

MEC support for Network Slicing: Status and Limitations from a Standardization Viewpoint

Luca Cominardi^{*,†}, Thomas Deiß[‡], Miltiadis Filippou[§], Vincenzo Sciancalepore[◇], Fabio Giust[⊥], Dario Sabella[§]
[†]ADLINK Technology, [‡]Nokia Germany, [§]Intel Germany, [◇]NEC Laboratories Europe, [⊥]Athonet

Edge computing and network slicing might be considered as main pillars of the upcoming 5G systems as they inject flexibility in the network management operations. While one prominent architectural framework for edge computing has been recently defined by the ETSI standard organization, namely Multi-access Edge Computing (MEC), network slicing has reached its momentum by fostering interest in different standardization bodies and fora. To better understand how such distinct network slicing definitions impact on the standardized MEC framework, ETSI has recently published a study on the matter. In this paper, we first overview with a comprehensive analysis the different network slicing concepts and their relationship. Then, we elaborate on the ETSI study to provide an integrated view of network slicing technology within the context of MEC. Finally, we report on the open challenges in the ETSI study and we propose two solutions to evolve the current MEC framework towards end-to-end multi-slice support and efficient multi-tenant inter-slice communication in 5G deployments.

Index Terms—Edge computing; virtualization; network slicing; orchestration; MEC; NFV; 3GPP;

I. INTRODUCTION

IN the last decade, the telco industry has clearly delineated the need for flexible network architectures capable of achieving apparently opposite objectives: supporting demanding use cases in terms of performance and reducing costs. To meet such requirements, the European Telecommunications Standards Institute (ETSI) targets dynamic network management and service provisioning via the introduction of virtualisation and computing capabilities into mobile operator networks. Such vision materializes in the architectural framework defined by the ETSI Industry Specification Group (ISG) on Network Function Virtualization (NFV) [1] and the ETSI ISG on Multi-access Edge Computing (MEC) [2].

To this end, the work in [3] is a remarkable study aiming at integrating the MEC and NFV technologies, which is the basis for the MEC-in-NFV architecture available in the second version of [2]. Fig. 1 shows the main MEC components and reference points (shown in red) from such new architecture, which is also the starting point for the discussions that will follow in the present article. The *Network Function Virtualization Infrastructure (NFVI)* provides compute, storage, and network resources to the *MEC Applications (MEC Apps)* and the *MEC Platform*. These are deployed as *Virtual Network Functions (VNFs)* and can be placed close to the users for

latency reduction purposes. The *MEC Platform* offers an environment (i.e., service registry and Domain Name System – DNS – handling), where the MEC Apps can discover, advertise, consume and offer MEC Services (e.g. the Radio Network Information Service – RNIS). The *MEC Platform Manager* is in charge of controlling and configuring the MEC Platform regarding MEC Apps and MEC Services authentication and authorization. Finally, the *MEC Application Orchestrator (MEAO)* is in charge of (i) maintaining an overall view of the MEC Applications and MEC Platforms, including MEC Services, and (ii) interacting with the NFV Orchestrator (NFVO) for resource orchestration. Additional details on the components and reference points can be found in [1], [2], [3].

In addition to network virtualization and edge computing, *network slicing* is another technology reckoned to bring benefits to network operators. It relies on a virtual network architecture and enables a mobile operator to provide dedicated self-contained networks with functionality tailored to the service or customer over a common infrastructure. This approach would hence allow to dynamically configure the same underlying network catering for efficient resource sharing among multiple customers and tenants. As a result, network slicing and multi-tenancy provide a cost effective way to heterogeneous customers to access and use the same shared physical network infrastructure tailored to their needs.

Several Standards Developing Organizations (SDO) and fora have been working on network slicing, which led to the proliferation of distinct definitions. In order to understand the relationship and scope of these concepts, ETSI NFV has published a study on network slicing [4] analysing the potential impact on the ETSI NFV architectural framework. A similar study has been recently published by ETSI MEC [5], with the goal of identifying any necessary extensions to existing components and interfaces. The study analyses the architectural framework shown in Fig. 1 and identifies a set of gaps in MEC that need to be filled to properly support network slicing.

The goal of this paper is twofold: (i) it provides an harmonized view of MEC and network slicing, going beyond the standalone analysis presented in the ETSI MEC study [5], and (ii) it proposes some reference solutions for the gaps left open in the study. Specifically, Section II provides an overview of the network slicing concepts, highlighting their relationship and scope. Next, Section III integrates all these different concepts in the MEC architecture and identifies the components and reference points requiring extension. Section IV proposes two solutions to evolve the current MEC framework towards end-to-end multi-slice and multi-tenant support in 5G deployments. Finally, Section V draws the conclusions of this article.

*Corresponding author: luca.cominardi@adlinktech.com

This work has been partially funded by the EU H2020 projects 5G-CARMEN (grant no. 825012), 5Genesis (grant no. 815178) and the H2020 collaborative EU/TW research project 5G-DIVE (grant no. 859881).

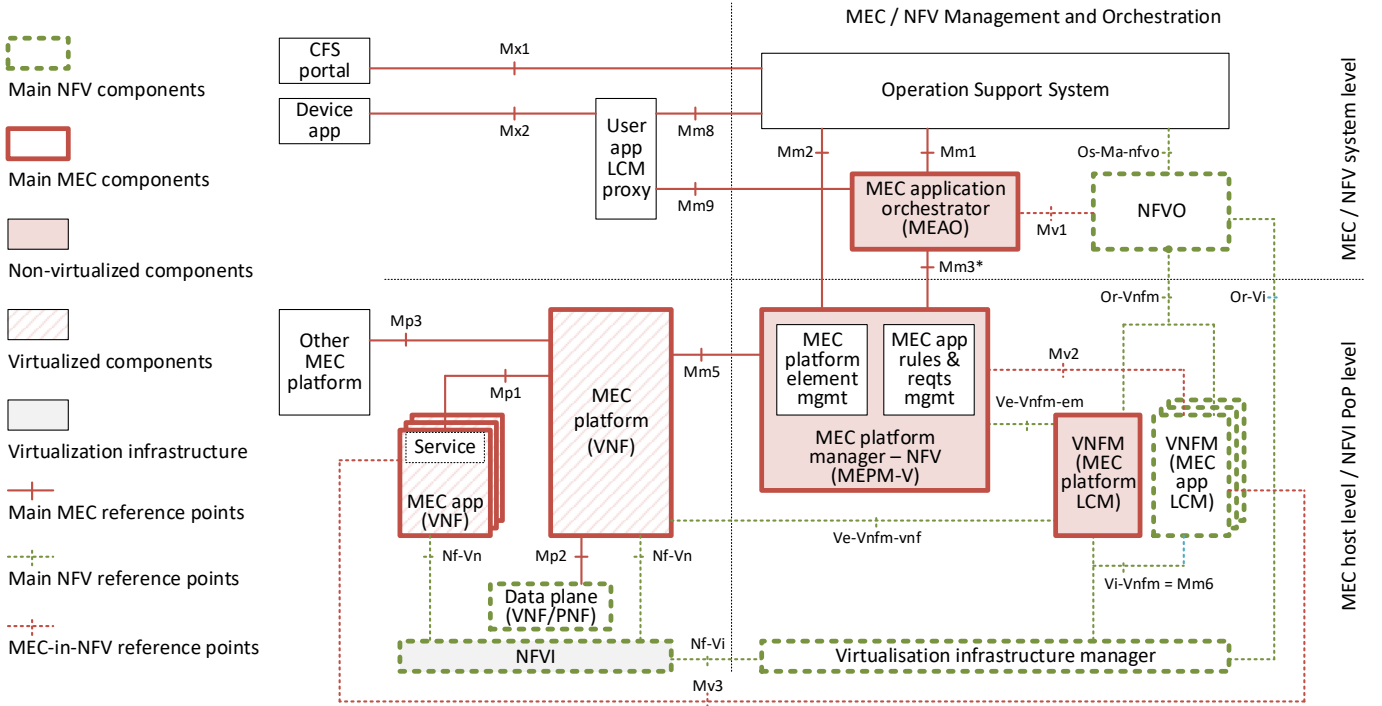


Fig. 1. MEC reference architecture in an NFV environment as proposed in ETSI GR MEC 017 [3].

II. NETWORK SLICING CONCEPTS

The following provides an overview of NGMN, ONF, ETSI NFV and 3GPP network slicing concepts and commonalities. A more extensive comparison can be found in [5]. The following overview is used next in Section III to understand how the various network slicing concepts can be integrated in the MEC framework.

A. NGMN

The Next Generation Mobile Networks (NGMN) white paper [6] first introduced the concept of network slice (i.e., “5G slice”), which is defined as a *collection of 5G network functions and configurations aimed at supporting the communication service of a particular connection type with a specific way of handling the service Control and User planes*. As a result, a 5G slice can span all the domains of the network: from cloud resources, to transport network and radio devices.

A second NGMN white paper [7] organizes the network slicing concepts in a three-layer architecture: Service Instance layer, Network Slice Instance layer, and Resource layer. The first layer encompasses one or more Service Instances (SIs) provided by the network operator or a 3rd party to end-users. The second layer comprises the Network Slice Instances (NSIs) which consist of several network functions forming a complete logical network tailored to certain service characteristics (e.g. ultra-low latency, ultra-high reliability). Finally, the third layer includes all the necessary *physical* and *logical* resources to support the NSIs. Finally, NGMN leverages technology-wise on Software Defined Networking (SDN) and NFV paradigms to provide the necessary flexibility for network slicing.

B. ONF

The Open Networking Foundation (ONF) is best known for standardizing the SDN architecture [8] and protocols (e.g. OpenFlow) [9] with the goal of decoupling the network control and data planes. The ONF architecture comprises three main components: (i) applications, (ii) SDN controller, and (iii) resources. The SDN controller is in charge of mapping the service requirements of the applications to the underlying resources according to certain policies.

In an effort to support the NGMN vision from an SDN perspective, ONF extended the SDN architecture in [10] to allow multiple network instances to share a common infrastructure. To that end, the SDN controller provides two types of resource views: one offered to the application through a *client context*, and a second one enabling the interaction with the resources through a *server context*. The scope of the client context is based on a business agreement between the client and the serving organizations. The client context is then enforced by the SDN controller that (i) allocates part of the resources and (ii) defines how these resources are exposed by creating a server context.

Summarizing, ONF addresses network slicing from a business perspective, in which clients request the network provider to fulfil their specific service needs, including the necessary set of resources and the supporting control logic. This is enabled by the client context in the SDN architecture, which can be directly mapped to a network slice.

C. ETSI NFV

ETSI NFV is best known for standardizing the Network Function Virtualization (NFV) architecture [1] whose goal is to

decouple the network functions from the underlying hardware. The ETSI NFV architecture comprises four main components: (i) Network Function Virtualization Infrastructure (NFVI), (ii) Virtual Network Functions (VNFs), (iii) Network Services (NSs), and (iv) MANagement and Orchestration (MANO). One exemplary instantiation of such architecture in the context of MEC is illustrated in Fig. 1 and described in Section I. In ETSI NFV, network slicing is considered as a means to run multiple logical networks on a common virtualisation infrastructure.

The ETSI NFV study on network slicing [11] mainly focuses on the functional and management isolation among network slices. To achieve management isolation, each network slice may include a tenant SDN controller, which is responsible for configuring the connectivity for the tenant VNFs in the infrastructure. Furthermore, to achieve a full management isolation, each network slice may contain a dedicated Operation and Support System (OSS), NFVO and VNF Managers (VNFM). Nevertheless, [11] allows other combinations of NFVO/VNFM, such as having an NFVO responsible for multiple network slices of a single tenant, following the different multi-domain models described in [4].

ETSI NFV on one hand leverages the various virtualisation technologies to support network slicing at functional level. On the other hand, it investigates the management aspects related to the deployment of a new network slice subsuming a limited set of resources. Isolation at management level is considered for both single and multi-domain scenarios.

D. 3GPP

The 3rd Generation Partnership Project (3GPP) approach to network slicing [?] is based on the NGMN concept and distinguishes between network slices and Network Slice Instances (NSI). A network slice is considered *a logical network that provides specific network capabilities and network characteristics* [?]. An NSI is architecturally defined as *a set of Network Function instances and the required resources (...) which form a deployed Network Slice* [?]. Distinct instances of the same network slice provide specific features based on different Slice/Service Type (SST). A Slice Differentiator (SD) may be used to differentiate amongst multiple network slices of the same SST (e.g. for different customers). 3GPP has defined three standardized SST values [?]: (i) enhanced Mobile Broadband (eMBB), (ii) Ultra-Reliable Low Latency Communications (URLLC), and (iii) Massive Internet-of-Things (MIoT).

Operational and management aspects of network slicing are described in [12]. Specifically, the network slice concept includes: (i) completeness of an NSI; an NSI is complete when it includes all functionalities and resources necessary to support certain communication services serving a business purpose; (ii) Network Slice Template (NST); an NSI is created using a NST containing instance-specific information like policies and configurations; (iii) isolation of NSIs; a NSI may be fully or partly isolated from another NSI. 3GPP defines in [12] an information model for network slices. Each network slice contains one or more network slice subnets, which, in their turn, contain one or more network functions. It should be noted

that a Network Slice Subnet Instance (NSSI) can be shared by multiple NSIs.

3GPP addresses network slicing from both architectural and management perspective by defining (i) an information model for network slices and NSIs and (ii) a set of control plane functions for network slices control and management. Finally, virtualization technologies are leveraged for isolation and flexibility purposes.

E. Commonalities in network slicing concepts

As can be seen from the above, the various network slicing concepts slightly differ in focus and scope. Nevertheless, they share important commonalities: (i) a network slice is seen as a logical network, (ii) a network slice is considered to include the resources to deploy such a logical network, (iii) SDN and NFV are considered key enablers to deploy network slices, (iv) isolation among NSIs is critical at functional and management level, especially in the case of common network functions being shared across multiple NSIs, (v) a network slice may simultaneously support multiple tenants.

Concluding, the different network slicing concepts are complementary and can be combined together in the context of MEC. Particular attention needs to be devoted to the use of either *dedicated* or *shared* network functions because of their impact on the MEC framework.

III. BRINGING NETWORK SLICING IN MEC

This section presents an integrated view of MEC in the context of network slicing, based on an elaboration of the use cases presented in the relevant ETSI MEC study [5]. Particularly, it describes the role of the MEC components and how they can be integrated to simultaneously support the various network slicing concepts. Finally, it reports on the identified open gaps.

A. MEC components and network slicing relationship

Fig. 2 shows a harmonized view of use cases [5], where two NSIs are considered, shown in blue and yellow colour, respectively. There is a number of points worth noting:

- 1) The system illustrated in Fig. 2 is based on the MEC-in-NFV architecture (see Fig. 1), resulting in the MEC Apps and MEC Platform being deployed as VNFs on the NFVI, which is composed of two NFVI-Points-of-Presence (NFVI-PoPs) for sake of simplicity. The MEC Platform Manager (MEPM-V) and the MEC Application Orchestrator (MEAO) are then part of the MANO system. The MEPM-V maintains the configuration of a MEP, e.g., by updating the traffic rules or DNS records.
- 2) The MEC components in the two NSIs are deployed as part of two separate NFV network services, which, in their turn, are mapped to two distinct 3GPP Network Slice Subnets.
- 3) The first NSI (in blue) comprises two MEC Platforms, each deployed on a different NFVI-PoP. The MEC Platform on NFVI-PoP#1 is internal to the NSI and dedicated to the MEC Apps belonging to the same NSI.

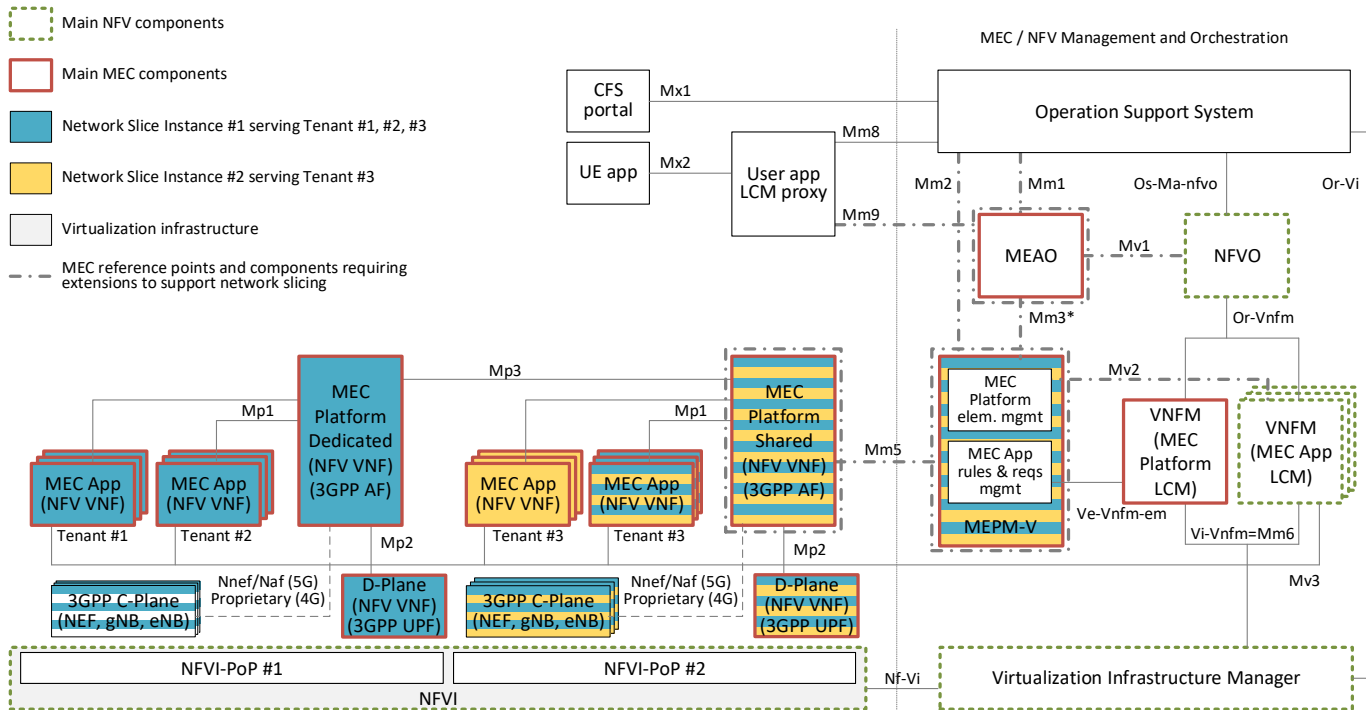


Fig. 2. Elaboration of MEC in NFV supporting network slicing based on the use cases presented in the ETSI GR MEC 024 [5]. Exemplifying scenario with two NFVI-PoPs and two Network Slice Instances serving three distinct tenants.

In this option, the data plane is realized via a dedicated VNF, which is also internal to the same NSI. By doing so, the MEC Platform can influence the traffic redirection via the Mp2 interface, which is kept as a MEC-internal reference point and it is agnostic to the way MEC is deployed. It is worth noting that a MEC Platform can be either shared across multiple NSIs or dedicated to a single NSI.

- 4) The first NSI considers the case where multiple tenants co-exist in the same NSI. For example, the NSI could be an automotive network slice where V2X features are offered to various tenants. Hence, each tenant could be represented by a car manufacturer that deploys its MEC Apps in the same NSI. Summarizing, a MEC Platform can be dedicated to a single NSI, whilst being shared among several MEC Apps belonging to different tenants.
- 5) The MEC Platform on the NFVI-PoP#2 is shared across two NSIs. This is aligned with the ETSI NFV view where VNFs may be shared across multiple NSIs [11]. As a result, the MEC Platform can serve distinct NSIs and offer different sets of features depending on the NSI. Similarly, MEC Apps may simultaneously belong to different NSIs to consume network slice-specific services. It is worth highlighting that MEC Apps can be only associated to one MEC Platform at a given time. In order to consume services across multiple MEC Platforms, these need to communicate over the Mp3 interface to share the services' data.
- 6) In the case of sharing MEC Apps and MEC Platforms among several NSIs, the MEAO and the MEPM-V need to perform different operations depending on the

NSI. For example, MEC Apps associated to a NSI may access only information/traffic of certain users who belong to the same network slice. This implies that MEC Apps requesting services from the MEC Platform, and the MEC Platform itself, need to share an authentication/authorization infrastructure that allows to accept/reject procedures on the granularity of the NSI. For instance, a MEC Platform may need to authenticate against a control plane function (e.g. 3GPP) to retrieve sensitive information of specific network slices or users, e.g., the status of the radio channel. This scenario becomes relevant, e.g., when handling users' mobility, and the consequent application mobility¹, which becomes yet another element to consider for the slice management.

- 7) Finally, in the specific case of a 5G network deployment, the MEC Platform may play the role of a 5G Application Function (AF) towards the 5G core network. In this role, the MEC Platform transmits the traffic offloading requirements to core network elements (e.g. Network Exposure Function – NEF) and the specific MEC Apps traffic to be offloaded by User Plane Functions (UPFs).

B. Open gaps and required extensions

The MEC components and reference points requiring extensions to support network slicing are highlighted in Fig. 2. Some of these gaps are derived from the procedures required for instantiating a network slice involving MEC components and NFV/3GPP systems. For instance, when the 3GPP Network

¹ETSI MEC handles mobility at application level, through the corresponding Application Mobility Service API defined in MEC GS 021.

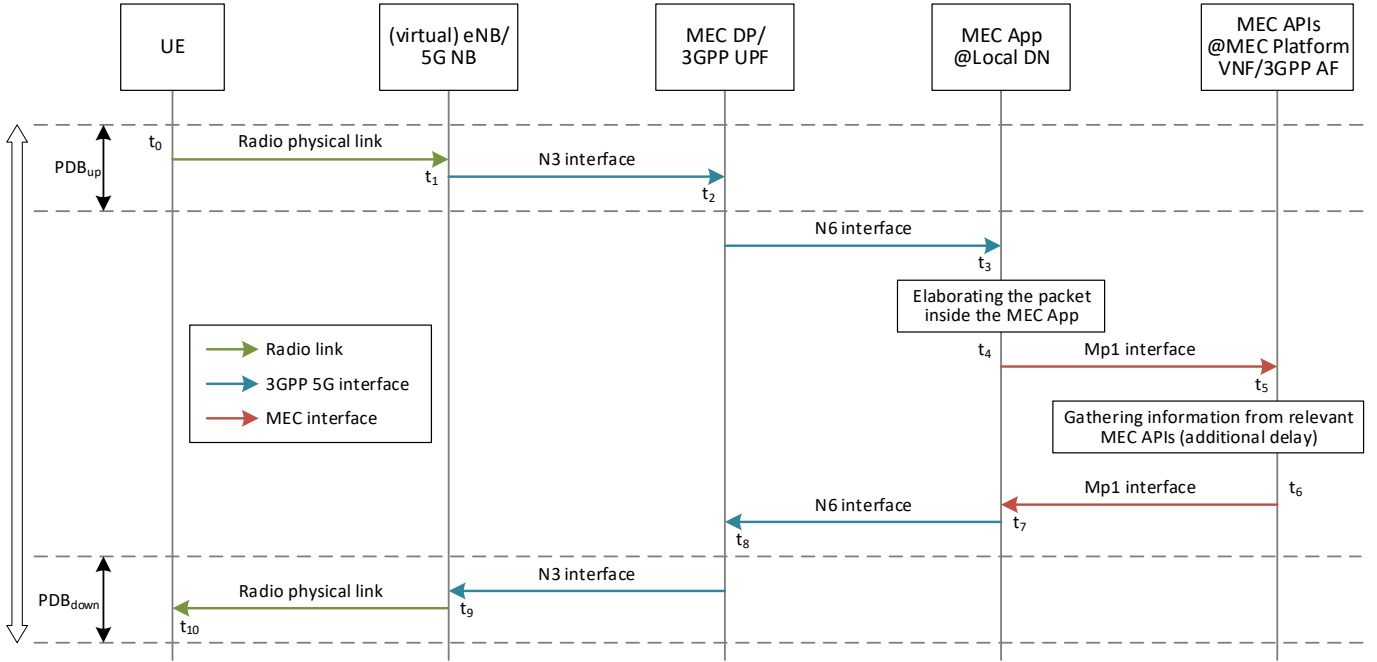


Fig. 3. Latency components involved in a direct communication between a UE client and a MEC App consuming a MEC service from a MEC Platform.

Slice Management Function (NSMF) triggers the creation of a new network slice via the Os-Ma-nfvo reference point, the NFVO may request the deployment of MEC Apps via an Application Descriptor (AppD) to the MEO. The AppD includes the type of traffic to offload and the MEC services to consume, which may belong to specific NSIs. In its turn, the MEO communicates the network slice requirements in terms of MEC services and traffic offloading to the MEPM-V, which, in their turn, are communicated to the MEC Platform. Finally, the MEC Platform enforces the network slices configuration for the MEC Apps, including authentication and authorization.

The *first* open gap in the above procedure resides in the lack of support of network slice-related information in the AppD, which is received by the MEO from clients and over-the-top systems. Additionally, the MEO needs to be slice-aware for enabling per-NSI operations based on their different requirements (e.g. bandwidth, latency, security, user mobility, etc.). Specifically, the Mv1, Mm1, Mm3*, and Mm9 reference points need to be extended to include a reference to the NSI so as to enable the orchestration of multiple network slices. The *second* open gap lies, therefore, in the missing support of network slice-related interfaces in the MEO.

Since a MEC Platform may be shared by several NSIs, this needs to ensure the isolation of the services and information available in a given NSI (or in a set of NSIs) from other network slices. For example, a MEC App may access only the information of the users connected to the same NSI. As a result, the *third* open gap involves the Mm2 and Mm5 reference points on the MEPM-V, which need to expose NSI-related information (e.g. ID, associated MEC Apps, etc.) to enable per-NSI management functionality (e.g. services authorization,

traffic rules, etc.). Finally, the *fourth* open gap resides on the Mv2 reference point between the MEPM-V and the VNFMs, which needs to expose network slice information to support the life cycle management of MEC Apps (e.g. service configuration upon MEC App migration).

As a consequence, to fill the above gaps, a joint effort between ETSI MEC, ETSI NFV, and 3GPP is necessary given the cross-domain nature of some of the involved reference points, as a “MEC-only” solution will not be able to satisfy the network slice requirements across multiple domains.

IV. SOLUTIONS AND EVOLUTION

Hereafter we present two solution proposals to fill the existing gaps of MEC when applied to network slicing scenarios spanning across multiple domains; the aim of these solution proposals is to enhance end-to-end performance and enable multi-tenancy support.

A. Efficient end-to-end multi-slice support for MEC-enabled 5G deployments

The first proposed solution focuses on how to achieve and guarantee the end-to-end latency requirement of a network slice, when instantiated in a virtualized 5G system [13] comprising a MEC system [2]. In the considered scenario, all logical functions, i.e., network functions (NFs) and application functions (AFs), are virtualized. The mapping of MEC components onto those of a 5G system is hence as follows: (i) the MEC Platform is implemented as an AF in 3GPP, (ii) the Data Plane (DP) in the MEC system corresponds to a UPF in 3GPP, and (iii) the MEC Apps are mapped to the local Data Network (DN) in 3GPP.

²The AppD can be transmitted over the Mv1, Mm1, and Mm9 reference points to request the instantiation of MEC Apps to the MEO.

1) End-to-end latency function modeling and evaluation

The starting point for the evaluation of the end-to-end performance, in terms of the total latency experienced by the User Equipment (UE), is the consideration of all entities involved in the Round Trip Time (RTT) of user traffic. As evident from Fig. 3, the end-to-end latency performance for a given network slice spans beyond the 3GPP 5G system domain, as the 5G Quality-of-Service (QoS) Class Identifier (5QI) only covers a part of the total RTT, thus, not accounting for the delays introduced by the MEC system.³ More specifically, the one-way latency introduced by the 3GPP system is composed of: (i) the latency corresponding to the UE-to-UPF communication i.e., the radio physical link followed by the N3 interface, which is characterized by a Packet Delay Budget (PDB), and (ii) the latency associated to the UPF-to-MEC App communication, i.e., involving the delay of the N6 interface along with the packet forwarding time in the UPF itself. For what concerns the MEC system instead, the latency is introduced by the time needed for the instantiated MEC App to gather/consume information from the MEC Platform via the Mp1 interface. From a deployment point of view, it is clear that having both the MEC Apps and the MEC Platform instantiated at the proximity of the 3GPP local DN is latency-efficient, assuming full local availability of the needed MEC services.

Finally, it is important to highlight that the end-to-end latency in a fully virtualized environment not only depends on the communication delay between the end points of the different components, but also on the processing delay within each component (or, functional entity). For instance, the computational operations performed by the MEC App upon receiving a packet from the local DN introduce additional delay which ultimately impacts the latency as perceived by the UE.

2) Slice-aware MEC App allocation policy

Having modeled and evaluated the different components of the network slice's end-to-end latency performance, as depicted in Fig. 3, an interaction between the involved system management entities is needed to instantiate the MEC Apps, whilst fulfilling the end-to-end latency requirement of the network slice the application belongs to. The proposed solution framework consists in the interaction between 3GPP (i.e., OSS of a mobile operator) and MEC system management entities (i.e., MEAO and NFVO). The goal of such interaction is to formulate and then implement a slice-aware MEC App allocation policy taking inputs from the 5G system (OSS), such as the Network Slice Template (NST), containing slice-specific inputs and attributes. The proposed interaction signaling may be performed iteratively and controlled by a new Slice Control Function (SCF); such a function could be implemented within the OSS, the NFVO, or even constitute an independent entity managed by a third party. Fig. 4 graphically illustrates an exemplary implementation of the signaling framework between cross-domain management entities, which is the main feature of the proposed solution.

The objective of the specific solution implementations will be to properly produce, expose and consume cross-domain latency measurements at the involved management functional entities,

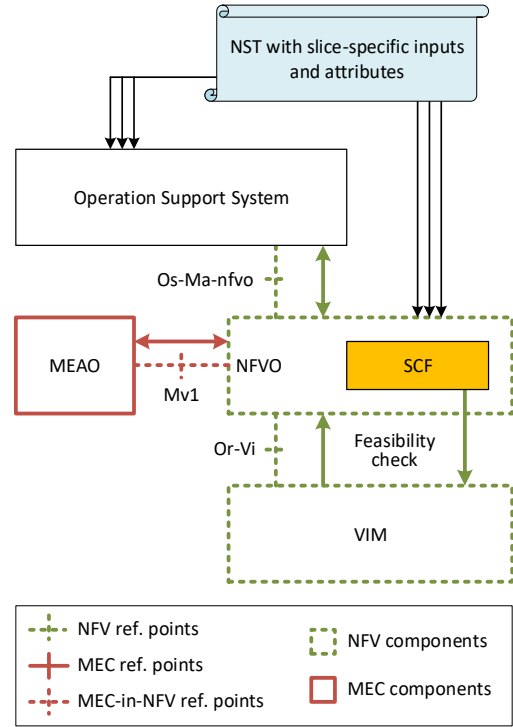


Fig. 4. Exemplary implementation of latency-aware end-to-end multi-slice support for a given QoS flow, addressing relevant recommendations of [5].

identify the latency *bottlenecks* and, based on a utility function, to design a MEC App allocation policy and recommend it to the VIM which is going to provide feasibility feedback. Examples of utility functions are the following:

- *5G system efficiency-centric utility*: The goal is the minimization of the end-to-end delay introduced by the 3GPP 5G system components affecting the slice QoS, subject to a maximum tolerable RTT constraint (with a certain confidence of satisfaction), and a fixed delay of the MEC system components;
- *MEC efficiency-centric utility*: The goal is the minimization of the end-to-end delay introduced by the MEC components affecting the slice QoS, subject to a maximum tolerable RTT constraint and a fixed delay of the 3GPP 5G system components.
- *Overall efficiency utility*: Assuming latency components are of the same priority across the domains, the goal is to minimize slice Service Level Agreement (SLA) breaches.

It should be noted that the described solution proposal addresses the use case recommendations provided in [5]. Nonetheless, its efficiency needs to be measured with regards to the new envisioned signaling between MEC and NFV management entities. This topic is left for future study.

B. Enabling inter-slices communication

When network slicing is in place, the new business model envisions multiple network tenants willing to pay for getting an isolated slice of the physical network. While this concept is generally developed to efficiently manage the network

³Note that UE traffic is terminated at the MEC Application.

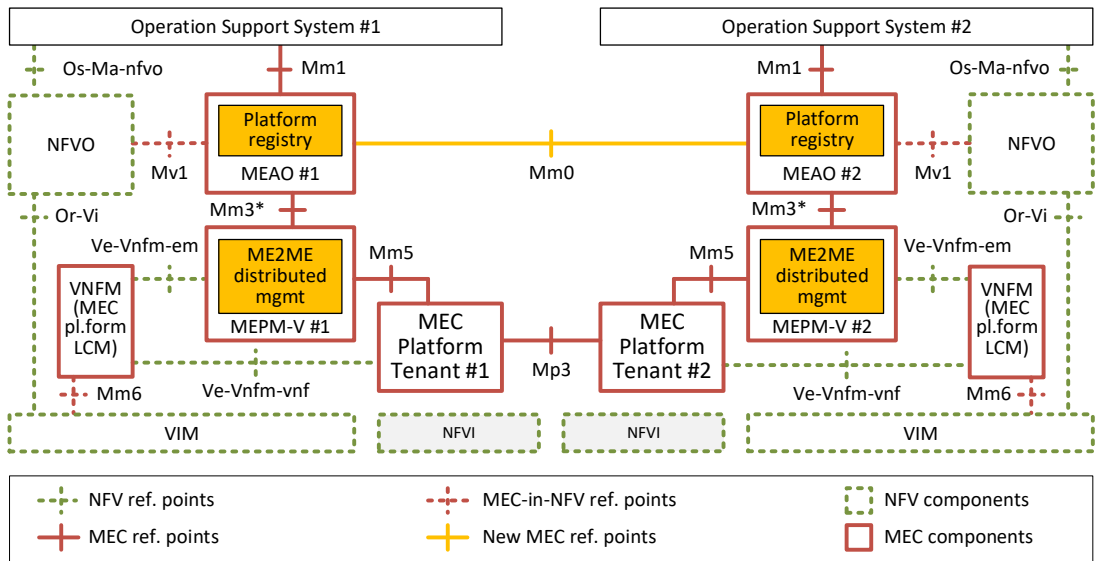


Fig. 5. Novel MEC functionality and components required for efficient inter-slices communication.

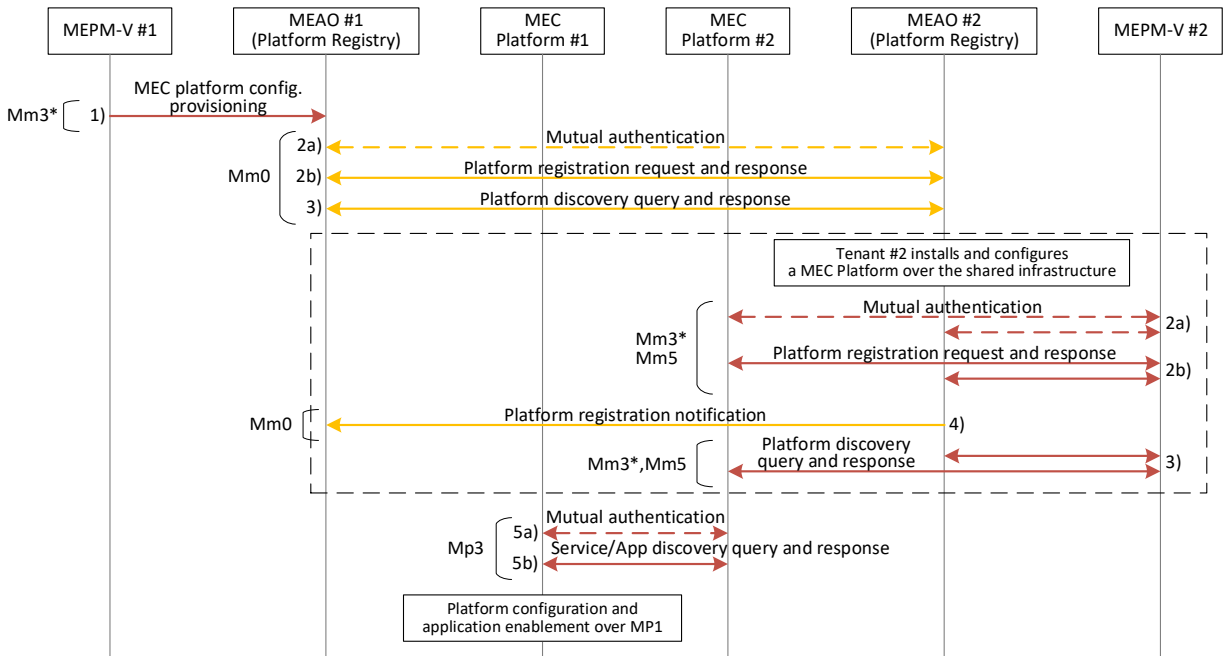


Fig. 6. Signaling chart for MEC Platform registration and configuration for efficient inter-slices communication.

resources, it might be also applied to edge computing premises and can be further supported by the ETSI MEC framework.

For this purpose, let us consider the following scenario: two different network tenants (#1 and #2) get an independent slice of the MEC deployment in order to deliver their own services. They share the same physical MEC deployment that is owned by the infrastructure provider, willing to guarantee isolation but at the same time to increase the system resource efficiency and, in turn, maximizing the return of investment. Physical resources are hardly split between those two tenants and two distinct MEC Platforms are installed onto the virtualized environments facilitated by the infrastructure provider as previously shown in Fig. 2.

MEC Apps and services installed on the MEC stack of tenant #1 will not be shared with end-users (and applications) of the MEC stack of tenant #2, as the physical infrastructure guarantees isolation between those two instances of MEC slices. However, in some cases the isolation property may be broken to facilitate an *inter-slice communication* only for specific resources. Let us consider the scenario where end-users associated to tenant #1 need to directly communicate with users associated to tenant #2. Users might be envisioned for instance as automotive customers, i.e., end-users driving their own car, and tenants can be envisioned as automotive industry [14]. Data to be exchanged between such users might be envisioned (but not limited to) as safety alert messages.

Such messages are time-constrained data messages, which must be delivered within short time deadline. If the inter-slice communication channel is not enabled, messages generated by users of tenant #1 will flow through the entire network stack (e.g. going through the core network of tenant #1) and will come back to the data network of tenant #2 that will instruct correctly the message towards the access network where users of tenant #2 – the recipient of such messages – are located.

Our proposed solution overcomes this problem by keeping affordable the delay of delivering such data by exploiting the overarching MEC capabilities. Specifically, the solution requires the introduction of a new component, namely Platform registry, as shown in Fig. 5. Such novel component is in charge of collecting all business-relation information between different MEC stacks and enabling their self-discovery. Information are collected by means of the reference point, Mm0, which allows the MEC Platform to directly communicate with the Platform registry. Therefore, MEC Platforms are instructed to communicate by means of the Mp3 interface as reported in Fig. 6.

In a nutshell, our proposal will be fully compliant with the ETSI MEC architecture, as per [2]. In particular, the network tenant installs and configures the ETSI-compliant MEC Platform in the shared infrastructure via Mm5 reference point. Once configured, the platform authenticates and registers itself towards the shared infrastructure’s platform registry. This operation is performed by exchanging control messages over the Mm0 reference point. Once successfully registered, the MEC Platform may start a MEC Platform discovery procedure over the Mm0 reference point. Should another MEC platform be installed in the system, it may perform the registration and discovery procedures as described above. After the successful registration, the platform registry may send a platform registration notification to all the MEC platforms that can access the newly installed one. To each recipient platform, the list of exposed services and applications is included as well. Then, a MEC Platform, upon discovering the presence of other MEC Platforms, can start a communication with one or more of them using the interfaces supported by reference point Mp3. Finally, after obtaining the list of services and applications available through other platforms, the requesting platform updates its service registry, DNS database and performs application enablement procedures as required following the standard procedures defined in [15].

V. CONCLUSIONS

This paper addresses the issue of integrating MEC with network slicing, two of the most promising technologies for mobile network operators to generate new revenue streams while reducing the network management complexity. Based on an elaboration of [5], an integrated view of network slicing in the context of MEC has been formulated and open gaps have been identified at functional and management level. By filling those gaps, it would be hence possible to create a multi-domain system capable of verifying the network slice requirements across distinct domains. Upon SLA breach, the system would be therefore able to react accordingly by exploiting multi-domain orchestration and management capability.

Two solutions have been proposed in this paper to evolve the current MEC framework towards end-to-end multi-slice support in 5G deployments: (i) a new Slice Control Function (SCF) to enable slice-aware MEC App allocation via the interaction of MEC, NFV and 3GPP systems and (ii) an inter-slice communication channel that automatically filters exchanged data between slices installed on the same MEC facilities. From a standardization point of view, a collaboration between ETSI MEC, ETSI NFV and 3GPP is required for extending the relevant reference points and procedures, like the ones proposed in this paper, to truly enable end-to-end multi-domain network slicing, including critical aspects like configuration, monitoring and charging.

REFERENCES

- [1] ETSI, “Network Functions Virtualisation (NFV); Architectural Framework,” European Telecommunications Standards Institute (ETSI), Group Specification (GS) NFV 002 v1.1.1, Oct. 2013.
- [2] —, “Mobile Edge Computing (MEC); Framework and Reference Architecture,” European Telecommunications Standards Institute (ETSI), Group Specification (GS) 003 v2.1.1, Jan. 2019.
- [3] —, “Mobile Edge Computing (MEC); Deployment of Mobile Edge Computing in an NFV environment,” European Telecommunications Standards Institute (ETSI), Group Report (GR) 017 V1.1.1, Feb. 2018.
- [4] —, “Network Functions Virtualisation (NFV); Management and Orchestration; Report on Architectural Options to Support Multiple Administrative Domains,” European Telecommunications Standards Institute (ETSI), Group Report (GR) NFV-IFA 028 v3.1.1, Jan. 2018.
- [5] —, “Multi-access Edge Computing (MEC); Support for network slicing,” European Telecommunications Standards Institute (ETSI), Group Specification (GS) 024 v2.1.1, Nov. 2019.
- [6] NGMN, “5G White Paper,” Next Generation Mobile Networks Alliance, White Paper v1.0, Feb. 2015.
- [7] —, “Description of Network Slicing Concept,” Next Generation Mobile Networks Alliance, White Paper v1.0, Feb. 2016.
- [8] ONF, “SDN architecture: Issue 1.1,” Open Networking Foundation, Technical Report (TR) 521, Jan. 2016.
- [9] —, “OpenFlow switch specification: Version 1.5.1,” Open Networking Foundation, Technical Specification (TS) 025, Mar. 2015.
- [10] —, “Applying SDN Architecture to 5G slicing,” Open Networking Foundation, Technical Report (TR) 526, Apr. 2016.
- [11] ETSI, “Network Functions Virtualisation (NFV) Release 3; Evolution and Ecosystem; Report on Network Slicing Support with ETSI NFV Architecture Framework,” European Telecommunications Standards Institute (ETSI), Tech. Rep. NFV 002 v3.1.1, Dec. 2017.
- [12] 3GPP, “Telecommunication management; Study on management and orchestration of network slicing for next generation network,” 3rd Generation Partnership Project (3GPP), Technical Report (TR) 28.801 v15.1.0, Apr. 2018.
- [13] —, “System Architecture for the 5G System,” 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 23.501 v15.8.0, Dec. 2019.
- [14] F. Giust, V. Sciancalepore, D. Sabella, M. C. Filippou, S. Mangiante, W. Featherstone, and D. Munaretto, “Multi-Access Edge Computing: The Driver Behind the Wheel of 5G-Connected Cars,” *IEEE Communications Standards Magazine*, vol. 2, no. 3, pp. 66–73, Sept. 2018.
- [15] ETSI, “Mobile Edge Platform Application Enablement,” European Telecommunications Standards Institute (ETSI), Group Specification (GS) 011 v1.1.1, Jul. 2017.