Mobile-Edge Computing

Objectives

This white paper is authored by the founders of the Mobile-edge Computing (MEC) industry initiative.

The objectives of this paper are to introduce the concept of Mobile-edge Computing and the related key market drivers, and to discuss the business, consumer and technical value/benefits that this technology offers. The paper discusses the enablers, the requirements and challenges for Mobile-edge Computing as well as the objectives of the MEC initiative.

This white paper presents the high-level architectural blueprint of Mobile-edge Computing which, together with the scope of work, will form the basis for the first release of the work in the initiative. In addition, it highlights the relationships between and the interfaces with other industry efforts.

The authors invite the various players in the value chain to actively participate in the work of the initiative.

Contributing Organizations and Authors

Huawei: Milan Patel, Yunchao Hu, Patrice Hédé

IBM: Jerome Joubert, Chris Thornton, Brian Naughton

Intel: Julian Roldan Ramos, Caroline Chan, Valerie Young, Soo Jin Tan, Daniel Lynch

Nokia Networks: Nurit Sprecher, Torsten Musiol, Carlos Manzanares, Uwe Rauschenbach

NTT DOCOMO: Sadayuki Abeta, Lan Chen, Kenji Shimizu

Vodafone: Adrian Neal, Peter Cosimini, Adam Pollard, Guenter Klas

Publication date: September 2014

Table of Contents

1 Executive summary

Mobile-edge Computing provides IT and cloud-computing capabilities within the Radio Access Network (RAN) in close proximity to mobile subscribers.

For application developers and content providers, the RAN edge offers a service environment with ultra-low latency and high-bandwidth as well as direct access to real-time radio network information (such as subscriber location, cell load, etc.) that can be used by applications and services to offer context-related services; these services are capable of differentiating the mobile broadband experience.

Mobile-edge Computing allows content, services and applications to be accelerated, increasing responsiveness from the edge. The mobile subscriber’s experience can be enriched through efficient network and service operations, based on insight into the radio and network conditions.

Operators can open the radio network edge to third-party partners, allowing them to rapidly deploy innovative applications and services towards mobile subscribers, enterprises and other vertical segments. Proximity, context, agility and speed can be translated into value and can create opportunities for mobile operators, service and content providers, Over the Top (OTT) players and Independent Software Vendors (ISVs), enabling them to play complementary and profitable roles within their respective business models and allowing them to monetize the mobile broadband experience.

This environment can create a new value chain and an energized ecosystem comprising application developers, content providers, OTT players, network equipment vendors and mobile operators. Based on innovation and business value, this value chain will allow all players to benefit from greater cooperation.

The intention is to develop favorable market conditions which will create sustainable business for all players in the value chain, and to facilitate global market growth. To this end, a standardized, open environment needs to be created to allow the efficient and seamless integration of such applications across multi-vendor Mobile-edge Computing platforms. This will also ensure that the vast majority of the customers of a mobile operator can be served. A new Industry Specification Group (ISG) is proposed to be set up in ETSI to allow the creation of industry specifications for Mobile-edge Computing (MEC). The ISG MEC will also work towards enabling and accelerating the development of edge applications across the industry, increasing the market scale and improving market economics.

Chapter 2 of this paper introduces the key market drivers for the evolution of the mobile base station as well as the business and technical benefits of Mobile-edge Computing. Use case scenarios are presented.

Chapter 3 raises the need for an industry initiative for Mobile-edge Computing.

Chapter 4 presents the deployment scenarios that will be supported in the first release and outlines scenarios that may be considered in future releases of the initiative.

Chapter 5 describes (a) the high-level architectural blueprint of Mobile-edge Computing that will form the basis of the work in the ISG MEC, and (b) the scope of the first release.

Chapter 6 outlines the enablers for Mobile-edge Computing.

Chapter 7 discusses related requirements and technical challenges.

Chapter 8 describes the relationships between and the interfaces with other industry efforts.

Chapter 9 calls for the active participation of the different players in the value chain: mobile operators, equipment vendors, platform providers, Application Service Providers (ASPs) and OTT players.

Chapter 10 presents the conclusion.

Chapter 11 provides the contact information of the authors of the white paper.

2 Introduction

2.1 Trends and market drivers

End users and businesses are demanding more from the telecommunication industry. While end users request personalized services, better performance, better user experience, businesses need to get more information about their consumers, easier and secured access to devices and greater flexibility for provisioning new services. There is a key role to play for Equipment providers, Service Providers and IT players together to make this a reality by providing converged IT and Network infrastructure.

The continuing growth of mobile traffic is well documented, driven mainly by consumer smart phones, streaming video, messaging and P2P applications. The growth in mobile traffic is set to increase dramatically as enterprises extend their business processes to smart mobile devices and as machine-to-machine solutions mature throughout vertical industries. Wireless sensors are key enablers to many mission-critical scenarios, from smarter traffic to video analytics. Wireless sensors are expected to grow in their numbers exponentially over the next 10 years. The cellular network is the ubiquitous platform for integrating these devices with vertical back office solutions.

The worlds of IT and Telecommunications Networking are converging bringing with them new possibilities and capabilities that can be deployed into the network (see Figure 1). A key transformation has been the ability to run IT based servers at network edge, applying the concepts of cloud computing. We define this as Mobile-edge Computing. Mobile-edge Computing can be seen as a cloud server running at the edge of a mobile network and performing specific tasks that could not be achieved with traditional network infrastructure. Machine-to-Machine gateway and control functions are one example, but there are many others.

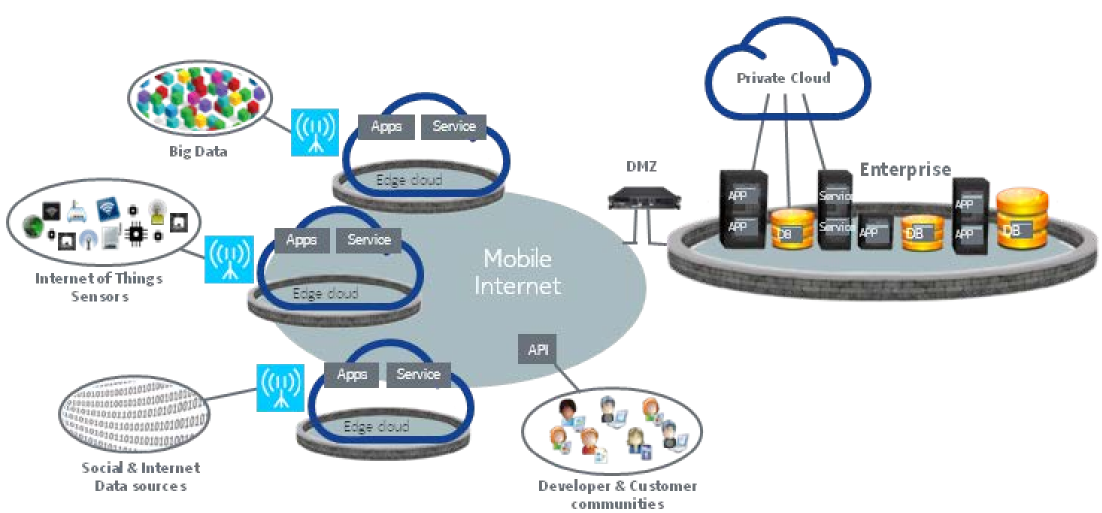

Figure 1: IT and Telecommunications networking convergence

Typically, Mobile-edge Computing is characterized by:

- **On-Premises**: The Edge is local, meaning that it can run isolated from the rest of the network, while having access to local resources. This becomes particularly important for Machine-to-Machine scenarios, for example when dealing with security or safety systems that need high levels of resilience.
- **Proximity:** Being close to the source of information, Edge Computing is particularly useful to capture key information for analytics and big data. Edge computing may also have direct access to the devices, which can easily be leveraged by business specific applications.
- **Lower latency**: As Edge services run close to end devices it considerably reduces latency. This can be utilized to react faster, to improve user experience, or to minimize congestion in other parts of the network.
- **Location awareness**: When a Network Edge is part of a wireless network, whether it is Wi-Fi or Cellular, a local service can leverage low-level signaling information to determine the location of each connected device. This gives birth to an entire family of business-oriented use cases, including Location Based Services, Analytics, and many more.
- **Network context information**: Real-time network data (such as radio conditions, network statistics, etc.) can be used by applications and services to offer context-related services that can differentiate the mobile broadband experience and be monetized. New applications can be developed (which will benefit from this real-time network data) to connect mobile subscribers with local points-of-interest, businesses and events.

2.2 Evolution of the mobile base station

In the past, the edge of a mobile network was a place where only specialist processing was done. It housed specialized computing that was designed from the ground up to perform a function in the overall architecture and was not able to be repurposed. Connectivity from the edge of the network back to the core was also a specific configuration, running over specialized protocols. The complete configuration was optimized in the pre smartphone era, where Voice quality was the key driver in network design and before the days where IP was the standard for network communications.

IP has spread from the internet, to enterprise networks and with widespread adoption of LTE, through the edge of networks to the end devices themselves. This has enabled new applications to emerge that have seen a transformation in telecommunication networks and their design. Single vendor radio network solutions are evolving into modular, open solutions that are able to integrate in an ecosystem of changeable components.

Edge computing in outdoor scenarios

There are different ways to implement Mobile-edge Computing, depending on the access technology. For outdoor, Macro cells vendors embed secured computing and virtualization capabilities directly into radio access network elements. This integration of applications with radio equipment allows operators to rapidly deliver innovative network features, accelerate over-the-top (OTT) services and enable a variety of new high value services. Such flexible services are executed at a very strategic location in the mobile network, making them much more essential than any other applications run at the core. This architecture is particularly relevant to:

- Improve mobile users' Quality of Experience (QoE), by reducing latency, improving quality of service or/and providing customized services.
- Improve infrastructure's efficiency, with more intelligent and optimized networks.
- Enable disruptive vertical services, particularly relevant for Machine-to-Machine scenarios, Big Data management, Analytics, Smart Cities and much more.
- Tight integration with radio equipment, making it easy to understand traffic characteristics and needs, deal with radio conditions, get device location information, etc.

Edge computing in indoor scenarios

When it comes to indoor, such as Wi-Fi and 3G/4G access points, edge clouds take the form of powerful on-premises gateways, where dedicated intelligence serves local purposes. Through lightweight virtualization, those gateways run multiple services applied to the particular location they are installed in, such as:

- **Machine-to-Machine scenarios**: Connecting to various sensors, Mobile-edge Computing services can deal with all sorts of monitoring activities (air conditioning, elevators, temperature, humidity, access control, etc.)
- **Retail Solutions:** Having the ability to locate and communicate with mobile devices, there is an opportunity to deliver higher value to the consumers and the malls. For example delivering content based on location, implementing augmenting reality, improving the overall shopping experience, or dealing with secured online payment.
- **Stadiums, Airports, Stations, Theatres:** Specific services can help manage other types of crowed places, in particular to deal with safety, security, evacuation, or to provide new kinds of services to the public. For example, stadiums could provide live content to the public, airports could guide passengers to their gate through an augmented reality service, and many more. All these

applications would leverage local content and conditions to be perfectly adapted to their audience.

- **Big data and Analytics**: Last but not least, the information gathered at this key point in the network, can be leveraged as part of a bigger analytics initiative to serve customers better.

2.3 Business and technical benefits

Mobile-edge Computing provides a new ecosystem and value chain and the opportunity for all players within it to collaborate and develop new business models they can each benefit from.

Mobile Network Operators (MNOs) can rapidly deploy new services for consumer and enterprise business segments which can help them differentiate their service portfolio. Adding new revenue streams from innovative services delivered from closer to the user can improve the MNOs bottom line whilst improving end user QoE. New applications which are aware of the local context in which they operate (RAN conditions, locality, etc.) can open up entire new service categories and enrich the offering to end users. Placing relevant applications on or near the base station not only offers advantages to consumer and enterprise end users. It also reduces the volume of signaling offloaded to the core network and could also reduce OPEX for the MNOs, compared to hosting in the core. The MNOs could increase their revenue by charging based on the resource usage (storage, NW bandwidth, CPU, etc.) of each content provider, if such resource usage could be obtained via specific APIs in MEC server.

Software and application providers can serve the new ecosystem by developing and bringing to the market innovative and ground breaking services and applications that can take advantage of the information on radio network capabilities and conditions available at the base station. The application space is open to anyone: software and application providers, infrastructure vendors and MNOs.

The use of open standards and Application Programming Interfaces (APIs), as well as the use of familiar programming models, relevant tools chain and Software Development Kits are key pillars to encourage and expedite the development of new disruptive applications or the adaptation of existing services and applications to the new Mobile-edge Computing environment.

2.4 Use cases

This subsection depicts some examples of use cases and benefits that have already been demonstrated. Many other use cases can be defined and enabled by the MEC open framework.

Use Case1: Active Device Location Tracking

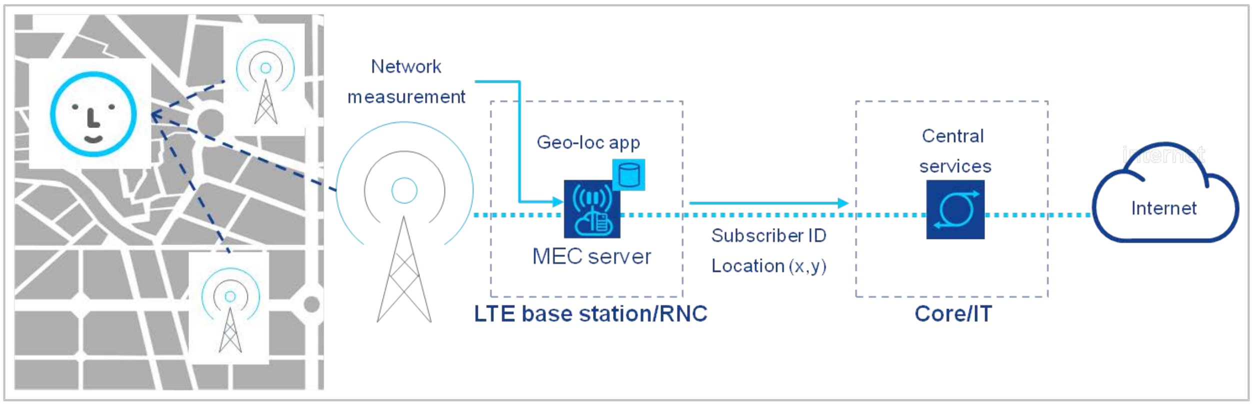

Figure 2: Example of active device location tracking

Figure 2 shows an example of the active device location tracking use case. This use case enables real-time, network measurement based tracking of active (GPS independent and network determined) terminal equipment, using 'best-in-class' third-party geo-location algorithms within a geo-location application hosted on the MEC server.

This provides an efficient and scalable solution with local measurement processing and event based triggers. It enables location based services for enterprises and consumers (e.g. on opt-in basis), for example in venues, retail locations and traditional coverage areas where GPS coverage is not available.

Services may include mobile advertising, 'Smart City', footfall analysis, campus management, etc.

Use case2: Augmented Reality Content Delivery

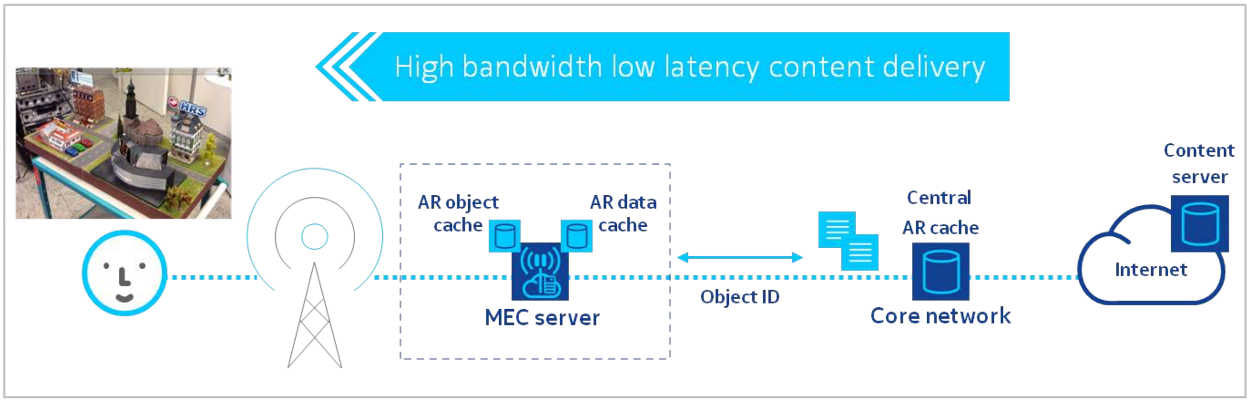

Figure 3: Example of augmented reality content delivery

An Augmented Reality (AR) application on a smart-phone or tablet overlays augmented reality content onto objects viewed on the device camera (as displayed in Figure 3). Applications on the MEC server can provide local object tracking and local AR content caching.

The solution minimizes round trip time and maximizes throughput for optimum quality of experience. It can be used to offer consumer or enterprise propositions, such as tourist information, sporting event information, advertisements etc.

Use Case 3: Video Analytics

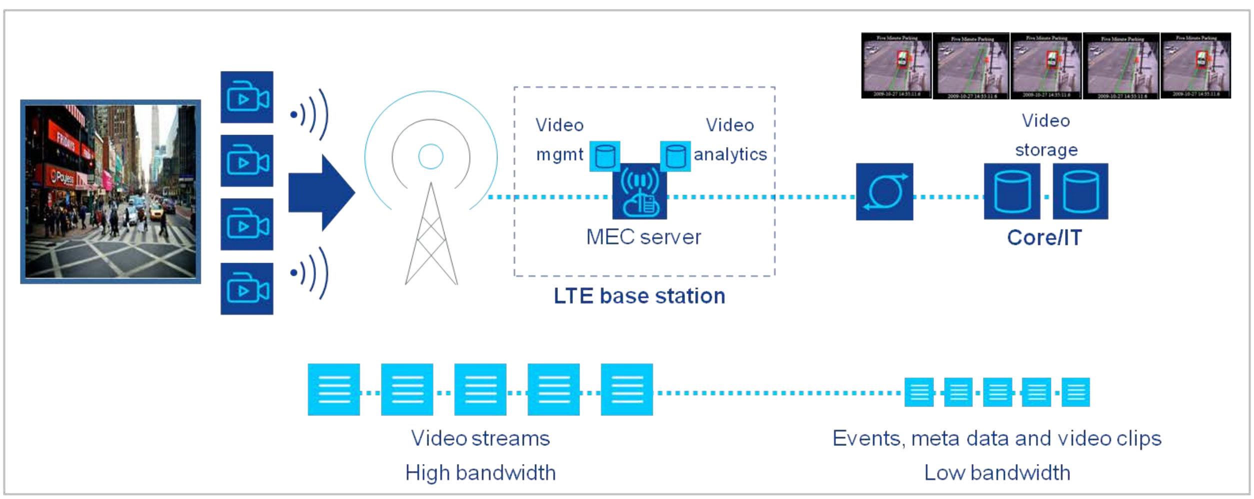

Figure 4: Example of video analytics

Figure 4 shows a distributed video analytics solution which provides an efficient and scalable mobile solution for LTE.

The video management application transcodes and stores captured video streams from cameras received on the LTE uplink. The video analytics application processes the video data to detect and notify specific configurable events e.g. object movement, lost child, abandoned luggage, etc. The application sends low bandwidth video metadata to the central operations and management server for database searches. Applications may range from safety, public security to smart cities.

Use Case 4: RAN-aware Content Optimization

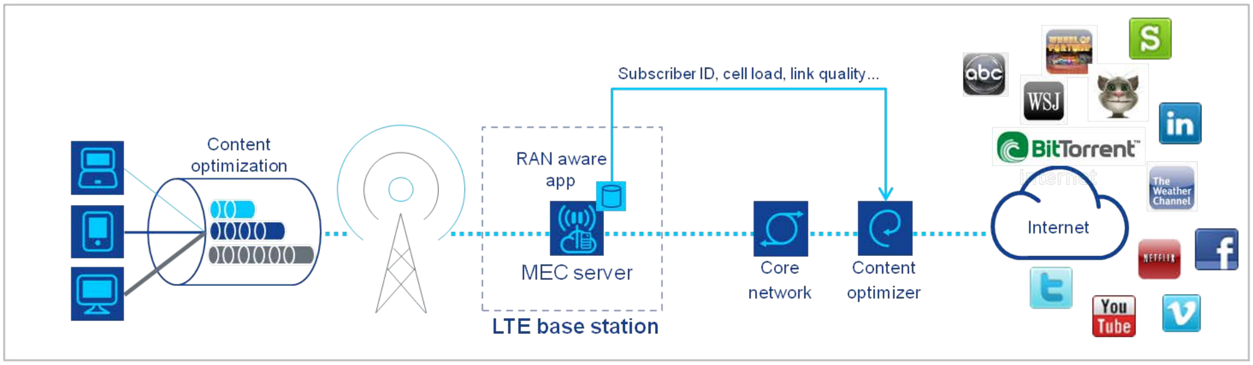

Figure 5: Example of RAN-aware content optimization

In this use case, the application exposes accurate cell and subscriber radio interface information (cell load, link quality) to the content optimizer, enabling dynamic content optimization, improving QoE, network efficiency and enabling new service and revenue opportunities. Dynamic content optimization enhances video delivery through reduced stalling, reduced time-to-start and 'best' video quality.

Figure 5 displays an example of RAN-aware content optimization.

The concept also enables enhanced use cases, such as promoted content delivery and subscriber throughput boosting.

Use Case 5: Distributed Content and DNS Caching

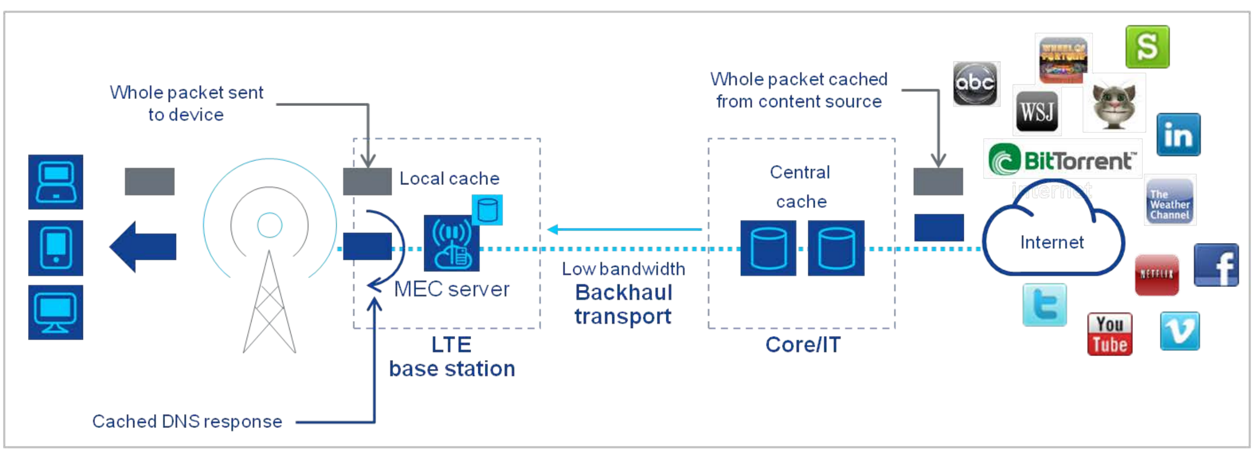

Figure 6: Example of distributed content and DNS caching

A distributed caching technology (as shown in Figure 6) can provide backhaul and transport savings and improved QoE. Content caching has the potential to reduce backhaul capacity requirements by up to 35%. Local Domain Name System (DNS) caching can reduce web page download time by 20%[1].

Use Case 6: Application-aware Performance Optimization

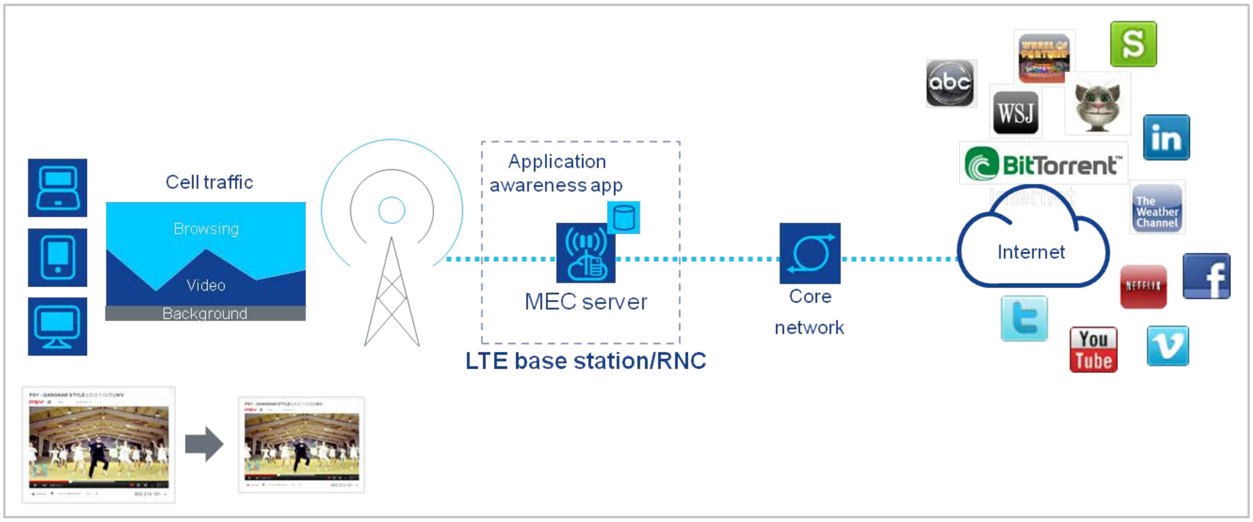

Figure 7: Example of Application-aware performance optimization

Application-aware cell performance optimization for each device in real time can improve network efficiency and customer experience (see Figure 7). It can reduce video stalling and increase browsing throughput. Latency may also be significantly reduced.

The solution can also provide independent metrics on application performance (video stalls, browsing throughput, and latency) for enhanced network management and reporting.

[1] Smart cells revolutionalize service delivery. Intel Corp. http://www.intel.co.uk/content/dam/www/public/us/en/documents/white-papers/smart-cells-revolutionize-service-delivery.pdf

3 The need for an industry initiative for Mobile-edge Computing

3.1 Rationale and objectives

In order to develop favorable market conditions which will create sustainable business for all players in the value chain, and to facilitate global market growth, it is proposed to set up a new Industry Specification Group (ISG) within ETSI for Mobile-edge Computing.

The purpose of the ISG MEC is to create a standardized, open environment which will allow the efficient and seamless integration of applications from vendors, service providers, and 3rd parties across multi-vendor Mobile-edge Computing platforms. This will ensure that the vast majority of the customers of a mobile operator can be served. In addition, the group will work to enable and accelerate the development of edge applications across the industry as well as to increase the market scale and to improve market economics. The group will also address compliance with regulatory and legal requirements.

The ISG MEC will produce interoperable and deployable Group Specifications (Standards Track Deliverables) that will allow the hosting of such applications in a multi-vendor Mobile-edge Computing environment.

The architectural blueprint presented in chapter 5 below forms the basis of the work in the ISG MEC and the scope of the first release.

3.2 Mobile-edge Computing work under the auspices of ETSI ISG

ETSI plays an active role in the development and implementation of telecommunication standards. ETSI has developed open mobile communication standards, such as GSM, UMTS and LTE, independently and with partners in other regions. These specifications have been successfully adopted worldwide.

In addition, ETSI is currently working on the development of standards relating to cloud and internet technologies as well as to end-to-end network architecture (which, among other aspects, also addresses Content Delivery Networks (CDN)).

The work of ETSI MEC aims to unite the telco and IT cloud worlds, providing IT and cloud-computing capabilities within the RAN (Radio Access Network). The ISG MEC will specify the elements that are needed to enable applications to be hosted in a multi-vendor Mobile-edge Computing environment. Proximity, context, agility and speed can be translated into value and utilized by mobile operators, service and content providers, OTT players and ISVs. These parties can play complementary and profitable roles within their respective business models, since they can monetize the mobile broadband experience.

Moreover, ETSI has an active ISG working on Network Function Virtualization (NFV) which is tasked with leveraging the standardized IT virtualization technology to consolidate multiple types of network equipment into standards-based, high-volume servers, switches and storage which could be located in data centers, network nodes and on the end user's premises. With NFV, entire classes of network node functions can be virtualized into building blocks that may be connected or chained together to create communication services. The work on MEC will complement the work on NFV and the scope will be highly focused (for the details of the scope, see section 5.2 below). MEC will enable applications and services (Layer 4 and above) to be hosted 'on top' of the base station or the RNC, i.e. above the network layer. These applications and services can benefit from being in close proximity to the customer and from receiving local radio-network contextual information.

The ISG MEC will align and liaise with ETSI ISG NFV, 3GPP and with other related initiatives (inside and outside of ETSI), reusing existing specifications where appropriate.

The lifetime of the ISG MEC is expected to be 18 months from setup. The Terms of References (ToR) for the ISG MEC present the expected deliverables of the ISG MEC and their target delivery time.

In addition to the specifications, the ISG MEC will produce informative reports in the form of white papers and tutorials that will be used to advance Mobile-edge Computing in the industry, accelerate the adoption of the concept and the specifications, and address legal and regulatory requirements. These reports will include descriptions of use case scenarios and facilitate exchanges of real-life stories, practices and insights.

4 Deployment scenarios

The first release of the ISG MEC will support (non-binding) deployment scenarios where the MEC server is deployed either at the LTE macro base station (eNB) site, or at the 3G Radio Network Controller (RNC) site, or at a multi-technology (3G/LTE) cell aggregation site. The multi-technology (LTE/3G) cell aggregation site can be located indoor within an enterprise (e.g. hospital, large corporate HQ), or indoor/outdoor for a special public coverage scenario (e.g. stadium, shopping mall) to control a number of local multi-technology (3G/LTE) access points providing radio coverage to the premises. This deployment option enables the direct delivery of locally-relevant, fast services from base station clusters. For more information on indoor and outdoor deployment scenarios, see section 2.2 above.

Figure 8 depicts the deployment scenarios of a Mobile-edge Computing server that are supported in the first release of the work in the ISG MEC.

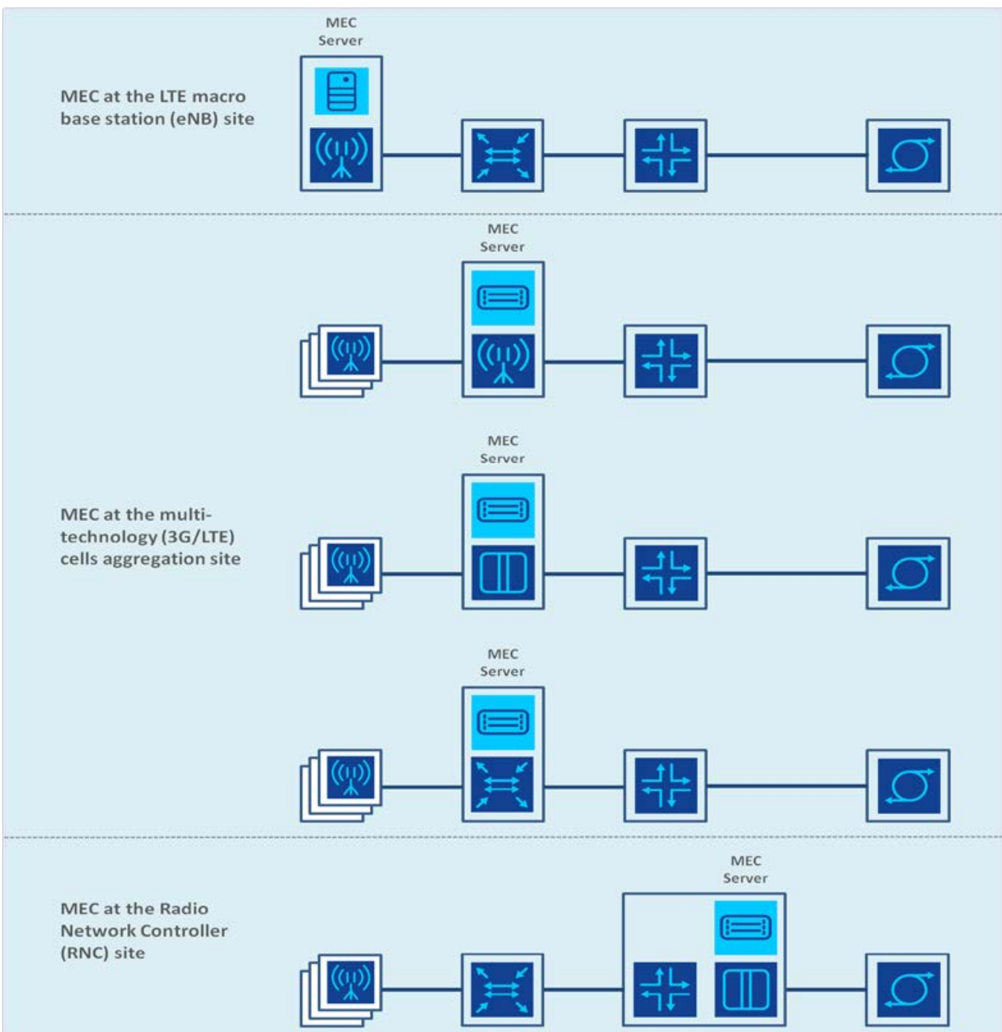

Figure 8: Deployment scenarios of the Mobile-edge Computing server

Note that for the sake of keeping the scope of the first release very focused in order to ensure the completion of the deliverables of the first release within 18 months, a multi-technology cell aggregation deployment scenario that includes WLAN access points will be considered once a more detailed analysis is available.

5 Architectural blueprint

Mobile-edge Computing provides a highly distributed computing environment that can be used to deploy applications and services as well as to store and process content in close proximity to mobile users. Applications can also be exposed to real-time radio and network information and can offer a personalized and contextualized experience to the mobile subscriber. This translates into a mobile-broadband experience that is not only more responsive, but also opens up new monetization opportunities. This creates an ecosystem where new services are developed in and around the base station.

The key element of Mobile-edge Computing is the Mobile-edge Computing (MEC) IT application server which is integrated at the RAN element (as described in chapter 4 above). The MEC server provides computing resources, storage capacity, connectivity, and access to user traffic and radio and network information.

The high-level architectural blueprint of Mobile-edge Computing, together with the scope of work that is described in this chapter, will form the basis of the first release of the work in the initiative. This chapter presents the abstract reference architecture relating to the MEC server; this server allows efficient and seamless integration of applications across multi-vendor platforms within the RAN.

The architecture includes components and functional elements that are key enablers for Mobile-edge Computing solutions in a multi-vendor environment. As enablers, they are capable of stimulating innovation and facilitating global market growth, while leaving room for differentiation and value creation.

5.1 MEC Platform

As depicted in Figure 9 below, the MEC server platform consists of a hosting infrastructure and an application platform.

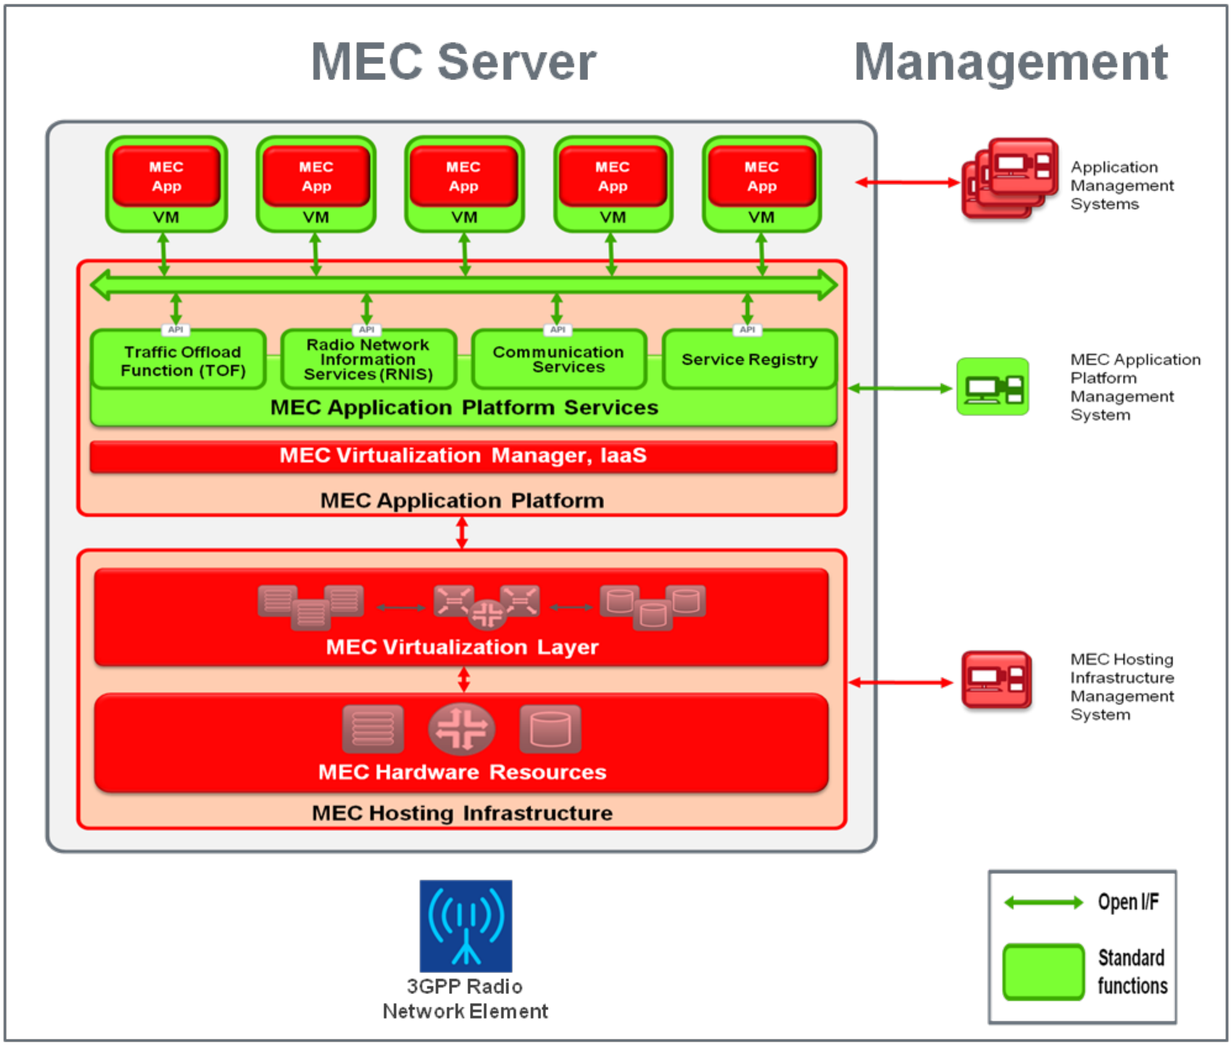

Figure 9: MEC server platform overview

The MEC hosting infrastructure consists of hardware resources and a virtualization layer. The details of the actual implementation of the MEC hosting infrastructure (including the actual hardware components) are abstracted from the applications being hosted on the platform.

The MEC hosting infrastructure, including the connectivity to the radio network element (eNB or RNC) and/or the network, is BEYOND the scope of the work of the MEC initiative. Multiple implementation options can be used to integrate the server within the RAN. Likewise, the interface towards the hosting infrastructure management system is BEYOND the scope of the work of the MEC initiative.

The MEC application platform provides the capabilities for hosting applications and consists of the application's virtualization manager and application platform services;

- The virtualization manager supports a flexible and efficient, multi-tenancy, run-time and hosting environment for applications by providing Infrastructure as a Service (IaaS) facilities. The IaaS controller provides a security and resource sandbox for the applications and the platform. Virtual-appliance applications run on top of an IaaS and are delivered as packaged-operating-system Virtual Machine (VM) images, allowing complete freedom of implementation.

  Note that in ***future releases*** of the initiative, Platform as a Service (PaaS) facilities may also be supported by the application platform services to allow application developers a different level of control over processing power, memory, storage and operating system support.

- The MEC application-platform services provide a set of middleware application services and infrastructure services to the applications hosted on the MEC platform. For details, see section 5.1.1 below.

MEC applications from vendors, service providers, and third-parties are deployed and executed within Virtual Machines. The MEC server and its services are application agnostic. The applications are managed by their related Application Management Systems which are application-specific components. Neither the applications nor their interfaces with the Application Management Systems are included in the scope of the work of the MEC initiative. The application management capabilities do not include application life-cycle management (e.g. start, stop, etc.), which is under the responsibility of the MEC application platform management system.

5.1.1 MEC application-platform services

The MEC application-platform services provide the following set of middleware services to the applications which are hosted on the MEC server:

- Infrastructure services:
  - Communication services;
  - Service registry;
- Radio Network Information Services (RNIS);
- Traffic Offload Function (TOF).

5.1.1.1 Infrastructure services

Communication between applications and services in the MEC server is designed according to the principles of Service-oriented Architecture (SOA).

The communication services allow applications hosted on a single MEC server to communicate with the application-platform services (through well-defined Application Programming Interfaces (APIs)) and with each other (through a service-specific API).

The service registry provides visibility of the services available on the MEC server. It uses the concept of loose coupling of services, providing flexibility in application deployment. In addition, the service registry presents service availability (status of the service) together with the related interfaces and versions. It is used by applications to discover and locate the end-points for the services they require, and to publish their own service end-point for other applications to use. The access to the service registry is controlled (authenticated and authorized).

For the communication services, a lightweight broker-based ‘publish and subscribe’ messaging protocol will be used. The ‘publish and subscribe’ capability provides one-to-many message distribution and application decoupling. Subscription and publishing by applications are access controlled (authenticated and authorized). The messaging transport should be agnostic to the content of the payload. Mechanisms should be provided to protect against malicious or misbehaving applications.

5.1.1.2 Radio Network Information Services (RNIS)

Mobile-edge Computing allows cloud application services to be hosted alongside mobile network elements and also facilitates leveraging of the available real-time network and radio information. The MEC RNIS provide authorized applications with low-level radio and network information. In the first release of the initiative, the information provided by RNIS can be used by the applications to calculate and present the following high-level and meaningful data: cell-ID, location of the subscriber, cell load and throughput guidance. Future releases of the initiative may provide additional information.

The RNIS deliver information from the radio network relating to users and cells.

The RNIS provide indications relating to the activation of User Equipment (UE) on a specific mobile network element. These include parameters on the UE context and the established E-UTRAN Radio Access Bearer (E-RAB), such as QoS, Cell ID for the E-RAB, identity of the UE-associated logical signaling connection, etc. The RNIS also indicate when a modification to an E-RAB or to a UE context occurs, or when an E-RAB is released. This information is based on 3GPP radio network-layer signaling messages (such as S1 Application Protocol (S1-AP), X2 Application Protocol (X2-AP) and Radio Resource Control (RRC)). There may be multiple implementation options regarding the way in which RNIS are exposed to the content arriving in the 3GPP signaling messages, but this is BEYOND the scope of the work of the MEC initiative.

RNIS also provide measurement and statistics information related to the user plane (at the cell and the E-RAB levels). This information is based on the 3GPP signaling messages and Performance Measurements (PM) defined by 3GPP and gathered by the mobile network element.

This initiative will produce gap analysis and may identify additional PM parameters that need to be standardized in 3GPP to provide greater value. Based on the results of the gap analysis, the MEC initiative will liaise with 3GPP and propose that these additional parameters be developed.

An authorized application can select the information to which it subscribes and derive some results based on suitable formulas. The application can use the results for its own purposes and/or publish them (partially or completely) for other applications.

5.1.1.3 Traffic Offload Function (TOF)

The TOF service prioritizes traffic and routes the selected, policy-based, user-data stream to and from applications that are authorized to receive the data.

The TOF service is supplied to applications in the following two ways:

- Pass-through mode where (uplink and/or downlink) U-plane traffic is passed to an application which can monitor, modify or shape it and then send it back to the original Packet Data Network (PDN) connection (3GPP bearer);
- End-point mode where the traffic is terminated by the application which acts as a server.

The traffic offloading policy sets filters at the E-RAB and the packet levels. The E-RAB policy filters are based on the Subscriber Profile ID (SPID), Quality Class Indicator (QCI) and Allocation Retention Priority (ARP). The packet filters are based on the 3-tuple (UE IP address, network IP address and IP protocol). In future releases of the initiative, additional filtering criteria may be supported.

The first release of the initiative does not handle service continuity during UE mobility; this is provided by the available end-to-end mechanisms.

5.1.2 MEC application platform management interface

The MEC application platform management interface is used by operators to manage the MEC application platform as well as the life cycle and operability of the applications and services which are hosted on the MEC platform. The management interface is application agnostic and supports the following functions:

- Configuration management of the application platform;
- Life cycle of the management application;

- VM Operations and Management (O&M).

Note that in addition to the MEC Application Platform Management System, there are application management systems which provide tools to manage the business logic of the respective applications. These are specific to each application and are BEYOND the scope of the work of the initiative.

5.1.2.1 Configuration management of the application platform

The configuration management of the application platform provides a standard way to package applications. It includes a descriptor of the application as well as system properties at the VM level relating to computing, storage and networking resources and configuration. The package also includes an access control mechanism which contains access policies to resources, services and applications for each virtual appliance.

The packaging design is aimed to address the portability and deployment of virtual appliances across multi-vendor MEC platforms, and to ensure software integrity protection.

In order to mitigate security threats, only authenticated and authorized applications can run on the MEC application platform. The MEC server must be able to verify that the application packages deployed in the server have a valid and trusted origin.

In addition, the MEC application platform management interface can be used by operators to configure policy for the applications. This includes traffic offloading and machine application policies, such as location in the service chain within the server when traffic is offloaded to multiple applications, recovery priority, etc.

5.1.2.2 Application life cycle

The MEC application platform management interface allows the network operator to manage the life cycle of the applications: deploy, start, stop and un-deploy.

5.1.2.3 VM O&M

The MEC application platform management interface supports fault management and performance measurements at the VM level.

5.2 Scope of the ISG MEC work

The architectural blueprint presented above forms the basis for the work in the ISG and the scope of the first release. The scope of the first release is restricted and includes the following elements:

- Platform services and APIs: communication services, service registry, RNIS, TOF;

- Virtual Machine (VM) SLA: description and negotiation of cloud services based on a standard unit of measurement; supervision mechanisms for the availability and the integrity of the VMs, platform services and APIs; Mechanisms for trouble shooting;
- MEC application platform management interface.

The ISG MEC will produce normative specifications for the requirements, framework and reference architecture, and specifications relating to the platform services and the APIs. The ISG will produce gap analysis to identify critical functional elements and techniques that need to be standardized to provide greater value.

As a general guideline, the ISG MEC will use and refer to existing specifications (both ETSI and external specifications) where appropriate.

Future releases of the ISG may include additional functionality, such as:

- Additional deployment scenarios;
- Integration with mobile core functions;
- Additional radio and network information exposure;
- Local mechanisms to handle service continuity during mobility of the UE;
- Orchestration functions with additional interfaces to the management entities;
- Support for PaaS;
- Support for network sharing;
- Etc.

6 Key enablers

In recent years, the telecommunications and IT industries have implemented several technological advancements that will serve as key enablers for the development and success of Mobile-edge Computing solutions. These key enabling technologies are presented in this section and their relevance is explained.

6.1 Cloud and virtualization

The separation of hardware and software and the enabling of horizontal, cloud-based solutions in the IT space have transformed the IT industry over the past decade. This transformation has fundamentally been made possible by the use of hypervisors which decouple the application and software environment in the Virtual Machine (VM) from the underlying hardware resources.

Multiple VMs can be deployed in a single platform and allowed to share the hardware resources in a controlled, efficient and flexible manner. Robust, effective and secured inter-VM communication is enabled by virtual switches. Traffic can be routed to a VM from a physical interface and subsequently routed from the VM back to the physical interface.

Cloud solutions make use of these technologies to provide computing and storage resources on-demand, introducing new levels of automation, flexibility and elasticity in network and service deployments as well as a faster innovation cycle for top-line growth.

Cloud and virtualization technologies have been leveraged by Telco Cloud and Network Functions Virtualization (NFV). They are currently transforming the communications industry in the way in which the IT industry was transformed over the past decade. These technologies are also key enablers for Mobile-edge Computing; they will allow applications to be deployed and run on top of the platform in a flexible, efficient and scalable way, independent of the lifecycle of the 3GPP network elements.

6.2 High-volume standard servers

High-volume IT hardware should be used to facilitate the commercial success of Mobile-edge Computing. The hardware platforms are built using mainstream and standardized IT components which can be provided in a competitive manner as key elements in the economy of scale of Mobile-edge Computing.

Standardized components can be changed within the server, enabling rapid and cost-effective maintenance and upgrades.

In recent years, general-purpose IT platforms are becoming increasingly adept at handling applications and services that consume vast amounts of hardware resources, such as packet processing. Ethernet controllers supporting 10 Gbps and even 40 Gbps are mainstream today, while the use of optimized drivers enables this high throughput, even in virtualized environments with general-purpose CPUs.

6.3 Enabling the application and service ecosystem

If the Mobile-edge Computing industry is to flourish and prosper, it is essential that software and application vendors develop and bring to market innovative and ground-breaking services and applications which can utilize Mobile-edge Computing functions and capabilities.

The use of open standards and APIs, as well as familiar programming models, a relevant tool chain and Software Development Kits (SDKs), are key pillars for encouraging and expediting the development of new cutting-edge applications or for adapting existing services and applications to the new Mobile-edge Computing environment.

To foster a healthy ecosystem of software vendors that thrives on developing applications for Mobile-edge Computing, it is imperative that support programs are put in place which will facilitate the development lifecycle of such applications and services. These support programs would enable the ISV community to develop and market advanced ideas, such as reference platforms, development tools or a laboratory infrastructure where newly developed applications can be optimized and verified in a test environment.

Finally, in addition to the support program mentioned above, once the necessary standards have been created and become mainstream, the industry would benefit from the establishment of a program aimed at ensuring that MEC applications not only comply with these standards, but that they are also capable of being adapted for use on platforms from different vendors. The creation of such a program is beyond the scope of this initiative.

7 Technical challenges and requirements

In order to promote and accelerate the advancement of Mobile-edge Computing, the community must overcome a diverse set of technical challenges. These challenges are described in the following subsections of this white paper.

7.1 Network integration

The introduction of a Mobile-edge Computing server is intended to be transparent to the 3GPP network architecture and the existing interfaces. User Equipment (UE) and Core Network elements that comply with the existing 3GPP specifications should not be affected by the presence of the MEC server as well as the applications being hosted on it. The 3GPP protocols and procedures should run and operate without affecting the SLAs.

7.2 Application portability

A fundamental requirement calls for applications to be seamlessly loaded and executed in the MEC platform that may be provided by different vendors.

This type of portability removes the need for dedicated development or integration efforts on a per-platform basis, which would unnecessarily encumber the software-application developers. It allows for the rapid transfer of applications (which may occur on the fly) between MEC servers, providing the freedom to optimize, without constraints, the location and required resources of the virtual appliances.

The precise but extensible definition of the services provided by the application platform is the key to ensuring application portability. The platform-management framework needs to be consistent across the different solutions to ensure that diverse management environments do not complicate the application developers' work on Mobile-edge Computing. The tools and mechanisms used to package, deploy and manage applications also need to be consistent across platforms and vendors. This will allow software application developers to ensure the seamless integration of their applications' management frameworks.

7.3 Security

The MEC platform poses a number of security challenges which, for the most part, stem from the introduction of IT applications into the telecom world; the environment and set of security requirements imposed by the radio network (resulting from both internal policies regarding operator security and regulatory bodies) is typically foreign to IT players. As a consequence, the MEC platform needs to simultaneously fulfil 3GPP-related security requirements while providing a secure sandbox for applications, i.e. isolating applications as much as possible from the burden of having to relate to all the implications of 3GPP security, operator security policies and local regulatory rules.

There are several mechanisms employed to provide such isolation / sandboxing:

- Ensuring isolation between Virtual Machines (and therefore, applications);
- Ensuring that the Virtual Machines only access platform resources and services for which they have authorization;
- Ensuring that the platform software and firmware, as well as the hosted application software, are not modified in any way by a malicious party;
- Ensuring that communication between applications, as well as between applications and the platform, is secured;
- Ensuring traffic isolation so that only the intended recipient(s) have access to traffic and data;
- Etc.

Another important aspect of MEC platform security is the difference between physical security constraints arising from deployment scenarios. With regard to the MEC platform deployed at the LTE macro base station, the available physical security is very poor in comparison with that at large data centers. This means that the MEC platform needs to be engineered in a way that will provide protection from both logical intrusions as well as physical intrusions. The MEC platform therefore needs to establish a trusted computing platform that is resilient to a multitude of attack vectors (including physical attacks).

In addition to trusting the platform itself, operators will be concerned about security implications with regard to the deployment of third-party software applications. In the particular case of Mobile-edge Computing, this concern may be exacerbated by the fact that these third-party software applications are deployed in very close proximity to their most valued asset, the Radio Access Network. The platform needs to provide the mechanisms to ensure that the virtual machines which contain the packaged application software (destined for deployment) come from a trusted source, are authenticated and authorized and are, therefore, safe to be incorporated in the platform.

The use of virtualization inherently creates a secure environment, since each virtual machine is isolated from all other machines; modern IT hardware platforms support several features which enforce this isolation. Thus, a software attack on a particular application is isolated from other applications running on different VMs. For this reason, it is critical that the hypervisor software is validated before it is launched.

The mechanisms for protecting these network interfaces will use other existing mechanisms (mutual authentication, integrity and confidentiality) which are robust and widely used in the industry.

7.4 Performance

Mobile network operators expect that Mobile-edge Computing solutions will have minimal impact on the KPIs (throughput, latency and packet loss to name a few) that they use to measure their networks' performance. Hence, Mobile-edge Computing platforms and the applications hosted on them need to be dimensioned and should have enough capacity to process the user traffic that is handled by the 3GPP Network Element with which they are paired. The mobile-edge applications shall be transparent to the UE and, at the same time, shall provide improved QoE.

Since the Mobile-edge Computing platform provides a virtualized environment where software applications are hosted, the main challenge facing the industry concerns the problem of extracting maximum performance while minimizing the impact of virtualization. The 'extra' layers introduced by virtualization should not degrade performance, particularly with regard to applications that require the intensive use of hardware resources (e.g. intensive I/O or computing) or require low latency.

The IT and cloud worlds have been dealing with these challenges for several years and solutions based on specific hardware-platform features that are widely supported by hypervisors have become mainstream. The communications industry, with Telco Cloud and NFV, is leveraging these platforms and hypervisor-technology advancements which maximize efficiency in virtualized environments. Mobile-edge Computing platforms will also make use of these innovations to ensure that hosted applications incur minimal overhead caused by virtualization, and that overall network performance is not impacted.

7.5 Resilience

The introduction of the Mobile-edge Computing platform should not affect network availability, hence MEC solution vendors should offer the level of resilience and address the high-availability requirements demanded by their network operators. In the event of a fault within the MEC platform that renders it inoperative, a failsafe mechanism is used to prevent it from adversely affecting the normal operation of the network.

In addition, the Mobile-edge Platform is designed to host applications that process user traffic; these hosted applications also need to be robust and resilient. To protect against any abnormality in the hosted software applications, the MEC Application Platform will have the necessary fault tolerance mechanisms to ensure that they operate without problems and within the established operational framework. If a fault is detected, or in the event that a hosted application is found to be operating outside the configured boundaries, the platform will ensure that corrective measures are put into place to prevent disruption to the user traffic being routed through the faulty or misbehaving application.

7.6 Operations

Mobile-edge Computing introduces three new management layers: one for the hosting infrastructure, one for the MEC application platform and one for the software applications and services which are hosted on the platform. The flexibility that virtualization and cloud technologies bring to Mobile-edge Computing enables different deployment scenarios where various organizations can be responsible for management at each level (e.g. the mobile operator manages the infrastructure and application platform layers while a third-party manages the application). Therefore, the implementation of the management framework should also consider the diversity of potential deployments.

In addition to this, the implemented management framework needs to take into account and complement the existing management framework of the radio access network. Moreover, it must ensure that operational and maintenance practices are not overly complex.

7.7 Regulatory and legal considerations

In the development of MEC platform regulatory and legal requirements will be taken into account. Potential examples include:

- Privacy - restricted information should not be passed to the application if the user hasn't given consent.
- Hosting an application on a network edge platform may provide certain advantages such as low latency for services like video streaming. Offering a type of "specialized service" that ensures sufficient quality of service for such applications to function, is appropriate and can be done consistently within the evolving principles of network neutrality. Additional technical analysis would be appropriate to determine the criteria to assess which applications and services would benefit from specialized treatment. Analysis could also examine how to configure networks to ensure sufficient capacity to accommodate demand for specialized services while maintaining suitable network conditions to support a robust user experience for non-specialized services. Transparency and non-discrimination within the specialized services category, and among all non-specialized traffic, could be observed as part of a net neutrality framework that would allow end users or the owner of the content, application, or service to pay for specialized treatment subject to reasonable consumer protections important to network neutrality.
- The introduction of a MEC server in the radio access network should not reduce the provision of lawful interception.
- Charging requirements for access to services during roaming scenarios. For example, in the case that the same services are available in the home network and visited access network, should

different charging requirements apply? Ongoing regulatory changes to roaming charges would need to be considered.

8 Relations and interfaces with other industry efforts (within and outside ETSI)

The MEC specification will refer to existing specifications where appropriate, hence the ETSI MEC ISG will establish liaisons with other groups within ETSI and with other Standards Definition Organizations (SDOs). If the ISG identifies additional, critical functional elements and techniques that need to be standardized to provide greater value, it will liaise with the relevant organizations and propose that they develop these standards. The MEC ISG will consider establishing a liaison relationship with the following ETSI Technical Bodies (TBs) and Partnership Project:

- **EP E2NA/TC NTECH**, which works on end-to-end network and service architecture;
- **ETSI ISG NFV**, which defines the requirements and architecture for the virtualization of network functions;
- **The 3rd Generation Partnership Project (3GPP)**, which covers cellular telecommunication network technologies, including radio access, the core-transport network and service capabilities. The work includes provisioning and management of the network and its services.

Depending on the way in which the work progresses, the ISG MEC may establish a liaison relationship with the following organizations:

- **Distributed Management Task Force (DMTF)**, which creates standards to enable interoperable cloud management (including elements of portability regarding appliance virtualization, service-level assurance, service lifecycle, etc.);
- **Cloud Security Alliance (CSA)**, which promotes the use of best practices for providing security assurance within Cloud Computing;
- **TM Forum**, which works on a common cloud framework and key service enablers (including Quality of Experience);
- **Organization for the Advancement of Structured Information Standards (OASIS)**, which, among other things, produces standards to enhance the portability and manageability of cloud applications, and the IT services comprising them, which run on a complex software and hardware infrastructure.

Liaison with additional organizations may become relevant in the future.

It is also expected that additional communication channels will need to be established with certain open-source software communities, such as OpenStack (which develops the cloud operating system).

9 Call for active participation

Mobile-edge Computing allows applications from different providers to run at the edge of the mobile network. The business and technical benefits it introduces are described in the preceding sections of this white paper.

Mobile-edge Computing will turn the mobile base station into a versatile computing platform, allowing the rapid deployment of new innovative services and the optimization of service delivery. Before this vision can become a reality in today's multi-vendor world, interoperability must be achieved so that applications from different providers can run on the mobile-edge platforms offered by the various vendors. This whitepaper introduces an architecture designed to implement this interoperability and lists the challenges that need to be overcome.

It is proposed that a new ISG be set up under the auspices of ETSI which will be tasked with developing the technical specifications for Mobile-edge Computing. These specifications will enable the implementation of the interoperability defined by the architecture and included in the scope of this whitepaper.

The different players in the value chain (mobile network operators, telecom equipment vendors, IT platform vendors, and potential mobile-edge service and application providers) are invited to actively participate in the ISG and to contribute to the development of the specifications based on industry consensus. The players are also requested to assist with the dissemination of the deliverables produced by the ISG in order to develop favorable market conditions for sustainable business for all players in the value chain. The participating players are encouraged to share best practices and demonstrate Proofs of Concepts (PoCs).

Companies requiring further information about the planned ISG and those interested in participating in discussions are welcome to contact the authors of this white paper (for the contacts, see section 11 below).

10 Conclusion

Mobile-edge Computing transforms base stations into intelligent service hubs that are capable of delivering highly personalized services directly from the very edge of the network while providing the best possible performance in mobile networks. Proximity, context, agility and speed can be translated into unique value and revenue generation, and can be exploited by operators and application service providers to create a new value chain.

Creating a new value chain and a refreshed ecosystem (based on innovation and business value) allows all players to benefit from greater cooperation. Mobile operators can play a pivotal role within the new value chain and attract OTTs, developers and Internet players to innovate over a new cutting-edge technology, while enabling context-aware applications to run in close proximity to the mobile subscriber. Mobile subscribers can enjoy a unique, truly gratifying and personalized mobile-broadband experience which is tailored to their needs and preferences.

This paper presents the concept of Mobile-edge Computing as well as the business and technical advantages that it offers. It explains how Mobile-edge Computing can help operators to meet the challenges posed by traffic explosion and increased user demands. Several use cases are presented that can benefit from Mobile-edge Computing.

The paper highlights the importance of promoting a healthy ecosystem in addition to the relevant toolkits for developers, all of which are aimed at driving the invention of new services and applications. In addition, the quest to create new revenue streams for all players in the MEC ecosystem, including operators, OTT players, vendors, etc., is discussed.

The motives behind the MEC initiative are described. Operators demand standards and seek an interoperable, multi-vendor approach. Moreover, the value of applications can only be appreciated if they cover the vast majority of the population which is served by a single operator. The applications must be deployed on top of Mobile-edge Computing servers originating from different vendors.

This paper introduces the high-level architectural blueprint that provides a framework for interoperability; this architecture will form the basis for the work in the initiative. The initiative will work to specify the MEC application-platform APIs which will be supported by multiple vendors and adopted by multiple operators. The open APIs will give the players in the value chain the opportunity to revolutionize, differentiate and create value in a multi-vendor mobile environment.

Various deployment scenarios are presented and a number of enabling technologies are considered. A number of technical challenges are described, including considerations regarding security, application portability, network integration, performance and resilience.

It is proposed to set up the initiative under the auspices of ETSI and to create a new ISG MEC. The architectural blueprint for Mobile-edge Computing, together with the scope of work that is described in this paper, will form the basis of the first release of the work in the initiative.

The relationships and the interfaces of the ISG MEC with other industry efforts (within and outside of ETSI) are highlighted. The MEC work is intended to complement and build upon the concepts defined in the ETSI ISG NFV.

It is also recognized that MEC will operate within the current regulatory frameworks and, if necessary, identify gaps where further regulatory considerations may be required. The work in the initiative will help to develop favourable market conditions for sustainable business for all players in the value chain.

The market players are invited to actively participate and contribute to the work in order to make Mobile-edge Computing a successful and developing part of the overall mobile broadband ecosystem.

The next steps for the MEC initiative are as follows:

- Establish the ETSI ISG MEC and approve its Terms of Reference;
- Promote the ETSI ISG MEC with the aim of building an ecosystem consisting of players representing the entire value chain;
- Consider the activities that fall outside the scope of specification development, including activities that demonstrate proof of concepts, market acceleration and further promotion of MEC towards network operators and applications developers.

11 Contact information

Huawei: Milan Patel, milan.Patel@huawei.com

IBM: Brian Naughton, naughton@uk.ibm.com

Intel: Caroline Chan, caroline.chan@intel.com

Nokia Networks: Nurit Sprecher, nurit.sprecher@nsn.com

NTT DOCOMO: Sadayuki Abeta, abeta@nttdocomo.com

Vodafone: Adrian Neal, adrian.Neal@vodafone.com