Software-driven Connectivity Orchestration for Multi-domain NFV Ecosystems

Borja Nogales*, Ivan Vidal*, Francisco Valera*, Victor Sanchez-Aguero*, Diego R. Lopez†
*Universidad Carlos III de Madrid
†Telefónica I+D

***Abstract*—Network Functions Virtualization (NFV) is a fundamental enabler in 5G networks, automating service deployment through softwarization and virtualization. However, challenges remain in providing network connectivity to service components, commonly referred to as Virtual Network Functions (VNFs), deployed on different domains. A usual approach to enable such connectivity is to rely on layer-3 routing from Internet service providers. However, this approach presents limitations to preserve isolation among remote VNFs, and may require undesirable network configurations. This paper proposes a novel connectivity orchestration service for multi-domain NFV ecosystems to address these limitations. The service is based on a set of interconnected software components that are deployed on every NFV domain. Connectivity is automatically provided to remote VNFs over the overlay network formed by these components using Software-Defined Network (SDN) technologies. The service has been prototyped using open-source software, being validated on a real NFV ecosystem.**

***Index Terms*—NFV, SDN, overlay networks, Inter-Domain Connectivity Orchestration.**

I. Introduction

THE last decade has witnessed the significant role of virtualization and softwarization technologies in the deployment of mobile networks [1], especially in the take-off of the 5^{th} generation (5G) of mobile communications. Softwarization enables the implementation in software of functionalities traditionally provided by specialized hardware (*e.g.*, routers, firewalls, proxies, etc.), whereas virtualization supports the flexible deployment of software functions on different locations, including cloud and edge facilities, using commercial off-the-shelf equipment [2]. Embracing both concepts, ETSI Network Functions Virtualization (NFV) [3] has been adopted as a fundamental enabler to realize the disruptive vision of a virtual, pervasive, autonomously scalable, and repairable (also referred to as *zero-touch*) 5G network.

NFV provides an enormous flexibility to deploy software functions, commonly referred to as Virtual Network functions (VNFs), on different NFV infrastructure domains, for instance on geographically dispersed cloud computing facilities or network edge locations. This flexibility is fundamental to meet the stringent performance requirements of services in 5G and beyond networks. However, the establishment of proper connectivity among VNFs deployed on different NFV domains has received little attention by both the research and standardization communities. Existing solutions normally assume the availability of link-level (layer-2) data paths between NFV domains [4], which may not always be realistic. Alternatively, communications among remote VNFs have been provided through network-level routing mechanisms (layer-3), using IP routers of Internet service providers (ISPs) that interconnect NFV infrastructure domains [5], [6]. However, adopting a layer-3 approach to support VNF communications poses significant challenges for NFV infrastructure providers, ISPs, software developers, telco operators, and vertical sector organizations involved in the development and operation of 5G services. This approach hinders the proper isolation between multi-domain services and may potentially involve additional, undesirable, network-layer configurations on VNFs and ISP networks after a service deployment.

In this context, this article presents a novel inter-domain connectivity orchestration service for NFV ecosystems. The service design carefully considers the lessons learned from our previous work [7], proposing a new approach to support VNF connectivity in multi-domain scenarios. Our solution uses NFV and SDN technologies to automate the creation of virtual networks to which VNFs deployed in different NFV infrastructure domains can be connected. Furthermore, it offers flexibility in applying distinct management policies to steer VNF data traffic between NFV domains, facilitating resilient and high-performance communications.

We implemented a prototype of this solution using open-source software and standard protocols, validating the suitability of our design with a proof-of-concept over a real multi-site NFV ecosystem, currently available at the 5G Telefonica Open Network Innovation Centre (5TONIC).

II. Related work

NFV was proposed to address the limitations of relying on proprietary hardware for network functionalities, by implementing them in software and deploying them as VNFs. It establishes an abstraction layer that decouples hardware from functionalities, enabling a greenfield environment for deploying network services. ETSI provides standardization activities for NFV, including the definition of an architectural framework [3] where management and orchestration (MANO) functions support the automated deployment of network services built by interconnected VNFs. On the other hand, SDN introduces a data plane with programmable network elements (*e.g.*, switches), and centralizes the control logic in an upper hierarchy central node referred to as controller. The SDN architecture, defined by the Open Networking Foundation [8], facilitates communication between the controller and programmable network devices, the exposure of state information to external applications, and the interoperation among different controllers.

While some research has explored the synergies of NFV and SDN to provide networking solutions, most studies have focused on conceptual aspects rather than on a practical implementation [9]. Only a reduced set of studies have experimentally evaluated the capacity of NFV platforms to deploy multi-domain 5G services. In [10], a streaming service is deployed over several geographically distributed NFV infrastructures, using layer-3 VPNs to support inter-domain VNF traffic. Another study by [11] formulates an algorithm for VNF placement and traffic steering in an NFV ecosystem, which is experimentally validated using NFV, SDN, and layer-3 VPNs for multi-domain communication. Other research studies conducted in European research projects [5], [6] showcase the prevalent approach, employing network-level routing to support inter-domain VNF communications.

However, this approach opens serious concerns that may potentially limit the correct operation of telecommunication and vertical sector services on multi-domain NFV environments. On the one hand, the use of network-level routing mechanisms hinders a proper isolation between multi-domain 5G services. That is, in the absence of specific mechanisms to prevent it, VNFs of one service could be reachable from VNFs of other services or by untrusted third parties, using the IP addresses of the target VNFs. On the other hand, a layer-3 approach entails the potential need for additional (undesirable) network-layer configurations on VNFs and their underlying ISP networks after a service deployment. For example, the next hop address of a VNF might not correspond to another VNF of the service, as would be expected according to the service descriptor, but to an edge router of the NFV infrastructure domain where the VNF is deployed. Moreover, the exchange of multicast and broadcast traffic among remote VNFs may require the installation of specific forwarding state on ISP routers, which may simply prohibit this traffic over their networks, preventing the proper execution of the service. A layer-2 approach based on the allocation of Virtual Local Area Networks (VLANs) to 5G services might be a natural alternative to address the aforementioned limitations, using link-level data paths from the underlying ISP networks or layer-2 VPN services [12]. However, this approach raises significant challenges in terms of scalability and automated provisioning.

All these aspects have been carefully studied in a prior research work [7], jointly developed between Universidad Carlos III de Madrid and Telefónica I+D. This work presented L2S, a platform to provide secure link-level connectivity among NFV domains. However, L2S operation is based on a set of VLANs whose configuration is relatively static at each NFV domain, presenting limitations to accommodate the dynamism of the communication requirements in multi-domain NFV ecosystems.

III. Design of the inter-domain connectivity orchestration service

The design of the connectivity orchestration service allows the creation of virtual networks between different NFV infrastructure domains. These networks are automatically provisioned on-demand, during the deployment process of a multi-domain 5G service. VNFs of the same service that are deployed on different NFV domains can be allocated an inter-domain virtual network, and attach to that network through one of their network interfaces. The orchestration service guarantees the provision of link-layer connectivity among all the VNFs that are connected to the same virtual network, regardless of the actual domain where each of the VNFs is executed. In practice, such connectivity is provided on top of the underlying networks of NFV infrastructure providers and ISPs. The inter-domain connectivity orchestration service is intended to support the complete lifecycle management of these virtual networks, including not only their creation, but also their modification (*e.g.*, extending a virtual network to new NFV domains), and termination. Figure 1 illustrates our abstraction of an inter-domain link-layer virtual network.

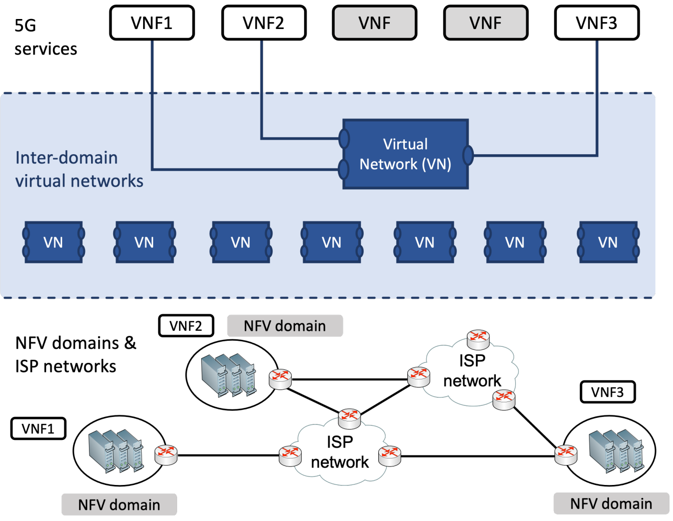

Fig. 1. Abstraction provided to VNFs by inter-domain virtual networks.

Our solution addresses the aforementioned limitations of layer-3 and layer-2 approaches to inter-domain VNF communication:

- It enables the isolated operation of VNFs of the same service, as only VNFs connected to a virtual network can transmit traffic over that network.
- VNFs can be interconnected honoring the NFV descriptor that defines a 5G service, avoiding undesirable network configurations by service providers and ISPs after a service deployment.
- Remote VNFs connected to a link-layer virtual network intrinsically share a broadcast domain. This allows the distribution of multicast and broadcast traffic among VNFs, regardless of their location and the network policies and configurations established on ISP networks.
- It uses a set of versatile software elements that support the automated and on-demand establishment of layer-2 VNF connectivity among NFV domains, without requiring link-level data paths on the ISP networks interconnecting them.

Figure 2 illustrates the design of the connectivity orches-

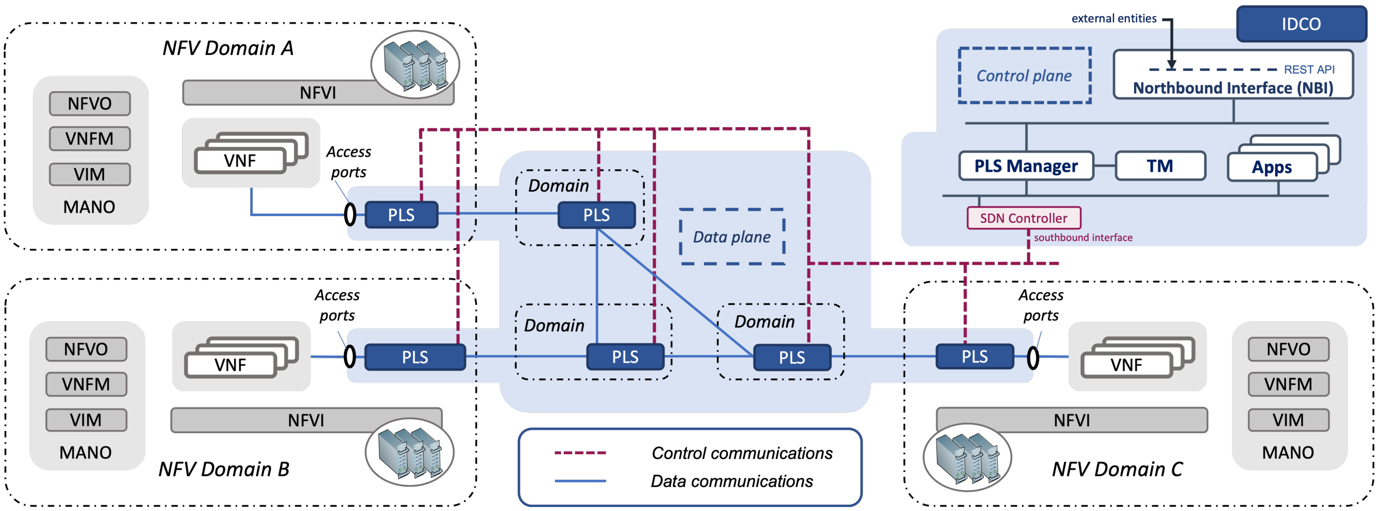

Fig. 2. Overview of the inter-domain connectivity orchestration service design.

tration service. It encompasses different logical components within a data and a control plane. Our design choices were guided by the following criteria: 1) automation in the creation of inter-domain virtual networks and the flexibility to implement different management policies for exchanging VNF data traffic between NFV domains; 2) ensuring that a service descriptor remains independent of the multi-domain nature of the NFV ecosystem (a 5G service should function correctly based on the service descriptor, regardless of the specific NFV domains involved in its deployment); 3) maintaining compatibility with existing initiatives and facilitating practical implementation by leveraging well-established and widely adopted NFV/SDN specifications and Internet protocols.

A. Data plane components

The service design incorporates data plane elements responsible for data forwarding functionalities. They are deployed at the network edge of every NFV domain to facilitate inter-domain VNF communications. These elements, referred to as **PLS** in Figure 2, operate as programmable layer-2 switches. PLS elements provide a number of access ports at every NFV domain to attach VNFs and support their inter-domain communications.

PLS elements can be implemented in software and be deployed as regular VNFs on NFV domains. This allows exploiting the inherent advantages of the NFV technology, such as the flexibility to change the allocation of resources to a PLS deployment, for instance by incorporating additional PLS instances or by scaling them vertically.

To support data plane communications across NFV domains, PLS elements are interconnected through point-to-point links. These links can be established over the underlying networks of NFV infrastructure providers and ISPs, using standard IP tunneling protocols, such as instance Virtual eXtensible Local Area Networks (VXLANs) or Generic Routing Encapsulation (GRE). Moreover, traffic exchanged between PLS entities may be protected through state-of-the-art security mechanisms, *e.g.*, IP security (IPsec).

The set of interconnected PLS elements creates an **overlay network** spanning all NFV infrastructure domains. This network can provide redundant links between domains and support multiple end-to-end communication paths for remote VNFs.

B. Control plane components

In our solution, the overlay network can be programmed by installing traffic forwarding rules on the PLS elements. This characteristic allows for the separation of inter-domain VNF communications into isolated virtual networks, which are built on top of the overlay network. It also provides flexibility in managing data communications among VNFs connected to the same inter-domain virtual network. This includes ensuring a minimum bandwidth, enabling shortest-path communications with different metrics, enhancing resilience in case of path failures, or performing load balancing in PLS elements.

In the architectural design of Figure 2, the **Inter-Domain Connectivity Orchestrator (IDCO)** is the control plane entity that exploits such programmability. It features an SDN-based control interface to interact with the PLS elements. This interface enables the management of forwarding rules on these devices. The IDCO comprises four modules:

- The **Northbound Interface (NBI)** module provides the point of contact for the connectivity orchestration service. It facilitates the creation, modification, and deletion of inter-domain virtual networks. The NBI also offers a user interface for managing the topology of the overlay network, creating and deleting links between NFV domains, and exposing information on its topology, statistics, and status. It may report significant events like point-to-point link failures between PLS elements.
- The **PLS Manager** module acts as an SDN application that manages traffic forwarding rules on PLS elements. It interacts with the Traffic Management (TM) module to identify the PLS elements involved in an inter-domain virtual network and determines the traffic forwarding

This article has been accepted for publication in IEEE Software. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MS.2023.3304854

rules for each PLS. These rules are installed through an SDN controller that provides a southbound interface to the programmable overlay network (*e.g.*, based on OpenFlow[13], or NETCONF [14]). The PLS Manager also manages the overlay network topology, collects statistics, and detects unexpected events like link failures, reacting to them.

- The **Traffic Management (TM)** module determines the network paths for VNF data flows on the overlay network. To this purpose, it may consider traffic management policies or the link status characterized by the PLS Manager, facilitating the application of traffic engineering principles.

The modular design of the IDCO allows for flexible integration of supplementary functionalities. For instance, to monitor data traffic transmitted over an inter-domain virtual network, or to temporarily deactivate and reactivate VNF communications for security purposes.

C. Operational aspects

Our solution considers a transport provider that has Service Level Agreements (SLA) with a set of NFV infrastructure providers. The transport provider can instantiate a multi-domain network service including a PLS element in each NFV domain. This service can be deployed as any other network service of the multi-domain NFV ecosystem, using the existing MANO facilities. Using the PLS elements, along with its own network infrastructure, the transport provider builds an overlay network that spans every NFV domain.

A service provider may now request the deployment of a service (referred to as the *target service*) from the MANO system of the NFV ecosystem. The *target service* may be built by multiple VNFs, which have to be deployed on different NFV domains (*e.g.*, to ensure appropriate latency for end-users). Some VNFs may require link-layer connectivity (*e.g.*, to isolate those VNFs from other services, or to share a broadcast/multicast domain).

As part of deploying the target service across the intended NFV domains, the MANO platform requests the creation of inter-domain virtual networks to the IDCO, identifying the point of attachment of each VNF to its corresponding PLS element. The IDCO then installs the necessary forwarding rules in the PLS elements to create the inter-domain virtual networks and support link-layer inter-VNF communications. Once the target service is available, the MANO system notifies the service provider accordingly.

IV. Validation of the solution

A. Prototype implementation

We have developed a basic prototype of each component of the architectural design shown in Figure 2 to validate the feasibility of our design. The PLS implementation is an enhancement of the prototype presented in [7]. It supports SDN-based operations using Open vSwitch, an open-source programmable layer-2 switch implementation. SDN programmability in Open vSwitch is supported through the OpenFlow protocol [13]. Point-to-point links between PLS elements are created using Linux VXLAN interfaces. These links are protected with standard IP security (IPsec). The PLS prototype has been packaged as a VNF, using a virtual machine with Linux Ubuntu Server 18.04 LTS, 1 GB RAM, 1 vCPU, and 20 GB disk storage.

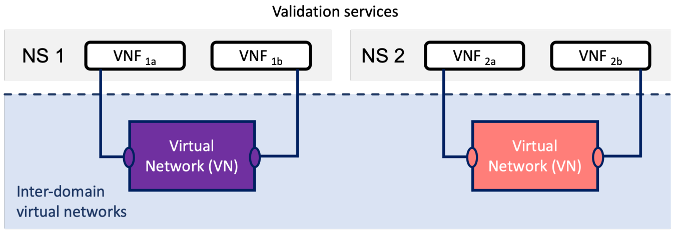

Fig. 3. Services defined to perform the experimental validation of the IDCO.

The PLS Manager and TM modules are implemented as a single SDN application using Ryu [15], an open-source SDN framework. The PLS Manager utilizes the NetworkX Python package to represent the overlay network topology as a graph, which can automatically be discovered and updated through the Ryu SDN controller. The TM module utilizes NetworkX to calculate the shortest paths on the graph, using link attributes such as the available bandwidth and latency and the Dijkstra's algorithm.

For validation purposes, our IDCO implementation supports the creation of point-to-point inter-domain virtual networks, enabling data communications between pairs of VNFs (to implement multi-point virtual networks is part of our future work). The NBI module acts as an HTTP application programming interface that accepts requests to create such networks between NFV domains. The IDCO prototype implementation has been installed on a virtual machine with a Linux Ubuntu Server 18.04 LTS distribution, 2 vCPUs, 4GB RAM, and 20 GB disk storage.

B. Proof-of-concept and results

We have accomplished a functional validation of our proposal through a proof-of-concept, using the prototypes of the IDCO and the PLS. The proof-of-concept encompasses the execution of two network services on a multi-domain NFV ecosystem (*NS 1* and *NS 2* in Figure 3).

NS 1 is built by two VNFs, VNF_{1a} and VNF_{1b}, that must be interconnected at layer-2 across different NFV domains. We will use our solution to create an inter-domain virtual network and support the connectivity of both VNFs. Thus, they will be able to communicate as if they were attached to a layer-2 switch, satisfying the connectivity requirements of the network service. *NS 2* also has two interconnected VNFs that must be deployed on separate NFV domains, VNF_{2a} and VNF_{2b}. Our solution will support their connectivity using an inter-domain virtual network. Inter-domain virtual networks will isolate data communications of *NS 1* from *NS 2*, and vice-versa.

The proof-of-concept has been realized on a multi-domain NFV ecosystem available at 5TONIC. Figure 4 represents the various components of the proof-of-concept and their

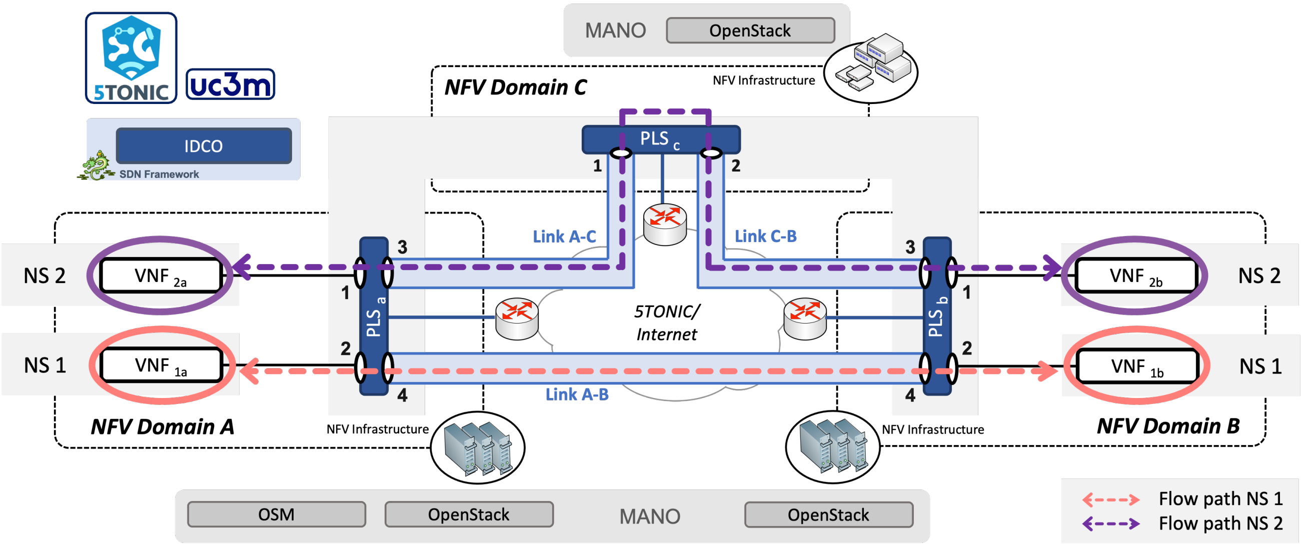

Fig. 4. Multi-domain scenario configured for experimentation.

deployment on the NFV ecosystem. The ecosystem comprises three distinct NFV domains. Domains A and B are built by commercial off-the-shelf server computers. Domain C is composed by a rack of resource-constrained devices (single board computers Raspberry Pi). Each NFV domain uses OpenStack as the virtual infrastructure manager (VIM). The management and orchestration of resources across the three NFV domains are facilitated by the ETSI Open Source MANO (OSM) software. The NFV ecosystem is presented in detail in [7]. The ecosystem was extended with an instance of the IDCO, hosted on a virtual machine collocated with the OSM software.

As an initial step, we created an NFV service descriptor with three PLS VNFs. The descriptor includes the definition of the point-to-point links to be established between PLS elements as VXLAN tunnels (links A-B, A-C, and C-B in the figure), and the necessary configuration actions to automatically set-up such links at instantiation time. The OSM software was then used to automatically deploy the service, instantiating a PLS VNF on each of the domains and creating their point-to-point links. The PLS VNFs configuration concluded with the registration of their respective programmable switching functions (provided by Open vSwitch) within the IDCO. Figure 5.a shows the registration of these three VNFs in the IDCO through OpenFlow (representing the OpenFlow directives exchanged between the PLS Manager module and the PLS VNFs). These registration events started at the beginning of the experiment, allowing the PLS Manager to discover the overlay network topology formed by the PLS VNFs.

The proof-of-concept proceeded with the deployment of *NS 1*. In the experiment, VNF_{1a} was deployed on domain A as a traffic generator, whereas VNF_{1b} served as a traffic sink on domain B. The OSM software was instructed to attach VNF_{1a} to port 2 of the PLS VNF in domain A (this was facilitated by the use of OpenStack provider networks). Similarly, VNF_{1b} was connected to the same port of the PLS VNF in domain B. Once both VNF were operational, approximately within 25 s of initiating the experiment, we made a request to the IDCO NBI to create an inter-domain virtual network between domains A and B. The request specified the allocation of port 2 in both PLS_a and PLS_b to the virtual network, along with the MAC addresses of VNF_{1a} and VNF_{1b}. These addresses were explicitly specified, although they can be obtained from OSM in a realistic situation.

In this proof-of-concept, the default traffic management policy prioritizes network paths based on Dijkstra's lowest cost algorithm, favoring links that do not accommodate additional flows. The costs assigned to each link are determined through a preliminary performance analysis, considering available bandwidth and round-trip-time (RTT) metrics. The analysis revealed the following results: link A-B offers an available bandwidth and RTT of 898 Mb/s and 1.716 ms, respectively; link A-C, 286 Mb/s and 4.492 ms; and link C-B, 288 Mb/s and 4.399 ms. Thus, we assign a lower cost to link A-B since it provides better performance.

Following the request to the IDCO, the TM module determines that data communications between VNF_{1a} and VNF_{1b} should utilize the network path conformed by link A-B. The PLS Manager then configures it using the SDN controller, installing traffic forwarding rules on PLS_a and PLS_b. These interactions take place over the VNFs management network, and are depicted in Figure 5.a with two events occurring at approximately 25 s. After creating the virtual network, we initiated a 10 Mb/s UDP traffic transmission from VNF_{1a} to VNF_{1b}, emulating the streaming of high-definition multimedia content. To confirm that the data traffic traversed the configured network path, we captured traffic on port 4 of PLS_b. The average traffic throughput observed on PLS_b is represented by a continuous blue-gray line in Figure 5.c. Likewise, Figure 5.b illustrates the impact of the established network path on the RTT of the transmission.

Using a similar procedure, we then deployed *NS 2*. The NFV domains and the PLS ports allocated to VNF_{2a} and VNF_{2b} are shown in Figure 4. Within 75 s of initiating the experiment, we performed a new request to the IDCO NBI to create an additional virtual network between domains A and B. According to the default traffic management policy, the TM selected an unused network path through links A-C and C-B, aiming to balance the traffic load in the overlay network. This resulted in the configuration of appropriate traffic forwarding rules on PLS_a, PLS_b, and PLS_c, as shown by the OpenFlow events in Figure 5.a. After the instantiation, we configured another 10 Mb/s UDP transmission from VNF_{2a} to VNF_{2b}. Figure 5.c shows the average traffic throughput traversing port 2 of PLS_c (dashed purple line), validating the utilization of the new data path for the inter-domain virtual network. Figure 5.b illustrates the different RTT of this network path, being aligned with our previous link performance analysis.

As next step, we tested the IDCO's ability to handle unexpected events within the overlay network. The link between PLS_c and PLS_b was brought down, triggering an event that was captured by the PLS Manager. The network topology was updated, and the TM module recalculated the network path for the inter-domain virtual network of *NS 2*. Consequently, the PLS Manager changed the traffic forwarding rules on the PLS VNFs, resulting in an increase of traffic throughput traversing port 4 of PLS_b, and a drop in traffic from PLS_c, as shown in Figure 5.c at 125 s.

Finally, the proof-of-concept involved the migration of VNF_{1a} from domain A to C. The migration event was manually triggered 176 s after the experiment initiation, as depicted in Figure 5.a. Accordingly, the PLS Manager changed the traffic forwarding rules associated with the virtual network of *NS 1*, so as to use links A-C and A-B. Once the migration process was completed, traffic arrives again at its destination (VNF_{1b}). Figures 5.b and 5.c show the RTT and the average throughput on the new network path, respectively.

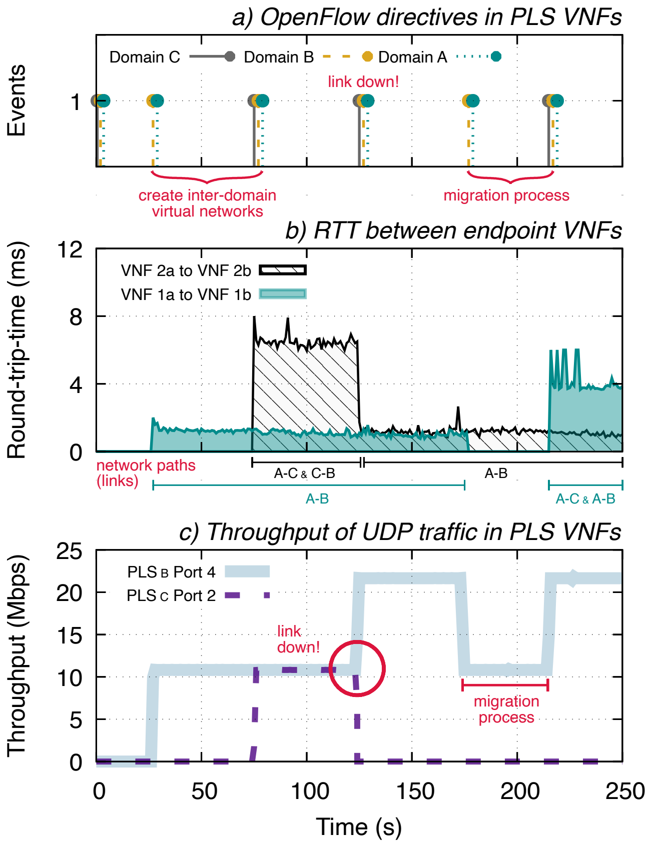

Fig. 5. Performance evaluation collected metrics.

V. Lessons Learned and Future Directions

This paper presents a connectivity orchestration service for multi-domain NFV ecosystems. Our research indicates that current NFV and SDN implementations are mature enough to support the flexible deployment of 5G services in different locations. However, they still present limitations in supporting VNF connectivity across different NFV domains. We have demonstrated the feasibility of leveraging open-source software technologies and standard protocols to address these limitations, creating new innovation opportunities for IT professionals and software practitioners in the field of 5G services. Additionally, we have conducted a thorough proof-of-concept in a realistic multi-domain NFV ecosystem, assessing the viability of our approach. Our orchestration service has demonstrated its capacity to reliably operate even under unexpected connectivity events. Its modular design, the prototype implementation, and the proof-of-concept may serve as a valuable reference of future research in this field.

In addition, our work unfolds novel and promising avenues for future developments. Regarding the design, we utilize a redundant overlay network across NFV domains. Whereas this enables multiple communication paths for inter-domain VNF communications, the overlay network topology and the path selection mechanisms are relatively static. Future research is required to adapt the overlay network to ISP network dynamism, considering bandwidth and delay variations, and implement traffic management policies based on the actual performance of overlay paths. We will also research the application of our connectivity orchestration service to cloud-native ecosystems, given their growing relevance in the NFV evolution.

Finally, we aim to contribute a full-featured implementation of our solution to the open-source software community. Consequently, further developments are needed, including multi-access virtual network creation and automated discovery of virtual network endpoints.

Acknowledgment

This article has partially been supported by the European H2020 5GZORRO Project (grant agreement 871533), and the TRUE5G project funded by the Spanish National Research Agency (PID2019-108713RB-C52/AEI/10.13039/501100011033).

References

[1] N. Slamnik-Kriještorac, H. Kremo, M. Ruffini, and J. M. Marquez-Barja, "Sharing distributed and heterogeneous resources toward end-to-end 5g networks: a comprehensive survey and a taxonomy," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 3, pp. 1592–1628, 2020.

[2] M. Condoluci and T. Mahmoodi, "Softwarization and virtualization in 5G mobile networks: Benefits, trends and challenges," *Computer Networks*, vol. 146, pp. 65–84, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128618302500

[3] B. Han, V. Gopalakrishnan, L. Ji, and S. Lee, "Network function virtualization: Challenges and opportunities for innovations," *IEEE communications magazine*, vol. 53, no. 2, pp. 90–97, 2015.

[4] R. V. Rosa, M. A. S. Santos, and C. E. Rothenberg, "MD2-NFV: The case for multi-domain distributed network functions virtualization," in *2015 International Conference and Workshops on Networked Systems (NetSys)*. IEEE, 2015, pp. 1–5.

[5] K. Mahmood *et al.*, "Design of 5G End-to-End Facility for Performance Evaluation and Use Case Trials," in *2019 IEEE 2nd 5G World Forum (5GWF)*, 2019, pp. 341–346.

[6] M. Gupta *et al.*, "The 5G EVE End-to-End 5G Facility for Extensive Trials," in *2019 IEEE International Conference on Communications Workshops (ICC Workshops)*, 2019, pp. 1–5.

[7] I. Vidal, B. Nogales, D. Lopez, J. Rodríguez, F. Valera, and A. Azcorra, "A Secure Link-Layer Connectivity Platform for Multi-Site NFV Services," *Electronics*, vol. 10, no. 15, p. 1868, 2021.

[8] J. A. Wickboldt, W. P. De Jesus, P. H. Isolani, C. B. Both, J. Rochol, and L. Z. Granville, "Software-defined networking: management requirements and challenges," *IEEE Communications Magazine*, vol. 53, no. 1, pp. 278–285, 2015.

[9] Q. Duan, N. Ansari, and M. Toy, "Software-defined network virtualization: An architectural framework for integrating SDN and NFV for service provisioning in future networks," *IEEE Network*, vol. 30, no. 5, pp. 10–16, 2016.

[10] J. Serrano *et al.*, "Design, Implementation, and Validation of a Multi-Site Gaming Streaming Service Over a 5G-Enabled Platform," *IEEE Transactions on Broadcasting*, 2022.

[11] F. Paganelli, P. Cappanera, and G. Cuffaro, "Tenant-defined service function chaining in a multi-site network slice," *Future Generation Computer Systems*, vol. 121, pp. 1–18, 2021.

[12] E. C. Rosen and L. Andersson, "Framework for Layer 2 Virtual Private Networks (L2VPNs)," RFC 4664, Sep. 2006. [Online]. Available: https://rfc-editor.org/rfc/rfc4664.txt

[13] Open Networking Foundation (ONF), *OpenFlow Switch Specification v1.0-v1.5*, [Online]. Available: https://opennetworking.org/software-defined-standards/specifications/ (accessed on August 8, 2023).

[14] R. Enns, M. Björklund, A. Bierman, and J. Schönwälder, "Network Configuration Protocol (NETCONF)," RFC 6241, Jun. 2011. [Online]. Available: https://www.rfc-editor.org/info/rfc6241

[15] Ryu, *component-based software defined networking framework*, [Online]. Available: https://ryu-sdn.org (accessed on August 8, 2023).

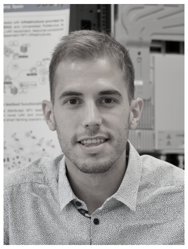

Borja Nogales (bdorado@pa.uc3m.es) received the Ph.D. in Telematics Engineering in 2022 from the Universidad Carlos III de Madrid (UC3M), where he is currently working as Post Doc. His research interests include Network Functions Virtualization (NFV), 5G networks, and Unmanned Aerial Vehicles (UAVs). He has been involved in several international and national research projects, including the H2020 5GZORRO, FISHY, 5GINFIRE, and TRUE5G.

Ivan Vidal (ividal@it.uc3m.es) received the Ph.D. in Telematics Engineering in 2008 from UC3M, where he works as Tenured Associate Professor. He has been involved in several international and national research projects, including the H2020 5GZORRO and TRUE5G, and has published more than 60 scientific papers in several conferences and international journals.

Francisco Valera (fvalera@it.uc3m.es) received the Telecommunication Engineering degree from the Technical University of Madrid (UPM), in 1998, and the Ph.D. degree in telecommunications from the UC3M, in 2002. He is currently a Tenured Associate Professor and the Deputy Director of the Telematics Engineering Department with UC3M.

Victor Sanchez-Aguero (visanche@it.uc3m.es) received the Ph.D. in Telematics Engineering in 2022 from the UC3M. His research interests include UAVs, SDN, NFV, and 5G networks. He has been involved in different international and national projects, including the H2020 LABYRINTH and TRUE5G.

Diego R. Lopez (diego.r.lopez@telefonica.com) joined Telefonica I+D in 2011, and is currently in charge of the Technology Exploration activities within the GCTIO Unit. He is especially focused on virtualization, data-driven management, new architectures, and security. Diego chairs ETSI ISG ZSM and ETSI ISG PDL.