

A Framework for Orchestration and Federation of 5G Services in a Multi-Domain Scenario

L. Valcarenghi¹, B. Martini^{1,2}
¹Scuola Superiore Sant'Anna, Pisa,
Italy
²CNIT, Pisa, Italy

K. Antevski, C. J. Bernardos
Universidad Carlos III de Madrid,
Spain

G. Landi, M. Capitani
Nextworks, Italy

J. Manges-Bafalluy, R.
Martinez, J. Baranda,
I. Pascual
Centre Tecnològic de
Telecomunicacions de Catalunya
CTTC/CERCA, Spain

A. Ksentini
EURECOM, France

C. F. Chiasserini, F. Malandrino
Politecnico di Torino, Italy

Xi Li
NEC Laboratories Europe,
Germany
Xi.Li@neclab.eu

D. Andrushko
KhNURE, Ukraine

Konstantin Tomakh
Mirantis, Ukraine

ABSTRACT

This paper presents the design of the Service Orchestrator (SO), which is one of the key components of the 5G-TRANSFORMER system for the deployment of vertical services. Depending on the requests from verticals, the SO offers service or resource orchestration and federation. These functions include all tasks related to coordinating and providing the vertical with an integrated view of services and resources from multiple administrative domains. In particular, service orchestration entails managing end-to-end services that are split into various domains based on requirements and availability. Federation entails managing administrative relations at the interface between the SOs belonging to different domains and handling abstraction of services. The SO key functionalities, architecture, interfaces, as well as two sample use cases for service federation and service and resource orchestration are presented. Results for the latter use case show that a vertical service is deployed in minutes.

CCS CONCEPTS

Networks, Network services, Programmable networks

KEYWORDS

5G, mobile transport, NFV, end-to-end service orchestration, federation

1 Introduction

5G networks are envisioned to expand the scope of current mobile networks services to support various vertical services, hence enriching the telecom network ecosystem. A wide range of vertical industries, such as eHealth, automotive, media, or cloud robotics, act as drivers to construct this ecosystem. The support of the diverse service requirements of different vertical industries is not only a question of providing broadband capacity, but also a matter of “ultra-reliable low-latency communications” and “massive density connections”.

To enable such vision, EU H2020 5G-PPP phase 2 5G-TRANSFORMER project [1] proposes a flexible and adaptable SDN/NFV-based design of the next generation Mobile Transport Networks. In this design, Network Function Virtualization (NFV), Network Slicing, Multi-access Edge Computing (MEC), and network federation are considered as key concepts to enable such mobile network transformation.

The 5G-TRANSFORMER solution consists of three novel building blocks, namely:

- 1) Vertical Slicer (VS) as the common entry point for all verticals into the system. The VS dynamically creates and maps the services onto network slices according to their requirements, managing their lifecycle. It also translates the vertical and slicing requests into NFV network services (NFV-NS) sent towards the SO. In this sense, a slice will be deployed as a NFV-NS instance.
- 2) Service Orchestrator (SO). It offers service or resource orchestration and federation, depending on the request coming from the VS. Orchestration entails managing end-to-end services or resources that were split into multiple administrative domains based on requirements and availability. Federation entails managing administrative relations at the interface between SOs belonging to different domains and handling abstraction of services and resources.
- 3) Mobile Transport and Computing Platform (MTP) as the underlying unified transport stratum, responsible for providing the resources required by the NFV-NS orchestrated by the SO. This includes their instantiation over the underlying physical transport network, computing, and storage infrastructure. It also may (de)abstract the MTP resources offered to the SO.

This paper focuses on the architecture, functionalities and interfaces of the SO of the 5G-TRANSFORMER system with a particular attention to resource orchestration and federation.

Current software implementing Network Function Virtualisation Orchestrators (NFVOs) (e.g., Cloudify, OpenBaton, Open Source MANO) often simply take in input Network Service Descriptors (NSDs) where the placement of the Virtual Network Functions (VNFs), their interconnections, and the utilized resources are already specified by the vertical/user.

The SO developed within 5G-TRANSFORMER goes beyond the current state of the art by autonomously making some of the aforementioned decisions. The SO takes decisions based on the slice requirements imposed by the different verticals' service requests and the network context (e.g., topology, available resources, etc.). When reaching the SO, vertical requests from the VS have been translated into NFV-NS requests. The SO takes decisions on end-to-end NFV-NS (de)composition along with the allocation of virtual resources from multiple administrative domains so that NFV-NS requirements are met. This decision implies not only the allocation of underlying network, computing and storage resources, and placement of virtual network functions, but also the interaction (federation) with the SOs of other administrative domains when, for instance, requirements cannot be met with services and resources of a single domain. In this way, the virtual resources and the services offered by multiple providers can be aggregated by federating them through their respective SOs.

2 State of the Art Solutions for the SO

Service orchestration and automated lifecycle management are mandatory components of the 5G network software vision. Multiple software implementations of the management and orchestration (MANO) platform are already available both as open source and proprietary solutions.

Open Network Automation Platform (ONAP) was founded in 2017 by merging AT&T driven OpenECOMP project and another NFV orchestration project called Open-O [2]. This project is hosted by Linux foundation and has one of the largest Telco and vendor community among all existing NFV orchestration projects. The first platform release, called Amsterdam, was delivered in November 2017. It was mostly focused on the integration of the software artifacts from the OpenECOMP and Open-O. Thus, just two use-cases are supported in the first release and one of them utilizes proprietary Virtual Network Functions (VNFs). The second platform release is about to emerge, however still at this moment ONAP cannot be considered as a generic orchestration platform suitable for arbitrary NFV use cases.

Open Source MANO (OSM) [3] is an ETSI-hosted project focused on the development of the software stack aligned with ETSI NFV specifications. At present the platform release 4 (OSMv4) is available. The platform has a flexible architecture and each component is both replaceable and pluggable. Overall OSMv4 was targeting production readiness of the OSM platform and multiple improvements were made in this direction. Thus, this platform might be used in pre-production scenarios.

Cloudify [4] is another quite mature open source solution that is already widely implemented in production environments. It has powerful workflow engine and flexible plugin mechanism and service templates are declared in a form of the TOSCA dialect. Despite of the open source nature of this product, Cloudify is a vendor solution and developed by a single company. Another project, OpenBaton [5] has only few developers behind and quite tiny community. OpenStack Tacker [6] project, even if it claimed to be an NFV Orchestrator, focuses on the OpenStack only.

In addition to the industry-driven communities, there are a number of R&D projects funded under the 5G-PPP/H2020 programs, which aim to deliver NFV orchestration platforms. For example, 5G-Crosshaul project [7] delivered an NFVO software prototype able to support the instantiation and termination of Network Services composed of VNFs over a Crosshaul transport network. Other 5G-PPP/H2020 projects like SONATA [8], 5GEx [9] focused on a various operation and use case specific scenarios, like VNF SDK development and multi-domain operation.

Upon a detailed analysis of the existing open source orchestration platforms, most of them do not support network slicing and they all lack of support service federation and MEC, which are the key technology enablers for 5G. Therefore, the SO proposed in the 5G-TRANSFORMER project advances the aforementioned solutions by providing: i) support for network slicing with isolated network service instance management and efficient resource orchestration; ii) support for MEC services and applications; iii) support for service and resource federation by developing novel algorithms, components and interfaces; iv) support of integrated orchestration of heterogeneous cloud and WAN technologies. The new features brought by the SO will pave the way to an autonomic service provisioning to support various vertical services in 5G. The SO implementation is based on Cloudify, while extensive extensions will be made to support the above new features. The selection of Cloudify was made due to its platform maturity, extendable TOSCA dialect to declare service templates and resource dependencies, powerful workflow engine and flexible plugin architecture.

3 SO Architecture and Key Functionalities

This section details the main functions and interworking of the SO subsystems and the related interfaces.

3.1 SO functional system architecture

Figure 1 presents the SO system architecture with a high level overview of the main functional modules and the interactions that need to be developed to achieve the SO operation. The SO system architecture is in line with ETSI NFV guidelines [10] and is inspired by 5GEx [9] and 5G-Crosshaul [7] design.

Northbound Interface (NBI): it offers Application Program Interfaces (APIs) to support requests for service on-boarding, service creation, service instantiation, service modification, service termination.

NS/VNF Catalogue DB and Manager: it is the repository of all usable Network Service Descriptors (NSDs) and VNF Descriptors (VNFDs) that can be accessed through the Catalogue Manager. The NSD is a deployment template for a Network Service the SO can provide (either on its own or by leveraging neighboring SOs). VNFD is a deployment template which describes a VNF in terms of its deployment and operational behavior requirements. The NSD/VNFD is used by the SO in the process of NS-VNF instantiation and their lifecycle management to obtain relevant information, e.g., configuration rules, out-scaling rules. The Catalogue Manager also takes care of advertising NSD to other domains for federation purposes.

NFV Orchestrator (NFVO): it is responsible for orchestration of virtual resources across multiple domains, fulfilling the Resource Orchestration (NFVO-RO) functions, as well as of coordinating the deployment of Network Services (NS) along with their lifecycle management, thereby fulfilling the Network Service Orchestration (NFVO-NSO) functions.




Figure 1: SO functional level architecture

More specifically, the NFVO-NSO coordinates all NS deployment operations including AAA and formal checks of service requests based on attributes retrieved from NSDs and VNFDs. Through the NS-Orchestration Engine (NS-OE), it decomposes the NSDs into several segments and decides to implement them either in the local administrative domain or by federating with the neighbor SOs. Finally, NFVO-NSO requests the NFVO-RO of the local domain or the NFVO-NSO of a neighbor domain to deploy the Network Service segment. The NFVO-RO maps the service segment into a set of virtual resources through the RO Orchestration Engine (RO-OE) by deciding where, in the virtual infrastructure exposed by the MTP, each VNF shall be placed based on specified computational, storage and networking (e.g., bandwidth, latency) requirements. Then, the NFVO-RO takes care of resource provisioning in the MTP domains directly by using So-Mtp/Southbound Interface (SBI) or indirectly by using So-So/east-westbound interface (EBI/WBI), respectively. In the latter case, the sharing of abstract views and the required coordination of correlated actions is carried out by the SO-SO Resource Federation element that executes/forwards requests to the SO NFVO-RO of other domains to request allocations.

VNF Manager (VNFM): it is in charge of the VNF deployment by using either local or remote resources (or a combination of thereof) and of the lifecycle management of the deployed VNFs. It receives relevant VNF lifecycle events from the local NFVO and provides reconfiguration according to specified counteractions decided by the NFVO through VNFDs/NSDs (e.g., auto-scaling).

NFVI Resource Repository: it stores consolidated abstract resource views received from the underlying MTP or from the SO-SO Resource Federation block as for the abstract resource views coming from other domains.

SO-SO Resource Advertisement: it is in charge of exchanging abstract resource views (e.g., abstract topologies, computing and storage capabilities) with other domains while feeding the SO-SO Resource Federation entity that consolidates inputs and stores federated resources into the NFVI Resource Repository.

NS/VNF Instance Repository: it stores the instances of VNFs and NSs that have been previously instantiated.

SO Monitoring Service: it provides the quality measurement reports for the SO to support SO monitoring management including performance monitoring and fault management, based on the collected monitoring data provided by the MTP.

Service Monitoring Data Consumer: it supports the lifecycle management of instantiated VNFs/NSs by collecting quality measurement reports from the SO Monitoring Service and reporting data to the NFVO, e.g., to trigger auto-scaling actions based on scaling rules in the NSD, or to the SLA Manager, e.g., to enable SLA on-line verification. Performance reports can be also used to trigger healing actions to recover from failures or service degradations. The aim is to adapt deployed services or provisioned resources while preventing service degradations due to the concurrent usage of resources from different services.

SLA Manager: it elaborates performance reports from the Service Monitoring Data Consumer during the service lifecycle and assures that the agreed SLAs are continuously met through on-line SLA verification. If a request's SLA parameters are not met, the SLA Manager may trigger scaling actions to prevent or recover from SLA violations.

3.2 Interfaces

The **SO Northbound interface (NBI)** enables the interaction between the VS and the SO. The implementation of the SO NBI is mostly based on the ETSI NFV IFA 13 [11], which defines the NBI of an NFVO. However, the SO NBI implements further methods to allow the SO to handle Network Services extended to include MEC Applications. The NBI is based on the following three reference points.

- **Vs-So(-LCM).** It deals with the NFV network service lifecycle management (LCM) and it is based on the IFA 013 NS Lifecycle Management Interface. It offers primitives to instantiate, terminate, query, update or re-configure network services or receive notifications about their lifecycle.
- **Vs-So(-MON).** It deals with the monitoring (MON) of network services and VNFs through queries or subscriptions and notifications about performance metrics, VNF indicators

and network services or VNFs failures. This interface is based on the IFA 013 NS Performance Management Interface and the IFA 013 NS Fault Management Interface.

- **Vs-So(-CAT)**. It deals with the catalogue (CAT) management of Network Service Descriptors (NSDs), VNF packages, including their VNF Descriptors (VNFDs), and MEC Application Packages, including their Application Descriptors (AppDs). This reference point offers primitives for on-boarding, removal, updates, queries and enabling/disabling of descriptors and packages. It is based on the IFA 013 NSD Management Interface, the IFA 013 VNF Package Management Interface, and the ETSI MEC 10-2 Application Package Management Interface [12], where the SO implements the role of the Mobile Edge Orchestrator (MEO).

The **SO Southbound Interface (SBI)** addresses the interworking between the SO and the MTP building blocks of the 5G-TRANSFORMER architecture. It is worth mentioning that the SO and MTP may follow a 1:N relationship. That is, a single SO may interact via multiple SBI instances with N MTPs which handle the configuration and programmability of a number of domains including heterogeneous virtualized resources for compute, storage and networking. The SO SBI is based on three reference points. They are mostly related to ongoing standard document produced within the ETSI NFV, namely, ETSI GS-NFV IFA 005 and 006 [13][14]:

- **So-Mtp(-RAM)**. It provides the Resource Advertisement Management (RAM) functions. That is, it allows feeding the SO's NFVI resource repository to have a view of the virtualized resources in the MTP that will accommodate the requested NFV-NS. Such information can be delivered by using different levels of detail/abstraction. Thus, the adopted abstraction and vision of the resources, will notably impact the SO NFVO-RO algorithms used for the VNF placement or networking computation.
- **So-Mtp(-RM)**. It encompasses the Resource Management (RM) operations over the virtualized resources. Basically, it contains the set of operations used for reserving, allocating, updating (in terms of scaling up or down) and terminating (i.e., release) the virtualized resources handled by each MTP. In short, and according to the abstracted information managed by the SO, the So-Mtp(-RM) coordinates the involved MTPs to manage the utilization of a selected set of virtualized resources. This entails the VNF placement on the selected compute and storage elements as well as triggering the reservation/ allocation of the network resources (i.e., ports, link bandwidth, etc.) forming the Virtual Network Forwarding Path (VNFP) of the targeted NFV-NS.
- **So-Mtp(-RMM)**. It provides the required primitives and parameters for supporting the SO Resource Monitoring Management (RMM) capability. This entails a twofold functionality: i) fault management to recover/restore the NFV-NS interrupted by a failure (e.g., link failure, VNF host crashes, etc.); ii) permanent performance monitoring providing selected measurements of parameters and metrics such as virtual CPU power consumption, virtual memory oversubscription, etc. which are crucial to ensure the demanded SLA for the existing NFV-NS.

The **SO Eastbound/Westbound Interface (EBI/WBI)** provides service and resource federation. The SO EBI/WBI enables

interaction between the local SO and the SOs of other administrative domains. This SO-SO interface enables an SO to request/offer Network Service as a Service (NSaaS) and NFVI as a Service (NFVIaaS). The SO EBI/WBI is based on six reference points.

- **So-So(-LCM)**, for the operation of NFV network services. The reference point is used to instantiate, terminate, query, update or re-configure network services or receive notifications for federated NFV network services. The SO NFVO-NSO uses this reference point.
- **So-So(-MON)**, for the monitoring of network services through queries or subscriptions/notifications about performance metrics, VNF indicators and network service failures. The SO NFVO-NSO uses this reference point.
- **So-So(-CAT)**, for the management of Network Service Descriptors (NSDs) flavours together with VNF/Vertical Application (VA) and MEC Application Packages, including their Application Descriptors (AppDs). This reference point offers primitives for on-boarding, removal, updates, queries and enabling/disabling of descriptors and packages. The SO NFVO-NSO uses this reference point.
- **So-So(-RM)**, for allocating, configuring, updating and releasing resources. The Resource Management reference point offers operations such as configuration of the resources, configuration of the network paths for connectivity of VNFs. These operations mainly depend of the level of abstraction applied to the actual resources. The SO NFVO-RO uses this reference point.
- **So-So(-RMM)**, for monitoring of different resources, computing power, network bandwidth or latency, storage capacity, VMs, MEC hosts provided by the peering administrative domain. The details level depends on the agreed abstraction level. The SO NFVO-RO uses this reference point.
- **So-So(-RAM)**, for advertising available resource abstractions to/from other SOs. It broadcasts available resources or resource abstractions upon capability calculation and periodic updates for near real-time availability of resources. The SO-SO Resource Advertisement uses this reference point.

The SO EBI/WBI implementation is based on the ETSI NFV IFA 013 and ETSI NFV IFA 005. NSaaS provides NFV network services through the reference points of the SO NFVO-NSO mainly by implementing and extending interfaces of ETSI NFV IFA-013 [11]. NFVIaaS provides resources or resource abstractions through the reference points of SO NFVO-RO and SO-SO Resource Advertising block, generally by implementing interfaces of ETSI NFV IFA-005 [13].

4 SO Use Cases

This section presents two specific use cases involving the SO: service federation and multicloud content delivery service.

4.1 Service Federation in 5GT

Orchestration decisions are triggered whenever the VS requests the creation of a NFV-NS or changes to an existing NFV-NS. In

either case, the SO NFVO must make the following decisions: i) how to decompose a complex NFV-NS (i.e., NSDs) into several segments for efficiency; ii) consequently, where to implement each segment, whether in the local domain or by leveraging neighbor SOs; iii) how to map an NFV-NS segment into a set of virtual resources, i.e., where to run the component VNFs in the virtual infrastructure based on specified resource demand.

This section focuses on the second decision. Service federation is a mechanism for integrating multiple administrative domains unifying them into an open platform for sharing services with a certain degree of trust between each other. Different administrative entities create a peer-to-peer network with pre-established business level terms translated into Service Level Agreements (SLAs). An administrative domain that requests federation procedure is a consumer domain whereas an administrative domain that provides federated resources and/or services is a provider domain.

In 5G-TRANSFORMER, the paradigm of Network Service as a Services (NSaaS) is used for service federation. All federation procedures are executed at the SO level through the E/WBI. NSaaS is a federation procedure where NFV-NS are provided to a consumer domain by a provider domain. The provider domain is responsible for instantiation and life-cycle management of a federated network service. The consumer is only a user of a federated network service. The NSaaS is established in two phases: pre-phase and instantiation phase.

The pre-phase of the NSaaS consists of business agreement between administrative domains, established “offline” on business meetings. Each domain can offer to other administrative domains a set of agreed and offered NFV-NS referred to as catalogue of services. Once the pre-phase is finished, peering catalogues are shared and stored into the NS Catalogue Database of each administrative domain.

Once the instantiation phase of NSaaS is triggered, SO NFVO-NSO checks the NS Catalogue for a desired NFV-NS. A consumer SO decides how to decompose the NFV-NS (i.e., NSDs) into several segments, for example based on service or resource availability, and for the most suitable provider domain. A request for NFV-NS creation is delivered over the E/WBI interface to the chosen provider’s SO NFVO-NSO. The chosen provider’s SO NFVO-NSO initiates service creation and instantiation procedure to its constituent MTP. Upon creation of a federated NFV-NS, positive feedback is sent to the consumer SO NFVO-NSO. Then the consumer SO starts using the federated NFV-NS or adds it to a set of nested services.

The provider SO is responsible for the Lifecycle Management (LCM) of the federated NFV-NS. The consumer SO only keeps track of the quality of the consumed federated NFV-NS and reacts accordingly in case there are violations of the SLA or there is a change in nesting of services.

An example of service federation is included in the workflow depicted in Figure 2. The workflow describes the federation of NFV-NSs (NSaaS) between different administrative domains. It is

assumed that the consumer 5GT-SO receives a request for instantiating a certain NFV-NS (with the NSD) and that it is not capable of satisfying it with its own resources or services. Thus, it exploits service federation to accommodate the requested service. The service federation steps are described here below.




Figure 2: SO Service Federation

1. The administrative domains (white and grey in Figure 2) create a white-list of agreed and offered NFV-NSs which is referred to as catalogue of services. The pre-phase is concluded when the catalogues from each peering administrative domain is stored into the Catalogue DB.
2. The NFV-NS instantiation request is received from the 5GT-VS. The 5GT-SO NFVO-NSO performs decomposition of the requested NFV-NS. In this workflow, a decision is made to consume one of the set of decomposed nested NFV-NSs through federation of NFV-NS or by consuming NFV-NSaaS. As mentioned, the decomposition process of NFV-NS is for further study.
3. NFVO-NSO checks the NFV-NS Catalogue to find out if the desired decomposed NFV-NS can be consumed from another administrative domain. The selection of the potential peering 5GT-SOs is based on either the parameters of the NSD or best-matching capabilities to enable the requested NFV-NS instance. The 5GT-SO takes the role of consumer 5GT-SO (white on Figure 2).
4. Consumer 5GT-SO NFVO-NSO sends the “check for availability” requests to multiple peering 5GT-SOs (or their NFVO-NSOs) that potentially could enable the NFV-NS.

5. The peering 5GT-SO NFVO-NSOs check their availability of the requested service and make decision if it can provide a service instance to the consumer 5GT-SO NFVO-NSO. The peering 5GT-SO NFVO-NSOs send feedback to the consumer 5GT-SO NFVO-NSO.
6. Depending on the received results, the consumer 5GT-SO NFVO-NSO selects the provider 5GT-SO NFVO-NSO or the procedure is terminated if all received responses are negative. The consumer 5GT-SO NFVO-NSO decides for the best-matching provider 5GT-SO NFVO-NSO (e.g., it provides lowest latency for a low-latency NFV-NS and/or at minimal cost). The consumer 5GT-SO NFVO-NSO sends a request for NFV-NS instantiation to the selected provider 5GT-SO NFVO-NSO. The selected provider 5GT-SO NFVO-NSO initiates the service creation and consequently the instantiation procedure involving its constituent 5GT-MTP.
7. Once the requested federated NFV-NS is instantiated, a positive feedback is sent back to the 5GT-SO NFVO-NSO which becomes the consumer of the federated NFV-NS instance.
8. The selected provider 5GT-SO NFVO-NSO updates its constituent Catalogue DB (grey). The consumer 5GT-SO NFVO-NSO consumes the federated NFV-NS by adding it to the set of nested services.

4.2 Multicloud Content Delivery Service

In this use case the implementation of a multicloud content delivery service is considered. The architecture of the implemented service is shown in Figure 3. Content delivery service is deployed in both Amazon Web Service (AWS) public cloud (i.e., SPR.1) and in a private Telco/NFV edge cloud (i.e., SPR.2). The SO, based on extensions of Cloudify, orchestrates the deployment of the service and resources both in the AWS public cloud and in the Telco/NFV edge cloud, including the IPsec tunnel between them with a predefined deployment.




Figure 3: Multicloud content delivery service

Performed deployment experiments show that the overall *service deployment time*, defined as the time elapsing between the service deployment request and the successful service deployment, is about 280s (i.e., less than five minutes).

Once the service is delivered, two performance parameters are measured: the data rate (both between the local cache and the end-user and the local cache and the global repository) and the round trip time (RTT) (in the two aforementioned scenarios). The traffic between the end-user and the local cache shows a bursty behavior with peaks of 10Mbps (the video content used has an average

bitrate of 2Mbps). When the video is not in the local cache the traffic between this virtual machine and the global repository exhibits the same traffic pattern. On average the measured RTT between the end-user and the local cache is about 40 μ s (notice that the end user is directly attached to the local cache) while the one between local cache and global repository is about 40ms.

5 Conclusions

In this paper, we presented the design of the 5G-TRANSFORMER Service Orchestrator (SO). The SO is a novel component responsible for end-to-end orchestration of services and resources across multiple administrative domains. In the paper, we reviewed state of the art solutions for service orchestration already available both as open source and as proprietary solutions, and under research of several EU H2020 5G-PPP projects. We presented in detail the SO functional system architecture, including the main function modules inside the SO, the required interfaces to other components of the 5G-TRANSFORMER system, namely the VS and the MTP, as well as the interfaces to the SOs of other administrative domains enabling multi-domain service and resource federations. We finally presented two use cases where the SO is utilized: service federation and deployment of a multicloud content delivery service involving resource orchestration. For the latter use case results, collected from an experimental testbed involving private and public cloud resources, a service deployment time in the order of minutes is achieved.

ACKNOWLEDGMENTS

This work has been partially funded by the EU H2020 5G-TRANSFORMER Project (grant no. 761536).

REFERENCES

- [1] 5G-transformer.eu
- [2] ONAP project - <https://www.onap.org/>
- [3] OSM project - <https://osm.etsi.org/>
- [4] Cloudify orchestration platform - <https://cloudify.co/>
- [5] OpenBaton project - <https://openbaton.github.io/>
- [6] OpenStack Tacker - <https://wiki.openstack.org/wiki/Tacker>
- [7] 5G-Crosshaul project - <http://5g-crosshaul.eu/>
- [8] SONATA NFV project <http://www.sonata-nfv.eu/>
- [9] 5G-Ex project <http://www.5gex.eu/>
- [10] ETSI GS NFV 002 v1.2.1, "Network Functions Virtualisation (NFV); Architectural Framework", Dec. 2014
- [11] ETSI GS NFV IFA 13, "Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Os-Ma-Nfv reference point – Interface and Information Model Specification" v2.3.1, August 2017
- [12] ETSI GS MEC 010-2, "Mobile Edge Computing (MEC); Mobile Edge Management; Part 2: Application lifecycle, rules and requirements management", v1.1.1, July 2017
- [13] ETSI GS NFV-IFA 005, "Network Function Virtualisation (NFV); Management and Orchestration; Or-Vi reference point – Interface and Information Model Specification", v2.1.1, Apr. 2016
- [14] ETSI GS NFV-IFA 006, "Network Functions Virtualisation (NFV); Management and Orchestration; Vi-Vnm reference point - Interface and Information Model Specification", V2.1.1, Apr. 2016