Service Orchestration and Federation for Verticals

Xi Li
NEC Laboratories Europe, Germany

Giada Landi, Francesca Moscatelli
Nextworks, Italy

Kiril Antevski, Carlos J. Bernardos
Universidad Carlos III de Madrid, Spain

Dmitriy Andrushko
NURE, Ukraine

Adlen Ksentini
EURECOM, France

Josep Mangues-Bafalluy, Iñaki Pascual
Centre Tecnològic de Telecomunicacions de Catalunya
CTTC/CERCA, Spain

Luca Valcarenghi, Barbara Martini
Scuola Superiore Sant'Anna, Pisa, Italy

Carla Fabiana Chiasserini, Claudio Casetti
Politecnico di Torino, Italy

Nicolás A. Serrano
Telefónica I+D, Spain

Abstract— **The next generation mobile transport networks shall transform into flexible and agile SDN/NFV-based transport and computing platforms, capable of simultaneously supporting the needs of different vertical industries, e.g., automotive, e-health and media, by meeting a diverse range of networking and computing requirements. Network slicing, has emerged as the most promising approach to address this challenge by enabling per-slice management of virtualized resources and provisioning and managing slices tailored to the needs of different vertical industries. Service orchestration is the key enabler for slicing that allows efficient placement of virtual network functions over the infrastructure as well as optimal allocation of virtual resources among all network slices to deliver guaranteed, reliable and scalable services of different verticals. Besides, due to the limited footprint of infrastructure operators, it is also required to enable the interconnection and federation of multiple administrative domains, to effectively allow services to span across several providers. This paper presents the design of Service Orchestrator (SO) in the 5G- TRANSFORMER system, which deals with service orchestration and federation across multiple domains.**

Keywords— network slicing, service orchestration, federation

I. INTRODUCTION

5G networks are not only envisioned to provide enhanced mobile broadband (eMBB), but also to support a wide range of vertical industries with very diverse and stringent service requirements. This paper dives into the design of the next generation Mobile Transport Networks, as understood by the EU H2020 5G-PPP phase 2 5G-TRANSFORMER project, to simultaneously support the needs of different vertical industries. This design brings the Network Slicing concept into mobile transport networks by provisioning and managing slices tailored to the needs of different vertical industries.

Network Function Virtualization (NFV), which enables infrastructure and function virtualization, is gaining an incredible momentum by mobile operators as one of the significant solutions to reduce costs and to dynamically optimize the resource allocation. Portions of network and compute resources (i.e., slices) provided by ditstributed data centers, together with the automated network provisioning, adpatation and data forwarding based on Software Defined Network (SDN) concept will provide verticals with the required virtualized services. Two technologies are being regarded as the key of our solutions: i) network slicing which is being studied in different standandard bodies (e.g., [1], [2], [3] [4]), and ii) Multi-access Edge Computing (MEC) [5]. While the main impact of slicing will be on the cost reduction side, MEC will enable low-delay or delay sensitive services that are not currently possible.

Orchestration is the key enabler in 5G-TRANSFORMER for slicing, efficient load distribution and arbitration among network slices to deliver guaranteed, realiable and scalable e2e services for different verticals with very diverse requirements. One major challenge in slicing is exposing the capabilities of the network including topologies and resources via proper abstraction to the orchestration layer. In this paper we present the orchestration layer of 5G-TRANSFORMER architecture, named as Service Orchestrator (SO), which enables multi-domain service orchestration (including federation of services and resources). This paper describes the functionality of the SO, and presents different architecture options and high-level workflows to illustrate its operaton. A sample vertical use case is also discussed to explore its use of the architecture.

The rest of the paper is organized as follows. In Section II we present an overview of the 5G-TRANSFORMER system and relate it in Section III with the relation to ETSI NFV MANO framework. Section IV introduces the concept of federation. In Sections V we present the high-level workflows for service instatiation, and in Section VI we present a smaple vertical use case. Finally, Section VII gives the conclusion and an outlook of the future work.

This work has been partially funded by the EU H2020 5G-TRANSFORMER Project (grant no. 761536). * Corresponding author email: xi.li@neclab.eu

II. 5G-TRANSFORMER SYSTEM

We envision the 5G-TRANSFORMER system consisting of 3 major components: Vertical Slicer (VS), Service Orchestrator (SO) and Mobile Transport and Computing Platform (MTP), as illustrated in Figure 1.

The VS is the common entry point for all verticals into the system, being part of the operating and business support systems (OSS/BSS) of the administrative domain of a given provider. The VS coordinates and arbitrates the requests for vertical services such as automative, eHealth, etc. Vertical services are offered to the verticals through a high-level interface focusing on the service logic and needs of vertical services. It allows composing vertical services from a set of vertical-oriented service blueprints, which along with instantiation parameters will result in a vertical service instantiation request. Then, the VS maps descriptions of vertical services and requirements at the vertical application level onto a network service descriptor (NSD), which is a service graph composed of a set of V(N)Fs chained with each other and fine-grained instantiation parameters (e.g., deployment flavor) that are sent to the SO. The VS and its functionalities are explained in [6].

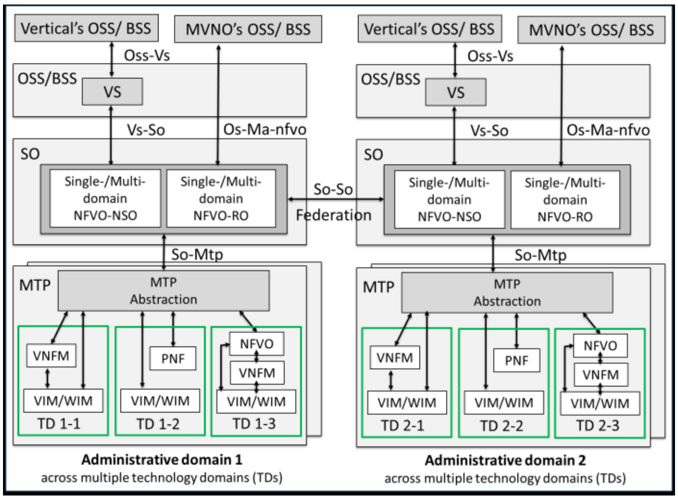

Figure 1: 5G-Transformer system architecture

The SO provides end-to-end orchestration of services, as composition of virtual (network) functions V(N)Fs, across multiple administrative domains. It acts as NFV orchestrator (NFVO). It receives requests from the VS or directly from the M(V)NO - Mobile (Virtual) Network Operator. Depending on the use case, both network service (NFVO-NSO) and resource (NFVO-RO) orchestration may be used for both single and multi-domains. In turn, based on the request, the SO may decide to create a new network slice instance or to reuse one previously created by the provider to be shared. Therefore, it manages the monitoring and allocation of virtual resources to network slices, be it for vertical services or for network services of an M(V)NO. If needed (e.g., not enough local resources), the SO interacts with the SOs of other administrative domains (federation) to take decisions on the end-to-end (de)composition of virtual services and their most suitable execution environment. Even if a service is deployed across several administrative domains, e.g., if roaming is required, a vertical still uses one VS to access the system, and so, the SO hides this federation from the vertical. The NFVO-RO functionality of the SO handles resources coming from the local MTP (real or abstracted) and from the SOs of other administrative domains (abstracted). The orchestration decision for creating or updating a network slice includes the placement of V(N)Fs over such virtual networks with virtual nodes and links, as well as the resources to be allocated. The SO will then create the slice by exploiting the interface exposed by the local MTP and also that exposed by peer SOs, which will eventually interact with their respective MTPs.

The MTP [7] is responsible for instatiation and realization of the slices by managing services and/or resources and their instantiation over the infrastructure under its control, as well as managing the underlying physical mobile transport network, computing and storage infrastructure. The computing and storage infrastructure may be deployed in central data centers as well as distributed, as in MEC [8]. The MTP provides support for slicing, enforces slice policies coming from the SO and provides physical infrastructure monitoring and analytics services. Depending on the use case, the MTP may offer different levels of abstraction to the SO via the MTP abstraction component, which in turn forwards the SO requests to the right entity accordingly (hence acting as a single logical point of contact): VIM (Virtual Infrastructure Manager)/WIM (WAN Infrastructure Manager), VNFM (VNF Manager) or PNF (Physical network Function) , or NFVO [9].

III. SO IN RELATION TO ETSI NFV

The SO exposes a northbound API from which it receives a Network Service Descriptor (NSD) along with instantiation parameters (e.g., deployment flavor) for the requested service. This API is used by the vertical slicer and M(V)NOs or other services. This service request is processed by the SO, which, in accordance with the ETSI NFV architecture [9], is in charge of the network service lifecycle management, including on-boarding, instantiation, termination, or update of network services, and VNF forwarding graph management.

Additionally, the SO is in charge of resource orchestration. The resources handled at the SO level may come not only from the MTP in its domain (abstracted or not), but also from those abstracted resources from other domains exposed through the SO-SO interface by peering SOs. Therefore, a key function of the SO is to decide whether the service can be offered within its own domain, which in general may be the preferred option, or jointly with federated MTP resources. More specifically, the functions at the resource orchestration level are [9]: 1) Validation and authorization of NFVI resource requests from VNF Manager (VNFM), 2) NFVI resource management across operator's Infrastructure Domains, 3) Supporting the management of the relationship between the VNF instances and the NFVI resources allocated, 4) Policy management and enforcement for the Network Service instances and VNF instances, and 5) Collecting usage information of NFVI resources by VNF instances or groups of VNF instances.

Another key feature of the SO is its support for federation which requires the exchange of resource and service-related offerings between administrative domains. According to [9], at least two types of administrative domains may be present (tenant and infrastructure domains), each featuring its own MANO. The tenant domain is the one having an end-to-end view of the services deployed (e.g., all the components for a collision avoidance service). This domain is service-aware in the sense that it understands the relationship between the logic of the VNFs that conform the service. The services of the tenant domain must be deployed over resources provided by one or various infrastructure domains. Therefore, consumer-provider MANO relationship naturally appear, which may have a hierarchical or peering nature. Architectural options of such multi-domain scenarios are developed in [15] for the use cases of NFVIaaS and composition of Network Services (NSs) consisting of constituent NSs provided by multiple domains. In the former, virtual resources are offered by the provider domain to other domains to be consumed and managed through their own MANO. In the latter, the provider domain offers complete constituent network services to the consumer domain.

A. NFVIaaS use case

For the NFVIaaS use case, the provider domain provides a set of virtual resources assigned to a given consumer domain, which will manage them according to its service logic. Hence, the consumer domain handles the NS lifecycle and generates resource management requests to the provider domain, which, in turn, provides a global view of the resources to the consumer domain, considers resource constrainsts from the consumer when executing placement decisions, and orchestrates the infrasture resources based on its requests.

There are various architectural options for offering the above functionalities. They can be classified based on two axes, namely 1) multiple logical points of contact (MLPOC) vs. single logical point of contact (SLPOC) and 2) direct vs. indirect mode. As for the former, the MLPOC approach allows NFVOs and VNFMs of the consumer domain to access the VIMs of the provider domain, whilst in SLPOC access is done through a single entity (the SLPOC) in the provider domain that hides the VIM-level granularity to the consumer domain. Such granularity is handled internally to the provider by forwarding the requests from the SLPOC to the appropriate VIM. As for the latter, direct mode allows both NFVOs and VNFMs of the consumer domain to directly access the entities of the provider domain (through MLPOC or SLPOC). On the other hand, in indirect mode, it is only the NFVO accessing them, and all requests from the VNFM require NFVO intervention.

The selection of architectural option and mapping of 5G-TRANSFORMER functionalities and interfaces is still under study. This paper explores the different possibile options. Figure 2 presents an example based on [15], but adding the SO as key building block for federating resources.

In this example, the tenant sends resource management requests on virtual resources to the SO. The SO embeds the NFVO of the provider, in charge of federating resources exposed by peer SOs and offering them in an aggregated manner to the tenant. Therefore, federation is transparent to the tenant. In turn, the SO embeds the SLPOC, which forwards the resource requests to the NFVOs of the MTPs providing the virtual resources to the tenant. After that, these requests are received by the SLPOC embedded in the NFVO of the MTP, which eventually forwards the resource management request to the corresponding VIM inside that particular MTP.

In terms of interfaces, the tenant exploits the SO-SAP (SO service access point) capabilities, which in accordance to [15] would be based on the new Or-Or interface, which is in turn based on the Or-Vi reference point [16]. Additionally, in case of direct mode access, Vi-Vnfm [17] would also be used between the VNFMs of the customer and the SLPOC of the SO. Besides, the additional level of indirection between the SO and the provider MTPs could also be based on IFA 005.

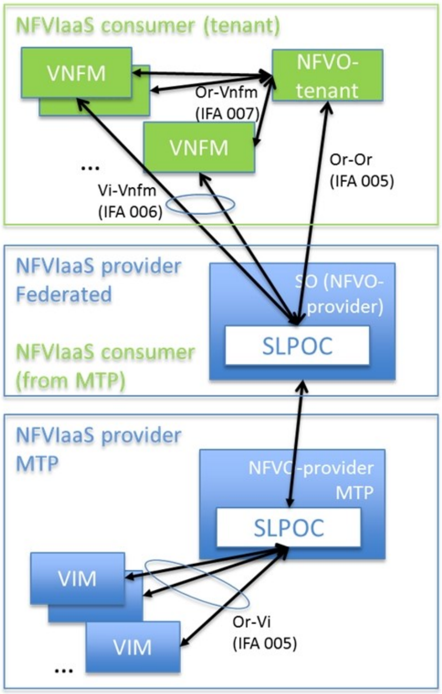

Figure 2: Possible architecture for NFVIaaS by the SO

B. Network services over multiple administrative domains

ETSI NFV IFA 028 [15] is also analyzing which implications the provisioning of Network Services across multiple administrative domains have on the ETSI NFV MANO architecture, identifying potential extensions in the MANO reference points to enable these scenarios. Each administrative domain is assumed to include all the components of the MANO stack, i.e. one or more NFVI PoPs with related VIMs, a set of VNFMs and an NFVO instance. The management of multiple VIMs, where needed, can be handled through M/SLPOC functions as analysed in section III.A and does not present further implications. Specific sets of NSs, called Constituent Nested NSs, are instantiated in each administrative domains; for example Figure 3 shows NS A and NS B instantiated in domain A and B. Their composition through a nesting procedure generates a new Composite NS that can be offered by one of the administrative domains providing the nested NS or even by another administrative domain (administrative domain C in the figure) that has established SLAs with the other ones.

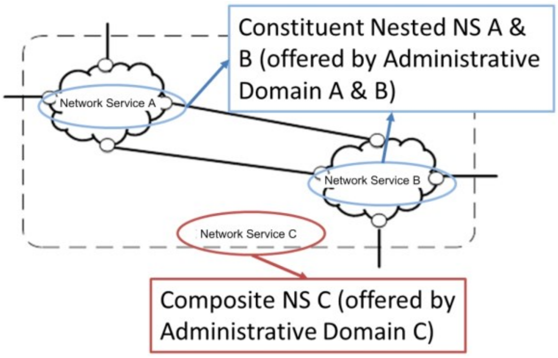

Figure 3: Composition of NSs across different adm. domains

In order to address this scenario, ETSI proposes a hierarchical architecture where the NFVO of the administrative domain offering the composite NS interacts with the NFVOs of the administrative domains offering the constituent nested NSs through a new reference point, called Or-Or (see Figure 4). It should be noted that this scenario is based on NS-level federation, where each domain offers complete NSs, not just pools of virtual resources like VMs or virtual links. The interface adopted through the Or-Or can thus re-use part of the Os-Ma-Nfvo interface specified in IFA 013 [10].

In particular, IFA 028 [15] currently proposes a subset of the NS Lifecycle Management (LCM) interface methods (NS ID creation and NS instantiation, scaling, healing, queries and lifecycle notifications) to enable the management of the constituent NSs and a subset of the NSD management interface to enable NSD queries about the services potentially offered by different administrative domains. Further specific Os-Ma-Nfvo interfaces should be supported on the Or-Or: for example a limited set of performance monitoring information about the Constituent NS should be exchanged between the NFVOs, depending on the established SLAs, to enable scaling decisions at the parent NFVO.

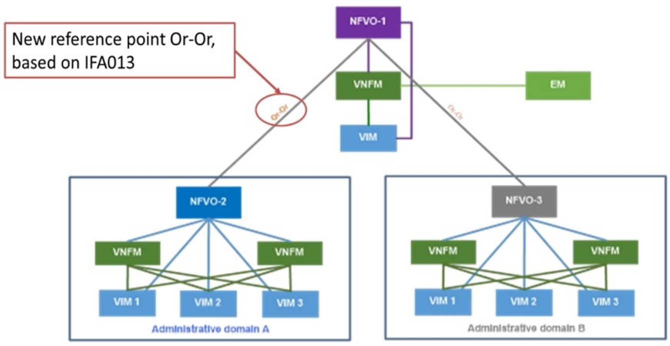

Figure 4: Hierarchical interaction between NFVOs

The architecture proposed by ETSI follows a typical hierarchical approach. 5G-TRANSFORMER is evaluating two different options. The former is based on the ETSI-like hierarchical approach and it is compatible with the architecture in described section II. The latter follows an alternative approach based on a peer-to-peer interactions between SOs and requires extensions to the architecture in section II.

In the hierarchical model, the 5G-TRANSFORMER architecture maps the SO to the NFVO orchestrating the composite NS while the administrative domains offering the constituent NSs are handled by an MTP (see Figure 5). In this case, each MTP should include its own NFVO, responsible to fully manage and orchestrate both lifecycle and resources for the NS instances deployed in its domain. The SO, on the other hand, is responsible to coordinate the NS nesting procedures and to take decisions about the placement of NSs across the available domains, based on the services offered by the "child" NFVOs at the MTPs. The Or-Or is thus mapped on the SO-MTP reference point of the 5G-TRANSFORMER architecture.

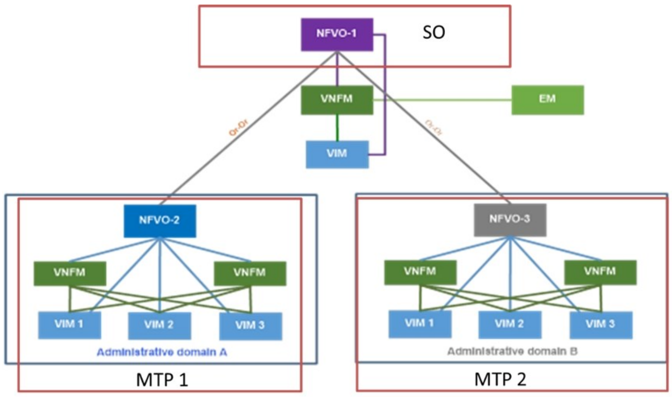

Figure 5: 5G-TRANSFORMER architecture based on hierarchical paradigm

The alternative approach follows a peer-to-peer paradigm where different SOs interact with each other for service federation (see Figure 6) through the SO-SO reference point of the 5G-TRANSFORMER architecture, which is mapped to the Or-Or reference point defined by ETSI NFV. In this option the MTP(s) are just responsible for resource allocation, based on the decisions taken by the SO; the SO-MTP reference point can thus be mapped to ETSI Or-Vi and Vi-Vnfm reference points.

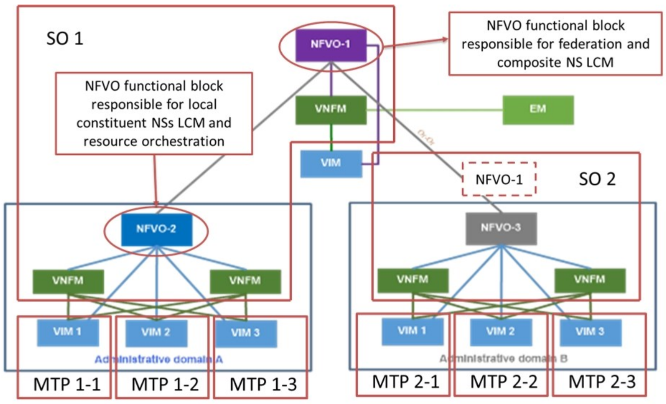

Figure 6: 5G-TRANSFORMER architecture based on peer-to-peer paradigm

On the other hand, the SO itself needs to extend the traditional NFV concept of NFVO. In particular, we can identify two relevant functional blocks within the SO. The former is responsible for the provisioning of constituent NSs in the local domain, it takes care of their LCM and resource orchestration and it can be considered as a standard NFVO. On top of that we need an additional NFVO-like component (the

NFVO-1 of the picture) that handles the logic of the service federation and takes care of the interaction with its peers placed in the other domains through the new Or-Or reference point.

Each administrative domain interested in joining the federation should include such additional component in its MANO architecture. Two main functionalities should be considered for that: (i) the exchange of NS offers between federated domains, as regulated through the active SLAs and (ii) the capability to split the request of a composite NS in several constituent NSs and coordinate their requests to other SOs along the entire LCM of the composite NS. It should be noted that the former function is required for all the administrative domains wanting to enter the federation with a provider role (i.e. domains offering NSs to other peer domains), while the latter function is required for domains with a consumer role (i.e. domains requesting and composing NSs offered by other domains to serve their own customers).

IV. FEDERATION

Federation procedure is defined on the SO-SO interface for sharing of resources and/or services between entities (mobile operators, service providers, infrastructure providers) with established terms and conditions of usage. Depending on the way of federation, the SO-SO interface is designed with different process and abstraction level details.

The federation can be formed as pre-established federation and open federation. Pre-established federation happens when the relationship agreement to share resources is previously established as a business agreement (offline) with well-defined sharing polices. The agreement is defined as Service Level Agreement (SLA) and each operator has to act accordingly to each pre-agreed peer SLA. The open federation is established between entities that advertise their resources publicly, with certain level of abstraction. Each operator maintains a list of publicly available abstractions of resources/services from all operators. Once an operator decides for the most suitable abstraction of resources/services, the negotiation continues bi-laterally or centralized. The operator directly initiates the agreement process with the operators that own the resources/services or this is managed by a trusted centralized entity. The key differences between pre-established and open federations are the level of abstraction applied to the advertised resources/services and the initial allocation method. Pre-established federation entities advertise their resources or services in a deterministic way towards the pre-agreed peering operators, e.g., using BGP-LS (update and keep-alive messages) [14]. The abstraction level of resources/services on each advertising connection depends of the SLA previously agreed. As an example, an Operator A will reveal more detailed information (type, capability, power, quantity, vendor, location) towards a premium peer Operator B compared to the information (type, power) shared with a silver peer operator C. Allocation of resources/services is done through a REST APIs.

Regarding open federation, recent advances in Blockchain network technologies enable fully distributed solution for advertisement of all abstractions, negotiation and management of connections between operators. Using Ethereum network [11], operators can have near real-time global view of the abstracted resources/services. Smart Contracts [12][13] can be used for negotiation and management of connections. Smart contract is a tiny piece of code that implements SLA logic and automatically can enforce terms and conditions as well as establishes and manages connections between operators. The maintenance and complexity levels are low, but the stability depends of the scalability of blockchain (Ethereum) network.

Vertical users gain access to resources/services using a User Interface (UI) portal of the Vertical Slicer (VS). Each operator can offer proprietary VS UI portal, populated with a view of the offered resources/services (including abstract external views). The resulting VS UI portal is unique for each operator. Another solution is when UI portal is generated by a federation. List of all available abstract views of the federation is generated.. In that case a vertical user connecting at any Vertical Slicer (any operator), connects directly to a federation. In that case, the boundaries of all operators in a federation are not visible from the vertical user perspective.

V. WORKFLOWS FOR SERVICE INSTATIATION

The high-level workflow of the SO is explained below. Upon reception of the NSD including the fine-grained instantiation parameters (e.g., deployment flavor) over NorthBound API, the SO performs the following actions:

1. It authenticates and authorizes the request of the VS.
2. It checks whether the requested network services (NS) or V(N)Fs are already on-boarding and available. If not, it informs the VS to ask for service on-boarding.
3. It performs service orchestration to decide the V(N)F placement and resource to be allocated/reserved for V(N)Fs and network connectivity to build a Network Service Instance (NSI), and then requests the MTP to allocate resources and/or instantiate V(N)Fs. If the local MTP cannot satisfy NS or V(N)F instantiation, SO initiates federation process to perform multi-domain service deployment. It shall guarantee the correct sequence of V(N)Fs and virtual links instantiation between different domains.
4. It notifies the VS of the acceptance/rejection of the NS instantiation request and it provides the status of the NS deployment progress.
5. After successful instantiation, it sends an Notification ACK to notify the VS and it monitors the NSI state and appropriate metrics. These metrics are made available through the NorthBound API to the VS and other tenants directly using the SO services.

VI. AUTOMOTIVE USE CASE

A sample vertical service use case is a collision avoidance (CA) service provided to vehicles travelling in urban areas (depicted in Figure 7). The CA service allows cars to avoid crashes at blind crossroads. Traditionally the CA service implementation is based in the vechicle, which collects and fuses data gathered by on board sensors. An alternative approach is the one here described, where the CA service is based on exchanging information between a CA-client application already installed in the cars and a CA-server located in the infrastructure (e.g., a Vehicle to Infrastructure --- V2I --- interaction, which is the one envisioned by 5G-TRANSFORMER) that makes decisions based on the received information. Clearly, the CA service requires a low end-to-end latency, especially when automated vehicles are involved. Based on the NSD provided by the VS and on the abstraction provided by the MTP, the SO needs to take a subset of the actions reported in section V to provide the service. This is a NFVaaS use case where the SO is responsible for orchestrating the service as a chained V(N)Fs according to the given NSD. For example, by referring to Figure 7, the SO shall orchestrate the CA-server application placement (it is assumed that the CA-client application is already installed in the car) and the virtual links to be utilized to connect the service components. Such components must meet the service latency constraints expressed by the VS in the NSD. Then the SO delegates the MTP to instantiate the service components.

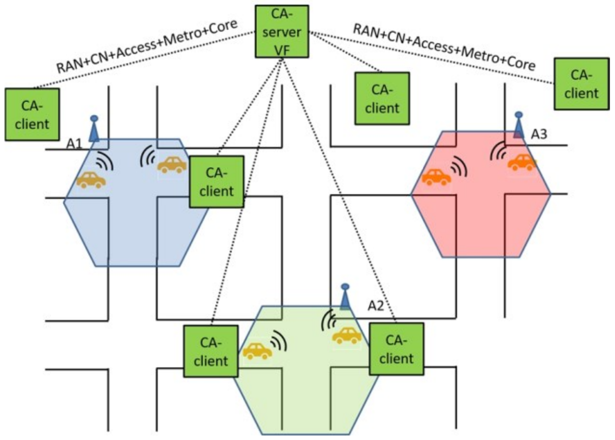

Figure 7: CA use case with coarse granularity communications service abstraction at the SO

Another possible scenario is the one depicted in Figure 8. In this scenario the SO is provided by the MTP with the visibility of the transport network components to be utilized for implementing the connectivity between the CA-client and the CA-server. Thus, the SO must orchestrate not only the CA-server placement and the virtual links selection but also the placement, selection, and interconnection of transport network components (e.g., eNB and EPC). Even in this case, the component instantiation as well as the type of functional split to be utilized is left to the MTP.

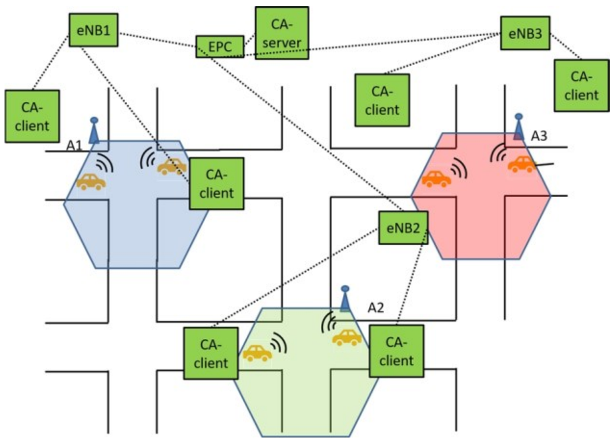

Figure 8: CA use case with fine granularity communications function abstraction at the SO

VII. Conclusion

In this paper we introduce the main orchestration layer of the 5G-TRANSFORMER architecture, namely the Service Orchestrator component. This component is responsible of end-to-end orchestration of services across multiple administrative domains. The SO manages the monitoring and allocation of virtual resources to network slices, and, if makes use of federation with other administrative domains to take decisions on the end-to-end (de)composition of virtual services. We have described the different architectural options being considered by the project, based on the current status of ETSI NFV ISG, where we plan to contribute with outcomes from the project.

We have also provided initial workflows as well as described a key 5G vertical use-case: the automotive one. Next steps include going further into identifying gaps from ETSI NFV specifications required to support 5G-TRANSFORMER features, the design of the required extensions and its validation and evaluation via prototype implementation.

References

[1] N. Alliance, “Description of network slicing concept,” NGMN 5G P, vol. 1, 2016.

[2] 3G-PPP TR 28.801, V15.0.0, Study on management and orchestration of network slicing for next generation network, September 2017

[3] ETSI: “DGR/NFV-EVE012” Work Item. [Online]. Available: https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=51391

[4] ETSI: “DGR/MEC-0024NWslicing” Work Item. [Online]. Available: https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=53580

[5] Y. C. Hu, M. Patel et.al.,, “Mobile edge computinga key technology towards 5g,” ETSI White Paper, vol. 11, 2015.

[6] Casetti et.al, “Network Slices for Vertical Industries”, IEEE WCNCW - Compass 2018.

[7] P. Iovanna, T. Pepe et.al., “5G Mobile Transport and Computing Platform for verticals”, IEEE WCNCW - Compass 2018.

[8] ETSI MEC, [Online]. Available: http://www.etsi.org/technologies-clusters/technologies/multi-access-edge-computing

[9] ETSI GS NFV-MAN 001, V1.1.1, Management and Orchestration, 2014

[10] ETSI GS NFV-IFA 013, V2.3.1, Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Os-Ma-Nfvo reference point - Interface and Information Model Specification, 2017.

[11] A. Anjum, M. Sporny and A. Sill, "Blockchain Standards for Compliance and Trust," in IEEE Cloud Computing, vol. 4, no. 4, pp. 84-90, July/August 2017.

[12] N. Szabo, “The idea of smart contracts,” 1997.

[13] V. Scoca, R. B. Uriarte and R. D. Nicola, "Smart Contract Negotiation in Cloud Computing," 2017 IEEE 10th International Conference on Cloud Computing (CLOUD), Honolulu, CA, 2017, pp. 592-599.

[14] A. Sgambelluri et al., “ Software prototype documentation and user manual,” D3.5 available at http://www.5gex.eu/.

[15] ETSI GR NFV-IFA 028 v0.7.0, “Architecture options to support the offering of NFV MANO services across multiple administrative domains,” June 2017.

[16] ETSI GS NFV-IFA 005, V2.3.1, Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Or-Vi reference point - Interface and Information Model Specification, Aug. 2017.

[17] ETSI GS NFV-IFA 006, V2.3.1, Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Vi-Vnfm reference point - Interface and Information Model Specification, Aug. 2017.