Scaling Federated Network Services: Managing SLAs in Multi-Provider Industry 4.0 Scenarios

Jorge Baranda, Josep Mangues-Bafalluy, Luca Vettori, Ricardo Martínez, Engin Zeydan
Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Castelldefels, Spain
{jbaranda, josep.mangues, lvettori, rmartinez, engin.zeydan}@cttc.cat

Abstract—Next generation mobile networks require flexibility and dynamicity to satisfy the needs of vertical industries. This may entail the deployment of slices instantiated in the form of composite network services (NSs) spanning multiple administrative domains through network service federation (NSF). In this way, different nested NSs of the composite service can be deployed by different service providers. But fulfilling the needs of verticals is not only needed during instantiation time but also during NS operation to honour the required service level agreements (SLAs) under changing network conditions. In this demonstration, we present the capabilities of the 5Growth platform to handle the scaling of federated NSs. In particular, we show the scale out/in of a nested NS deployed in a federated domain, which is part of a composite NS. These scaling operations, triggered to maintain the NS SLAs, imply a set of coordinated operations between involved administrative domains.

I. Introduction

The flexibility and dynamicity required in next generation mobile networks may involve the deployment of slices using NFV network services (NSs) spanning multiple administrative domains (ADs) owned by different service providers due to business (e.g., cost) or technical reasons (e.g., service availability and coverage). This is possible thanks to the network service federation (NSF) concept [1], which enables the deployment of composite NS across different ADs and the interconnection of its constituent parts (the so-called nested NSs) to build such end-to-end slices. However, the mentioned capabilities are not only needed at deployment time but also during operation time to maintain the requested service level agreements (SLAs) of the different deployed NSs over the same shared infrastructure despite dynamic network conditions. Indeed, this is one of the objectives of the EU 5Growth (5Gr) project [2].

Up to now, most of previous work involving NSF focuses on the instantiation phase, being the report presented by ETSI NFV SDO in [3] one of the sole works covering the scaling of NSs deployed in multiple ADs. But this approach is only restricted to service orchestration aspects at the composite NS level, not considering the direct application of scaling actions at the nested NS level to react to changing network conditions and enforce requested SLAs. This demonstration showcases the capabilities of the 5Growth platform, and specially of its 5Gr-Service Orchestrator (5Gr-SO) module to broaden this scope and support auto-scaling operations at the nested

NS level implying NSF and also considering the resource orchestration perspective handling the update of inter-nested connectivity between involved ADs, which is essential for the NSF concept. The scenario under demonstration considers an emulated Industry 4.0 scenario where a standalone Non-Public Network (NPN) [4] regular NS can be shared among composite NSs acting as *industrial virtual services*. To the best of our knowledge, this is the first operational demonstration involving the (auto-)scaling (out/in) of a nested NS deployed in a federated domain, which is part of a composite NS. Additionally, this composite NS is sharing a previously deployed regular NS with another deployed composite NS.

II. Testbed Architecture

Fig. 1 presents the multi-AD setup deployed at the CTTC 5G Lab, where each AD runs its own instance of the 5Gr platform. The aim of this demonstration is to show the capabilities of the 5Gr-SO to handle the scaling of composite NSs, whose deployment implied the use of NSF [1]. The 5Gr-SO is the module in the 5Gr platform that oversees the end-to-end orchestration and the lifecycle management of NSs in single and multi-AD scenarios by interacting with the infrastructure manager (5Gr-RL) and/or with peering 5Gr-SOs. For doing so, the 5Gr-SO relies on its hierarchical (parent-child) service orchestrator engine (SOE) and its resource orchestrator modules [1].

In this setup, both 5Gr-SOs use an instance of Open Source MANO (OSM) as Core MANO platform and communicate through the depicted *So-So NSF* interface. At the NFVI level, AD1 presents three NFVI-PoPs managed each by a dedicated Virtual Infrastructure Manager (VIM) based on OpenStack (Devstack Queens Release). These NFVI-PoPs are interconnected by the depicted transport network emulated with GNS3 [5], which is managed by an ONOS SDN controller acting as Wide Area Infrastructure Manager (WIM). These VIMs and WIM are controlled by the 5Gr-RL. AD2 counts with two NFVI-PoPs, each managed by an Openstack-based VIM, as in AD1. The 5Gr-RL of AD2 also uses an ONOS SDN controller to interconnect these NFVI-PoPs through the GNS3 emulated transport network represented in Fig. 1. The transport network of both ADs are connected through an Inter-AD link, implemented with a VPN tunnel.

III. Demonstration

We demonstrate the capabilities of the 5Gr platform to handle the scaling of federated NSs. Considering an Industry
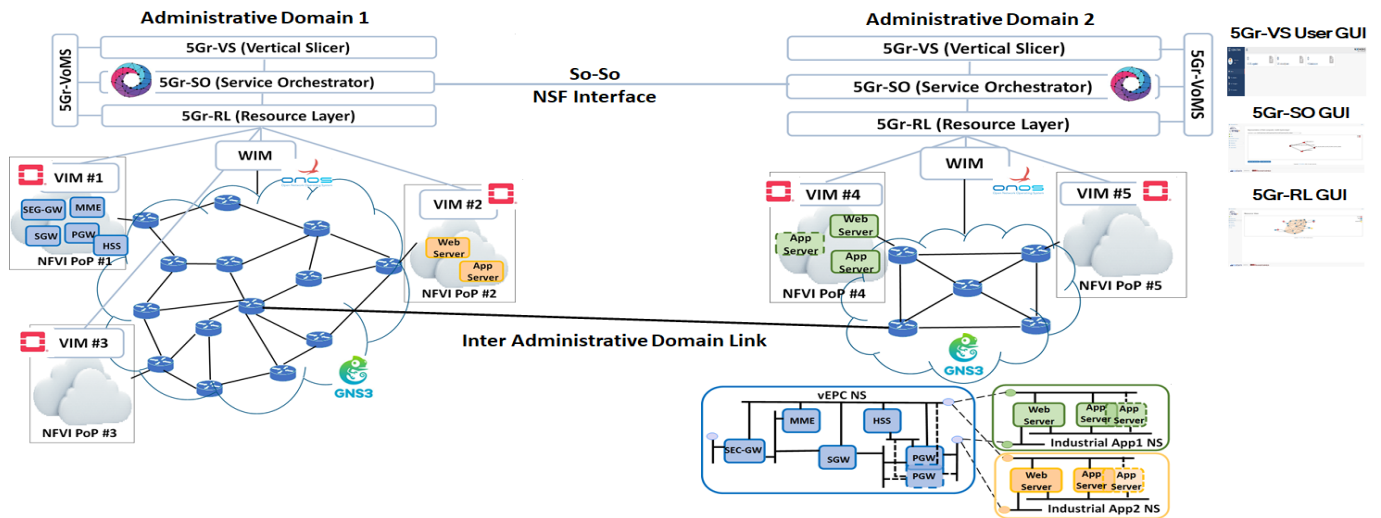
Fig. 1. Multi-administrative domain experimental setup under demonstration

4.0-like scenario, we focus on the auto-scaling of a federated nested NS part of a composite NS, which shares components with other composite NSs. The mentioned composite NSs have two nested NSs. The *common* nested NS is an emulated standalone NPN network [4] vEPC NS (represented by the blue VNFs in Fig. 1) providing connectivity to different emulated industrial application NSs (represented by green and yellow VNFs in Fig. 1). The composite NSs define multiple instantiation levels (ILs) covering the different ILs of its associated nested NSs to be able to adapt the structure of the whole service to cope with changes in network conditions and be able to satisfy SLA agreements. The main steps of the demonstration are:

1) The vEPC NS implementing the standalone NPN is instantiated in AD1 as explained in [6]. The five VNFs of the NS (SEC-GW, SGW, MME, PGW and HSS) are deployed in NFVI-PoP#1, placed at the edge of the network (e.g., the factory premises).

2) Then, the two composite NSs are instantiated. Both composite NSs rely on and share the previously instantiated vEPC NS. However, the remaining nested NS is deployed in different ADs. In the first case, the *Industrial App1 NS* is deployed in NFVI-PoP#4 of AD2 as part of the agreement between ADs to share its network service catalogue by means of NSF, as explained in [1]. In the second case, the *Industrial App2 NS* is deployed locally in NFVI-PoP#2 of AD1 because this NS is available in the AD1 catalogue and the characteristics of the composite NS require the nested industrial app to be deployed close to the vEPC NS.

3) Due to growing activity in the factory, the load on the App server VNF in *Industrial App1 NS* grows and starts jeopardizing the SLA requirements expressed in the NS descriptor (NSD) of the nested NS. The 5Gr-SO of AD2 receives an (auto-)scaling operation request after processing the alert sent from the 5Gr Vertical-Oriented Monitoring System (5Gr-VoMs). It requests to change the IL of the *Industrial App1 NS* to increase the number of instances (scale out) of the App

server VNF, hence coping with the increasing load.

4) When processing the scaling request, the 5Gr-SO at AD2 realises that the *Industrial App1 NS* is part of a composite NS federated with AD1. The 5Gr-SO of AD2 notifies the change in the nested federated NS to AD1 through the *So-So* interface. The 5Gr-SO at AD1 checks the implications of the auto-scaling being done at AD2 in the NSs deployed in AD1 and after confirming the execution of the scaling operation at AD2, proceeds to update the inter-nested connections established though the available Inter-AD Link. In this case, new inter-nested connections need to be established between ADs to communicate the new App server VNF with the vEPC NS.

5) When the traffic load decreases, the *Industrial App1 NS* can turn back to its initial IL with a single instance of the App server VNF. The 5Gr-SO at AD2 is notified about this situation and receives an auto-scaling operation to change the IL (scale in) of *Industrial App1 NS*. Then, the 5Gr-SO performs again step 4) to adapt to the new situation. In this case, the previously established inter-nested connections are removed and the system turns back to the state after step 2). During the demonstration, all steps are shown through the GUIs of the different building blocks of the 5Gr platform at the involved ADs (e.g., NFV-NS deployment information status, selected/updated network paths), as depicted in Fig. 1. A demonstration video is available online[1].

REFERENCES

[1] J. Baranda et al., "Realising the Network Service Federation vision," in IEEE Vehicular Tech. Mag., June 2020.
[2] EU 5G-PPP 5Growth Project: 5G-enabled Growth in Vertical Industries, Available at https://5growth.eu/, [Accessed in January 2021].
[3] ETSI GR NFV-IFA 028, "Report on architecture options to support multiple administrative domains", Jan 2018.
[4] 5G Alliance for Connected Industries and Automation, "5G Non-Public Networks for Industrial Scenarios White Paper", July 2019.
[5] The Graphical Network Simulator-3. Available at: https://www.gns3.com/, [Accessed in January 2021].
[6] J. Mangues et al., "5G-TRANSFORMER Service Orchestrator: Design Implementation and Evaluation", EUCNC'19, June 2019.

[1]Available at https://youtu.be/RhO9ronLhTw