

This is a postprint version of the following published document:

Fonseca, J., Alegria, J., Cunha, V. A., Quevedo, J., Santos, D., Gomes, D., Barraca, J. P, Corujo, D. & Aguiar, R. L. (9-11 Nov. 2021). *Dynamic Interdomain Network Slicing for verticals in the 5Growth project* [proceedings]. 2021 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), Heraklion, Greece.

DOI: [10.1109/NFV-SDN53031.2021.9665037](https://doi.org/10.1109/NFV-SDN53031.2021.9665037)

© 2021, IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Dynamic Interdomain Network Slicing for Verticals in the 5Growth Project

João Fonseca, João Alegria, Vitor A. Cunha, José Quevedo, David Santos, Diogo Gomes, João P. Barraca, Daniel Corujo, Rui L. Aguiar

Instituto de Telecomunicações and Universidade de Aveiro, Portugal

Email: {*jpedrofonseca, joao.p, vitorcunha, quevedo, dauidasantos, dgomes, jpbarraca, dcorujo, ruilaa*}@av.it.pt;

Abstract—This paper proposes and validates a Interdomain Network Slicing framework for verticals, allowing them to directly participate in the establishment and control of end-to-end Communication Services deployment across multiple inter-operator domains. The framework progresses the means made available by different standards and research initiatives to enhance service requesting and provisioning interfaces for the stakeholders involved, namely operators and verticals. The framework is validated under two different use cases, showcasing effective end-to-end service instantiation and a first assessment towards dynamic service modification capability.

Index Terms—interdomain, 5g, slicing, network, virtualization

I. INTRODUCTION

One of the key contributions of 5G Fifth Generation of Mobile Networks (5G), beyond enhanced radio performance, was the coupling of cloud-based capabilities for more flexible and dynamic network service operation. This coupling allows Mobile Network Providers (MNOs) to leverage Network Functions Virtualization (NFV) and Software-Defined Networks (SDN) mechanisms to efficiently provide turn-key solutions towards complex communication scenarios with service assurance and added functionality such as security. As a result, such simplification attracted verticals whose strict communication requirements traditionally mandated dedicated and isolated network deployments, optimizing expenditure. However, these newly integrated vertical sectors are also capable of imprinting new utilization considerations. For example, in the transportation or energy sectors, it is common for their assets to be geographically widespread and fall under the coverage domain of different (and separate) operators. In this way, End-to-End (E2E) service provisioning poses an added degree of complexity to existing telecommunication mechanisms and standards, which require solutions for such cases. On the one hand, standards-based best practices and assurances are essential to these verticals due to the critical nature of their inherent communications. On the other hand, despite the technical ability of one MNO to act as an intermediate with secondary ones for coverage extension purposes, the economic benefit of the end-user (i.e., the vertical) might not be optimal.

This paper aims to contribute by proposing and validating an Inter-domain Network Slicing (INS) framework that leverages the contributions of the H2020 5Growth (5Growth) project [1] to deploy E2E communication services for verticals, across multiple domains, and through different operators. It also

provides an initial assessment of the framework’s capabilities towards dynamic service modification. The remainder of this document is structured as follows. Section II presents standardization and research efforts on the area, followed by the system architecture definition in Section III and the validation in Section IV. Finally, the paper concludes in Section V.

II. STATE OF THE ART

In this section, we will start by introducing the standards by 3GPP and then ETSI. We will then discuss the related research projects and their contributions to our proposal.

A. 3GPP

The 3rd Generation Partnership Project (3GPP) under 3GPP TSG SA Working Group 5 (SA5) specified the roles for the entities involved in providing a Communication Service (CS), which uses 5G networks and network slicing. These entities are Communication Service Consumer (CSC), Communication Service Providers (CSP), Network Operator (NO), Virtual Infrastructure Manager (VIM) provider, and the Data Center Service Provider.

CSCs in a 3GPP system can also be a CSP, which leverages CSs offered by other CSPs. No matter the type of CS, CSCs are dependent on the existence of a CSP. A CS can be mapped as a Communication Service Instance (CSI), which a CSC requests to CSPs. In [2], the 3GPP specifies several categories for these CSs: Business to Business (B2B), Business to Business to everything (B2B2X), Business to Consumer (B2C), and Business to Household (B2H).

According to the CS requirements, a Network Slice can be delivered as a CS, exposing a Network Slice Instance (NSI) that uses the resources available in the CSP domain. Furthermore, the management of these resources is made available to the CSC. The lifecycle of an NSI is defined in 3GPP Specifications [3] and [4], occurring in four phases:

- 1) The **Preparation Phase** corresponds to the design of the network slice, creating a Network Slice Template (NST) that is then onboarded and the associated supportive network environment prepared.
- 2) On the **Commissioning Phase**, the CSP requests the instantiation of a network slice as an NSI. Sometimes it is followed by the instantiation of possible Network Slice Subnet Instances (NSSIs) associated with it.

- 3) The next phase is the **Operation**. In it, the NSI can be activated, modified, and deactivated.
- 4) The last phase is **Decommissioning**, where an NSI and the related resources are released and cease to exist if they are not shared with other NSIs. This phase also triggers the termination of NSSI when they are not being used.

3GPP defines Network Slice as a Service (NSaaS) as a CS, enabling CSCs to choose the type of access the Network Slice: CSC, end-user, or manager using the management interfaces exposed by the CSP [2]. The latter option enables the CSC to provide the Network Slice to other entities such as a normal CSP. This CS can be seen as a B2B2X service.

B. ETSI

In contrast to the definitions by 3GPP, the European Telecommunications Standards Institute (ETSI) NFV WG only considers the existence of Network Services (NSs) and Network Functions (NFs) [5]. It is worth noting that an NS can itself contain another NS, in what is called a nested NS [6]. An NS is considered a resource-centric view of a Network Slice when an NSI contains at least one Virtual Network Function (VNF). A corollary of this definition was also applied to the Network Slice Subnets and NSSI.

An NSSI can be shared by several NSI [3] [4], and a nested NS can be shared with a parent NS. This fact led ETSI to recommend the 3GPP Network Slice Subnet connectivity to physical resources [7], which a nested NS can represent. This definition sets the possibility for a Network Slice to use an ETSI VNFs or Physical Network Functions (PNFs) attached to an NFV-NS.

Considering the need for management interfaces, 3GPP also recommends adopting Management Functions related to CS [3]: Communication Service Management Function (CSMF), Network Slice Management Function (NSMF), Network Slice Subnet Management Function (NSSMF).

C. Research Projects

All these recommendations by Standards Development Organizations (SDOs) led to the proposal of several network slice management platforms: Open Baton¹, Open Source MANO (OSM)², SliMANO [8], 5G-TRANSFORMER (5G-TRANSFORMER)³, 5GTANGO⁴. While OSM, Open Baton, and 5GTANGO provide network slice manager platforms that interact with ETSI NFV based Network Function Virtualization Orchestrator (NFVO), thus providing just access to network resources, both SliMANO and 5G-TRANSFORMER exceed those capabilities aiming at providing the management of SDN controllers and Radio Access Network (RAN).

From all the possibilities mentioned before, only the 5G-TRANSFORMER project provided a component that allowed Verticals to create CSs that fit their needs while mapping

the CSs to NSIs, the 5G-TRANSFORMER Vertical Slicer (5G-TRANSFORMER-VS). This component also maps an existing NSI into an NFV-NS [7], which is then requested to the NFVO. Although the 5G-TRANSFORMER's Vertical Slicer (VS) allows the CSI and NSI requests, it did not follow the recommendations by 3GPP on the division of the Management Functions. Further, this monolithic architecture prevented the request of slices that could be obtained by leveraging NSaaS. These problems have been addressed in 5Growth.

5Growth⁵ [1] is a 5G Infrastructure Public Private Partnership (5GPP) Phase-3 project funded under H2020-ICT-2019. It builds on top of the breakthrough architecture of 5G-TRANSFORMER, a Phase 2 project, using its platform as a reference. The main focus of 5Growth is the automation of processes for supporting several industry verticals. A vertical portal is inherited from 5G-TRANSFORMER and is available to help in this process, where CS requests are made. 5Growth leverages Artificial Intelligence (AI) to enable the deployment of E2E network solutions across multiple technologies and domains. The project aims at improving the architecture of the 5G-TRANSFORMER platform by adding: support to RAN segments in the network slices, vertical-service monitoring extensions, service slice orchestration monitoring, control loops stability, AI/ML support, federation and inter-domain capabilities, support for the next generation RANs. The 5Growth framework also uses security and auditability mechanisms and CI/CD capabilities by applying new algorithmic innovations that enable smart orchestration and resource control, anomaly detection capabilities, forecasting, and inference based on previous data analytics. All these innovations provide the Verticals, like Mobile Virtual Network Providers (MVNOs) and other Communication Service Providers (CSP), more autonomy from the NOs while providing the latter with tools for providing a better service and fulfill their Service Level Agreements (SLAs).

Focusing on the 5Growth Vertical Slicer (5GR-VS), separating the CSMF and the NSMF components enables new CS requested by verticals to different administrative domains. At the CSMF level, a CSMF federation was explored between domains to request an NFVO level federation. On the other hand, it is also possible to request a CSI mapped into an NSI. This NSI can incorporate network slice resources from different NSMF domains. This possibility is the main focus of this document, as we propose a solution for the use of inter-domain network slice resources. We present a possible way of moving forward to guarantee the dynamic change of the CS throughput quality.

III. SYSTEM ARCHITECTURE

An example use case for the use of the INS concept is presented in Figure 1. It pictures, in three different views (management, service, and network), a solution provider for verticals (i.e., CSP) that mediates the access to a CS, which is

¹OPEN BATON: <https://openbaton.github.io/>

²OSM: <https://osm.etsi.org/>

³5G-TRANSFORMER: <http://5g-transformer.eu/>

⁴5GTANGO: <https://www.5gtango.eu/>

⁵5Growth: <https://5growth.eu/>

built by stitching two independent NSaaS solutions provided by two different CSPs. As shown in the Service View, the main objective allows the vertical to monitor assets deployed on different Administrative Domain (AD) in real-time. Following the recommendations by 3GPP on network slicing, this will create an E2E network slice with access to the different resources. As we have two different NO providing access to their NSaaS solutions, the slice will have two Network Slice Subnets corresponding to those ADs. The only thing left to decide is the correct way to stitch these network parts. Let us assume the solution provider for verticals networks can be the interconnection network. A tunnel between the two Points of Presence (PoP) is available to stitch the different slices.

The 5Growth project introduced an Moving Target Defense (MTD) function [9] to protect these interdomain interfaces against reconnaissance, delivery, and undiscovered exploits in the exposed service stack. The MTD function works alongside the secure tunnel that interconnects the PoPs, leveraging a known Two-Factor Authentication (2FA) protocol to defend the secure tunnel [10]. Orchestrating the secret keys distribution, function deployment, configuration, and synchronization across the PoPs is critical for the MTD function.

Considering the scenario described before, the Management View shows the relations between all the entities in the environment. We can see how a CSC (Vertical) requests a service to a CSP, the solution provider for verticals. A CSP operator then uses the 5Growth portal, a platform designed to onboard and request CSs or Vertical Services. In this case, the requested Vertical Service is an E2E CSI over two ADs. This request is then forwarded to the 5GR-VS. In the 5GR-VS, the CSI is mapped into an E2E NSI by the CSMF. This management function is responsible for identifying the associated NSSIs needed by the NSI, requesting the resources it needs from both domains, and using the NSMF domain-specific drivers. In this case, both CSP A and B provide NSaaS solutions to the E2E CSP. In this example, each CSP uses OSM in slice manager + NFVO mode, which means the domain drivers are the same (however, our design is generic enough to encompass other orchestrators). At 5GR-VS, the translator will check each domain catalog and map the instantiation and configuration requests to different domains over the same driver.

The NFVOs tries to control the network resources of the respective ADs. In the Network View, we present an overview of the infrastructure to be used in this scenario. The network services/slices provided by CSP A and B can leverage cloud resources available in Openstack⁶ clusters, which act as the VIMs. It is inside these clusters that the NFs will be deployed. In this scenario, OSM will then be used to control and deploy the NF over the VIMs in each of these AD. The network topology of the interconnection network can be described as an Non-Public Network (NPN) owned by the solution provider for verticals, or it can also be the Internet. This topology results from the focus of this paper being the orchestration of cloud-based resources over different AD. We opted for

the use of an NPN as the interconnection network. However, the network resources managed by this solution provider for verticals, which makes use of the interfaces provided by the CSPs in their NSaaS solutions, can be easily integrated into a slice that makes use of radio resources from Public Land Mobile Network (PLMN) if required.

A. NFVO and Network Slice Resources

OSM Release 9 was used, which can deploy several types of NF compliant with SOL006 [11] and slices compliant with SOL005 [12] and IFA014 [6]. It was necessary to create custom NFVO entities deployed by OSM to enable the intended inter-domain E2E service, namely new VNFs, NSs, and NSIs. For that reason, all of these entities follow the guidelines and requirements defined in OSM documentation.

1) *VNFs*: In the OSM environment, the VNF is the lowest entity managed by the NFVO. Using a Virtual Network Function Descriptor (VNFD) and Juju charm, it is possible to define the function's topology and behavior. Two VNFs were created with the E2E service in mind, one for the inter-domain mechanism and another for the MTD functionality. Both VNFs had the same topology, consisting of only one Virtual Deployment Unit (VDU) and two interfaces. The final inter-domain VNF provides a Virtual Private Network (VPN) tunnel peer and several operations to configure it: get tunnel peer information, add tunnel peer, remove tunnel peer, and modify tunnel quality parameters. The final MTD VNF provides an MTD agent and operations to configure it, such as: get MTD information and activate MTD.

2) *NSs*: With functional VNFs, it is up to the NS to operate them. An Network Service Descriptor (NSD) defines which VNFs compose the service and how they should be connected. With the inter-domain E2E service in mind, we developed two NSDs based on the VNFs previously defined. The first NS contains only the inter-domain VNF, supporting the inter-domain mechanism and providing a tunnel peer ready to be used. This service establishes a simple inter-domain connectivity and serves as a baseline for more complex services and scenarios. The second service chains together the MTD and the inter-domain VNFs. These two NFs are internally connected, where the MTD function serves as a gateway for the inter-domain function. This service was created to elaborate over the previous one, proving the inter-domain mechanism still works despite the MTD mechanism restrictions.

3) *Network Slice Resources*: Finally, we created one NST to test the inter-domain mechanism using the network slicing capabilities of OSM. This NST has internally one subnet, which entails one of the NSs previously defined. Therefore, this network slice is a viable option for an AD that provides NSaaS solutions. Its dedicated subnet contains a VPN tunnel peer needed to secure the inter-domain environment.

B. Signaling Diagram

The Inter-domain Network Slicing (INS) CS can be requested over the 5GR-VS after the onboarding of a Vertical Service Blueprint (VSB). This VSB defines the composition

⁶Openstack: <https://www.openstack.org/>

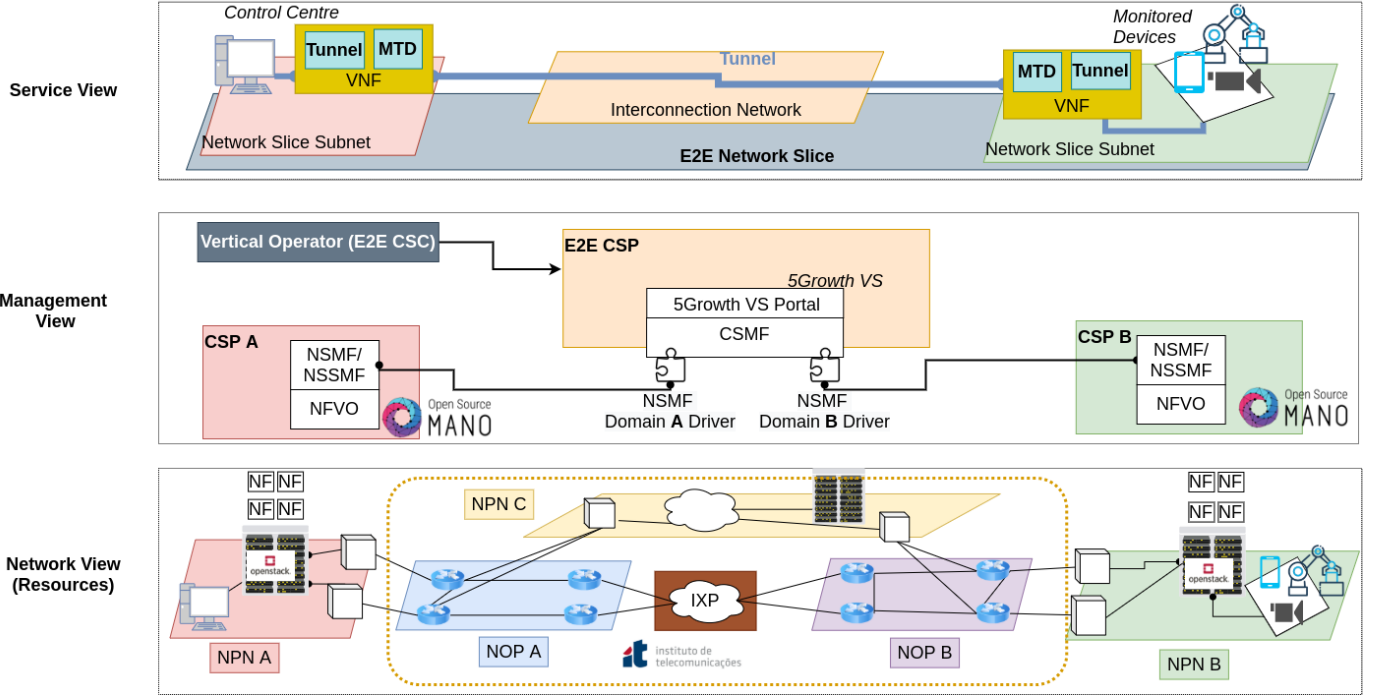


Fig. 1: Inter-domain Network Slicing

of the Vertical Service Instance (VSI), the associated E2E slice, and Key Performance Indicators (KPIs). The process of instantiation/termination of the VSI and associated slices are represented in the signaling diagram of Figure 2. Knowing the 5GR-VS is a platform responsible for communicating the requests to the slice managers of the different AD, it is of interest to understand the delays created by its use.

The instantiation of the desired CS or Vertical Service is represented in Stage 1. It can be divided into steps to achieve the E2E network slice across domains. Depending on the number of NFs being used, this stage can have several operation steps and sub-steps associated with their instantiation/configuration (day-0 and day-1 operations). The first step corresponds to the assignment of a unique identifier to each of the network slice subnets. The second step refers to the delay of the instantiation of the network slice subnets over the 5GR-VS until the OSM slice managers receives the request. The third step refers to the configuration of the VPN tunnel peers. Likewise, the configuration of the MTD mechanism happens during this time, if it is being used. When this step finishes, assuming everything goes as expected, the Vertical Service is ready to be used by the Vertical.

Upon the availability of the Vertical Service, the only operations available and supported at the time of writing for changing the status of the Vertical Service and the associated E2E NSI are updating and terminating the Vertical Service. In both cases, this leads to the termination of the service. That creates an issue for the Industry Vertical because the service could be interrupted while changing the properties of the Vertical Service (e.g., changing the quality of the service, following the KPIs).

Stage 2 focuses on the termination of the Vertical Service and the teardown of the E2E inter-domain network slice. In this context, step 4 specifies the requests at the 5GR-VS level for terminating the network slice subnets. The termination request for the E2E network slice subnets and the respective network resources at the different ADs leads to the release of the NFVO resources at the VIMs. Thus, the network slice subnet is terminated, leading to the termination of the E2E slice when the last subnet is released. Lastly, the termination of the E2E slice leads to the termination of the Vertical Service.

In the diagram, we show our current view for the Service Level Modification of the characteristics of the Network Slice resources deployed at the different AD, see VSI Modification section. This procedure should follow the KPIs listed during the CS onboarding. It is expected that a solution provider for verticals can request the change of the characteristics of the offered service at a given time without needing to tear the whole service down. Therefore, we propose a new flow of actions, signalled with red arrows, to be executed in the 5GR-VS, that allows a CSP to request the change of a CS whenever needed. These actions will be later forwarded to the slice managers and respective NFVOs can change the correct NF, using day-2 primitives. This concept is not yet supported by the 5GR-VS, which means that later in the article, we will only focus on the modification of the service at the slice managers level.

IV. EVALUATION

In this section, we will evaluate our proposal through a Proof-of-Concept (PoC) deployment in our laboratory data center. We will start by describing the test environment.

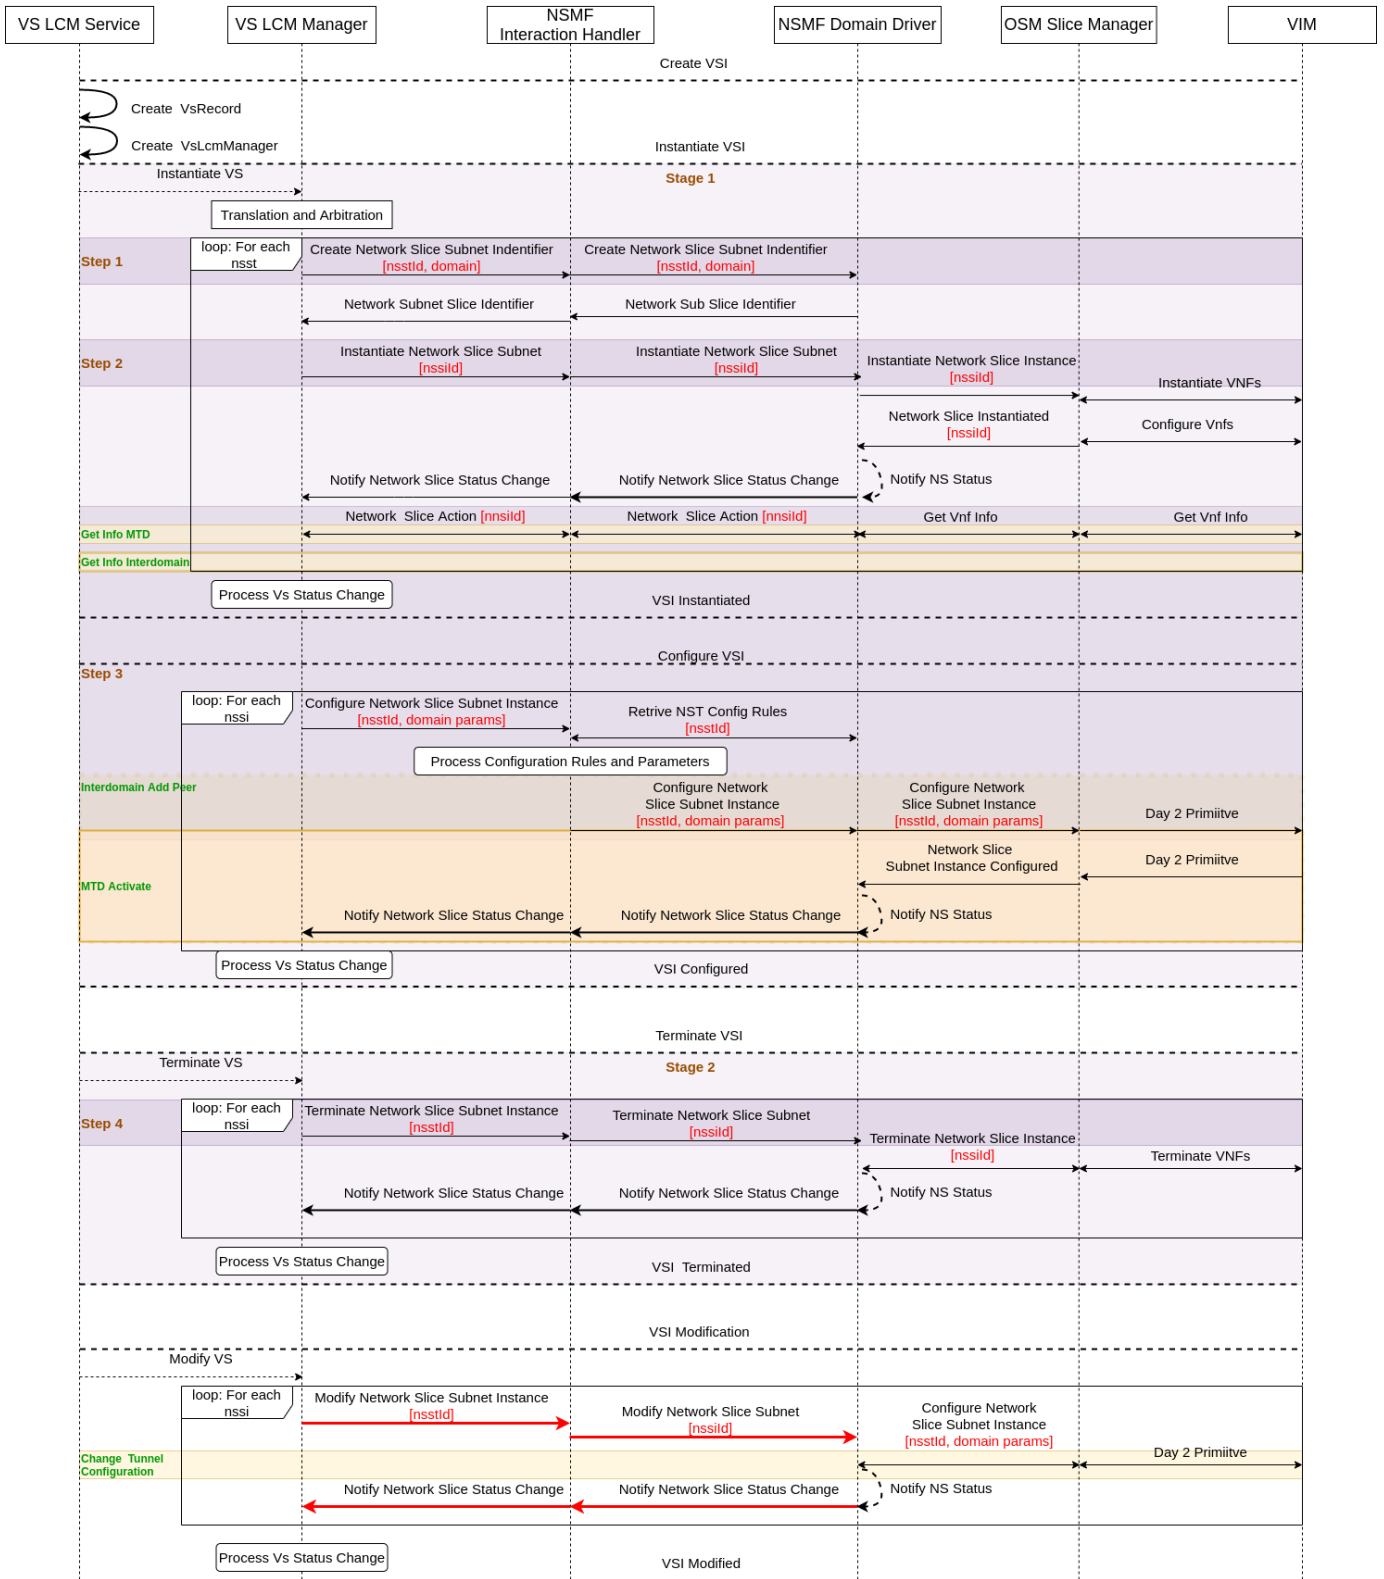


Fig. 2: Inter-domain signaling diagram

Then, we will discuss the relevant results gathered with two representative scenarios.

A. Testing Environment

To test this PoC we created two independent OSMs, one for each domain, and each one deployed in a Virtual Machine (VM) with 12 GB of RAM, 4 vCPUs, 150 GB of storage,

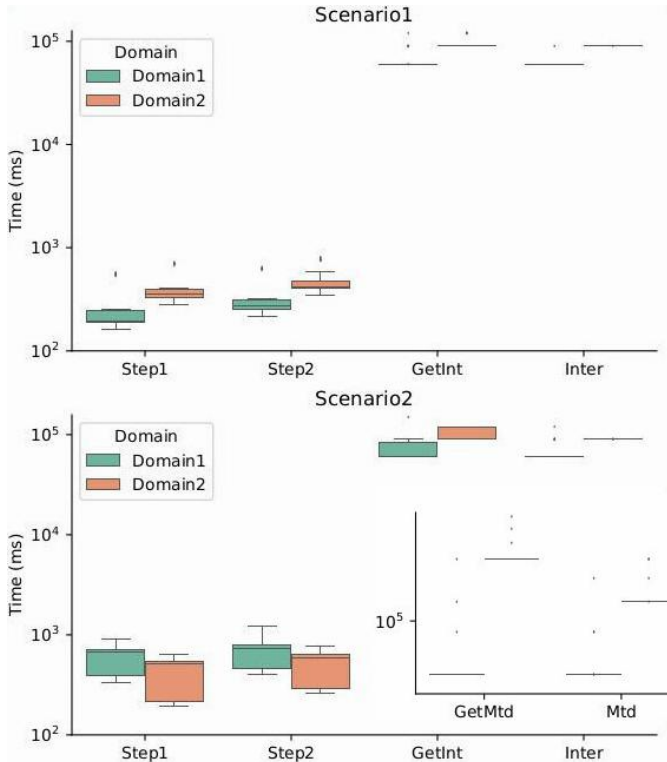


Fig. 3: Instantiation stages delays of both scenarios

and running *Ubuntu 18.04*. The first OSM integrates with a production OpenStack deployment, where the project defined had the limitations of 30 VM instances, 60 vCPUs, 70 GB of RAM, 1 TB of storage, and 10 networks. The second OSM integrates with a DevStack, an OpenStack version focused on the quick deployment of development infrastructures. The project assigned had the limitations of 10 VM instances, 20 vCPUs, 50 GB of RAM, 1 TB of storage, and 100 networks. Additionally, we used a dedicated VM with 8 GB of RAM, 4 vCPUs, 32 GB of storage, and running *Ubuntu 18.04* to deploy the 5GR-VS system.

B. Results

In this PoC, we focused on the VSI orchestration aspects, namely the delays and number of signaling bytes associated with each step in its instantiation and termination phases. Those phases are detailed in Figure 2. Two testing scenarios were conducted: the first consisted of a simple inter-domain service (Scenario 1). The second consisted of a service composed of the inter-domain and MTD functions (Scenario 2). The delays of all steps related to the instantiation and termination of each scenario's service are presented in Table I. Figure 3 presents, specifically, the instantiation phase delays for both scenarios.

Concerning Scenario 1, the E2E instantiation and termination delays are on average 9 minutes and 46 seconds, respectively. This difference was expected since the instantiation phase is considerably more complex than the termination one. By analyzing the instantiation phase, after deploying each

Scenario	Stage	Min	Max	Avg	StDev
1	Step1-1	162.0	556.0	248.14	124.64
	Step1-2	283.0	705.0	393.86	126.88
	Step2-1	214.0	625.0	315.82	128.03
	Step2-2	342.0	773.0	467.5	130.65
	GetInter-1	60070.0	90103.0	64219.45	10534.95
	GetInter-2	90086.0	90140.0	90108.31	14.35
	AddPeer-1	60094.0	60141.0	60111.03	13.20
	AddPeer-2	90111.0	90167.0	90134.72	16.12
	Stage1	496856.0	591173.0	538047.79	19188.84
	Step4-1	38.0	138.0	83.55	31.21
	Step4-2	44.0	142.0	82.97	29.50
	Stage2	22034.0	83027.0	45706.9	23486.80
2	Step1-1	336.0	896.0	580.5	185.47
	Step1-2	194.0	646.0	403.3	171.51
	Step2-1	397.0	915.0	644.93	181.25
	Step2-2	259.0	775.0	489.8	184.06
	GetMtd-1	60070.0	90093.0	61157.64	5670.80
	GetMtd-2	180133.0	210173.0	181227.14	5672.87
	GetInter-1	60071.0	90112.0	67329.10	13069.58
	GetInter-2	90088.0	120152.0	103128.17	15117.42
	Addpeer-1	60100.0	90150.0	65295.24	11539.65
	Addpeer-2	90121.0	90190.0	90144.18	17.76
	Mtd-1	60073.0	90146.0	62158.76	7750.61
	Mtd-2	120108.0	150168.0	121205.79	5676.08
	Stage1	938127.0	1174247.0	1027347.64	65086.07
	Step4-1	48.0	146.0	89.38	32.97
	Step4-2	49.0	147.0	91.93	28.98
	Stage2	20790.0	84359.0	56500.7	21225.13

TABLE I: Orchestration stages delays of both scenarios

subnet in the respecting ADs, the E2E VSI needs to be configured, which in this scenario corresponds to activating the inter-domain mechanism, namely fetching the necessary information and configuring the tunnel peers. This mechanism needs to be triggered in each AD, taking on average 2.5 minutes per domain. Concerning the number of bytes exchanged in each stage, the instantiation phase needed approximately 93.02 KB and the termination phase needed 12.08 KB.

Compared to the first scenario, the E2E instantiation and termination delays in Scenario 2 are higher, averaging on 17 minutes and 57 seconds, respectively. This significant increase in the instantiation phase is due to the higher complexity of the service, combining the MTD and inter-domain functions. Furthermore, adding the MTD function required more service configuration actions. After deploying each subnet in the respecting ADs, the E2E service was configured by triggering both domains via the inter-domain mechanism. This scenario took on average 2.7 minutes per domain, and the MTD mechanism took on average 3.5 minutes per domain. When compared with Scenario 1, the E2E service configuration phase caused more delay in the instantiation process, taking on average 6.2 minutes per domain instead of the previous 2.5 minutes. Being a service with higher complexity and orchestration delays also affected the number of bytes exchanged in both stages, needing approximately 139.03 KB for instantiating and 15.30 KB for terminating the service.

One aspect directly influencing the orchestration delays obtained for these scenarios is the virtualization infrastructure itself. Depending on the technologies and hardware used in each AD, there will be resources limitations determining the services that are deployed in that infrastructure. These constraints also mean that the instantiation of the same service in different infrastructures can generate distinct values. This

	Mean	StdDev	Min	Max
Domain1	8.3706	1.4918	6.3437	11.3395
Domain2	11.2389	3.4270	6.9619	19.60

TABLE II: Tunnel Bandwidth Modification Delay (Mb/s)

phenomenon is illustrated in the difference of delays between Domain 1 and Domain 2 in both tested scenarios. Domain 2 generated higher delays in every instantiation step, ranging from 25 seconds to 2 minutes of difference.

The services were deployed through the 5GR-VS orchestrator in these tests, which has its pros and cons. Its advantages are the VSI abstraction, enabling the vertical user to focus only on the service itself and not on the network and infrastructure complexities, the separation between CSMF and NSMF, enabling the Vertical Slicer (VS) to connect to different ADs, and finally, the support of instantiation configuration operations, allowing the VS to trigger those actions automatically after the service is instantiated. On the other hand, its disadvantages are the limited support of runtime configuration, restricting the possibility of triggering Day-2 operations over the services, and the sequential approach when managing and orchestrating VSIs, its corresponding E2E NSIs, and associated NSSIs, meaning that only one action from one subnet is processed at any given time. This last disadvantage directly affected the delays measured in our two test scenarios. Given that in an inter-domain environment, several independent AD are coordinated to provide the E2E service, each domain's subnet could be processed in parallel, which would reduce the VSI instantiation delay. In the scenarios tested, by using that parallelization, the instantiation delays could decrease 2 minutes in scenario 1 and 4 minutes in scenario 2, considering the worst domain delays. Another implementation aspect that impacted our results was the polling approach used to update the service status, influencing the number of bytes obtained from the signaling stages. Many bytes are exchanged regularly by constantly polling the service, meaning that the longer the instantiation and termination phases take, the higher the number of bytes exchanged is.

Following the tests done using the MTD and the inter-domain NFs, we decided to change the characteristics of the bandwidth of the inter-domain tunnel. For that we, requested the applicability of a day-2 primitive in each AD for changing the tunnel bandwidth. Given the nonexistent support at the time for these operations at the 5GR-VS CSMF, we focused on the results obtained by requesting this change directly to the slice managers. The results for the modification delay for each domain are listed in Table II (note that these results only address the signalling exchange involving the elements beyond the 5gr-vs). As pointed out before, the delay is far more significant in Domain 2.

The tests we performed for the change of the characteristics of the tunnel are available in Table III. They show the correspondence between the maximum bandwidth available and the correspondent measured bandwidth. The tunnel occupancy seems to increase with the decrease of the available bandwidth.

Available Bandwidth	Mean	StdDev	Min	Max
1000	705.1379	111.14332	462.0	848.0
500	469.0345	5.5580	453.0	476.0
250	238.7778	0.9740	237.0	241.0

TABLE III: Available Throughput inside the Scenario2 after Service Modification (Mb/s)

V. CONCLUSION

We have successfully implemented and demonstrated two scenarios that required inter-domain communications across different Point of Presences (PoPs). Our solution shows one of the first efforts integrating a CSMF with several NSMF and stitching the network slice resources into a E2E slice. The evaluation showed promising results that validated the usefulness of the solution for a practical deployment concerning direct intervention from verticals in reaching geographically widespread assets (such as the 5Growth pilots). The proposed solution for the dynamic change of the Vertical Services enables on-demand shaping of the network resources conditions following the KPIs defined at the CS onboarding.

ACKNOWLEDGMENT

This work has been supported by EC H2020 5GPPP 5Growth project (Grant 856709).

REFERENCES

- [1] X. Li, A. Garcia-Saavedra, X. Costa-Perez, C. J. Bernardos, C. Guimarães, K. Antevski, J. Mangués-Bafalluy, J. Baranda, E. Zeydan, D. Corujo, P. Iovanna, G. Landi, J. Alonso, P. Paixão, H. Martins, M. Lorenzo, J. Ordonez-Lucena, and D. R. López, "5Growth: An end-to-end service platform for automated deployment and management of vertical services over 5g networks," *IEEE Communications Magazine*, vol. 59, no. 3, pp. 84–90, 2021.
- [2] 3GPP, "Aspects; Management and orchestration; Concepts, use cases and requirements," 3rd Generation Partnership Project, Technical Specification 28.530, 2019.
- [3] —, "Study on management and orchestration of network slicing for next generation network," 3rd Generation Partnership Project, Technical Specification 28.801, 2018.
- [4] —, "Management and orchestration; Provisioning," 3rd Generation Partnership Project, Technical Specification 28.531, 2019.
- [5] S. e. a. Clayman, "The NECOS approach to end-to-end cloud-network slicing as a service," *IEEE Communications Magazine*, vol. 59, no. 3, pp. 91–97, 2021.
- [6] "ETSI GS NFV-IFA 014 V2.4.1, Management and Orchestration; Network Service Templates Specification," 2 2018.
- [7] "ETSI GS NFV-EVE 012 V3.1.1; Network Functions Virtualisation (NFV) Release 3; Evolution and Ecosystem; Report on Network Slicing Support with ETSI NFV Architecture Framework," 12 2017.
- [8] F. Meneses, M. Fernandes, D. Corujo, and R. L. Aguiar, "Slimano: An expandable framework for the management and orchestration of end-to-end network slices," in *2019 IEEE 8th International Conference on Cloud Networking (CloudNet)*, 2019, pp. 1–6.
- [9] V. A. Cunha, D. Corujo, J. P. Barraca, and R. L. Aguiar, "TOTP Moving Target Defense for sensitive network services," *Pervasive and Mobile Computing*, vol. 74, p. 101412, 2021.
- [10] V. A. C. et al., "5Growth: Secure and reliable network slicing for verticals," in *2021 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit): Network Softwareisation (NET) (2021 EuCNC & 6G Summit - NET)*, Porto, Portugal, Jun. 2021.
- [11] "ETSI GS NFV-SOL 006 V2.7.1; Network Functions Virtualisation (NFV) Release 2; Protocols and Data Models; NFV descriptors based on YANG Specification," 12 2019.
- [12] "ETSI GS NFV-SOL 005 V2.4.1; Network Functions Virtualisation (NFV) Release 2; Protocols and Data Models; RESTful protocols specification for the Os-Ma-nfv Reference Point," 2 2018.