



Multi-access Edge Computing (MEC); Framework and Reference Architecture

Disclaimer: This DRAFT is a working document of ETSI ISG MEC. It is provided for information only and is still under development within ETSI ISG MEC. ETSI and its Members accept no liability for any further use/implementation of this Specification.

Non-published MEC drafts stored in the ["Open Area"](#) are working documents, these may be updated, replaced, or removed at any time

Do not use as reference material.

Disclaimer

Do not cite this document other than as "work in progress".

The present document has been produced and approved by the Multi-access Edge Computing (MEC) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG. It does not necessarily represent the views of the entire ETSI membership.

Approved and published Specifications and reports for implementation of the MEC system shall be obtained via the ETSI Standards search page at:
<http://www.etsi.org/standards-search>

Reference

RGS/MEC-0003v221Arch

Keywords

architecture, MEC

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

The present document can be downloaded from:
<http://www.etsi.org/standards-search>

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format at www.etsi.org/deliver.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at <https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx>

If you find errors in the present document, please send your comment to one of the following services:
<https://portal.etsi.org/People/CommitteeSupportStaff.aspx>

Copyright Notification

Reproduction is only permitted for the purpose of standardization work undertaken within ETSI.
The copyright and the foregoing restrictions extend to reproduction in all media.

© ETSI 2020.
All rights reserved.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members.
3GPP™ and **LTE™** are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.

oneM2M™ logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners.

GSM® and the GSM logo are trademarks registered and owned by the GSM Association.

Contents

Intellectual Property Rights	5
Foreword.....	5
Modal verbs terminology.....	5
1 Scope	6
2 References	6
2.1 Normative references	6
2.2 Informative references.....	6
3 Definition of terms, symbols and abbreviations.....	7
3.1 Terms.....	7
3.2 Symbols.....	7
3.3 Abbreviations	7
4 Overview	7
5 Multi-access Edge Computing framework	8
6 Reference architecture.....	9
6.1 Generic reference architecture.....	9
6.2 Reference architecture variant for MEC in NFV	10
6.2.1 Description.....	10
6.2.2 Architecture diagram	10
7 Functional elements and reference points	11
7.1 Functional elements.....	11
7.1.1 MEC host.....	11
7.1.2 MEC platform.....	11
7.1.3 MEC application.....	12
7.1.4 MEC system level management.....	12
7.1.4.1 Multi-access edge orchestrator.....	12
7.1.4.2 Operations Support System (OSS).....	12
7.1.4.3 User application lifecycle management proxy	12
7.1.5 MEC host level management	13
7.1.5.1 MEC platform manager.....	13
7.1.5.2 Virtualization infrastructure manager	13
7.1.6 Device application	13
7.1.7 Customer facing service portal	14
7.1.8 Specific functional elements in the MEC in NFV architecture variant.....	14
7.1.8.1 Overview.....	14
7.1.8.2 MEC application orchestrator	14
7.1.8.3 MEC platform manager - NFV	14
7.1.8.4 NFVO.....	14
7.1.8.5 VNFM (MEC Platform LCM)	14
7.1.8.6 VNFM (MEC App LCM)	14
7.2 Reference points	14
7.2.1 Reference points related to the MEC platform	14
7.2.2 Reference points related to the MEC management.....	15
7.2.3 Reference points related to external entities	15
7.2.4 Reference points related to the MEC in NFV architecture variant	15
8 MEC services	16
8.1 General	16
8.2 Radio Network Information	16
8.3 Location.....	17
8.4 Traffic Management Services.....	17
9 Inter-MEC system communication	17

Annex A (informative):	Key concepts	18
A.1	MEC host selection	18
A.2	DNS support	18
A.3	Application traffic filtering and routing	19
A.4	Support of application and UE mobility	19
A.4.1	Background: UE mobility	19
A.4.2	MEC application scenarios for UE mobility	19
A.4.2.1	MEC applications not sensitive to UE mobility	19
A.4.2.2	MEC applications sensitive to UE mobility	19
A.4.2.2.1	Maintaining connectivity between UE and MEC application instance	19
A.4.2.2.2	Application state relocation	19
A.4.2.2.3	Application instance relocation within the MEC system	20
A.4.2.2.4	Application instance relocation between the MEC system and an external cloud environment	20
A.5	Void	20
A.6	Data Plane	20
A.7	API gateway support	20
History		21

Draft

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: "*Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards*", which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (<https://ipr.etsi.org/>).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification Group (ISG) Multi-access Edge Computing (MEC).

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the [ETSI Drafting Rules](#) (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The present document provides a framework and reference architecture for Multi-access Edge Computing that describes a MEC system that enables MEC applications to run efficiently and seamlessly in a multi-access network. The present document also describes the functional elements and the reference points between them, and a number of MEC services that comprise the solution. It finally presents a number of key concepts related to the multi-access edge architecture.

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at <https://docbox.etsi.org/Reference/>.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are necessary for the application of the present document.

- [1] ETSI GS MEC 002: "Multi-access Edge Computing (MEC); Phase 2: Use Cases and Requirements".
- [2] ETSI GS NFV 002: "Network Functions Virtualisation (NFV); Architectural Framework".
- [3] ETSI GS NFV-IFA 013: "Network Functions Virtualisation (NFV); Management and Orchestration; Os-Ma-nfvo reference point - Interface and Information Model Specification".
- [4] ETSI GS NFV-IFA 008: "Network Functions Virtualisation (NFV); Management and Orchestration; Ve-Vnfm reference point - Interface and Information Model Specification".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

- [i.1] ETSI GS MEC 001: "Multi-access Edge Computing (MEC); Terminology".
- [i.2] Void.
- [i.3] OpenStack®: "OpenStack++ for Cloudlet Deployment".

NOTE 1: Available at <http://reports-archive.adm.cs.cmu.edu/anon/2015/CMU-CS-15-123.pdf>.

NOTE 2: The OpenStack® Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. ETSI is not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

- [i.4] Void.

- [i.5] ETSI GR MEC 017: "Mobile Edge Computing (MEC); Deployment of Mobile Edge Computing in an NFV environment".
- [i.6] ETSI TS 123 501: "5G; System architecture for the 5G System (5GS) (3GPP TS 23.501 Release 15)".

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the terms given in ETSI GS MEC 001 [i.1] apply.

3.2 Symbols

Void.

3.3 Abbreviations

For the purposes of the present document, the abbreviations given in ETSI GS MEC 001 [i.1] and the following apply:

CFS	Customer Facing Service
LCM	LifeCycle Management
MEAO	MEC Application Orchestrator
MEPM-V	MEC Platform Manager - NFV
NFV	Network Functions Virtualization
NFVI	Network Functions Virtualization Infrastructure
NFVO	NFV Orchestrator
NS	Network Service
OSS	Operations Support System
VIM	Virtualization Infrastructure Manager
VNF	Virtualized Network Function
VNFM	VNF Manager

4 Overview

The present document presents a framework and a reference architecture to support the requirements defined for Multi-access Edge Computing in ETSI GS MEC 002 [1].

The framework described in clause 5 shows the structure of the Multi-access Edge Computing environment.

The reference architecture described in clause 6 shows the functional elements that compose the multi-access edge system, including the MEC platform and the MEC management, as well as the reference points between them.

The functional elements and reference points listed in clause 7 describe the high-level functionality of the different functional elements and reference points.

Clause 8 describes the high-level functionality of a number of MEC services, comprising the solution for Multi-access Edge Computing.

Annex A describes at a high-level a number of key concepts that underlie the principles used to develop the framework and reference architecture described in the present document.

5 Multi-access Edge Computing framework

Multi-access Edge Computing enables the implementation of MEC applications as software-only entities that run on top of a virtualization infrastructure, which is located in or close to the network edge. The Multi-access Edge Computing framework shows the general entities involved. These can be grouped into system level, host level and network level entities.

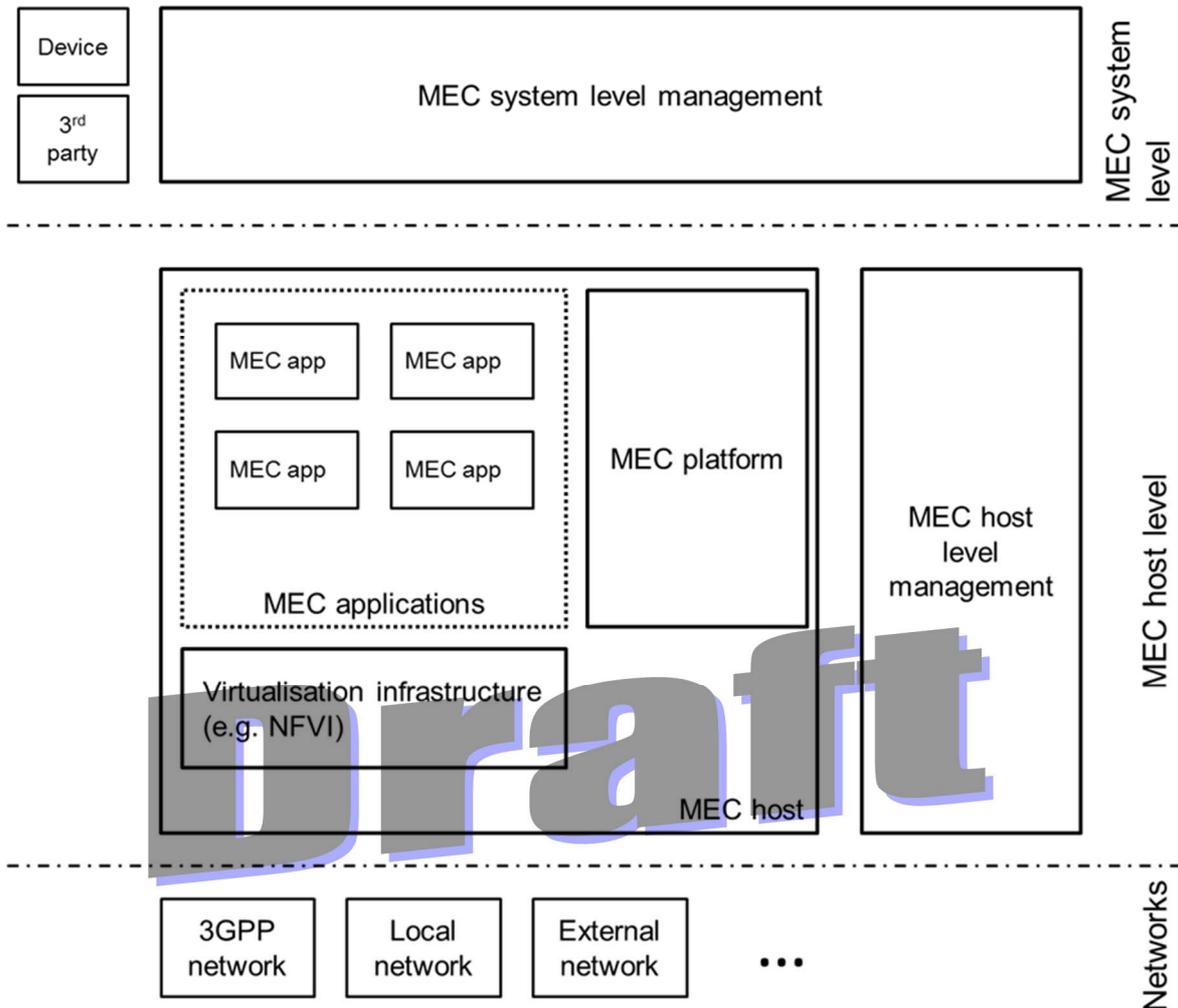


Figure 5-1: Multi-access Edge Computing framework

Figure 5-1 illustrates the framework for Multi-access Edge Computing consisting of the following entities:

- MEC host, including the following:
 - MEC platform;
 - MEC applications;
 - virtualization infrastructure;
- MEC system level management;
- MEC host level management;
- external related entities, i.e. network level entities.

6 Reference architecture

6.1 Generic reference architecture

The reference architecture shows the functional elements that comprise the multi-access edge system and the reference points between them.

Figure 6-1 depicts the generic multi-access edge system reference architecture. There are three groups of reference points defined between the system entities:

- reference points regarding the MEC platform functionality (Mp);
- management reference points (Mm); and
- reference points connecting to external entities (Mx).

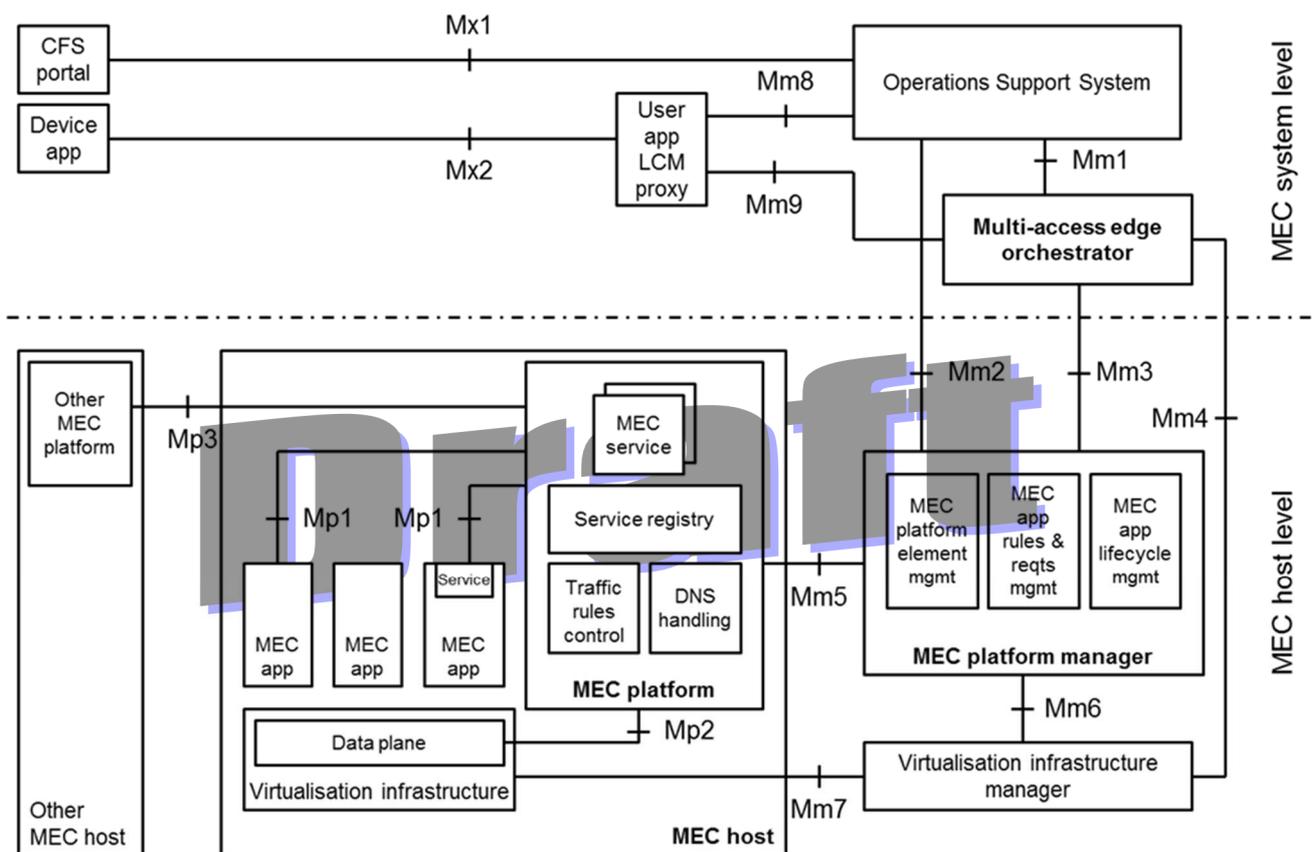


Figure 6-1: Multi-access edge system reference architecture

The multi-access edge system consists of the MEC hosts and the MEC management necessary to run MEC applications within an operator network or a subset of an operator network.

The **MEC host** is an entity that contains a MEC platform and a virtualization infrastructure which provides compute, storage, and network resources, for the purpose of running MEC applications. The MEC host is further described in clause 7.1.1.

The **MEC platform** is the collection of essential functionality required to run MEC applications on a particular virtualization infrastructure and enable them to provide and consume MEC services. The MEC platform can also provide services. The MEC platform is further described in clause 7.1.2.

MEC applications are instantiated on the virtualization infrastructure of the MEC host based on configuration or requests validated by the MEC management. MEC applications are further described in clause 7.1.3.

The MEC management comprises the MEC system level management and the MEC host level management.

The MEC system level management includes the **Multi-access edge orchestrator** as its core component, which has an overview of the complete MEC system. The MEC system level management is further described in clause 7.1.4.

The MEC host level management comprises the **MEC platform manager** and the **virtualization infrastructure manager**, and handles the management of the MEC specific functionality of a particular MEC host and the applications running on it. The MEC host level management is further described in clause 7.1.5.

6.2 Reference architecture variant for MEC in NFV

6.2.1 Description

Multi-access Edge Computing (MEC) and Network Functions Virtualization (NFV) are complementary concepts. The MEC architecture has been designed in such a way that a number of different deployment options of MEC systems are possible. A dedicated Group Report, ETSI GR MEC 017 [i.5], provides an analysis of solution details of the deployment of MEC in an NFV environment.

In clauses 6.2.2, 7.1.8 and 7.2.4 of the present document, a MEC architecture variant is specified that allows to instantiate MEC applications and NFV virtualized network functions on the same virtualization infrastructure, and to re-use ETSI NFV MANO components to fulfil a part of the MEC management and orchestration tasks.

6.2.2 Architecture diagram

Figure 6-2 depicts a variant of the multi-access edge system reference architecture for the deployment in a Network Functions Virtualization (NFV) environment [2].

In addition to the definitions for the generic reference architecture in clause 6.1, the following new architectural assumptions apply:

- The MEC platform is deployed as a VNF.
- The MEC applications appear as VNFs towards the ETSI NFV MANO components.
- The virtualization infrastructure is deployed as an NFVI and is managed by a VIM as defined by ETSI GS NFV 002 [2].
- The MEC Platform Manager (MEPM) is replaced by a MEC Platform Manager - NFV (MEPM-V) that delegates the VNF lifecycle management to one or more VNF Managers (VNFM).
- The MEC Orchestrator (MEO) is replaced by a MEC Application Orchestrator (MEAO) that relies on the NFV Orchestrator (NFVO) for resource orchestration and for orchestration of the set of MEC application VNFs as one or more NFV Network Services (NSs).

The new reference points shown in figure 6-2 are further described in clause 7.2.4.

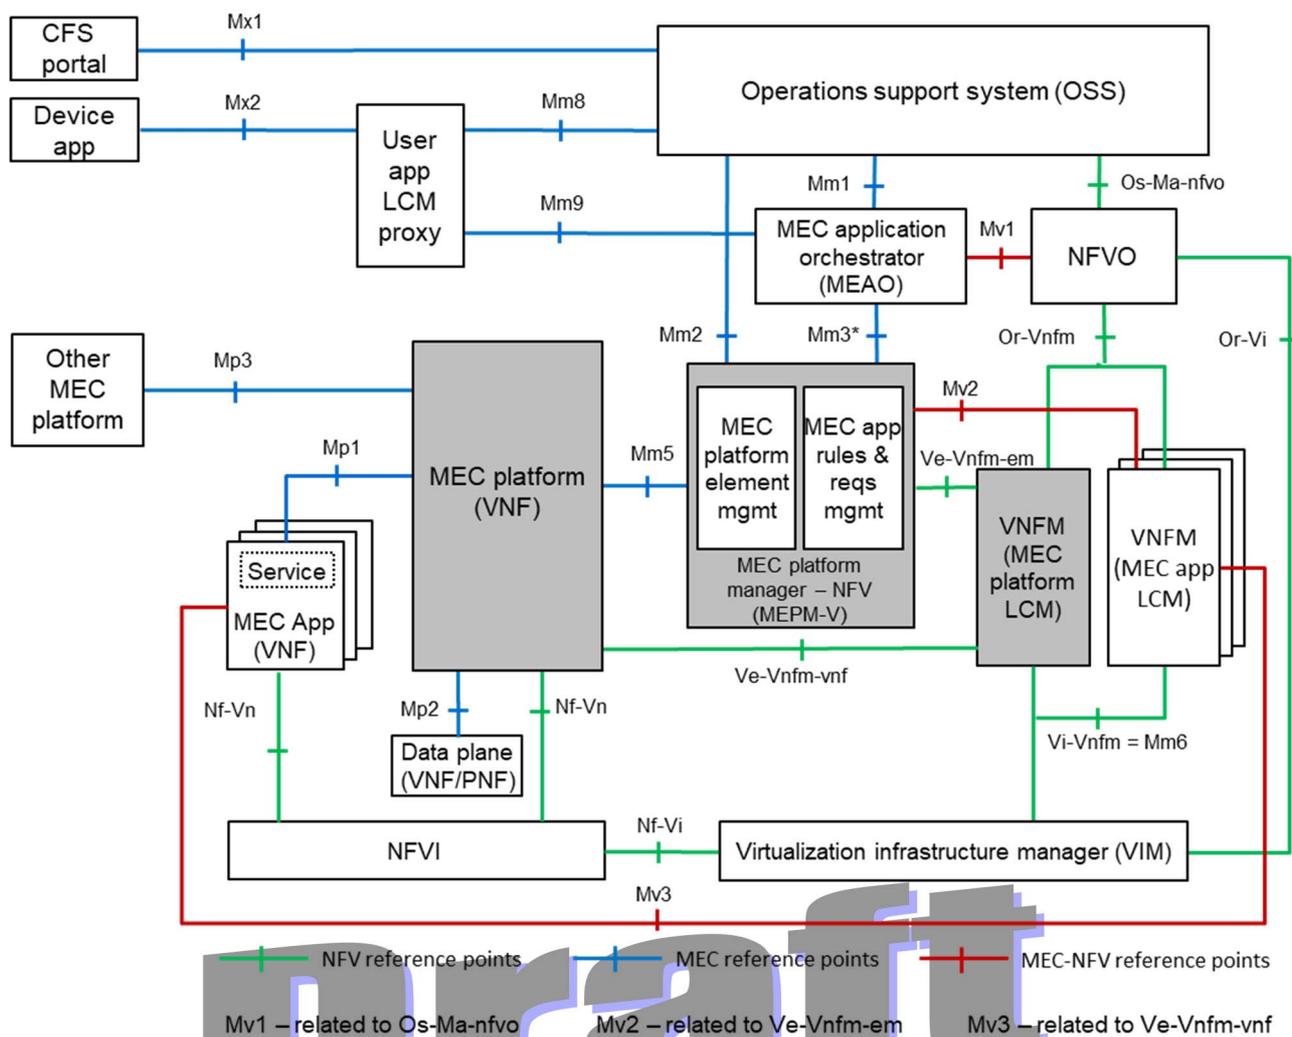


Figure 6-2: Multi-access edge system reference architecture variant for MEC in NFV

7 Functional elements and reference points

7.1 Functional elements

7.1.1 MEC host

The MEC host is an entity that contains the MEC platform and a virtualization infrastructure which provides compute, storage, and network resources for the MEC applications. The virtualization infrastructure includes a data plane that executes the traffic rules received by the MEC platform, and routes the traffic among applications, services, DNS server/proxy, 3GPP network, other access networks, local networks and external networks.

7.1.2 MEC platform

The MEC platform is responsible for the following functions:

- offering an environment where the MEC applications can discover, advertise, consume and offer MEC services (see clause 8), including, when supported, MEC services available via other platforms (that may be in the same or a different MEC system);

- receiving traffic rules from the MEC platform manager, applications, or services, and instructing the data plane accordingly. When supported, this includes the translation of tags representing UEs in the traffic rules into specific IP addresses;
- receiving DNS records from the MEC platform manager and configuring a DNS proxy/server accordingly;
- hosting MEC services, possibly including services that are described in clause 8;
- providing access to persistent storage and time of day information.

The MEC platform may be accompanied by an API gateway functionality for MEC applications to access the MEC service APIs.

7.1.3 MEC application

MEC applications are running as Virtual Machines (VM) on top of the virtualization infrastructure provided by the MEC host, and can interact with the MEC platform to consume and provide MEC services (described in clause 8).

In certain cases, MEC applications can also interact with the MEC platform to perform certain support procedures related to the lifecycle of the application, such as indicating availability, preparing relocation of user state, etc.

MEC applications can have a certain number of rules and requirements associated to them, such as required resources, maximum latency, required or useful services, etc. These requirements are validated by the MEC system level management, and can be assigned to default values if missing.

7.1.4 MEC system level management

7.1.4.1 Multi-access edge orchestrator

The multi-access edge orchestrator is the core functionality in MEC system level management.

The multi-access edge orchestrator is responsible for the following functions:

- maintaining an overall view of the MEC system based on deployed MEC hosts, available resources, available MEC services, and topology;
- on-boarding of application packages, including checking the integrity and authenticity of the packages, validating application rules and requirements and if necessary adjusting them to comply with operator policies, keeping a record of on-boarded packages, and preparing the virtualization infrastructure manager(s) to handle the applications;
- selecting appropriate MEC host(s) for application instantiation based on constraints, such as latency, available resources, and available services;
- triggering application instantiation and termination;
- triggering application relocation as needed when supported.

7.1.4.2 Operations Support System (OSS)

The Operations Support System (OSS) in figure 6-1 refers to the OSS of an operator. It receives requests via the CFS portal and from device applications for instantiation or termination of applications, and decides on the granting of these requests. Granted requests are forwarded to the multi-access edge orchestrator for further processing.

When supported, the OSS also receives requests from device applications for relocating applications between external clouds and the MEC system.

7.1.4.3 User application lifecycle management proxy

A user application is a MEC application that is instantiated in the MEC system in response to a request of a user via an application running in the device (device application).

The user application lifecycle management proxy allows device applications to request on-boarding, instantiation, termination of user applications and when supported, relocation of user applications in and out of the MEC system. It also allows informing the device applications about the state of the user applications.

The user application lifecycle management proxy authorizes requests from device applications in the device (e.g. UE, laptop with internet connectivity) and interacts with the OSS and the multi-access edge orchestrator for further processing of these requests.

The user application lifecycle management proxy is only available when supported by the MEC system.

7.1.5 MEC host level management

7.1.5.1 MEC platform manager

The MEC platform manager is responsible for the following functions:

- managing the life cycle of applications including informing the multi-access edge orchestrator of relevant application related events;
- providing element management functions to the MEC platform;
- managing the application rules and requirements including service authorizations, traffic rules, DNS configuration and resolving conflicts.

The MEC platform manager also receives virtualized resources fault reports and performance measurements from the virtualization infrastructure manager for further processing.

7.1.5.2 Virtualization infrastructure manager

The virtualization infrastructure manager is responsible for the following functions:

- allocating, managing and releasing virtualized (compute, storage and networking) resources of the virtualization infrastructure;
- preparing the virtualization infrastructure to run a software image. The preparation includes configuring the infrastructure, and can include receiving and storing the software image;
- when supported, rapid provisioning of applications, as described in "Openstack++ for Cloudlet Deployments" [i.3];
- collecting and reporting performance and fault information about the virtualized resources;
- when supported, performing application relocation. For application relocation from/to external cloud environments, the virtualization infrastructure manager interacts with the external cloud manager to perform the application relocation, for example using the mechanism described in "Adaptive VM Handoff Across Cloudlets" [i.3], possibly through a proxy.

The functionality provided by the virtualization infrastructure manager in the present document and the functionality provided by the virtualized infrastructure manager described in ETSI GS NFV 002 [2], clause 7.2.5, overlap to a large extent.

7.1.6 Device application

Device applications as defined in the present document are applications in the device (e.g. UE, laptop with internet connectivity) that have the capability to interact with the MEC system via a user application lifecycle management proxy, as defined in clause 7.1.4.3.

7.1.7 Customer facing service portal

The customer facing service portal allows operators' third-party customers (e.g. commercial enterprises) to select and order a set of MEC applications that meet their particular needs, and to receive back service level information from the provisioned applications.

7.1.8 Specific functional elements in the MEC in NFV architecture variant

7.1.8.1 Overview

The MEC in NFV architecture variant introduces two specific functional blocks which replace the MEO and the MEPM in the generic architecture. These specific blocks realize MEC-specific functionality and delegate tasks that can be fulfilled by ETSI NFV MANO to the related MANO functional blocks NFVO and VNFM.

7.1.8.2 MEC application orchestrator

The MEAO has the same responsibilities as the MEO (see clause 7.1.4.1), however, it delegates the management of the set of MEC apps to an NFVO which manages these as part of one or more NFV network services. Similarly, it can delegate the management of the MEC app VNF packages to the NFVO.

7.1.8.3 MEC platform manager - NFV

The MEPM-V has the same responsibilities as the MEPM (see clause 7.1.5.1), however, it does not perform LCM actions itself, but delegates these to a VNFM. Also, it does not receive virtualized resources fault reports and performance measurements directly from the VIM, but these are routed via the VNFM.

7.1.8.4 NFVO

This functional block, as defined by ETSI GS NFV 002 [2], is responsible for managing the lifecycle of the network services that include the MEC app instances, treating each MEC app instance as a VNF instance.

7.1.8.5 VNFM (MEC Platform LCM)

This functional block, as defined by ETSI GS NFV 002 [2], is responsible for managing the lifecycle of the MEC platform using standard NFV LCM procedures.

7.1.8.6 VNFM (MEC App LCM)

This functional block, as defined by ETSI GS NFV 002 [2], is responsible for managing the lifecycle of the MEC apps, treating each MEC app instance as a VNF instance. There may be more than one instance of this block in a deployment. Likewise, this block may be the same as the VNFM (MEC Platform LCM) as defined in clause 7.1.8.5.

7.2 Reference points

7.2.1 Reference points related to the MEC platform

- Mp1: The Mp1 reference point between the MEC platform and the MEC applications provides service registration, service discovery, and communication support for services. It also provides other functionality such as application availability, session state relocation support procedures, traffic rules and DNS rules activation, access to persistent storage and time of day information, etc. This reference point can be used for consuming services and providing service specific functionality.
- Mp2: The Mp2 reference point between the MEC platform and the data plane of the virtualization infrastructure is used to instruct the data plane on how to route traffic among applications, networks, services, etc. This reference point is not further specified.
- Mp3: The Mp3 reference point between MEC platforms is used for control communication between MEC platforms.

NOTE: Optionally control signalling can be exchanged between two MEC platforms in different MEC systems in order to facilitate feature specific inter-MEC system coordination. Such features include application mobility support and V2X support.

7.2.2 Reference points related to the MEC management

- Mm1: The Mm1 reference point between the multi-access edge orchestrator and the OSS is used for triggering the instantiation and the termination of MEC applications in the MEC system.
- Mm2: The Mm2 reference point between the OSS and the MEC platform manager is used for the MEC platform configuration, fault and performance management.
- Mm3: The Mm3 reference point between the multi-access edge orchestrator and the MEC platform manager is used for the management of the application lifecycle, application rules and requirements and keeping track of available MEC services.
- Mm4: The Mm4 reference point between the multi-access edge orchestrator and the virtualization infrastructure manager is used to manage virtualized resources of the MEC host, including keeping track of available resource capacity, and to manage application images.
- Mm5: The Mm5 reference point between the MEC platform manager and the MEC platform is used to perform platform configuration, configuration of the application rules and requirements, application lifecycle support procedures, management of application relocation, etc. This reference point is not further specified.
- Mm6: The Mm6 reference point between the MEC platform manager and the virtualization infrastructure manager is used to manage virtualized resources e.g. to realize the application lifecycle management.
- Mm7: The Mm7 reference point between the virtualization infrastructure manager and the virtualization infrastructure is used to manage the virtualization infrastructure. This reference point is not further specified.
- Mm8: The Mm8 reference point between the user application lifecycle management proxy and the OSS is used to handle device applications requests for running applications in the MEC system. This reference point is not further specified.
- Mm9: The Mm9 reference point between the user application lifecycle management proxy and the multi-access edge orchestrator of the MEC system is used to manage MEC applications requested by device application. This reference point is not further specified.

7.2.3 Reference points related to external entities

- Mx1: The Mx1 reference point between the OSS and the customer facing service portal is used by the third-parties to request the MEC system to run applications in the MEC system. This reference point is not further specified.
- Mx2: The Mx2 reference point between the user application lifecycle management proxy and the device application is used by a device application to request the MEC system to run an application in the MEC system, or to move an application in or out of the MEC system. It is only available when supported by the MEC system.

7.2.4 Reference points related to the MEC in NFV architecture variant

- Mm3*: The Mm3* reference point between MEAO and MEPM-V is based on the Mm3 reference point. Changes cater for the split between MEPM-V and VNFM (MEC applications LCM).
- Mv1: The Mv1 reference point between the MEAO and the NFVO is related to the Os-Ma-nfvo reference point, as defined in ETSI GS NFV-IFA 013 [3].

- Mv2: The Mv2 reference point between the VNF Manager that performs the LCM of the MEC application VNFs and the MEPM-V allows LCM related notifications to be exchanged between these entities. It is related to the Ve-Vnfm-em reference point as defined in ETSI GS NFV-IFA 008 [4].
- Mv3: The Mv3 reference point between the VNF Manager and the MEC application VNF instance allows the exchange of messages e.g. related to MEC application LCM or initial deployment-specific configuration. It is related to the Ve-Vnfm-vnf reference point, as defined in ETSI GS NFV-IFA 008 [4].

The following reference points are used as they are defined by ETSI NFV:

- Ve-Vnfm-em: This reference point connects the VNF Manager (VNFM) that manages the lifecycle of the MEC platform with the MEC Platform Manager - NFV (MEPM-V). It is the Ve-Vnfm-em reference point as defined in ETSI GS NFV-IFA 008 [4].
- Ve-Vnfm-vnf: This reference point connects the VNFM that manages the lifecycle of the MEC platform with the MEC Platform VNF.
- Nf-Vn: This reference point connects the MEC Platform VNF with the NFVI.
- Nf-Vi: This reference point connects the NFVI with the VIM.
- Os-Ma-nfvo: This reference point connects the OSS with the NFVO. It is primarily used to manage Network Services, i.e. a number of VNFs connected and orchestrated to deliver a service.
- Or-Vnfm: This reference point connects the NFVO with the VNFM that manages the lifecycle of the MEC platform. It is primarily used for the NFVO to invoke VNF LCM operations.
- Vi-Vnfm: This reference point connects the VIM with the VNFM that manages the lifecycle of the MEC platform. It is primarily used by the VNFM to invoke resource management operations to manage the resources that are needed by the VNF.
- Or-Vi: This reference point connects the NFVO with the VIM. It is primarily used by the NFVO to manage resources capacity.

8 MEC services

8.1 General

A MEC service is a service provided and consumed either by the MEC platform or a MEC application. When provided by an application, it can be registered in the list of services to the MEC platform over the Mp1 reference point (see clause 7.2.1).

A MEC application can subscribe to a service for which it is authorized over the Mp1 reference point.

A certain number of MEC services are necessary in order to fulfil the requirements defined in ETSI GS MEC 002 [1] and are described in clauses 8.2 to 8.4.

8.2 Radio Network Information

The Radio Network Information service, when available, provides authorized applications with radio network related information.

It exposes information to applications, such as:

- appropriate up-to-date radio network information regarding radio network conditions;
- measurement and statistics information related to the user plane;

- information (e.g. UE context and radio access bearers) related to UEs served by the radio node(s) associated with the MEC host;
- changes on information related to UEs served by the radio node(s) associated with the MEC host.

The radio network information is provided at the relevant granularity (e.g. per User Equipment (UE) or per cell, per period of time).

8.3 Location

The Location service, when available, provides authorized applications with location-related information.

It exposes information to applications, such as:

- the location of specific UEs currently served by the radio node(s) associated with the MEC host;
- information about the location of all UEs currently served by the radio node(s) associated with the MEC host;
- optionally, information about the location of a certain category of UEs currently served by the radio node(s) associated with the MEC host;
- a list of UEs in a particular location;
- information about the location of all radio nodes currently associated with the MEC host.

NOTE: Location can be geolocation, Cell ID, etc.

8.4 Traffic Management Services

The following optional Traffic Management services are supported:

- BandWidth Management (BWM) service.
- Multi-access Traffic Steering (MTS) service.

The BandWidth Management (BWM) service, when available, allows allocation of bandwidth to certain traffic routed to and from MEC applications and the prioritization of certain traffic.

The Multi-access Traffic Steering (MTS) service, when available, allows seamlessly steering/splitting/duplicating application data traffic across multiple access network connections.

9 Inter-MEC system communication

Inter-MEC system communication addresses the following high level requirements:

- 1) A MEC platform should be able to discover other MEC platforms that may belong to different MEC systems.
- 2) A MEC platform should be able to exchange information in a secure manner with other MEC platforms that may belong to different MEC systems.
- 3) A MEC application should be able to exchange information in a secure manner with other MEC applications that may belong to different MEC systems.

To enable the inter-MEC system communication, the following hierarchical inter-MEC system discovery and communication framework is assumed:

- MEC system level inter-system discovery and communication.
- MEC host level inter-system communication between the MEC platforms.

NOTE: It is for further study if MEC platforms in different MEC systems should be able to discover each other without the involvement of the MEC system level functional elements.

Annex A (informative): Key concepts

A.1 MEC host selection

In order to run a MEC application in the MEC system, the multi-access edge orchestrator receives requests triggered by the OSS, a third-party, or a device application.

These requests provide information about the application to run, and possibly other information, such as the location where the application needs to be active, other application rules and requirements, as well as the location of the application image if it is not yet on-boarded in the MEC system.

The information considered by the multi-access edge orchestrator when selecting a MEC host(s) for a MEC application can include:

- deployment model of the application (e.g. whether it is one instance per user, one instance per host, one instance on each host, etc.);
- required virtualized resources (compute, storage, network resources, including specific hardware support);
- latency requirements (e.g. how strict the latency constraints are, latency fairness between users);
- requirements on location;
- required MEC services that are needed for the MEC application to be able to run;
- MEC services that the MEC application can take advantage of if available;
- connectivity or mobility requirements (e.g. application state relocation, application instance relocation);
- required MEC features, such as VM relocation support or UE identity;
- required network connectivity (e.g. connectivity to applications within the MEC system, connectivity to local networks, or to the Internet);
- information on the operator's MEC system deployment or mobile network deployment (e.g. topology, cost);
- requirements on access to user traffic;
- requirements on persistent storage.

The multi-access edge orchestrator considers the requirements and information listed above and information on the resources currently available in the MEC system to select one or several MEC hosts within the MEC system, and requests the selected host(s) to instantiate the application.

NOTE: The actual algorithm used to select the hosts depends on the implementation, configuration, and operator deployment and is not intended to be specified.

Under certain circumstances (e.g. UE mobility events resulting in increased latency, load balancing decisions), and if supported, the multi-access edge orchestrator could decide to select a new host and initiate the transfer of an application instance or application-related state information from a source host to a target host, as described in clause A.4 "Support of application and UE mobility".

A.2 DNS support

The MEC platform provides access to DNS, which includes a name server and a proxy/cache function. The MEC platform receives the application DNS rules from the MEC management. Based on configuration or following an activation request from the MEC application, the MEC platform configures the mapping between an IP address and its FQDN into the DNS based on these rules.

A.3 Application traffic filtering and routing

The MEC platform allows the activation, update, and deactivation of the MEC application traffic rules, and applies these rules to the underlying data plane. This allows IP traffic routing or tapping to the MEC applications or to locally accessible networks (e.g. enterprise network, Internet access, etc.) according to the configured traffic rules.

Within the constraints set by the MEC management, an authorized MEC application can request the activation, update and deactivation of the MEC application traffic rules dynamically. For example, this allows the redirection of traffic of a certain UE to an enterprise network after the UE has been authenticated and authorized by the MEC application.

A.4 Support of application and UE mobility

A.4.1 Background: UE mobility

UE mobility is supported by the underlying 3GPP networks and can result in the UE moving to a radio node associated with a different MEC host due to handover events. The likelihood for MEC host changes due to UE mobility depends on the deployment options and the topology of the network.

MEC applications can be impacted by UE mobility events and the possible scenarios are described in clause A.4.2.

A.4.2 MEC application scenarios for UE mobility

A.4.2.1 MEC applications not sensitive to UE mobility

Some MEC applications remain unaffected by UE mobility events, for example if their purpose is to only process traffic that goes through the local MEC host, or if the UE related state can be rapidly rebuilt after a handover.

A.4.2.2 MEC applications sensitive to UE mobility

A.4.2.2.1 Maintaining connectivity between UE and MEC application instance

Some MEC applications expect to continue serving the UE after a location change of the UE in the mobile network. In order to provide continuity of the service, the connectivity between the device application and the MEC application needs to be maintained. As the UE moves around the network, the traffic between the UE and the MEC application is routed so that it reaches the intended destination.

As the UE moves further away from the location of the MEC application, there could be an increased latency between the UE and the MEC application. Due to this reason or others (e.g. network congestion), for some MEC applications, it might become necessary to relocate the application state or application instance in order to satisfy the latency requirements.

A.4.2.2.2 Application state relocation

Application state relocation requires MEC application support. As the UE moves around in the mobile network, another (target) MEC host might be identified as being a more appropriate location for the MEC application to serve the UE instead. If the application does not run yet on the target MEC host, the application can be instantiated before starting the application state relocation. Interaction between the two application instances on the source and target MEC hosts is enabled, e.g. to move the application state from one instance to the other. The exact data, the procedures, and the pacing of this interaction are dependent on the application design. At some point in time, the application state will have been fully transferred, the source instance will stop serving the UE, and the target instance will take over serving the UE.

A.4.2.2.3 Application instance relocation within the MEC system

In some cases, and when it is supported, it is preferable for some applications to move the serving application instance itself from the source host to the target host. This is done by having the multi-access edge orchestrator trigger the application instance relocation. In that case, the application instance runs in the source host, while it is copied over to the target host. When the process is complete, the instance then runs in the target host and takes over the traffic from the UE. This concept is exemplified by the "VM handoff" described in "Adaptive VM Handoff Across Cloudlets" [i.3].

A.4.2.2.4 Application instance relocation between the MEC system and an external cloud environment

In some cases, and when it is supported, the UE can request the MEC system to move application instances out of the MEC system to an external cloud environment, or from an external cloud environment to the MEC system. In that case, the application instance relocation is triggered between the MEC system and the external cloud environment under the supervision of the multi-access edge orchestrator.

A.5 Void

A.6 Data Plane

In the MEC reference architecture the Virtualization infrastructure in the MEC host includes a data plane that executes the traffic rules received by the MEC platform, and routes the traffic among applications, services, DNS server/proxy, 3GPP network, local networks and external networks (as described in clause 7.1.1). The Data plane described in the present document is only a representation of the execution environment for the traffic rules and routing. Mapping of all or part of MEC data plane functionality to any functional element(s) of a real network architecture, e.g. UPF in ETSI TS 123 501 [i.6], implies a specific deployment option of MEC in such a network architecture.

A.7 API gateway support

The MEC platform may be accompanied by and API gateway functionality supporting e.g., the following:

- Receiving the service API request from MEC application via a stable service connection endpoint. The IP address change of MEC service instances will not lead to the update of the service connection endpoint information.
- Load balancing for the multiple backend MEC service instances.
- Throttling API requests for better throughput based on the configuration.
- Monitoring the API requests. It can be used for statistics and charging.

History

Document history		
V1.1.1	March 2016	Publication
V2.1.1	January 2019	Publication
V2.1.2	June 2019	Implements document MEC(19)000206
V2.1.3	September 2019	Implements document MEC(19)000330r1
V2.1.4	May 2020	Implements documents MEC(20)000029, MEC(20)000132r2
V2.1.5	July 2020	Clean-up done by <i>editHelp!</i> E-mail: mailto:edithelp@etsi.org

Draft