ETSI White Paper No. 30

MEC in an Enterprise Setting: A Solution Outline

First edition – September 2018

ISBN No. 979-10-92620-25-2

Authors:

Alex Reznik, Editor (HPE), Anthony Sulistio (Bosch), Alexander Artemenko (Bosch), Yonggang Fang (ZTE), Danny Frydman (Saguna), Fabio Giust (Athonet), HuaZhang Lv (China Unicom), Saad Ullah Sheikh (STC), Yifan Yu (Intel), Zhou Zheng (Huawei)

ETSI
06921 Sophia Antipolis CEDEX, France
Tel +33 4 92 94 42 00
info@etsi.org
www.etsi.org

Contents

Introduction

Multi-access Edge Computing (MEC) complements the corporate data centre by providing compute, storage, networking and data analytics at locations closer to the data source (e.g. Internet of Things (IoT) devices, , workers, operators, etc.) and points of consumption [1]. For example, MEC solutions enable enterprises to manage their security and compliance requirements effectively, since data analysis can be performed in a more deterministic manner and the data are kept within a specific premises or political region compared to the centralized data centre. Additionally, enterprise MEC solutions are closely “integrated” with access network(s) to provide a fully-converged enterprise access and compute environment. Such environments often include enterprise WiFi. However, they may also include mobile access, especially when large campus and outdoor coverage is required. An example is a “private LTE” network which can use licensed spectrum, unlicensed spectrum or new technologies such as Citizens Broadband Radio Service (CBRS).

Having a well-defined and structured set of functionalities, a MEC-enabled enterprise infrastructure supports both scaling options: (i) Small scale application scenarios in small enterprises benefit from easy deployment, setup and scalability advantages of edge computing solutions; (ii) middle to large scale scenarios will boost the enterprise size by using flexible and full-automated management and orchestration functions.

The environment of a MEC deployment can differ for each location depending on the use cases and the digital services that are offered to end users, which are sensitive to network latency and require a high level of performance. However, most of the existing solutions are done on a rather ad-hoc and proprietary basis for a specific environment, with minimal adherence to standards and/or interoperability. To avoid building a MEC solution from scratch for each location, the ETSI work on MEC aims to address this problem.

Due to historical reasons, many companies utilize a large set of heterogeneous technologies in different domains, including communication networks and data processing. By introducing MEC, the unification of the applied interfaces will enable a new level of interoperability for various components from different vendors and facilitate a natural reuse of underutilized capacity elements (like network, storage, processing, etc).

The purpose of this white paper is to give a solution overview of MEC deployments in the enterprise environment. Firstly, this paper presents several use cases and MEC deployment options. It then highlights key challenges when trying to deploy these use cases in an existing enterprise infrastructure. In addition, it demonstrates how the ETSI MEC APIs help to overcome these challenges.

Use Cases and Deployment Examples

Enterprise Solution Use Cases

Use Case A: Smart Enterprise Building

Enterprises already invest a lot in smart buildings since they are increasingly valuable for a company's facility management and its employees. Some of the advantages gained from smart buildings are:

- Reduced energy and utility costs and thus lowering the carbon footprint,
- Improved building operations from sensors data collected for predictive maintenance,
- Increased occupant comfort and productivity by having personalization of the working environment, such as room ambience, temperature and lightning,
- More efficient use of space by providing live data of available desks and meeting rooms,
- New level of working experience and collaboration opportunities.

A smart building contains many IoT devices and sensors, such as motion, light, temperature, humidity, infrared, video, etc. A typical application is to continuously monitor energy consumption, room occupancy, parking spaces, temperature, coffee machines, etc. The information is shown in a central dashboard for facility management. The building's occupants can access relevant information via a smartphone app for personalized workspaces, finding empty desks, booking a meeting room at a short notice, etc.

As these sensors generate gigabytes of data (depending on the size of the building). The facility management deals with an operational challenge in capturing, processing and storing data (also for historical data). A local processing in the edge (e.g. per floor) rather than a centralized data centre reduces network latency and analyses the data more quickly. Thus, MEC provides a better user experience.

Another challenge is that the IoT devices and sensors are heterogeneous, coming from various manufacturers. The devices need to communicate through standardized protocols and the smartphone app needs to interact with the local edge or IoT gateway via open Application Programming Interfaces (APIs). The ETSI work on MEC aims to address this challenge.

Use Case B: Data Analysis and Security

In regulated industries, such as finance and healthcare, data need to be stored and analysed on-premises in order to comply with local regulations. For example, financial institution branches can find non-compliant transactions in real-time and stop them more quickly, compared to sending the data to a central data centre [2]. In addition, the local branches can provide online digital services to their customers by transforming ATMs with a video capability into interactive tellers [3], [4]. Using MEC, enterprises with sensitive digital assets can protect the security and integrity of their data by providing real-time security monitoring for traffic anomalies [5].

Use Case C: Augmented Reality Conferencing

Telephone conferences represent a vital communication means in every enterprise giving a high level of flexibility and location independence to workers. A transition from pure voice to video plus voice conferences increases a feeling of close human presence which remains an important factor of successful collaboration within working teams. The next level can be reached using augmented reality (AR). With AR, an immersive user experience helps to reach a feeling of real presence. However, higher requirements on latency and image quality present a showstopper for today's mobile devices. Using offloading of image processing into the edge, almost unlimited possibilities arise.

Use Case D: Location-restricted BYOD access

A large enterprise with a large mixed indoor/outdoor campus (e.g. an automobile test and assembly facility) is providing Bring Your Own Device (BYOD) access to numerous enterprise applications to its employees. The access is available only when employees are on-site, but in such cases it is automatic. The employees' devices and mobile identities are mapped to appropriate enterprise identities allowing for proper application of enterprise access policies. This happens over an indoor WiFi network as well as an outdoor LTE-based access. Enterprise traffic does not leave the enterprise premises. Non-employees are able to use LTE-based data access, however this happens over a separate "network slice" – none of the enterprise network assets are visible to non-employees.

Use Case E: Streaming media and entertainment in Enterprise

An end-to-end streaming media solution suite is deployed locally within each enterprise location, which could mean a building or a portion of a building. This solution supports 4K/8K video playback, in-campus video conferencing, and a number of applications which include AR/VR. Minimizing the amount of streaming media that has to leave the boundaries of the location improves the overall experience, while reducing the traffic on the enterprise WAN – thus reducing both costs to the enterprise and to the carrier.

An Example of an Enterprise MEC Deployment

Enterprise ICT services target a clear and usually well-controlled set of subscribers, i.e., staff members and other authorized collaborators. In addition, they are inherently localized within the premises of the enterprise and/or where the core operations are carried out, and are designed for very specific purposes, sometimes implemented by tailor made solutions. For these reasons, MEC appears as a natural partner technology to provide edge computing and communication infrastructure to enterprises.

In this section, we highlight what such an implementation may look like using an example that captures much of the complexity associated with an enterprise deployment, keeping in mind that many enterprise deployments may actually be simpler than this example. Consider the overall system shown in Figure 1 and let's take what happens at the AR/VR terminal as an example. The VR terminal uses the 5G NR air interface (gNB) to access local application content on the MEC edge business platform. Video transcoding processing and cloud game graphics calculation and rendering are all performed at the edge site, avoiding the need to upload the business flow to the centralized cloud in the Internet. Because the MEC edge business platform is an extension of the cloud platform in the internet, it does not require customized development of apps, but rather can run application components of well-designed apps "as-is.[1]" This enables rapid deployment and iteration of applications.

The MEC management platform is deployed in a local or regional data centre, which enables the coordination and management of MEC business platforms across the enterprise.

[1] Please see [6] for some discussion of what a well-designed application is

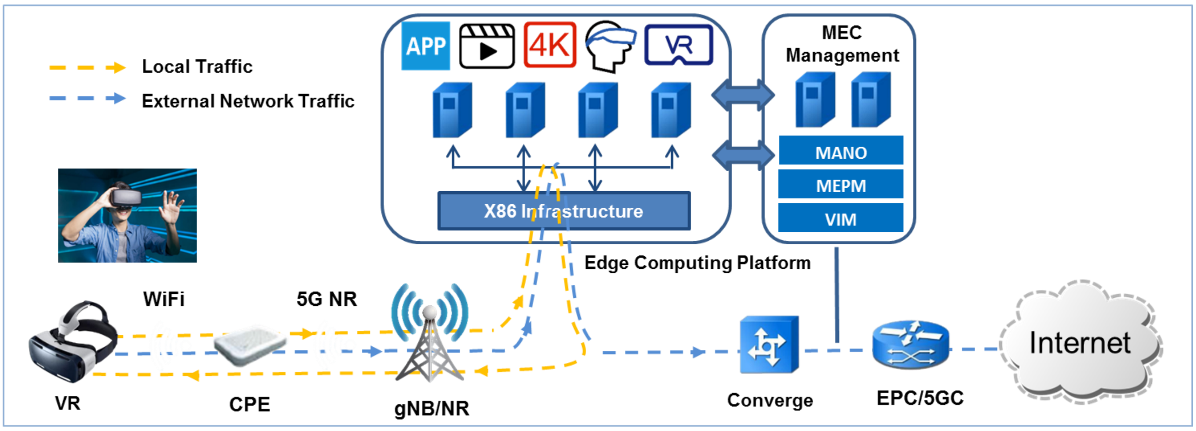

Figure 1: Illustration of an in-building streaming media system

Focusing now on each local MEC site, we note that it is a miniaturization of a full data centre; a detailed diagram is shown in Figure 2. A typical scale involves 10 to 20 enterprise-grade x86-based compute nodes with built-in storage. These are used for general computing, as well as network functions (thus distinguishing them from a traditional enterprise-owned cloud). A dedicated cluster, e.g. for storage, video transcoding or AI can be added within the framework as necessary.

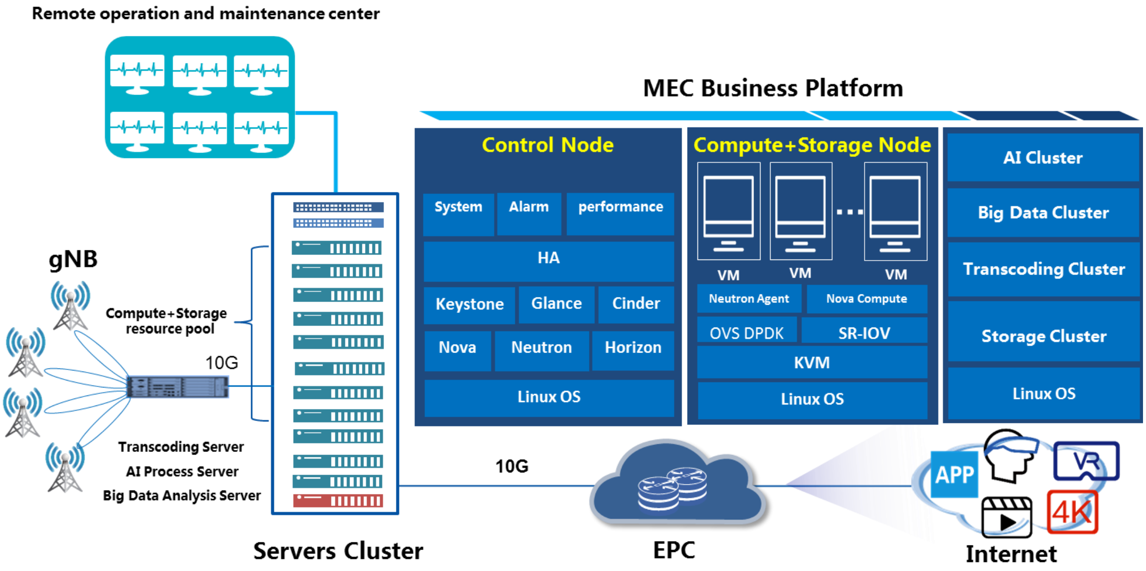

Figure 2: Detailed diagram of a MEC site

Service Enablement Challenges in MEC

It is widely accepted that the enterprise market, as enabled by MEC, is going to be a major component of the emerging 5G mobile ecosystem. Moreover, many in our industry hold the opinion that enterprise will be the lead market for 5G. However, integrating mobile networks and enterprise networks continues to involve a number of challenges. Some of the key ones, especially as related to MEC, are listed below.

Supporting Enterprise-Grade MEC Applications

Enterprise IT application development has increasingly been shifting towards a micro-services based application architecture that heavily utilizes containers. Such an approach holds additional benefits for MEC applications, as pointed out in our earlier white paper on software development for MEC [6]. However, existing telecommunication network architectures, including MEC specifications, often assume (implicitly or explicitly) that a virtual machine-based NFV management environment is in place. For true enterprise applications, container-based services must be supported in MEC.

Unity of experience across all networks (Fixed /3G /4G /5G / Wi-Fi)

MEC will deliver services to many enterprises that require applications to be delivered in a manner agnostic of network access type. Having said this, each wireless network type has its unique delay/latency characteristics. One of the main challenges for enterprises will be to ensure the same experience across different access technologies.

A related challenge concerns enterprises that operate a large number of remote locations and/or mobile employees. Here, a unity of experience must also be provided over a large number of distinct types of public internet accesses, e.g. public mobile access, public WiFi, etc.

Integration of Access Control

In order to properly support enterprise applications, operator-offered MEC solutions must be able to "understand" the world of enterprise identity and access management. Unfortunately, enterprise identity and access management is based on completely different technologies to those used in mobile networks. Mobile networks utilize 3GPP-defined SIM-based approaches, while enterprise networks are built around systems such as LDAP, IdAM, etc.

A key challenge will be to devise techniques for identity and access management that are able to "connect" an operator's, in particular a mobile operator's subscriber management systems and an enterprise access and identity management system in a way that is acceptable to both – i.e. taking into account that the MEC system (i.e. the MEC platform and applications running on it) may not be considered a trusted entity by either party.

Furthermore, this creates an opportunity for operators to offer common authentication /identification in both mobile and enterprise segments "as-a-service" providing additional added value to enterprises.

High Bandwidth Content Optimization and QOE enablement

Many of the enterprise use cases will rely on video and similar high bandwidth content. In the case of mid to large size organizations, the cost associated with sufficient throughput to support such applications is a significant challenge. Translating this into an operator-provided MEC system supporting multiple applications means solving this issue potentially for multiple enterprises at a time – while keeping the traffic of each enterprise fully differentiated and separate from each other and from public network traffic.

Enterprise Operations & Maintenance requirements

Every enterprise requires the ability to monitor and control its assets. The same will apply to enterprise applications running in MEC clouds. This means that MEC clouds will require O&M tools and solutions like those in use today for other public clouds – but adapted to the unique nature of MEC as a highly distributed collection of smaller mini and micro clouds. This issue is particularly acute in those cases where enterprise locations include hard to reach areas. For example, oil/gas pumping sites are often remote, poorly connected and many – an important and highly challenging case of highly distributed infrastructure where effective O&M is critical.

Addressing the MEC Enterprise Challenges

Supporting Enterprise-Grade MEC Applications

Enterprise deployments of MEC are expected to require co-location of operator-managed network functions and enterprise-managed IT applications on a shared infrastructure. Such deployments need to satisfy several requirements.

Co-existence with NFV management framework, such as ETSI NFV

Virtualizing network functions is subject to very different performance and management requirements than virtualization of traditional IT applications, see [8] for a detailed list of NFV requirements. The upshot of this different set of requirements is the development of an understanding that virtualizing network functions represents a different type of virtual application and that the infrastructure for enabling such functions must be different as well. This different approach to virtualization is now called Network Functions Virtualization (NFV).

However, this does mean that co-located NFV and enterprise applications on the same infrastructure means that the infrastructure management framework must be able to deal with an additional layer of complexity. Specifically:

- Because network functions and enterprise applications belong to two different domains of trust (the carrier's domain and the enterprise domain), they must be located in well-separated tenant spaces (and, if possible on separate physical resources). Communication between these domains must be enabled using separate LANs (virtual and physical) with appropriate security infrastructure deployed (firewalls, policy-based routing, etc.) For management of virtual network functions, a useful resource is ETSI MEC's report on integration with the NFV management framework [9].
- Notwithstanding the above requirement, we must recognize that the shared physical infrastructure ultimately has a single owner (enterprise or carrier) and that this owner must be able to manage the infrastructure as a whole – preferably in a fairly dynamic fashion so as to be able to realize the benefits resulting from the flexibility of virtualization.

This leads to a layered approach towards management of enterprise-based MEC deployments. Each entity maintains its own management framework which has control over one or several tenant spaces allocated to it. The enterprise management framework can be based on traditional enterprise tools, while the carrier management framework can be based on traditional NFV management tools. However, in both cases, the management framework must limit its "scope" to management of virtual infrastructure (vCPUs, volumes, vLANs, etc.) assigned to it. In addition, the owner of this physical infrastructure maintains a third management framework for the physical infrastructure. Its scope must be limited to management of physical infrastructure and allocation of virtual resources to each of the tenants using the physical infrastructure. A well-defined simple API framework is required for this infrastructure and a good practice is never to integrate it with one of the virtual management frameworks – even when the two happen to be operated by the same entity (e.g. both are operated by the carrier).

Support of multiple approaches to virtualization

As we all know, modern approaches to virtualization have evolved from a single, Virtual Machine-based approach, to several, notably including containerization and serverless compute. It is widely believed that both the NFV and MEC management frameworks are mostly agnostic to the virtualization type. However, it is likely that each could benefit from certain optimizations that are specific to a virtualization approach and both ETSI MEC and ETSI NFV are in the process of studying this topic.

The ongoing study in ETSI MEC includes the gaps in currently defined MEC functionalities when running MEC applications as containers. There are several use cases considered in the study, which includes containerized MEC application packaging, on-boarding, instantiating etc. The MEC study is expected to take into account the requirements of application developers and identify gaps in existing MEC specifications. The study is expected to be published in the first quarter of 2019.

Unity of experience across all networks (Fixed /3G /4G /5G / Wi-Fi)

An enterprise network is a private network. It could be large and geographically distributed across multiple cities/countries in the world, or it could be as small and localized as a single office. Each enterprise may choose a different approach to build up its networks based on its size and business requirements.

Figure 3 illustrates an example of enterprise network deployment, consisting of several "zones":

- The headquarters, where the core business services are located.
- Satellite offices, with local enterprise networks being inter-connected with the headquarters cloud through secured backhaul networks provide accessibility for enterprise employees to access the enterprise services. An enterprise network may use 4G/5G small cells for the outdoor coverage, Wi-Fi networks for the indoor coverage and fixed access for static devices.
- Remote employees, which access enterprise services using VPN over the public Wi-Fi or cellular networks.

A MEC host provides a computing environment with networking interfaces for running applications. In order to meet its service requirements, an enterprise may distribute services from its central cloud to MEC hosts deployed at the edges of networks so that the timing critical or bandwidth consuming applications can be run very close to the device's location. Some applications may only need to run over local area networks. This may require the MEC system to support all the access network connection types.

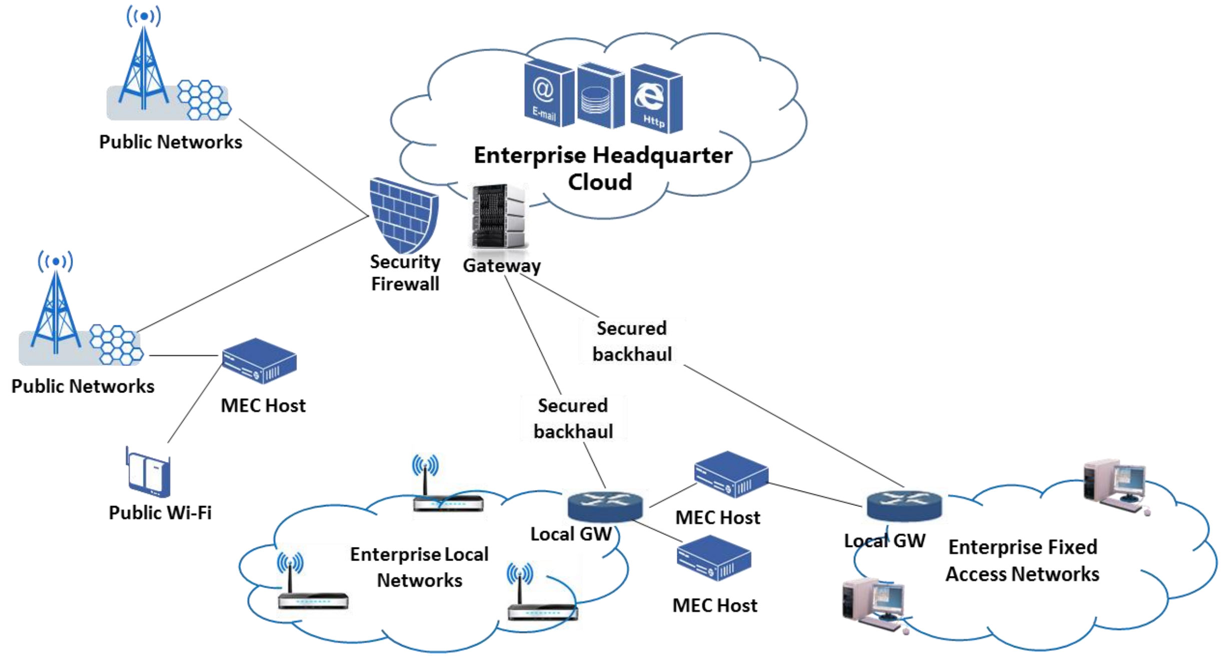

Figure 3: MEC deployment across different enterprise networks

Orchestration and optimization of enterprise applications

Different enterprise applications may have different characteristics and requirements. For example, augmented reality conferencing may be latency critical and bandwidth consuming, while location-restricted BYOD may require a seamless service when employees move in the enterprise campus across different access networks. Multiple types of enterprise networks provide an opportunity for MEC to orchestrate and optimize the performance of enterprise applications. In order for the MEC system to deliver a unity of experience over a large number of distinct types of access networks, it may require information on enterprise application requirements and information on access networks.

Table 1 shows an example of characteristics of different types of access network. Table 2 provides an example of characteristics for enterprise use cases described in above and possible access networks for running applications.

Table 1: A example of characteristics of access networks for enterprise

	4G LTE	5G NR	Wi-Fi	Fixed Access
Frequency band	Licensed	Licensed	License exempt	NA
Max data rate*	>100Mbps (DL) >50Mbps (UL)	10Gbps	6.9 Gbps (802.11ac) 9.6 Gbps (802.11ax)	Variable
Min latency*	10ms (air interface)	1ms (air interface)	Varies as the access loading	Variable
Session continuity	It supports “make before break”	It supports “make before break” inter-cell handover.	It supports the “break before	NA

	inter-cell handover.		make” for ESS inter-AP handover.	
Note 1: the maximum data rate depends on the bandwidth of the operational channel. Note 2: the minimum latency refers to one-way physical layer latency.				

Table 2: An example of enterprise application characteristics

Enterprise Usecases	Characteristics		
	Data rate	Latency	Possible access
Smart enterprise building	Variable	Variable	Wi-Fi or Fixed
Data analysis and security	> 20Mbps	Variable	Wi-Fi or Fixed
Augmented reality conferencing	100Mbps – 9.4Gbps	< 5ms	Wi-Fi or Fixed
Location-restricted BYOD access	Variable	Variable	Wi-Fi, Small Cell, Fixed
Video streaming	> 25Mbps	Variable	Wi-Fi or Fixed

Based on enterprise application requirements and access network characteristics, the MEC management could optimally orchestrate and schedule enterprise applications running on a MEC host close to devices’ locations over one or multiple appropriate access network connections. For enterprise applications like augmented reality conferencing, MEC management may choose a MEC host with wide bandwidth WiFi or fixed access to instantiate the application for delivery of the service to enterprise users. For the location-restricted BYOD access, enterprise applications may only be on-boarded to the MEC host at a specified location. Therefore only on-site employees can receive the services produced from those enterprise applications over local enterprise networks.

Unity of MEC APIs across enterprise networks

ETSI ISG MEC is developing a series of API specifications for different access networks:

- GS MEC 012 [14] specifies the APIs for radio network information service (RNIS). This specification defines an API that provides access to a large amount of network information for a 3GPP-defined network.
- GS MEC 028 is a specification under development that will specify the APIs for WLAN information service (WIS) and which will serve a purpose for WiFi networks that is similar to that of [14] for 3GPP-based access.

- GS MEC 029 is a specification under development that will specify the APIs for Fixed Access Information Service (FAIS) and which will serve a purpose for Fixed-Access networks that is similar to that of [14] for 3GPP-based access.

These APIs provide facilitate the unity of service interfaces by providing a common standardized service access to enterprise applications. Moreover, by providing information on the status of the access network, these APIs assist applications and service orchestrators in properly configuring and mapping applications across available access networks.

Integration of Access Control

A fundamental MEC operation is the ability to forward traffic between the access network and an application instance on a MEC host. ETSI MEC specifications enable this operation by specifying a traffic filtering service that a MEC Platform (MEP) must provide. The service is specified in ETSI GS MEC 011 [10]. The most common approach to indicating which traffic to forward is to use the IP 5-tuple: the transport protocol (TCP/UDP/etc.) and the source and destination IP addresses and port numbers. However, ETSI MEC recognizes that a number of other means of traffic filtering may be of use. In particular, when the access network is a 3GPP mobile access network and the traffic is encapsulated in GTP tunnels, filtering by GTP tunnel parameters may be of use as well. The TrafficFilter data type ([10], clause 6.5.6) supports all these capabilities. Additionally, [10] effectively enables filtering by web names (Fully Qualified Domain Name - FQDN) by defining a DNS service which returns a set of IP address for an FQDN and these IP addresses can then be used to define a traffic rule.

However, in the case of enterprise services, an additional filtering capability is required – filtering by an "enterprise user" – i.e. a service where all traffic associated with a particular enterprise user is directed between the access network and the enterprise application. A simple solution would be to use some service to look up all the IP flows associated with a particular user and set up IP-based traffic filters for all such flows. Indeed, in the mobile network such a service is readily available – e.g. the MME in 4G networks. Unfortunately, the association of IP flows is made to the mobile identity (e.g. IMSI) and not the enterprise identity – which highlights the need (as noted previously) for a way to associate such identities.

A naïve solution would be to create a table mapping mobile identities and enterprise identities. If a mobile number (MS-ISDN) is sufficient and the enterprise is willing to maintain a mapping of user identities and their MS-ISDNs then this is a sufficient approach. However, in some instances the MS-ISDN cannot be used and a mapping to IMSI (the actual mobile network identity) is required. This creates a problem: IMSI is a critical "identity asset" within the mobile network, much as an enterprise user identity (e.g. an LDAP identity) is within an enterprise network. Neither can be expected to share its identity with the other – doing so would be a major violation of standard security practices and expose both the mobile network and the enterprise to significant potential security risks. Unfortunately, this means that a naïve direct mapping is not a feasible solution.

The problem we are describing can be viewed as a special case of a well-known single-sign-on (SSO) problem and SSO techniques can be used to solve this problem. Essentially, a mutually trusted entity generates a stand-in value – a *token*, which is used for the following purposes:

- The MEC System is able to associate the *token* to an access network identity; however the *token* does not reveal the access network identity to any entity that is not "trusted" by the access network operator

- The enterprise application is able to associate the *token* to an enterprise identity; however the *token* does not reveal the enterprise identity to any entity that is not “trusted” by the enterprise.

Assuming that such a *token* can be defined, ETSI MEC has specified the means to use it for traffic filtering as follows:

- ETSI MEC 014 [11] defines an API by which such a *token* can be made available to the application.
- ETSI MEC 011 [10] supports traffic filtering by *token* (see clause 6.5.6).

What remains, therefore is how to define such a token. While the specific approach is left up to the design of each system, several well-known and standardized means are available. Below are a few examples:

- **Using MS-ISDN.** As noted above, the public phone number (or, more broadly, MS-ISDN) is one means to identify a user. Although not fully secure – as it is not secret – it may be sufficient for some applications. In this case, the MS-ISDN becomes the *token*; the enterprise maintains the mapping between the MS-ISDN and enterprise identities and the MEP is able to associate MS-ISDN and IP flows (typically by invoking services provided by mobile network entities such as the MME).
- **Using SIM-based authentication, such as EAP-SIM.** EAP-SIM [12] is a well-known protocol developed for the purposes of integrating authentication and access control mechanism between WiFi and 3GPP systems. The protocol is in use in HotSpot 2.0 systems (marketed under the Passpoint® brand) as defined by the WiFi Alliance. It uses a mobile device, which includes both a UE functionality with a SIM module for 3GPP access and WiFi functionality which relies on the AAA server and EAP protocol for access. As part of the EAP-SIM access procedure, a SIM-based “key” is generated to be used as the “master key” in the WiFi keying system. Because it is SIM-based, it can be generated by the mobile network and made available to MEP. Thus, it can be used as a *token*. The enterprise application is provided this token by an appropriate enterprise agent on the client device using secure means that are enterprise specific. Note that for the purposes of MEC *token* generation, the procedure can be ran without any WiFi based interaction – it is just being re-used for a different purpose.
- **Using 3GPP’s Generic Authentication Architecture / Generic Bootstrapping Architecture (GAA/GBA).** GAA/GBA are 3GPP defined mechanisms allowing a mobile operator to become a provider of SIM-based SSO services. Should an enterprise take advantage of these services, any of the GAA/GBA generated identities/keys available to the enterprise can be used as a *token*.
- **3rd party SSO Providers, such as OAuth.** Authentication by-products of 3rd party SSO providers can be used as *tokens* – provided that both the enterprise and the access network operator agree to use the same SSO provider. As with GAA/GBA, any of a number of by-products of the SSO access procedure known to both entities can be used as tokens.

High Bandwidth Content Optimization and QOE enablement

High-bandwidth applications, such as video conferencing and video streaming, continue to suffer from the mismatch between the network design and the application demand. For example, it well known that traditional TCP congestion control is designed with a view towards wired networks and highly heterogeneous traffic. While cross-layer optimization across architectural boundaries remains the wrong approach for broad-use public applications and for consumer products, in the context of an enterprise network it is perfectly acceptable and can bring about significant efficiencies in network performance, resulting in both user QoE improvement and IT cost reductions. The challenge is then, to achieve such vertical cross-layer optimization given a system designed from general purpose consumer components (HW and SW) – and to do it in the proper location within the system.

MEC can facilitate this by helping an application to identify the type of service and user profile in accordance with user data packets, and then by defining the QoS information for appropriate information flows that can be propagated into the access network (e.g. an LTE eNB). By integrating OTT service information and radio access network information, the MEC platform can use the advantages of intelligent channels to guarantee QoS of key users and services. Figure 4 shows an example of this.

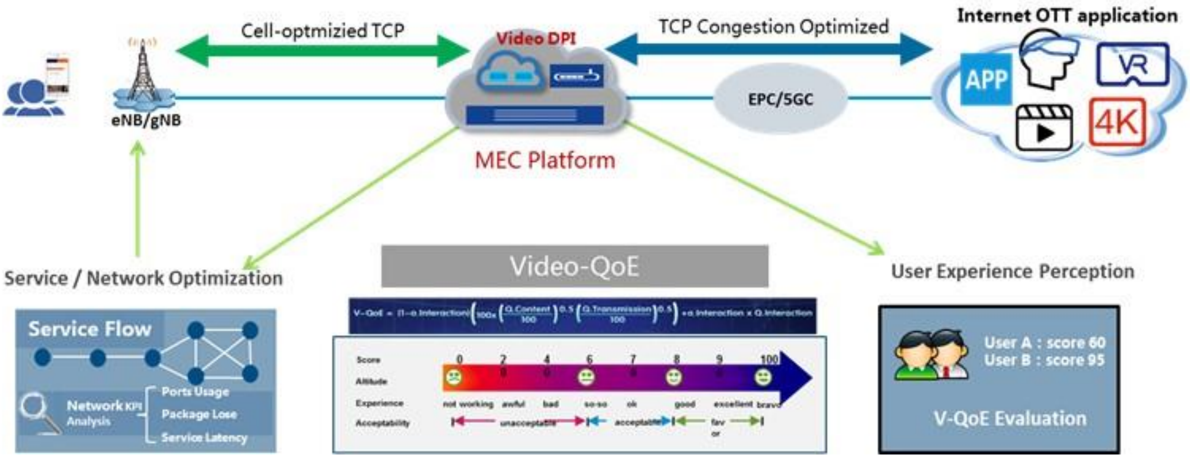

Figure 4: Example of a MEC-based Video-QoE optimizing application

How does ETSI MEC help achieve this? First the “Video DPI” component of the MEC service requires information about the state of the access network, which it can obtain using RNIS [14] for a 3GPP-based network (as shown) or using the upcoming WiFi and Fixed-Access information service APIs. Furthermore, user information in a particular edge site can be obtained using the Location APIs [15] and mapped to enterprise identities as discussed above. These can then be used to filter by specific user traffic, which allows such traffic to be operated on. Finally, the BW Management set of APIs, as defined in ETSI GS MEC 015 [20] enables the definition of QoS parameters to the various traffic flows and thus achieving the necessary goals as defined by the “computation” block in Figure 4.

Enterprise O&M requirements

As noted above, a key concern with operation and management of edge clouds is the highly distributed nature of these clouds and the fact that the communication links on which the O&M operations rely may be unreliable, or latency and throughput constrained. This issue is widely recognized, see e.g. [19]. This means that a successful O&M approach requires the following:

- A partitioning between a centralized system-wide orchestration and a localized on-edge-site management entity for implementation of decision.
- A well-defined secure set of APIs between these entities that is designed to be robust to communication links that may be unreliable, or latency and throughput constrained.

ETSI MEC facilitates the implementation of such systems, firstly by defining a reference architecture [16] that defines an on-the-host MEP Management Entity (MEPM) and a centralized, system-wide orchestration function (MEO). Additionally, REST-ful APIs for the management of the MEC platform [17] and the applications running at a MEC site [18] are defined with the requirement of robustness to imperfect communication links (none of the APIs are latency sensitive nor require significant throughput).

Summary and Conclusions

Enterprise is a key focus area for edge computing and represents most of the early deployments of edge computing. However, integration of edge computing in an access network presents a number of challenges that go beyond the typical issues that enterprises deal with. In this paper we have highlighted some of the key such challenges and outlined approaches to solutions. Clearly, it is not possible to provide detailed solutions in a short white paper, moreover a good solution should always take the specific needs and characteristics of each enterprise into account. However we do hope that this paper helps its readers in designing an appropriate solution. Additionally, we hope that by illustrating how ETSI MEC specifications enable such solutions in a simple, industry-standard interoperable way, we can encourage enterprises, operators and vendors to think of enterprise edge as a highly scalable market where much can be reused and duplicated despite the need to design to the specifics of each customer.

References

All ETSI MEC Specifications listed below can be accessed via: https://www.etsi.org/technologies-clusters/technologies/multi-access-edge-computing

[1] "Edge Computing Will Decentralize but Complement Traditional and Cloud-Based Data Center Architectures," [Online]. Available: https://itcblogs.currentanalysis.com/2018/04/09/edge-computing-will-decentralize-but-complement-traditional-and-cloud-based-data-center-architectures/ . [Accessed April 2018].

[2] "Why Edge Computing Is Here to Stay: Five Use Cases," [Online]. Available: https://www.rtinsights.com/why-edge-computing-is-here-to-stay-five-use-cases/ . [Accessed April 2018].

[3] "A Local Edge Lifecycle to Create Competitive Differentiation," [Online]. Available via: https://blog.schneider-electric.com/datacenter/2018/04/04/local-edge-lifecycle-competitive-differentiation/

[4] "An A.T.M., With a Real Teller on the Screen," [Online]. Available via: https://bucks.blogs.nytimes.com/2013/04/04/an-a-t-m-with-a-real-teller-on-the-screen/

[5] "Which Data Center Use Cases Are Best For a Value Added Reseller Business Plan," [Online]. Available via: http://www.ingrammicroadvisor.com/data-center/which-data-center-use-cases-are-best-for-a-value-added-reseller-business-plan

[6] ETSI, "Developing Software for Multi-Access Edge Computing," [Online]. Available via: http://www.etsi.org/images/files/ETSIWhitePapers/etsi_wp20_MEC_SoftwareDevelopment_FINAL.pdf

[7] "The Edge Is the Greenest, Most Intelligent Building in the World," [Online]. Available via: https://www.bloomberg.com/features/2015-the-edge-the-worlds-greenest-building/

[8] ETSI GS NFV-IFA 010, "Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Functional requirements specification," v. 2.4.1, 02/2018.

[9] ETSI GR MEC 017, "Mobile Edge Computing (MEC); Deployment of Mobile Edge Computing in an NFV environment," v. 1.1.1, 02/2018.

[10] ETSI GS MEC 011, "Mobile Edge Computing(MEC); Mobile Edge Platform Application Enablement," v. 1.1.1, 07/2017.

[11] ETSI GS MEC 014, "Mobile Edge Computing (MEC); UE Identity API," v. 1.1.1, 02/2018.

[12] IETF RFC 4186, "EAP-SIM Authentication," 01/2006. Available via: https://tools.ietf.org/html/rfc4186

[13] Wi-Fi Alliance, "Hotspot 2.0 (Release 2) Technical Specification," v. 1.2, 2016. Available via https://www.wi-fi.org/discover-wi-fi/passpoint

[14] ETSI GS MEC 012, "Mobile Edge Computing (MEC); Radio Network Information API," v. 1.1.1. 07/2017.

[15] ETSI GS MEC 013, "Mobile Edge Computing (MEC); Location API," v. 1.1.1, 07/2017.

[16] ETSI GS MEC 003, “Mobile Edge Computing; Framework and Reference Architecture,” v. 1.1.1, 03/2016.

[17] ETSI GS MEC 010-1, “Mobile Edge Computing (MEC); Mobile Edge Management; Part 1: System, host and platform management,” v. 1.1.1, 10/2017.

[18] ETSI GS MEC 010-2, “Mobile Edge Computing (MEC); Mobile Edge Management; Part 2: Application lifecycle, rules and requirements management,” v. 1.1.1, 07/2017.

[19] OpenStack, “Cloud Edge Computing: Beyond the Data Center,” Available via https://www.openstack.org/assets/edge/OpenStack-EdgeWhitepaper-v3-online.pdf

[20] ETSI GS MEC 015, “Mobile Edge Computing (MEC); Bandwidth Management API,” v. 1.1.1, 10/2017.

ETSI
06921 Sophia Antipolis CEDEX, France
Tel +33 4 92 94 42 00
info@etsi.org
www.etsi.org

This White Paper is issued for information only. It does not constitute an official or agreed position of ETSI, nor of its Members. The views expressed are entirely those of the author(s).

ETSI declines all responsibility for any errors and any loss or damage resulting from use of the contents of this White Paper.

ETSI also declines responsibility for any infringement of any third party's Intellectual Property Rights (IPR), but will be pleased to acknowledge any IPR and correct any infringement of which it is advised.