Baranda, J., Mangues, J., Martínez, R., Vettori, L., Antevski, K., Bernardos, C.J. y Li, X. (2020). Realising the Network Service Federation vision. *IEEE Vehicular Technology Magazine, Future Networks Initiative Special Issue on 5G Technologies and Applications*.

Realising the Network Service Federation vision

Jorge Baranda ,CTTC/CERCA; Josep Mangues-Bafalluy, CTTC/CERCA; **Ricardo Martínez, CTTC/CERCA; Luca Vettori, CTTC/CERCA; Kiril Antevski,** University Carlos III; **Carlos J. Bernardos,** University Carlos III; **Xi Li,** NEC Laboratories Europe

The 5G-TRANSFORMER project proposes an NFV/SDN-based architecture to manage the end-to-end deployment of composite NFV network services, which may involve multiple administrative domains, hence, requiring network service federation capabilities. At the architectural level, this article presents the service federation functionality of the 5G-TRANSFORMER service orchestrator. It covers the gaps identified in ETSI NFV reports and specifications (e.g., IFA028). Some recommendations are also presented based on this experience, particularly on the relevance of multi-domain resource orchestration. Experimental results show that the federated service under evaluation is deployed in less than 5 minutes. Time profiling of the various processing federation-related operation shows its reduced impact in the experienced deployment time. A comparison of service deployments of increasing complexity also offers valuable insights.

Keywords – *Network Service Federation, End-to-End service orchestration, NFV/SDN, Experimental evaluation, 5G mobile transport*

1. Introduction

5G brings a deep revolution in the network ecosystem. This revolution implies a redesign of the overall network architecture towards a *network of services*. Complexity, heterogeneity and dynamicity will be the rule to serve the diverse requirements imposed by multiple vertical industries (e.g., automotive, eHealth) over the same shared infrastructure, using network slices, or spanning through different administrative domains (ADs). In this context, network automation and programmability become essential to support these objectives. It is widely accepted that Software Defined Networking (SDN) and Network Function Virtualization (NFV) concepts represent the foundations to effectively provide these network capabilities.

To enable such a vision, the 5G-TRANSFORMER (5GT) project [1] is developing an NFV/SDN 5G mobile transport network platform – based on ETSI NFV procedures, specifications and reports – capable of managing the deployment of end-to-end (E2E) network services (NSs) spanning across complex and heterogenous transport networks, using computing resources deployed in different Points of Presence (PoPs) and involving multiple ADs (each one managed by a different 5GT platform instance).

We refer to the NS deployment in different ADs, each one using its own management and orchestration (MANO) platform as *network service federation (NSF)*. The 5GT platform can perform NSF thanks to its capabilities to orchestrate composite NS descriptors (NSDs) [2].

Composite NSs are composed by a set of constituents nested NSs. These different nested NSs may need to be instantiated in different ADs due to different reasons like shortage of resources, NS availability or the needs to deploy an NS while satisfying certain requirements (e.g, coverage of a given geographical area). Then, composite NSs allows designing NSs not only as integral units that work on their own, but also as sets of NSs that can also be dynamically grouped to create more complex and tailored offerings to vertical industries. Hence, NSF can provide enlarged service coverage and connectivity, seamless service continuity and flexibility to vertical industries. This enables the service providers (SPs) to offer a broader spectrum of NSs to vertical industries, opening the door to new business models. For instance, vertical industries may set up a business relationship with a single SP satisfying its requests, while this SP establishes agreements with other SPs to extend its service offerings, thus increasing its profit by avoiding the rejection of service deployment requests.

In the literature, the NSF problem has been split into two. First, there is the problem of mapping the different parts of the NSs across multiple ADs. The work in [3] presents a solution based on Integer Linear Programming and compares it with other previous solutions through simulation. However, these solutions do not delve into the required procedures and interfaces to truly deploy such NSs in different ADs, which is the second part of the problem and the focus of this work.

In that sense, different reports of Standard Development Organizations, like the Open Networking Foundation [4], the Metro Ethernet Forum [5] and the ETSI NFV [6] propose recommendations for procedures and interfaces to achieve

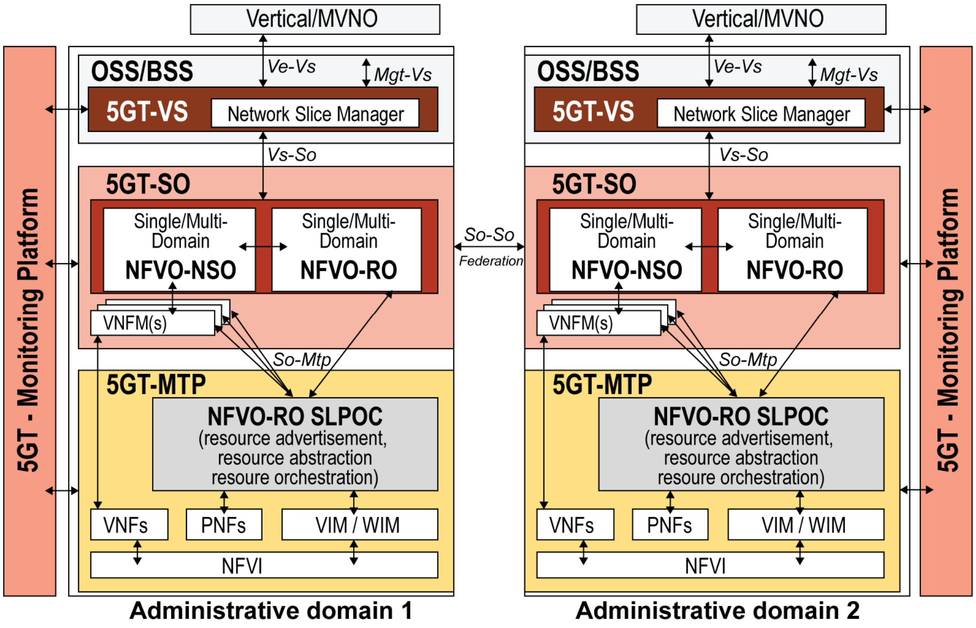

Figure 1. 5G-TRANSFORMER Platform Architecture

NSF. The ETSI NFV work, which is the base of the 5GT platform, goes beyond and defines an interface in the ETSI-NFV IFA030 specification [7] based on the ETSI-NFV IFA028 report [6]. The IFA030 specification presents a description of the interface between orchestrator platforms (*Or-Or*) to perform MANO operations among multiple ADs, which is the new interface reference point identified by IFA028. Basically, IFA030 considers using, for the *Or-Or* interface, a subset of the operations defined in ETSI NFV-IFA 013 [8] for the NSD Management and NS Lifecycle Management interface. The 5GT platform adopts these operations as a baseline to perform NSF. However, IFA030/013 are a service level interfaces (and do not handle resources). The potential architectural options to support NSF are only reported at high level by IFA028, which does not address properly an essential aspect of NSF: the interconnection of nested NSs running in different ADs. This is an often disregarded functionality that is fundamental to make E2E multi-domain orchestration and network slicing a reality in practice. In fact, a slice is not only the logical construct defining a E2E virtual network, but also the corresponding underlying resources assigned in exclusivity to a given tenant over a shared infrastructure [9]. This gap is tackled in this work, where we propose enhancements to the work of IFA028 by defining a more detailed workflow, which includes a set of additional interactions between the different Service Orchestrators to achieve the interconnections of nested NSs at the resource level. To the best of our knowledge, this is the first work realising and evaluating a complete NSF procedure, including the deployment of a composite NS involving different SPs in a real multi-PoP, multi-technology and multi AD experimental infrastructure.

The main contributions of this work are: (*i*) providing a set of recommendations at the Service Orchestrator architectural level to translate high-level ETSI NFV report guidelines into a feasible operational implementation of NSF, including the two main aspects of the E2E orchestration functionality, namely service and resource orchestration; (*ii*) presenting the complete workflow that implements NSF to explain what should be coordinated among peering domains at service and resource levels; and, (*iii*) evaluating experimentally the automatic deployment of a composite NS in two ADs managed by different SPs, each running its own 5GT platform.

The rest of the article is structured as follows. Section 2 describes the service orchestration solution of 5GT, by first introducing the high-level 5GT architecture, and then focusing on the 5GT Service Orchestrator (5GT-SO), providing further details about its internal architecture and the followed workflow to support NSF. Section 3 presents the multi-technology transport network, multi-pop, and multi-AD setup where the experimental evaluation (in Section 4) is carried out. The aim of this evaluation is to measure the *service deployment time* and the impact of the federation-related operations when deploying a composite NS designed for the eHealth use case of the 5GT project. Finally, Section 5 concludes this work.

2. 5G-TRANSFORMER Service Orchestration

This section first introduces the general 5GT architecture, and then presents a more detailed description (software architecture and proposed workflow to support NSF) of the 5GT-SO building block.

2.1. 5G-TRANSFORMER architecture description

The 5GT architecture is based on the ETSI NFV work and aims at providing a platform with flexible and dynamic management capabilities to accommodate multiple and heterogeneous services derived from different vertical industries. This platform concurrently maps such services into a shared infrastructure combining multiple heterogeneous resources in terms of computing, storage and networking. The source code of the 5GT platform is available in the 5GT repository [15] as open source code under Apache 2.0 licence.

Figure 1 presents the main building blocks of the 5GT platform. It is worth mentioning that the 5GT platform includes a certain level of security in its operations, although this is not its main scope. All blocks within the 5GT architecture keep a register of allowed associated blocks and requesters for each action to avoid not authorised third-party entities to launch or modify ongoing operations.

The Vertical Slicer (5GT-VS) is the frontend for vertical industries to the 5GT platform. It offers a high-level interface allowing users to easily request vertical services. 5GT-VS only focuses on the service and business demands without taking over how services are eventually deployed at the resource level. The 5GT-VS offers a catalogue of vertical services, which are particularised by the vertical users with their requirements. The internal logic of the 5GT-VS translates business-oriented service requirements into slice-related requirements to manage the lifecycle of network slices, which are deployed as NFV Network Services (NFV-NSs).

The Service Orchestrator (5GT-SO) oversees the E2E orchestration and the lifecycle management of the NFV-NSs based on the available resources advertised by the underlying Mobile Transport and Computing Platform (5GT-MTP). The 5GT-SO embeds the network service orchestrator (NFVO-NSO) and the resource orchestrator (NFVO-RO), whose functionalities are equivalent to those performed by typical NFV orchestrators (NFVO) [10].

In 5GT, NFV-NSs may embrace one or multiple ADs. It means that the 5GT-SO, besides interacting with its local 5GT-MTP (within a single AD), may interwork with other 5GT-SOs governing remote domains through the *So-So* interface, hence performing NSF.

The 5G-MTP is the unified controller responsible for orchestrating resources: VNF instantiation and, connectivity management at the underlying physical transport network. The 5G-MTP has full control of the underlying resources by interacting with different managers: the Virtual Infrastructure Manager (VIM), for managing compute resources, or the WAN Infrastructure Manager (WIM), for enabling network connectivity among different PoPs. For the sake of scalability and to simplify 5GT-SO operations, the 5GT-MTP applies abstraction mechanisms when exposing the resource view towards the 5GT-SO.

Finally, the 5GT Monitoring Platform (5GT-MON) provides metrics to the 5GT platform so it can perform reactive actions to continuously ensure targeted service level agreements. More specifically, the monitoring service at 5GT-MTP collects data about the local physical and virtual resources; the 5GT-SO monitoring service collects data about the managed VNFs and NFV-NSs; and the 5GT-VS monitoring service collects data about network slices and vertical services.

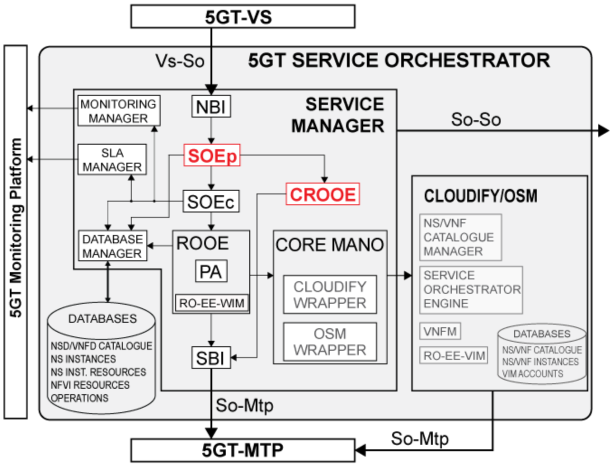

Figure 2. 5GT-Service Orchestrator software architecture

2.2. 5GT-SO Internal Software Architecture

Figure 2 presents the internal software architecture of the 5GT-SO and its relationship with the 5GT-VS, the 5GT-MTP and the 5GT-MON blocks.

As explained in [11], the 5GT-SO introduces the concept of Service Manager (SM), which allows the exploitation of open-source production-level MANO platforms (e.g., OSM and Cloudify) by means of the use of wrappers. This approach allows focusing on innovative functionalities, such as the network service federation, being easily supported regardless of the specific MANO platform.

The architecture of the 5GT-SO is organised in four main building blocks. First, the SM, which represents the "*brain*" of the 5GT-SO. It embeds the service, resource and monitoring orchestration logic, dispatching relevant tasks to the other building blocks according to the corresponding operational workflow. Second, the Core MANO platform (i.e., OSM or Cloudify), which interacts with the 5GT-MTP to handle the operations related to computing resources. Third, the placement algorithm (PA), which runs as a separate process of the 5GT-SO. The PA interacts with the resource orchestrator module of the 5GT-SO through a well-defined REST API, allowing the exchange of PA algorithms [12]. Finally, the databases maintain the state of the 5GT-SO system in terms of service offering, instantiated network services, 5GT-MTP used resources, etc.

2.3. Network Service Federation Workflow

The 5GT-SO handles NSF in a transparent way to the 5GT-VS (and to the vertical user) thanks to two additional blocks introduced in the SM with respect to the baseline architecture presented in [11]. In the same way any single-domain SO has two main functionalities (service and resource orchestration), a multi-domain-capable SO, like the 5GT-SO, must introduce the corresponding service and resource orchestration entities for coordination with peering domains. These blocks (highlighted in Figure 2), are a hierarchical (parent-child) Service

Orchestration Engine (SOE), and a Composite Resource Orchestrator Engine (CROOE). In the hierarchical SOE, the SOE parent (SOEp) focuses on the multi-domain service orchestration aspects, by coordinating with peering SOEps. In this sense, it analyses the incoming NS requests, and orchestrates the required operations at the different underlying sub-modules depending on such request. It also coordinates the interaction with the resource orchestration block (CROOE), which handles the interconnections between nested NSs. In case of NSF, CROOEs of peering domains jointly handle the resources at both sides to establish the required connection. This is the fundamental step that enables a full E2E multi-domain orchestration.

The NSF workflow is explained next. It assumes that there has been a previous *"offline"* agreement process between peering domains. In this offline process, service providers: (*i*) establish physically the control plane and data plane links between ADs and update their 5GT platforms accordingly to incorporate information about federated domains and the available links; and (*ii*) onboard composite and nested NSs to the 5GT platforms of the different peering ADs, so it is known in advance in which AD a nested NS can be deployed. This workflow follows a similar approach to the *top-down* solution presented in the IFA028 report. However, this work enhances the IFA028 solution by explicitly dealing with the interconnection of nested NSs. Achieving this interconnection requires a set of new additional steps where peering CROOEs exchange information related to instantiated VNFs through the defined *So-So* interface depicted in Figure 2. The specification of these exchanged messages is available at [15].

For a single NS, the SOEp relies on the SOE child (SOEc) to perform the NS instantiation operation following the original orchestration logic explained in [11] using ETSI NFV IFA 014, 011, 013 specifications.

For a composite NS, the set of operations depend on where the different nested services in a composite NS can be instantiated (either at the consumer or the provider domain).

Initially, the SOEp decomposes the composite NS between nested NSs to be instantiated locally (consumer domain), and nested NSs to be instantiated in a federated domain (provider domain). Then, the SOEp asks the CROOE block to determine how the different nested NSs in the composite NS are interconnected based on the NSD. Next, the SOEp starts the instantiation of the different *local* nested NS using the SOEc as explained previously for a single NS. Differently to the workflow defined in IFA028, our proposal starts dealing with operations at the consumer domain since it considers this domain as the coordinator of the process (as it is the receiver of the NS instantiation request).

When finishing the instantiation of *local* nested NSs, the SOEp launches the instantiation of *federated* nested NSs contacting the 5GT-SO of the corresponding provider domain through the northbound interface (NBI). At the federated domain, the NBI first validates that the request comes from an AD allowed to solicit NSF operation. Due to the presence of additional parameters in the instantiation request, the provider SOEp realises that this is a request from a consumer domain and asks this domain about required information to coordinate the instantiation of the nested NS at the provider domain. This interaction between 5GT-SOs is done between its respective CROOE modules, through the *So-So* interface. Then, the provider SOEp relies on its SOEc to instantiate the *federated* nested NS. While the operation progresses, the consumer domain polls periodically the provider domain to check the instantiation status of the *federated* nested NS. Once the consumer domain has confirmed that the *federated* nested NS has been correctly instantiated, its CROOE module asks its counterpart at the provider domain about the instantiation details of the *federated* nested NS. This information is needed to setup the inter-nested NS connections between the different VNFs and coordinate the instantiation of other possible subsequent *federated* nested NSs while avoiding address collisions between ADs. After the different *federated* nested NSs have been instantiated, the SOEp manages the required interconnections between the different nested NSs as expressed in the composite NSD. First, it asks its CROOE to determine and set the interconnections among VNFs of different nested NSs deployed locally at the multiple PoPs under its control. Then, the CROOE block interacts with the local 5GT-MTP through the southbound Interface (SBI) to set these

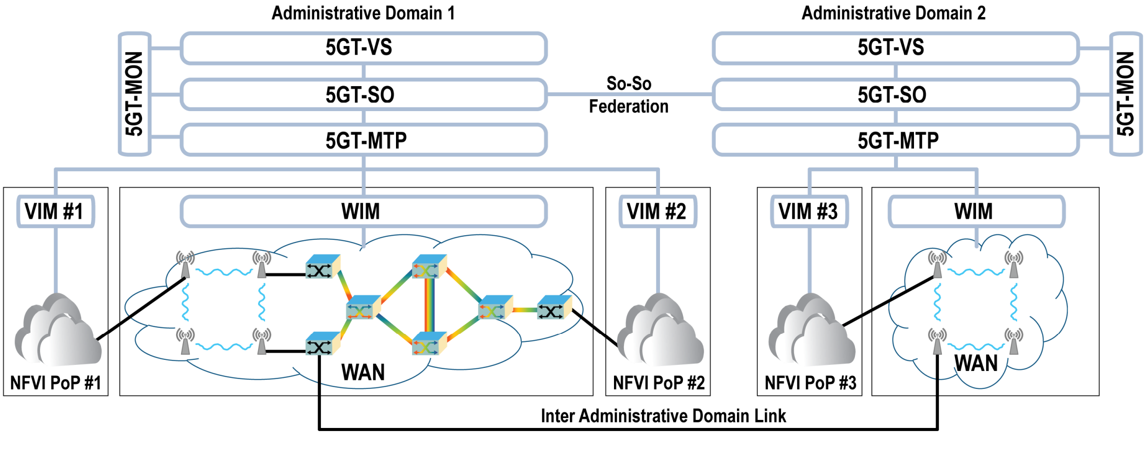

Figure 3. Multi administrative domain 5GT platform setup under evaluation

connections.

Second, the SOEp asks the CROOE to interconnect VNFs of different nested NSs deployed in PoPs of provider/federated domains with the ones instantiated at the local PoPs of the consumer domain. To do so, for each combination of instantiated *"consumer-provider"* nested NSs, peering CROOEs exchange the interconnectivity information (the addresses of pairs of VNFs at each AD requiring interconnection) through the *So-So* interface. Afterwards, each CROOE contacts its 5GT-MTP to set their segment of the inter-domain inter-nested path, that is, from the different PoPs to the point in its transport network where the inter administrative domain link is.

Finally, it is worth mentioning that the design of the presented workflow considers the instantiation of composite NSs having more than two nested NSs, and that the consumer domain could interact with more than one provider domain for the same composite NS. Thus, this approach is functionally more scalable than the approach proposed by IFA028 [6], which only considers a single provider domain and a composite NS constituted by two nested NSs.

In summary, based on this work, there are some key takeaways in terms of architectural design. First, a correct modular design allows adding the NSF/multi-domain functionality on top of already available single-domain orchestration. Second, a multi-administrative domain SO must handle both service and resource orchestration between peering domains. This includes introducing the corresponding entities in both domains and defining the corresponding interfaces and information models at both levels. Third, the service orchestration part requires minimum modifications compared to IFA013/030 for offering the NSF functionality. However, the fundamental component to offer full E2E control is the resource orchestration part, which exchanges relevant resource information between domains and interacts with the resource layer (5GT-MTP) to setup the inter-domain interconnection according to the requirements and to avoid potential resource conflicts in IDs of different domains (e.g., IP addresses).

3. 5G-TRANSFORMER experimental setup

Figure 3 presents the experimental setup, in which, for each administrative domain (AD), we map the 5GT architecture presented in Figure 1. At the control plane level, ADs communicate each other through the *So-So* interface, which traverses over a Virtual Private Network (VPN) connection. At the data plane level, the inter administrative domain links interconnect the transport networks of peering ADs. In this setup, there is a single inter AD link, implemented as well via a VPN connection, hence ensuring the security in the communications between ADs.

Next subsections present the different deployed ADs following a bottom-up approach, that is, from the available physical resources to the MANO system. This experimental setup has been built to satisfy the requirements of the eHealth use case defined within the 5G-TRANSFORMER project.

3.1. Administrative Domain 1

At the AD1, two different PoPs are deployed and managed by respective OpenStack instances. Both PoPs provide storage, computing and networking resources for the NSs to be deployed. PoPs are placed at both ends of the AD to emulate distributed computing resources at the different and remote segments of the transport network, i.e. edge and core.

At the transport network level, AD1 contains two different technological domains. At one domain, a ring of four forwarding elements (FE) use wireless mmWave/WiFi (IEEE 802.11ad/802.11ac) links, representing the edge packet-switched domain of the transport network, as depicted in Figure 3. The other domain represents the *core* part, which relies on a multi-layer network combining three packet switches with an optical wavelength division multiplexing mesh-network counting with two colorless Reconfigurable Optical Add/Drop Multiplexer (ROADM) and two Optical Cross-Connect (OXC) nodes.

This transport network is managed by a hierarchy of SDN controllers. Each of the aforementioned domains is controlled by its dedicated technology aware SDN controller, acting as child controller. On top of them, a parent controller handles the E2E SDN transport orchestration. This parent controller is based on the IETF Application-Based Network Operation architecture [14]. This architecture supports hierarchical deployments thanks to the use of a unified interface both at the Southbound/Northbound Interface (S/NBI): the Control Orchestration Protocol (COP). Finally, child-controllers interact with the underlying FEs. They are configured using the OpenFlow (OF) protocol in the wireless domain and in the packet part of the *core* domain. The optical infrastructure is controlled via an Active Stateful Path Computation Element using the Path Computation Element Protocol. A complete description of this multi-technology transport network and its control plane solution is available at [13].

The 5GT platform interacts with this infrastructure through the 5GT-MTP. The latter considers the parent controller as the WIM, whilst the different OpenStack instances are the VIMs. The 5GT-MTP is governed by the 5GT-SO, which in turn communicates with its associated 5GT-VS and the 5GT-SO of the peering AD2. The MANO platform associated to the AD1 5GT-SO is an instance of OSM Release 3.

3.2 Administrative Domain 2

The AD2, which represents another edge domain, is connected to the AD1 at the output of the core domain, through the Inter AD Link represented in Figure 3. In this case, the depicted ring of four available FEs use WiFi (IEEE 802.11ac) links. These FEs are configured using OF. The AD2 consists of a single PoP

Table 1. NS deployment schemes under evaluation

Deployment scheme	**Description**	**Label in Figure 5**
Single Nested NS MB	MB NS in NFVI PoP #2 of AD1	Nested-MB
Single Nested NS vEPC	vEPC NS in NFVI PoP #2 of AD1	Nested-vEPC
Composite NS Single-Pop	Composite NS (vEPC+MB), both Nested NSs in NFVI PoP #1 of AD1	Compo S-Pop
Composite NS Multi-Pop	Composite NS (vEPC+MB) in NFVI PoP #2 (vEPC) and NFVI PoP #1 (MB) of AD1	Compo M-Pop
Federated NS	Composite NS (vEPC+MB) in NFVI PoP #2 (vEPC) of AD1 and NFVI PoP #3 (MB) of AD2	Federated

managed by a different instance of OpenStack. The transport network uses its own SDN controller, which also uses COP protocol to handle connection requests. The SDN controller and the OpenStack instance are coordinated by the corresponding AD2 5GT-MTP. Likewise, in AD1, the 5GT-MTP is controlled by a 5GT-SO instance, which relies on another instance of OSM Release 3 as MANO platform. This 5GT-SO communicates with its own 5GT-VS and with the peering 5GT-SO at AD1.

4. 5G-TRANSFORMER experimental results

We next report on an experimental analysis aimed at validating the network service federation capabilities of the 5GT platform. To do so, we evaluate the *service deployment time* (SDT) using the setup described in Section 3. We define the SDT as the elapsed time between the consumer 5GT-SO receives the NS instantiation request until the 5GT-SO declares the NS as correctly instantiated.

Five different deployment schemes are evaluated, as summarised in Table 1. These experiments cover the full complexity range: simple ones: instantiation of a single NS; intermediate: composite NS services (both single and multi-PoP); and complex: federation, involving multiple ADs. Each experiment is repeated 10 times.

As depicted in Figure 4, the composite NS under evaluation consists of two nested NSs. This composite NS has been defined within the scope of the eHealth use case defined in the 5GT project. The first nested NS emulates a Virtualised Evolved Packet Core (vEPC), and consists of four VNFs, namely MME, HSS, SGW and PGW. The second nested NS is a monitoring backend (MB) NS, and consists of two VNFs, namely a load balancer (LB) and a processing server. Thanks to the capabilities of the 5GT platform, the eHealth vertical could request the deployment of this second nested NS as part of a composite NS *when* needed (re-using an already deployed nested NS (vEPC NS) in the case of an emergency, for instance) and *where* needed (using NSF to deploy in different ADs to meet certain service constraints).

Figure 5 shows the statistical behavior of the SDT for each of the considered deployment schemes. The box stretches from the 20th and 80th percentiles, including the mean and the median values. The whiskers represent the maximum and minimum values.

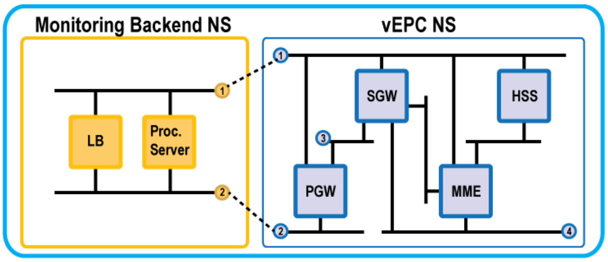

Figure 4. Composite NS under evaluation. Dashed lines show the inter-nested NS connections and numbers represent the NS service access points

As it can be observed in Figure 5, an increase in the number of operations at each deployment scheme has an impact in the total SDT. The maximum measured SDT corresponds to the "*Federated*" scheme (around 270 seconds), which is in line with the 5G target of achieving SDT in the order of minutes. Looking closer at the results, we can conclude that the most time-consuming operations are *(i)* the creation of virtual networks to support virtual links defined in the nested NSs and the instantiation of the VNFs (virtual machine booting up) (see "*Nested-vEPC*" vs "*Nested-MB*"), and *(ii)* the setting of the required connections through the transport network (see "*Compo M-Pop*" and "*Federated*" schemes compared to the others).

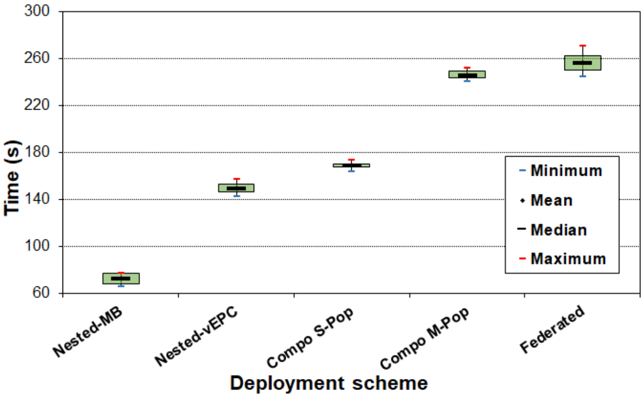

Figure 5. Boxplots of Service Deployment Time for each of the considered deployment schemes

It is also worth highlighting two more learnings from Figure 5. One is the reduced difference between schemes "*Compo S-Pop*" and "*Nested-vEPC*" (around 13% on average), which is due to the difference in time experienced when deploying with VIM#2 (slower) in comparison to VIM#1. The other one, is the

Table 2. Description of processing and operations associated to the network service federation procedure

Component	Description	Label in Figure 6
SOEp processing	Decomposition between local/federated nested NSs of the composite NSD, iteration over the different nested NSs	SOEp proc
CROOE analysing inter-nested NS connections	Analysis of the composite NSD to determine inter-nested NS connections at the virtual link level	CROOE nets
SOEp getting Federated Nested NS ID	Before instantiating in the provider domain, SOEp needs to ask for an ID according to IFA013 procedure	SOEp get Fed ID
SOEp getting Federated instantiation result	After instantiation of Federated nested NS, consumer domain asks the result to the provider domain to update its database registries	SOEp get Fed Res
CROOE determining inter-nested VNFs connections	After instantiating all nested NSs and based on the available information in the DB's, CROOE at consumer domain determines the interconnection between VNFs of different deployed nested NSs	CROOE interNS

reduced difference between "*Federated*" and "*Compo M-Pop*" schemes (around 5% on average), which is due to the proximity between ADs (both placed in the same laboratory premises) and more importantly, the lower complexity of the AD2 transport network and its managing entity (WIM#2) with respect to AD1. Although setting twice the number of connections in the "*Federated*" scheme with respect to the "*Compo M-Pop*" scheme (due to the involvement of both AD transport networks), the lower complexity of the AD2 transport network keeps the experienced SDT in similar values.

Figure 6 shows the statistical behavior (using the same representation as in Figure 5) of the elapsed time in the processing operations at the entities introduced at the 5GT-SO to handle NSF operations and the interactions between ADs.

Table 2 presents a description of such operations. Figure 6 reflects the time values experienced from the consumer domain perspective (i.e., the one driving all the instantiation process). As we can observe, these operations do not contribute significantly to the SDT attending to the time values presented in Figure 5, being the processing carried out at the SOEp the most significant component. However, as mentioned previously, the impact of all these interactions is low due to the proximity between the ADs.

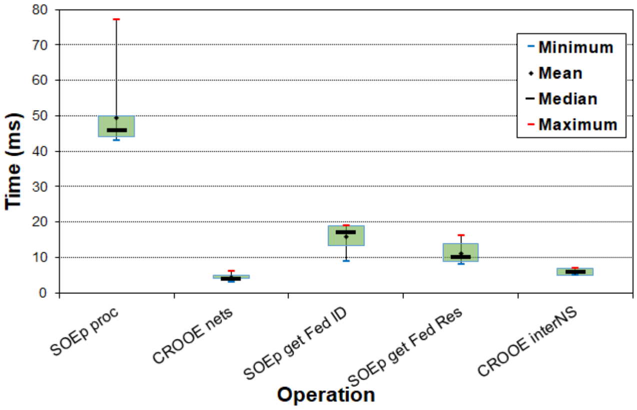

Figure 6. Boxplots of additional operations involved in the network service federation process

Nevertheless, with longer distances, e.g., one 5GT platform in Europe and another in Asia, the impact would be in the order of milliseconds, not having a significant relevance in the final experienced SDT.

5. Conclusions and future work

The 5G-TRANSFORMER project designed an NFV/SDN-based architecture to manage the deployment of composite NSs across multiple AD in what is referred to as network service federation. This article describes the new entities and workflow introduced in the service orchestrator. The proposed enhancements cover the gaps identified in the work of ETSI-NFV [6][7]. General recommendations on NSF architectural design are also provided based on experience, most notably, the importance of multi-domain resource orchestration if full E2E connectivity is to be provided.

To the best of our knowledge, our solution is the first NFV NSF solution to be designed and implemented. Furthermore, it is evaluated using a real infrastructure with multiple technological domains. This evaluation compares various deployment schemes of increasing complexity (ranging from regular NS deployment, to composite NS deployment, and finally, federation) to identify the impact in SDT of associated operations. The main components contributing to the SDT are the intrinsic operations associated to the instantiation of NSs, like booting up virtual machines or setting connections through the transport network interconnecting PoPs. However, the impact of processing federation-related operations and the latency in the link between orchestrators is limited. Ranges of values for each of these deployment schemes can also be taken as reference. For instance, in our setup, federated service creation is in the order of minutes (lower than five).

There are two main ways in which this work can be extended. First, at the architectural level, by introducing novel Distributed Leger Technology solutions to automate the exchange of NSD catalogues between peering ADs. Second, at the algorithmic level, once the framework is in place, it can be extended by adding intelligence (e.g., through machine learning) to the decision making processes to best select the provider domains offering the required services.

References

[1] H2020 5G-TRANSFORMER, 5G Mobile Transport Platform for Verticals. http://5g-transformer.eu/

[2] J. Baranda et al., "Demo: Composing Services in 5G-TRANSFORMER", in Proceedings of ACM Int. Symp. on Mobile Ad Hoc Net. and Computing, 2-5 July, Catania, Italy

[3] G. Sun et al, "Service Function Chain Orchestration Across Multiple Domains: A Full Mesh Aggregation Approach", IEEE Transaction on Network and Service Management, Sept. 2018

[4] Open Networking Foundation, "SDN Architecture ONF TR-521 Technical Report", 2016

[5] Metro Ethernet Forum, "Service Operation Specification (MEF 55), Lifecycle Service Orchestration, Reference Architecture and Framework", 2016

[6] ETSI GR NFV-IFA 028, Network Functions Virtualisation (NFV) Release 3; Management and Orchestration; Report on architecture options to support multiple administrative domains, January 2018

[7] ETSI GS NFV-IFA 030, Network Function Virtualisation (NFV) Release 3; Management and Orchestration; Multiple Administrative Domain Aspect Interfaces Specifications, May 2019

[8] ETSI GS NFV-IFA 013, Network Function Virtualisation (NFV) Release 3; Management and Orchestration; Os-Ma-Nfvo reference point – Interface and Information Model Specification

[9] NGMN Alliance, Description of Network Slicing Concept, January 2016.

[10] ETSI GS NFV-MAN 001, "Network Functions Virtualisation (NFV); Management and Orchestration", v1.1.1, 2014

[11] J. Mangues et al, "5G-TRANSFORMER Service Orchestrator: Design Implementation and Evaluation", in Procs of the 28th European Conf. on Net. and Comms. (EUCNC'19), 18-21 June 2019, Valencia, Spain

[12] K. Antevski et al, "Resource Orchestration of 5G Transport Networks for Vertical Industries", in Procs of the IEEE Int. Symp. on Personal, Indoor and Mobile Radio Comms (PIMRC), 9-12 September 2018, Bologna, Italy

[13] J. Baranda et al.,"Orchestration of end-to-end network services in the 5G-crosshaul multi-domain multi-technology transport network," IEEE Communications Magazine 56(7), July 2018

[14] D. King and A. Farrell, "A PCE-Based Architecture for Application-Based Network Operations", IETF RFC 7491, March 2015

[15] 5G-TRANSFORMER project code repository, Available at: http://github.com/5g-transformer

Acknowledgements

This work has been partially funded by the EC H2020 5G-Transformer Project (grant no. 761536), by MINECO grant TEC2017-88373-R (5G-REFINE) and Generalitat de Catalunya grant 2017 SGR 1195.

Jorge Baranda Hortigüela *(jbaranda@cttc.es) is currently a Senior Researcher in the Mobile Networks Department at CTTC in Barcelona. He received an Msc (2008) degree in Electrical Engineering from the Technical University of Catalonia. At CTTC, he has participated in several European, national and industrial projects related with management and orchestration of SDN/NFV mobile networks, efficient routing strategies for mobile network backhauling and novel wireless communication systems. He has co-authored over 30 different peer-reviewed journal and conference papers. His current research interests include management and orchestration of SDN/NFV mobile networks, wireless communications, wireless backhaul, network routing protocols, and network optimization.*

Josep Mangues-Bafalluy *(josep.mangues@cttc.cat) received the MSc (1996) and PhD (2003) in Telecommunications from the Technical University of Catalonia (UPC). He is senior researcher and head of the communication networks division (networks.cttc.cat) of the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC) and coordinated its IP Technologies Area from 2003 to 2012. He was also assistant professor at UPC. He regularly participates in EU (FP6, FP7, H2020), industrial (e.g., Cisco, Orange), Spanish, and Catalan research projects, of which six of them were led by him. His research interests include software-defined networks (SDN) and Network Functions Virtualization (NFV), including its application to self-organized mobile networks.*

Ricardo Martínez *(ricardo.martinez@cttc.es) is a graduate and PhD in Telecommunications Engineering of the Universitat Politècnica de Catalunya (UPC) in 2002 and 2007, respectively. He has been actively involved in several National and EU- founded (FP6, FP7 and H2020) R&D projects as well as technology transfer projects with industrial telecom operators. He has co-authored 1 book chapter and over 175 journal and conference papers. His research interest areas include control/orchestration and service management architectures, protocols and traffic engineering mechanisms for future packet and optical transport networks within aggregation/metro and core infrastructures as well distributed cloud (NFV) architectures.*

Luca Vettori *(luca.vettori@cttc.cat) received his Ms. Degree in Telecommunication Engineering from the University of Florence (Italy) in 2008. After graduated, he worked for three years as a Research Engineer in LENST Laboratory (currently LESC) for Department of Information Engineering of University of Florence. In July 2017, he joined CTTC as Research Assistant in the Optical Networks and Systems Department of Communication Networks Division. His current research activity is focused on Software Defined Network, Network Virtualization and Network Orchestration.*

Kiril Antevski *(kantevsk@pa.uc3m.es) received the B.S. degree in telecommunication engineering from The Saints Cyril and Methodius University of Skopje, Macedonia in 2008 and the M.S. degree in telecommunication engineering from the Politecnico di Milano, Milan, Italy in 2016. He is currently pursuing the Ph.D. degree in telematics engineering at University Carlos III Madrid (UC3M), Spain. His research interest includes the development of mechanisms to integrate and enhance NFV and MEC for 5G Networks in dynamic and* *heterogeneous environments. Mr. Antevski's has won the best-paper award on the 9th IFIP Wireless and Mobile Networking Conference (WMNC) in 2016.*

Carlos J. Bernardos *(cjbc@it.uc3m.es) received a Telecommunication Engineering degree in 2003, and a PhD in Telematics in 2006, both from the University Carlos III of Madrid, where he worked as a research and teaching assistant from 2003 to 2008 and, since then, has worked as an Associate Professor. His research interests focus on network virtualization and wireless mobility management. He has published over 100 scientific papers in international journals and conferences. He has worked in several EU funded projects, being the technical manager of the FP7 MEDIEVAL and 5GEx projects. He is the Project Coordinator of the 5GPPP 5G-TRANSFORMER and 5Growth projects.*

Xi Li *(xi.li@neclab.eu) received her M.Sc. in 2002 from TU Dresden and her Ph.D. in 2009 from the University of Bremen, Germany, where she worked as a research fellow and lecturer. In 2014 she worked as a solution designer at Telefonica, Germany. Since 2015 she has been with NEC Laboratories Europe, working as a senior researcher on 5G networks. Her research interest includes 5G architecture design, network slicing management, orchestration of SDN/NFV mobile networks, wireless fronthaul and backhaul, and network optimization.*