required to complete the three-way handshake of TCP between the Android application user and the MCX server. The latter reports the time taken to complete an HTTP request from the user to the server. The values obtained depend on the Round Trip Time (RTT) of the data connection. We can see that both fields are approximately 200 ms larger in the connection towards the MCX core than to the MCX edge. That is perfectly in line with the additional latency introduced in the path towards the Core DC, which is

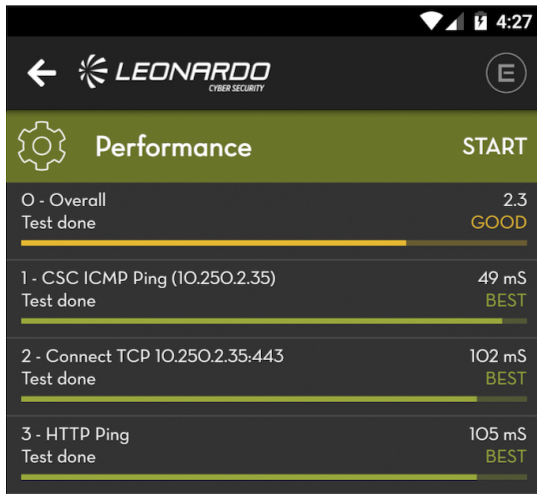


Fig. 10. Screenshot of the MC smartphone application executing performance measurements while communicating with the MCX in the edge.




Fig. 11. Screenshot of the MC smartphone application executing performance measurements while communicating with the MCX in the core.

200 ms in this experiment. Therefore, we can conclude that in the case of a real call the RTT of the media flows (voice and video) would be significantly lower than the RTT of the signaling towards the MCX in the core. This is one of the expected advantages of the CUPS approach.

C. Inter-DC Quality of Service Management

The QoS management in the data centers interconnection network will depend on the features made available by the network owner and/or provider. In this work, we assumed Software Defined Networking capabilities to test the possibility of managing the QoS in an integrated way with the network slice management.

The scenario considered is again the one sketched in Figure 4:

- network slices *A* and *B* are deployed, serving customers *C1* and *C2* respectively;
- the network slices are split in two sections and share an interconnection link at 10 Gbit/s;




Fig. 12. Guaranteed bandwidth per slice interconnection when a background traffic able to saturate the link is applied. The background traffic starts when the inter slice traffic flows are already established.

- the transport link determines the end-to-end bandwidth availability (given that inter-VNF bandwidth inside the same data center is typically larger);
- *C1* and *C2* negotiated the following service level agreements:
 - Slice *A* minimum guaranteed interconnection capacity $C_A = 1$ Gbit/s
 - Slice *B* minimum guaranteed interconnection capacity $C_B = 3$ Gbit/s

The interconnection network is emulated as a virtual switch (implemented with Open vSwitch) controlled by ONOS SDN controller. In the switch, token bucket queues at the minimum guaranteed capacity of the network slices are implemented with a higher priority over a standard FIFO queue used by other traffic flows. Traffic forwarding rules were set by exploiting the ONOS intent framework [42], forcing the switch to push the packets of the two slices into their specific queue. With reference to the slice lifecycle presented in Section III-D, these forwarding rules can be prepared during the network slice design and can be instantiated when the network slice is created. The QoS control is reactive and safeguards the minimum required bandwidth of the slices when a network overload happens. We forced these overload events by generating a very high bandwidth background traffic into the interconnection link.

In Figures 12 and 13 we show results proving the effectiveness of the QoS management strategy. In both cases, 2 minutes (120 seconds) of communication are shown. Slices *A* and *B* generate traffic trying to saturate the available bandwidth. The greedy background traffic causing the link overload is at 10Gbit/s and causes link congestion with high traffic losses in two different ways, as described below. In Figure 12, the background traffic starts when the one generated by slices *A* and *B* is already active with an almost even share between them. When the background traffic starts and congestion arises the traffic control strategy safeguards the minimum guaranteed bandwidth requested by the two network slices. After 90 seconds, the background traffic stops, and the two network




Fig. 13. Guaranteed bandwidth per slice interconnection when a background traffic able to saturate the link is applied. The inter slice traffic flows are started sequentially when the background traffic is already established.

slices can take over. Finally, when slice *B* stops, slice *A* can consume the whole link capacity. In Figure 13 the background traffic is already active and saturates the link for the complete duration of the experiment. Instead, slice *B* and slice *A* start generating traffic at about 10s and 20s, respectively. As before, the traffic control strategy throttles the background traffic to enforce the minimum capacity required by the two slices.

VI. CONCLUSION

In this manuscript, we have addressed the problem of designing network architectures for Public Protection and Disaster Relief forces, following the 3GPP standards for Mission Critical communications and adopting the Network Function Virtualization technology. The goal of the paper is to design and demonstrate the implementation of a network slice including 5G core network elements and Mission Critical communication services, to support and integrate the communication of PPDR forces in different European countries.

We addressed both qualitatively and quantitatively two main principles of 5G and network slicing, namely the separation between the control and data plane and the distribution of the key network slice components where they best fit the purpose. The manuscript describes the design principles of a network slice architecture, which is split into a core and an edge section. The edge section is designed to keep the media servers as close to the user as possible to provide optimal communications performance.

The manuscript reports the results of the practical implementation of the network slice deployment and management, automated according to the ETSI NFV MANO standard. We have shown that the entire network slice can be instantiated in a few minutes, thus allowing maximum flexibility for infrastructure deployment in case of emergency events. Moreover, we have shown in practice that the deployment of the media servers in the edge provides faster access to the Mission Critical services and low latency communications. Finally, we have also shown that the quality of service of the interconnection between the network slice sections can

be managed in an integrated way with the slice management thanks to Software Defined Networking control.

Overall the manuscript demonstrates the feasibility and the effectiveness of the architectural approach proposed for future PPDR networks and for their integration at the European level.

ACKNOWLEDGMENT

This work was partially funded in the framework of the PPDR4EU pre-competitive tender.

REFERENCES

- [1] Mission Critical Services in 3GPP. Accessed: Jun. 28, 2022. [Online]. Available: https://www.3gpp.org/news-events/1875-mc_services
- [2] "An Introduction to Network Slicing," GSMA, Tech. Rep., Sep. 2020.
- [3] "5G; System architecture for the 5G System (5GS)," 3GPP, Tech. Spec. 23.501 version 16.6.0, Oct. 2020.
- [4] "Network Functions Virtualisation (NFV); Management and Orchestration, Report on Management and Orchestration Framework," ETSI, Group Rep. NFV-MAN 001 version 1.2.1, Dec. 2021.
- [5] Regulation (EU) 2015/2120 of the European Parliament and of the Council. Accessed: Jun. 28, 2022. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32015R2120&from=EN>
- [6] D. Lund (ed.), "EU Interoperable Broadband Communication Applications and Technology for Public Safety: Final Definition of the Transition Roadmap and PCP Specification," Broadmap project deliverable D5.2, Apr. 2017. [Online]. Available: <http://www.broadmap.eu>
- [7] "Release 13 analytical view," 3GPP, RP-151569, Sep. 2015.
- [8] "Technical Specification Group Services and System Aspects; Release 14 Description; Summary of Rel-14 Work Items," 3GPP, Tech. Rep. 21.914 version 14.0.0, May 2018.
- [9] "Technical Specification Group Services and System Aspects; Release 15 Description; Summary of Rel-15 Work Items," 3GPP, Tech. Rep. 21.915 version 15.0.0, Sep. 2019.
- [10] "Technical Specification Group Services and System Aspects; Release 16 Description; Summary of Rel-16 Work Items," 3GPP, Tech. Rep. 21.916 version 16.2.0, Jun. 2022.
- [11] 3GPP Specification series 22. Accessed: Jul. 15, 2022. [Online]. Available: <https://www.3gpp.org/DynaReport/22-series.htm>
- [12] 3GPP Specification series 23. Accessed: Jul. 15, 2022. [Online]. Available: <https://www.3gpp.org/DynaReport/23-series.htm>
- [13] 3GPP Specification series 24. Accessed: Jul. 15, 2022. [Online]. Available: <https://www.3gpp.org/dynareport/24-series.htm>
- [14] "Framework of network virtualization for future networks," ITU-T, Standard Series Y: Global information infrastructure, Internet protocol aspects and next-generation networks, Jan. 2012.
- [15] "Description of Network Slicing Concept," NGMN Alliance, Deliverable version 1.0.8, Sep. 2016.
- [16] E. Obiodu and M. Giles, "The 5G era: Age of boundless connectivity and intelligent automation," GSMA, Tech. Rep., 2017.
- [17] A. Kaloxylas, "A Survey and an Analysis of Network Slicing in 5G Networks," *IEEE Communications Standards Magazine*, vol. 2, no. 1, pp. 60–65, 2018.
- [18] X. Foukas, G. Patounas, A. Elmokashfi, and M. K. Marina, "Network Slicing in 5G: Survey and Challenges," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 94–100, 2017.
- [19] I. Afolabi, T. Taleb, K. Samdanis, A. Ksentini, and H. Flinck, "Network Slicing and Softwareization: A Survey on Principles, Enabling Technologies, and Solutions," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 2429–2453, 2018.
- [20] A. A. Barakabitze, A. Ahmad, R. Mijumbi, and A. Hines, "5G network slicing using SDN and NFV: A survey of taxonomy, architectures and future challenges," *Computer Networks*, vol. 167, p. 106984, Feb. 2020. [Online]. Available: <https://doi.org/10.1016/j.comnet.2019.106984>
- [21] A. Matencio-Escobar, Q. Wang, and J. M. A. Calero, "SliceNetVSwitch: Definition, Design and Implementation of 5G Multi-Tenant Network Slicing in Software Data Paths," *IEEE Transactions on Network and Service Management*, vol. 17, pp. 2212–2225, 2020.
- [22] L. U. Khan, I. Yaqoob, N. H. Tran, Z. Han, and C. S. Hong, "Network Slicing: Recent Advances, Taxonomy, Requirements, and Open Research Challenges," *IEEE Access*, vol. 8, pp. 36 009–36 028, 2020.

- [23] S. Abe, G. Hasegawa, and M. Murata, "Effects of C/U Plane Separation and Bearer Aggregation in Mobile Core Network," *IEEE Transactions on Network and Service Management*, vol. 15, no. 2, pp. 611–624, 2018.
- [24] E. Schiller, N. Nikaein, R. Favraud, K. Kostas, D. Stavropoulos, S. Alyafawi, Z. Zhao, T. Braun, and T. Korakis, "Network Store: Exploring Slicing in Future 5G Networks," Sep. 2015.
- [25] T. Taleb, B. Mada, M.-I. Corici, A. Nakao, and H. Flinck, "PERMIT: Network Slicing for Personalized 5G Mobile Telecommunications," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 88–93, 2017.
- [26] F. Esposito, D. Di Paola, and I. Matta, "On Distributed Virtual Network Embedding With Guarantees," *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 569–582, 2016.
- [27] Y. Wang, N. Li, P. Yu, W. Li, X. Qiu, S. Wang, and M. Cheriet, "Intelligent and Collaborative Orchestration of Network Slices," *IEEE Transactions on Services Computing*, pp. 1–14, 2022.
- [28] K. Abbas, T. A. Khan, M. Afaq, and W.-C. Song, "Ensemble Learning-based Network Data Analytics for Network Slice Orchestration and Management: An Intent-Based Networking Mechanism," in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*, 2022, pp. 1–5.
- [29] "From Vertical Industry Requirements to Network Slice Characteristics," GSMA, Tech. Rep., Aug. 2018.
- [30] "Aspects; Management and orchestration; Concepts, use cases and requirements," 3GPP, Tech. Spec. 28.530 version 15.3.0, Dec. 2019.
- [31] "Study on management and orchestration of network slicing for next generation network," 3GPP, Tech. Rep. 28.801 version 15.1.0, Dec. 2020.
- [32] "Network Functions Virtualisation (NFV); Architectural Framework," ETSI, Group Spec. NFV-MAN 002 version 1.2.1, Dec. 2014.
- [33] OpenStack. Accessed: Jun. 28, 2022. [Online]. Available: <https://www.openstack.org>
- [34] Open Source MANO. Accessed: Jun. 28, 2022. [Online]. Available: <https://osm.etsi.org/>
- [35] "Network Functions Virtualisation (NFV) Release 3; Protocols and Data Models; NFV descriptors based on YANG Specification," ETSI, Group Spec. NFV-SOL 006 V3.6.1, Mar. 2022.
- [36] Open vSwitch. Accessed: Jun. 28, 2022. [Online]. Available: <https://www.openvswitch.org>
- [37] ONOS. Accessed: Jun. 28, 2022. [Online]. Available: <https://wiki.onosproject.org/display/ONOS/Downloads>
- [38] Open5GS. Accessed: Jun. 28, 2022. [Online]. Available: <https://open5gs.org/>
- [39] LTE broadband solutions. Accessed: Jun. 28, 2022. [Online]. Available: [https://www.leonardo.com/documents/15646808/16735768/CSP-MCX+broadband+MCC+platform+LQ+\(mm09006\).pdf](https://www.leonardo.com/documents/15646808/16735768/CSP-MCX+broadband+MCC+platform+LQ+(mm09006).pdf)
- [40] "LTE; Mission Critical Communication Interworking with Land Mobile Radio Systems," ETSI, Tech. Spec. 123 283 version 15.1.0, Jul. 2018.
- [41] UERANSIM. Accessed: Jun. 28, 2022. [Online]. Available: <https://github.com/aligungr/UERANSIM>
- [42] ONOS Intent Framework. Accessed: Jun. 28, 2022. [Online]. Available: <https://wiki.onosproject.org/display/ONOS/Intent+Framework#IntentFramework-Intents>