



Mobile Edge Computing (MEC); Framework and Reference Architecture

Disclaimer

The present document has been produced and approved by the Mobile 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.

Reference

DGS/MEC-003Arch

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 only prevailing document is the print of the Portable Document Format (PDF) version kept on a specific network drive within ETSI Secretariat.

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/CommiteeSupportStaff.aspx>

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI.
The copyright and the foregoing restriction extend to reproduction in all media.

© European Telecommunications Standards Institute 2016.
All rights reserved.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are Trade Marks of ETSI registered for the benefit of its Members.
3GPP™ and **LTE™** are Trade Marks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.
GSM® and the GSM logo are Trade Marks 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 Definitions and abbreviations.....	6
3.1 Definitions.....	6
3.2 Abbreviations	7
4 Overview	7
5 Mobile Edge Computing framework.....	8
6 Reference architecture.....	9
7 Functional elements and reference points	10
7.1 Functional elements.....	10
7.1.1 Mobile edge host.....	10
7.1.2 Mobile edge platform.....	10
7.1.3 Mobile edge application.....	10
7.1.4 Mobile edge system level management	11
7.1.4.1 Mobile edge orchestrator.....	11
7.1.4.2 Operations Support System (OSS).....	11
7.1.4.3 User application lifecycle management proxy	11
7.1.5 Mobile edge host level management.....	11
7.1.5.1 Mobile edge platform manager	11
7.1.5.2 Virtualisation infrastructure manager.....	12
7.1.6 User equipment application	12
7.1.7 Customer facing service portal	12
7.2 Reference points	12
7.2.1 Reference points related to the mobile edge platform.....	12
7.2.2 Reference points related to the mobile edge management	12
7.2.3 Reference points related to external entities	13
8 Mobile edge services	13
8.1 General	13
8.2 Radio Network Information	13
8.3 Location.....	14
8.4 Bandwidth Manager	14
Annex A (informative): Key concepts.....	15
A.1 Mobile edge host selection	15
A.2 DNS support.....	15
A.3 Application traffic filtering and routing	16
A.4 Support of application and UE mobility.....	16
A.4.1 Background: UE mobility	16
A.4.2 Mobile edge application scenarios for UE mobility	16
A.4.2.1 Mobile edge applications not sensitive to UE mobility	16
A.4.2.2 Mobile edge applications sensitive to UE mobility	16
A.4.2.2.1 Maintaining connectivity between UE and mobile edge application instance	16
A.4.2.2.2 Application state relocation.....	16
A.4.2.2.3 Application instance relocation within the mobile edge system	17

A.4.2.2.4	Application instance relocation between the mobile edge system and an external cloud environment	17
A.5	Integrating MEC and NFV in the same network deployment	17
History	18

Intellectual Property Rights

IPRs essential or potentially essential to the present document 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.

Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification Group (ISG) Mobile 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 Mobile Edge Computing that describes a mobile edge system that enables mobile edge applications to run efficiently and seamlessly in a mobile network. The present document also describes the functional elements and the reference points between them, and a number of mobile edge services that comprise the solution. It finally presents a number of key concepts related to the mobile 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 <http://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: "Mobile Edge Computing (MEC); Technical Requirements".

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: "Mobile-Edge Computing (MEC); Terminology".

[i.2] ETSI GS NFV 002: "Network Functions Virtualisation (NFV); Architectural Framework".

[i.3] "OpenStack++ for Cloudlet Deployment".

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

[i.4] "Adaptive VM Handoff Across Cloudlets".

NOTE: Available at <http://reports-archive.adm.cs.cmu.edu/anon/2015/CMU-CS-15-113.pdf>.

[i.5] ETSI GS MEC 017: "Mobile-Edge Computing (MEC); Deployment of Mobile Edge Computing in an NFV environment".

3 Definitions and abbreviations

3.1 Definitions

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

3.2 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
NFVI-PoP	Network Functions Virtualisation Infrastructure Point of Presence
OSS	Operations Support System

4 Overview

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

The framework described in clause 5 shows the structure of the Mobile Edge Computing environment.

The reference architecture described in clause 6 shows the functional elements that compose the mobile edge system, including the mobile edge platform and the mobile edge 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 mobile edge services, comprising the solution for Mobile 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 Mobile Edge Computing framework

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

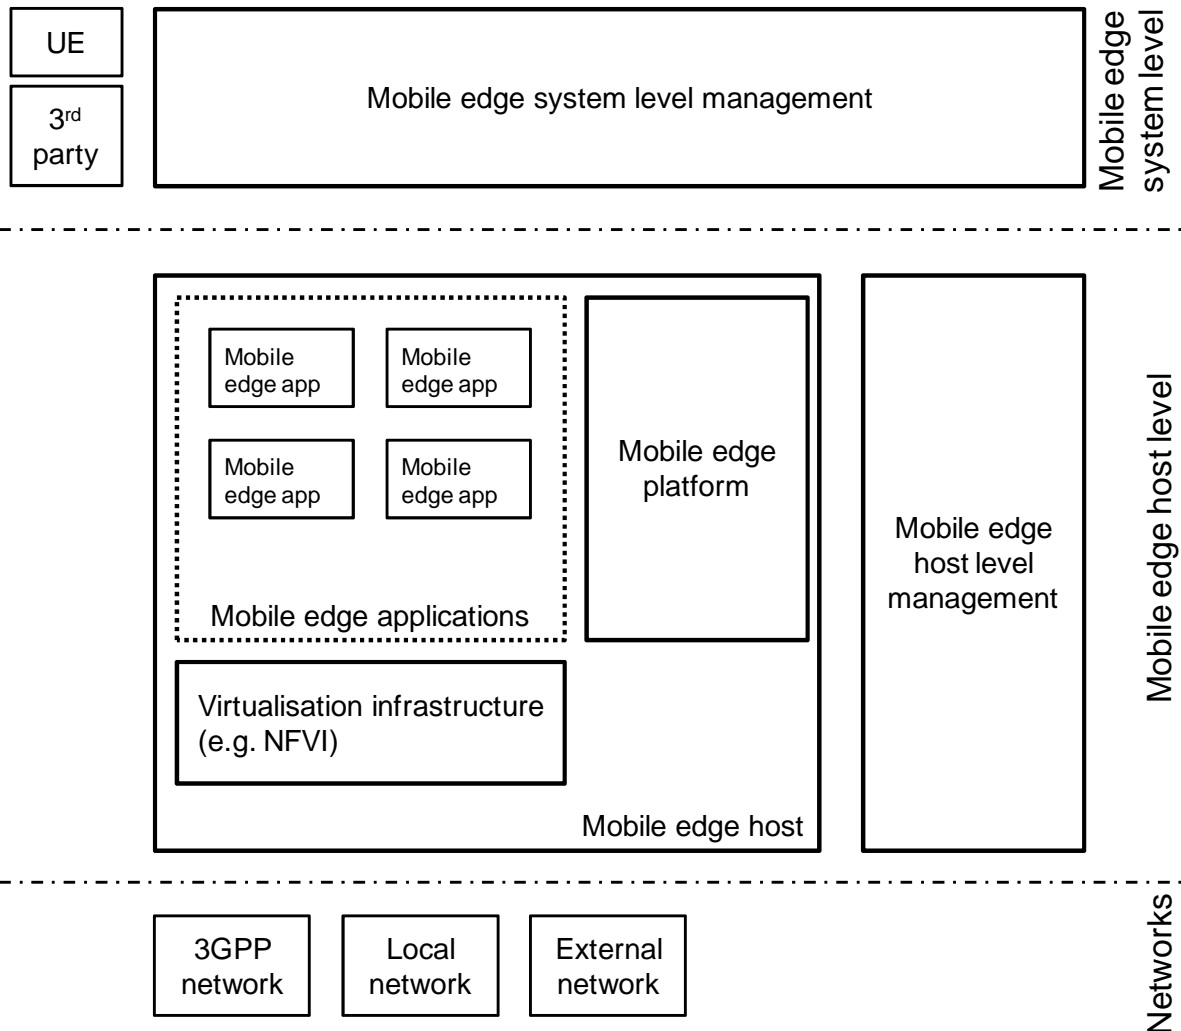


Figure 5-1: Mobile Edge Computing framework

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

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

6 Reference architecture

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

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

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

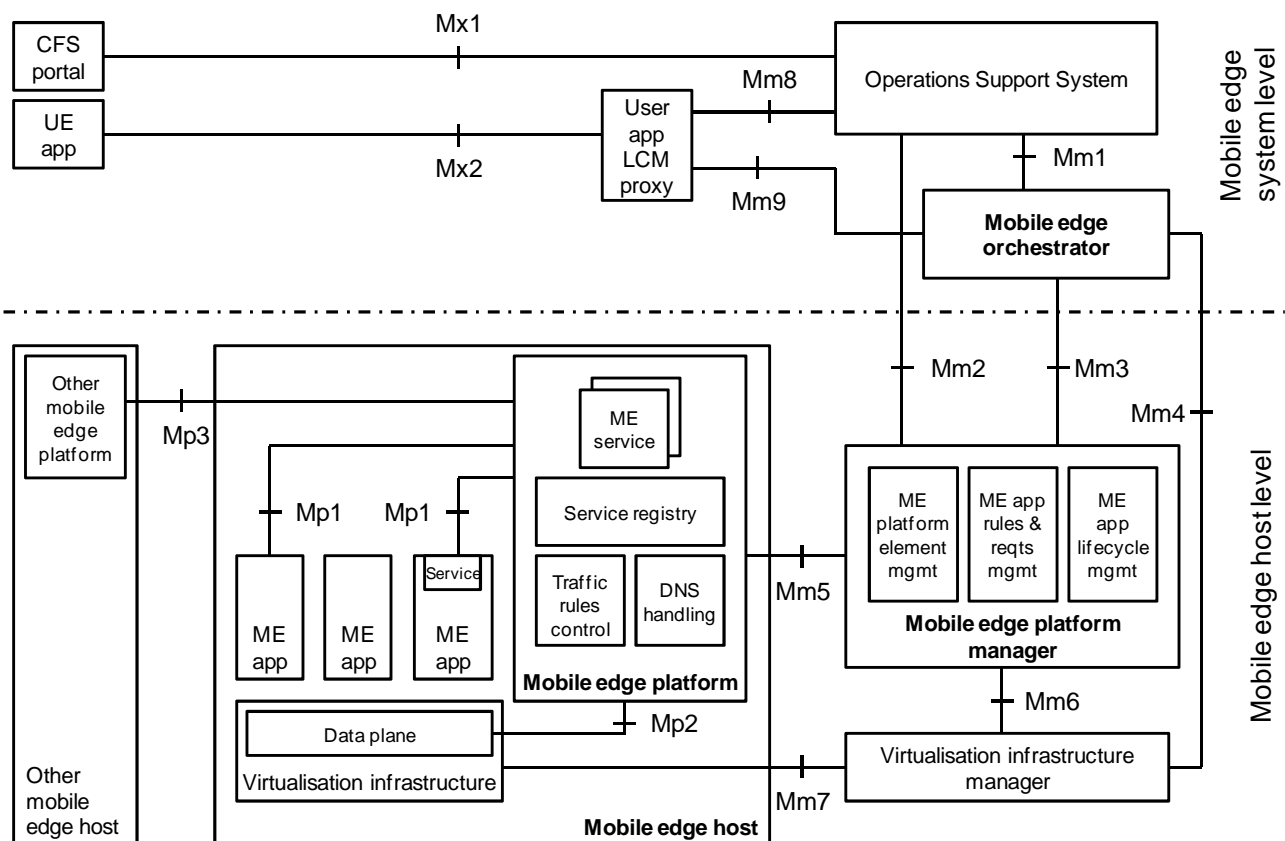


Figure 6-1: Mobile edge system reference architecture

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

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

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

Mobile edge applications are instantiated on the virtualisation infrastructure of the mobile edge host based on configuration or requests validated by the mobile edge management. Mobile edge applications are further described in clause 7.1.3.

The mobile edge management comprises the mobile edge system level management and the mobile edge host level management.

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

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

7 Functional elements and reference points

7.1 Functional elements

7.1.1 Mobile edge host

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

7.1.2 Mobile edge platform

The mobile edge platform is responsible for the following functions:

- offering an environment where the mobile edge applications can discover, advertise, consume and offer mobile edge services (see clause 8), including, when supported, mobile edge services available via other platforms;
- receiving traffic rules from the mobile edge platform manager, applications, or services, and instructing the data plane accordingly. When supported, this includes the translation of tokens representing UEs in the traffic rules into specific IP addresses;
- receiving DNS records from the mobile edge platform manager and configuring a DNS proxy/server accordingly;
- hosting mobile edge services, possibly including services that are described in clause 8;
- providing access to persistent storage and time of day information.

7.1.3 Mobile edge application

Mobile edge applications are running as virtual machines (VM) on top of the virtualisation infrastructure provided by the mobile edge host, and can interact with the mobile edge platform to consume and provide mobile edge services (described in clause 8).

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

Mobile edge 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 mobile edge system level management, and can be assigned to default values if missing.

7.1.4 Mobile edge system level management

7.1.4.1 Mobile edge orchestrator

The mobile edge orchestrator is the core functionality in mobile edge system level management.

The mobile edge orchestrator is responsible for the following functions:

- maintaining an overall view of the mobile edge system based on deployed mobile edge hosts, available resources, available mobile edge 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 virtualisation infrastructure manager(s) to handle the applications;
- selecting appropriate mobile edge 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 the figure 6-1 refers to the OSS of an operator. It receives requests via the CFS portal and from UE applications for instantiation or termination of applications, and decides on the granting of these requests. Granted requests are forwarded to the mobile edge orchestrator for further processing.

When supported, the OSS also receives requests from UE applications for relocating applications between external clouds and the mobile edge system.

7.1.4.3 User application lifecycle management proxy

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

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

The user application lifecycle management proxy authorizes requests from UE applications in the UE and interacts with the OSS and the mobile edge orchestrator for further processing of these requests.

The user application lifecycle management proxy is only accessible from within the mobile network. It is only available when supported by the mobile edge system.

7.1.5 Mobile edge host level management

7.1.5.1 Mobile edge platform manager

The mobile edge platform manager is responsible for the following functions:

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

The mobile edge platform manager also receives virtualised resources fault reports and performance measurements from the virtualisation infrastructure manager for further processing.

7.1.5.2 Virtualisation infrastructure manager

The virtualisation infrastructure manager is responsible for the following functions:

- allocating, managing and releasing virtualised (compute, storage and networking) resources of the virtualisation infrastructure;
- preparing the virtualisation 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 virtualised resources;
- when supported, performing application relocation. For application relocation from/to external cloud environments, the virtualisation 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.4], possibly through a proxy.

The functionality provided by the virtualisation infrastructure manager in the present specification and the functionality provided by the virtualised infrastructure manager described in ETSI GS NFV 002 [i.2], clause 7.2.5, overlap to a large extent.

7.1.6 User equipment application

UE applications as defined in the present specification are applications in the UE that have the capability to interact with the mobile edge 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 mobile edge applications that meet their particular needs, and to receive back service level information from the provisioned applications.

7.2 Reference points

7.2.1 Reference points related to the mobile edge platform

- Mp1: The Mp1 reference point between the mobile edge platform and the mobile edge 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 and providing service specific [i.3] functionality.
- Mp2: The Mp2 reference point between the mobile edge platform and the data plane of the virtualisation 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 mobile edge platforms is used for control communication between mobile edge platforms.

7.2.2 Reference points related to the mobile edge management

- Mm1: The Mm1 reference point between the mobile edge orchestrator and the OSS is used for triggering the instantiation and the termination of mobile edge applications in the mobile edge system.
- Mm2: The Mm2 reference point between the OSS and the mobile edge platform manager is used for the mobile edge platform configuration, fault and performance management.
- Mm3: The Mm3 reference point between the mobile edge orchestrator and the mobile edge platform manager is used for the management of the application lifecycle, application rules and requirements and keeping track of available mobile edge services.
- Mm4: The Mm4 reference point between the mobile edge orchestrator and the virtualisation infrastructure manager is used to manage virtualised resources of the mobile edge host, including keeping track of available resource capacity, and to manage application images.
- Mm5: The Mm5 reference point between the mobile edge platform manager and the mobile edge 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 mobile edge platform manager and the virtualisation infrastructure manager is used to manage virtualised resources e.g. to realize the application lifecycle management.
- Mm7: The Mm7 reference point between the virtualisation infrastructure manager and the virtualisation infrastructure is used to manage the virtualisation 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 UE applications requests for running applications in the mobile edge system. This reference point is not further specified.
- Mm9: The Mm9 reference point between the user application lifecycle management proxy and the mobile edge orchestrator of the mobile edge system is used to manage mobile edge applications requested by UE 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 mobile edge system to run applications in the mobile edge system. This reference point is not further specified.
- Mx2: The Mx2 reference point between the user application lifecycle management proxy and the UE application is used by a UE application to request the mobile edge system to run an application in the mobile edge system, or to move an application in or out of the mobile edge system. This reference point is only accessible within the mobile network. It is only available when supported by the mobile edge system.

8 Mobile edge services

8.1 General

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

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

A certain number of mobile edge 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 mobile edge host;
- changes on information related to UEs served by the radio node(s) associated with the mobile edge 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 mobile edge host;
- information about the location of all UEs currently served by the radio node(s) associated with the mobile edge host;
- optionally, information about the location of a certain category of UEs currently served by the radio node(s) associated with the mobile edge host;
- a list of UEs in a particular location;
- information about the location of all radio nodes currently associated with the mobile edge host.

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

8.4 Bandwidth Manager

The Bandwidth Manager service, when available, allows allocation of bandwidth to certain traffic routed to and from mobile edge applications and the prioritization of certain traffic.

Annex A (informative): Key concepts

A.1 Mobile edge host selection

In order to run a mobile edge application in the mobile edge system, the mobile edge orchestrator receives requests triggered by the OSS, a third-party, or a UE 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 mobile edge system.

The information considered by the mobile edge orchestrator when selecting a mobile edge host(s) for a mobile edge 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 virtualised 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 mobile edge services that are needed for the mobile edge application to be able to run;
- mobile edge services that the mobile edge application can take advantage of if available;
- connectivity or mobility requirements (e.g. application state relocation, application instance relocation);
- required mobile edge features, such as VM relocation support or UE identity;
- required network connectivity (e.g. connectivity to applications within the mobile edge system, connectivity to local networks, or to the Internet);
- information on the operator's mobile edge system deployment or mobile network deployment (e.g. topology, cost);
- requirements on access to user traffic;
- requirements on persistent storage.

The mobile edge orchestrator considers the requirements and information listed above and information on the resources currently available in the mobile edge system to select one or several mobile edge hosts within the mobile edge 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 mobile 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 mobile edge platform provides access to DNS, which includes a name server and a proxy/cache function. The mobile edge platform receives the application DNS rules from the mobile edge management. Based on configuration or following an activation request from the mobile edge application, the mobile edge 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 mobile edge platform allows the activation, update, and deactivation of the mobile edge application traffic rules, and applies these rules to the underlying data plane. This allows IP traffic routing or tapping to the mobile edge 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 mobile edge management, an authorized mobile edge application can request the activation, update and deactivation of the mobile edge 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 mobile edge 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 mobile edge host due to handover events. The likelihood for mobile edge host changes due to UE mobility depends on the deployment options and the topology of the network.

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

A.4.2 Mobile edge application scenarios for UE mobility

A.4.2.1 Mobile edge applications not sensitive to UE mobility

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

A.4.2.2 Mobile edge applications sensitive to UE mobility

A.4.2.2.1 Maintaining connectivity between UE and mobile edge application instance

Some mobile edge 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 UE application and the mobile edge application needs to be maintained. As the UE moves around the network, the traffic between the UE and the mobile edge application is routed so that it reaches the intended destination.

As the UE moves further away from the location of the mobile edge application, there could be an increased latency between the UE and the mobile edge application. Due to this reason or others (e.g. network congestion), for some mobile edge 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 mobile edge application support. As the UE moves around in the mobile network, another (target) mobile edge host might be identified as being a more appropriate location for the mobile edge application to serve the UE instead. If the application does not run yet on the target mobile edge host, the application can be instantiated before starting the application state relocation. Interaction between the two application instances on the source and target mobile edge 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 mobile edge 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 mobile 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.4].

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

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

A.5 Integrating MEC and NFV in the same network deployment

Mobile Edge Computing (MEC) and Network Functions Virtualisation (NFV) are complementary concepts that can exist independently. The mobile edge architecture has been designed in such a way that a number of different deployment options of mobile edge systems are possible. With respect to NFV, a mobile edge system can be realized independently from the presence of an NFV environment in the same network, or can be coexisting with it. As both MEC and NFV are based on the use of virtualisation technology, MEC mobile edge applications and NFV virtualised network functions can be instantiated partly or entirely over the same virtualisation infrastructure.

The reference architecture described in clause 6 reuses the concept of a virtualisation infrastructure manager similar, with some enhancements, to the Virtualised Infrastructure Manager of the NFV framework, as well as the concept of virtualisation infrastructure, which corresponds roughly to an NFVI-PoP as described in ETSI GS NFV 002 [i.2].

The reference architecture has been designed in such a way that further synergy between MEC and NFV can be achieved.

Multiple scenarios for deployments are possible, depending on operators' preferences for their networks and their migration strategy, e.g. fully virtualised environments or mixed environments, with MEC being deployed first, or NFV being deployed first, with different levels of integration between the two technologies, taking into account secondary aspects such as multi-tenancy, etc. The way MEC and NFV management and orchestration components relate to each other (e.g. integration, interworking, coexistence) is an important aspect of integrated MEC-NFV deployments [i.4].

A dedicated Group Specification, ETSI GS MEC 017 [i.5], describes the topic in more detail.

History

Document history		
V1.1.1	March 2016	Publication