White Paper

Mobile Edge Computing Use Cases & Deployment Options

Prepared by

Gabriel Brown
Senior Analyst, Heavy Reading
www.heavyreading.com

on behalf of

Juniper Networks

www.juniper.net

July 2016

What Is MEC? Where Is the Value?

Mobile edge computing (MEC) aims to place compute and storage resources in the 4G radio access network (RAN) to improve the delivery of content and applications to end users. The technology enables operators to better adapt traffic to the prevailing radio conditions, optimize service quality and improve network efficiency. This white paper discusses the benefits of edge computing in mobile networks and investigate the use cases that are deployable in the nearer term. It provides an overview of the MEC architecture and deployment options, and evaluates the longer-term role of MEC in evolved 4G and 5G contexts.

Services in the RAN: State of the Market

MEC enables the deployment of services within the RAN to improve user experience and provides a capability that enables operators to better handle latency-sensitive services. To improve network efficiency, it can help adapt service delivery according to the load on the radio link and can reduce the need for long distance backhaul through the use of local content caching and local Internet breakout.

The MEC concept originated soon after 3G smartphones became popular with the idea to develop mobile content delivery networks (CDNs) specifically for mobile networks (mobile CDNs). Some operators and vendors took the concept a step further with "pre-MEC" solutions that added compute capability to 4G base stations. This option to deploy services at the base station is still present in MEC and may yet become important; however, other options, such as deployment at aggregation points in the IP transport network, are now generally considered more cost-effective. Most recently, the MEC concept has evolved to draw on network functions virtualization (NFV) technologies to allow virtual network functions (VNFs) to run on this distributed MEC platform. This inserts MEC into a broader, more strategic discussion about network architecture evolution and distributed cloud in the journey to 5G.

MEC, in its current guise, was approved as a formal Industry Specification Group (ISG) by the European Telecommunications Standards Institute (ETSI) in 2014. The ISG will specify a functional architecture around the mobile edge platform function and its interactions with MEC applications in a multi-vendor environment. It has already created a document on service scenarios and initiated a proof-of-concept (PoC) program to inform specification development and drive market interest. "Foundation specifications" were released in April 2016, with normative specifications expected by the end of the year.

To align with broader developments in service provider networking, and to increase the strategic value of MEC, the ISG has adopted the ETSI management and orchestration (MANO) model for service orchestration and management. Over time, this means that a MEC hosting location will, in effect, be a distributed "micro" data center capable of running generic NFV and cloud services under one operational model. The group now counts over 40 member or participant companies and is backed by high-profile operators such as Docomo, Vodafone, TIM, SK Telecom, EE, Telefónica and Orange. Other large operators with global influence, such as Verizon, AT&T and China Mobile, are also interested in MEC and its relationship to distributed cloud.

Strategic Value of Distributed Cloud

One of the most interesting aspects of MEC is its influence on future network architectures – notably in distributed cloud and to meet the needs of low-latency 5G services.

Distributed cloud infrastructure, in combination with software-defined networking (SDN), enables operators to take advantage of physical assets, compute and storage located close to customers.

An important aspect of MEC deployment is to balance low latency and mobility. This is important as operators start to distribute mobile packet core gateways (PGWs) for reasons of performance and economics. The presence of MEC can reduce the number distributed PGWs needed, thus reducing inter-PGW handoffs. With the correct deployment of MEC nodes in the RAN – typically at an IP aggregation site – operators are able to reduce excessive application-layer handovers and meet the performance targets needed for latency-critical applications.

The new 5G system architecture, and associated new radio (NR) and next-generation core (NG Core), is under development in 3GPP, with deployments expected in the next three to five years. To meet high bandwidth and low-latency requirements with this architecture, and to enable innovative use cases at the fixed and mobile edge, 5G will require one-hop access to the cloud. Today's 4G operators looking to deploy virtualized "cloud native" core networks can therefore incorporate MEC technology to future-proof their architecture to be "5G-ready."

MEC High-Level Architecture

MEC is specified to provide a multi-tenant hosting environment for 4G RAN edge applications. The hosting environment consists of hardware resources and virtualization infrastructure (virtual compute, storage and network resources) and associated management services for MEC applications.

Figure 1: MEC Logical Architecture in Mobile Operator Network

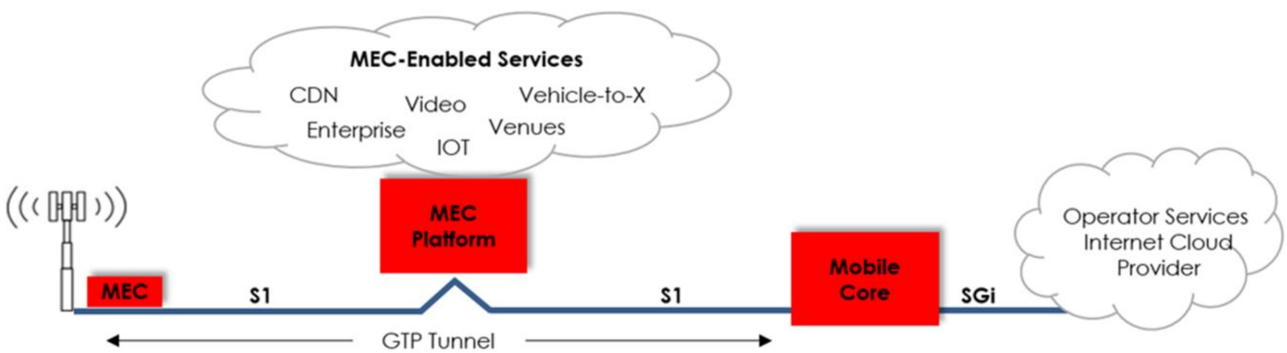

Source: Heavy Reading

In terms of logical network architecture, ETSI specifies that MEC can be a part of the eNodeB or be run on an external server that can be deployed between the radio base station and the mobile core on the S1 interface, as shown in **Figure 1**. The control plane over S1-MME is SCTP-based, and the user plane over S1-U is GTP-based. S1 traffic, in some cases, is also encapsulated in IPsec over the IP backhaul network. The MEC platform should, therefore, be deployed inline and be transparent to GTP, such that it should open the tunnel and re-encapsulate traffic, once services have been applied via locally-hosted MEC applications, without impacting the core network. In cases where S1 traffic is encapsulated in IPsec, the MEC platform must additionally support distributed security gateway functions.

The major components of the MEC architecture are shown in **Figure 2** below and can be described as follows:

- The **Mobile Edge Platform** sets the policy and configuration rules for forwarding user plane traffic to MEC applications. It also provides a set of services that expose radio network data and other real-time context information to authorized MEC applications.
- The **Mobile Edge Orchestrator** maintains an overall view of the deployed MEC Servers to determine the optimum location(s) for instantiating a MEC application.
- The **Mobile Edge Platform Manager** is responsible for lifecycle management of the MEC applications and management of the MEC Application Platform.
- The **Virtualized Infrastructure Manager** is responsible for managing the resources of the virtualized infrastructure, which also includes preparing the infrastructure for running a software image.
- **Mobile Edge Applications** – MEC-enabled services provided independently of the platform are what make MEC valuable. They must use the MEC application programming interfaces (APIs) (definition in progress) and be manageable within the NFV framework.

Figure 2: MEC Architecture

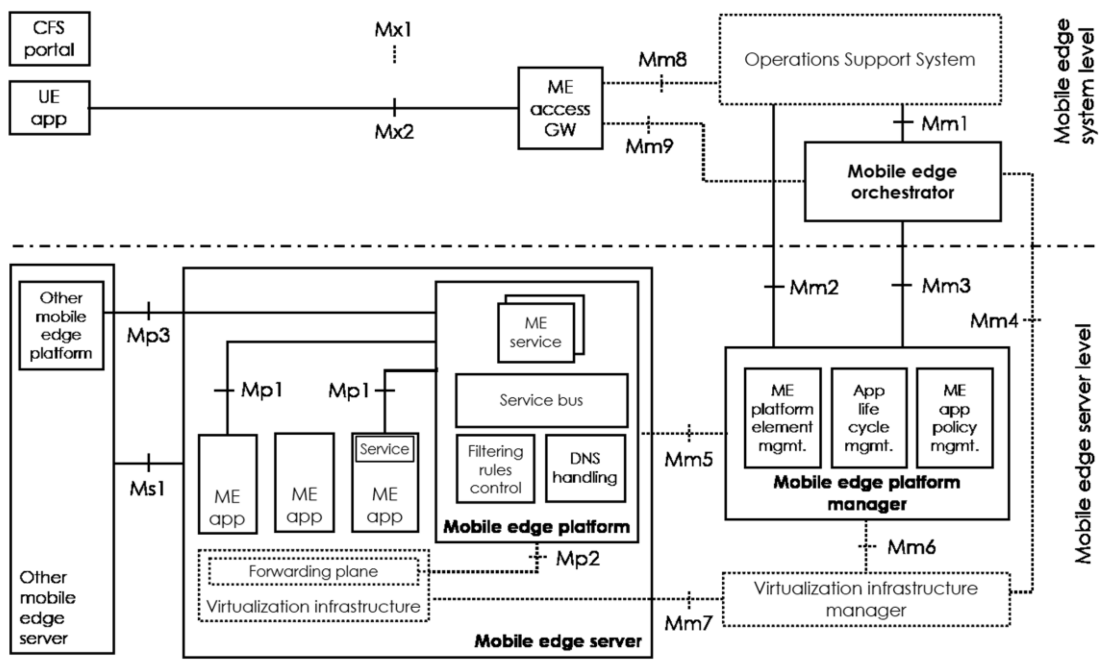

Source: ETSI MEC ISG, July 2016

Adopting the ETSI MANO model makes the MEC infrastructure an open environment that will allow the operators to deploy applications and services from third-party vendors. This makes interoperability important and, because it is a distributed deployment, places a great emphasis on low cost of operations. A standardized NFV Infrastructure (NFVI) platform connected by an SDN-enabled network is important to this realizing this goal.

Leading MEC Use Cases

There are many services that could benefit from being hosted in the distributed cloud close to customers. In combination with visibility into prevailing radio conditions via MEC's Radio Network Information Service (RNIS) function, applications can adapt content delivery in real time to ensure a consistent quality of experience (QoE) to the end user. Some of the leading use cases are discussed in this section.

RAN-Aware Video Optimization

Video is currently around half of mobile network traffic and set to exceed 70 percent of traffic over the next couple of years. Providing throughput guidance information to a video server is one of the classic MEC use cases.

In a mobile network, available capacity can vary by an order of magnitude within seconds as a result of changes in radio channel conditions. Transmission Control Protocol (TCP) may not be able to adapt fast enough to rapidly varying conditions, leading to underutilization of radio resources and/or to a sub-optimal user experience, such as buffering. In this use case, the proposed solution is to use MEC technology to inform the video server of the optimal bit rate to use given the radio conditions for a particular video stream or user. This is shown in **Figure 3**.

Figure 3: RAN-Aware Video Optimization Using MEC

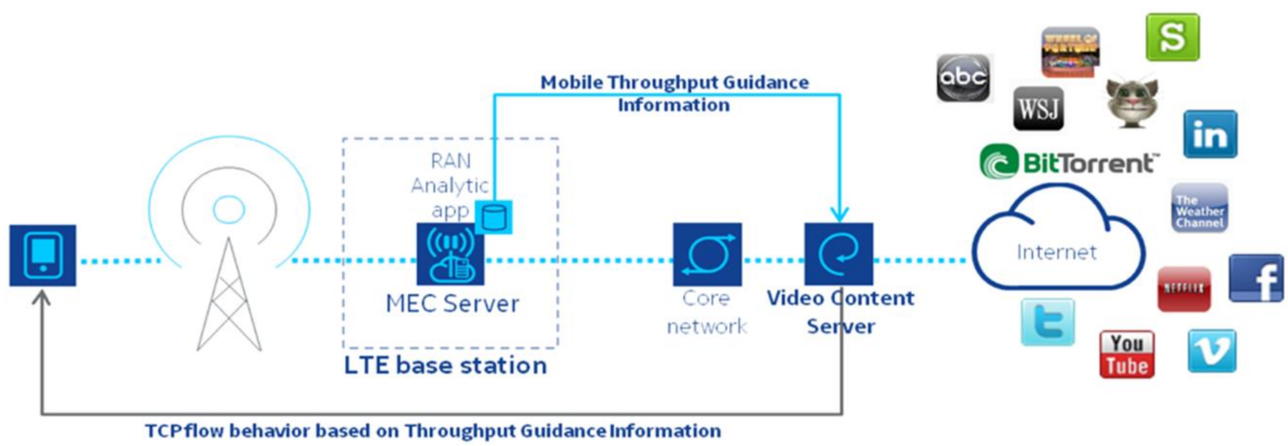

Source: ETSI

The idea is to use a RAN analytics application to determine/estimate the throughput likely to be available at the radio downlink interface for a user, and then use packet headers to convey that information to the video server, so that it can adapt the stream accordingly. Field trials have shown noticeable performance improvements when operators communicate RAN conditions to the video server in this way.

Stadiums & Public Venues – Edge Video Orchestration

Large public venues are good candidates for MEC – especially where localized venue services are important, such as stadiums. In this use case, video created at a sports event or concert is served to on-site consumers from a MEC server, running appropriate applications located on the stadium premises. This video traffic is locally stored, processed and delivered directly to users at the event and does not require backhaul to a centralized core network and to then be returned to the user at the venue. This model potentially enables unique stadium services, such as the ability to view in-game "player cams" that enrich the fan experience or can support fan and venue services, such as food ordering, ticketing queue management, etc.

This use case is the subject of an approved ETSI PoC and is being investigated by the largest U.K. mobile operator, EE, in association with several vendors. EE is the sponsor of the U.K. national football (soccer) stadium, which helped drive enthusiasm for the trial and to sharpen the commercial focus. Typically, this type of service is linked with a dedicated RAN deployment at the venue such that MEC servers can be co-located with RAN controllers and backhaul equipment.

Enterprise & Campus Networks

In large enterprise organizations, where premises-specific services are used, there is a desire to process users locally rather than backhaul traffic to the centralized mobile core and only to send it back again. This could be for services as simple as access to the corporate Intranet – for example, to offer a 4K corporate training video to a mass of employees at the same time would be difficult from the macro network – or more advanced services, such as security policy, location tracking and asset tracking services that can run as local MEC applications.

This use case is associated with the deployment of enterprise small cell networks because organizations that would make use of MEC on-premises would also typically have dedicated 4G networks deployed by an operator. This same operator may also provide enterprise virtual private network (VPN) services and an interesting deployment option is to run the MEC platform and services on the same customer premises equipment (CPE) device used to deliver the VPN and associated services.

New-generation CPE devices offer a compute and virtualization environment that can support local deployment of multiple VNFs, including MEC platform services and applications. There is a potential opportunity, therefore, for MEC and VPN services to be orchestrated by the same integrated service provider. Manufacturing facilities, hospitals, warehouses or even large outdoor facilities, such as mines or campuses, are examples of sectors where "private" 4G RANs are being deployed, and where applications related to voice over IP (VoIP), video, Internet of Things (IoT), etc., could run on the local MEC server.

Other MEC Services – IoT, Augmented Reality, Vehicle-to-X

MEC is a platform technology and is not prescriptive about services. There are many potential applications that can run on this infrastructure. Some interesting ones that are perhaps some time from being deployed commercially include:

- **Smart Building Services:** Much of the data generated in smart buildings is inherently local and involves device to device communication that could be processed by a locally-hosted IoT gateway. Examples of services that can benefit from local computing and control capability include security, tracking, climate control, smart signage, entry-control, etc.
- **Internet of Things (IoT):** MEC can be used to process and aggregate the small packets generated by IoT services before they reach the core network. This will be important for scalability as the number of IoT connections

increase and may be crucial for battery-powered IoT devices. A shorter transmission time between device and application server reduces drain on the battery and, therefore, can increase the life of the device and improve the business case for the service.

- **Vehicle Services:** Vehicle-to-infrastructure (V2I) use cases require local processing and low latency, which can be offered by MEC. As such, they are also an example of how MEC concepts are expected to be important to 5G services. This may extend to the ultra-low latency requirements of autonomous vehicles and other mission critical applications of the future.
- **Augmented Reality:** To overlay information from the phone camera (or, perhaps, eyewear in the future), requires localized content to be rendered very quickly on the viewing surface. Ergo, it would benefit from local processing and is an ideal candidate for MEC.

MEC Deployment Options

MEC infrastructure elements can be deployed at multiple places in the network, often determined by the use case – for example, public venues, such as stadiums or airports, may favor an on-premises deployment. In other cases, the operator must decide on the MEC hosting location based on its existing physical footprint and on its operations model. Selecting the correct location is critical because capabilities deployed at the edge may have a higher operational overhead than centralized deployments.

MEC Deployment at Aggregation Sites

An important factor is the "hit ratio" or, when in the distributed content use case, the "cache ratio." This refers to the point in the network where the MEC infrastructure can serve a sufficient number of users to make the deployment worthwhile. There may not be sufficient users of a service at single cell site, for example, to justify the investment in MEC infrastructure and the operator may need to look deeper in the network where multiple cell sites are aggregated for the optimal location. This is shown in **Figure 4**.

Figure 4: MEC Deployment in Service Provider Network

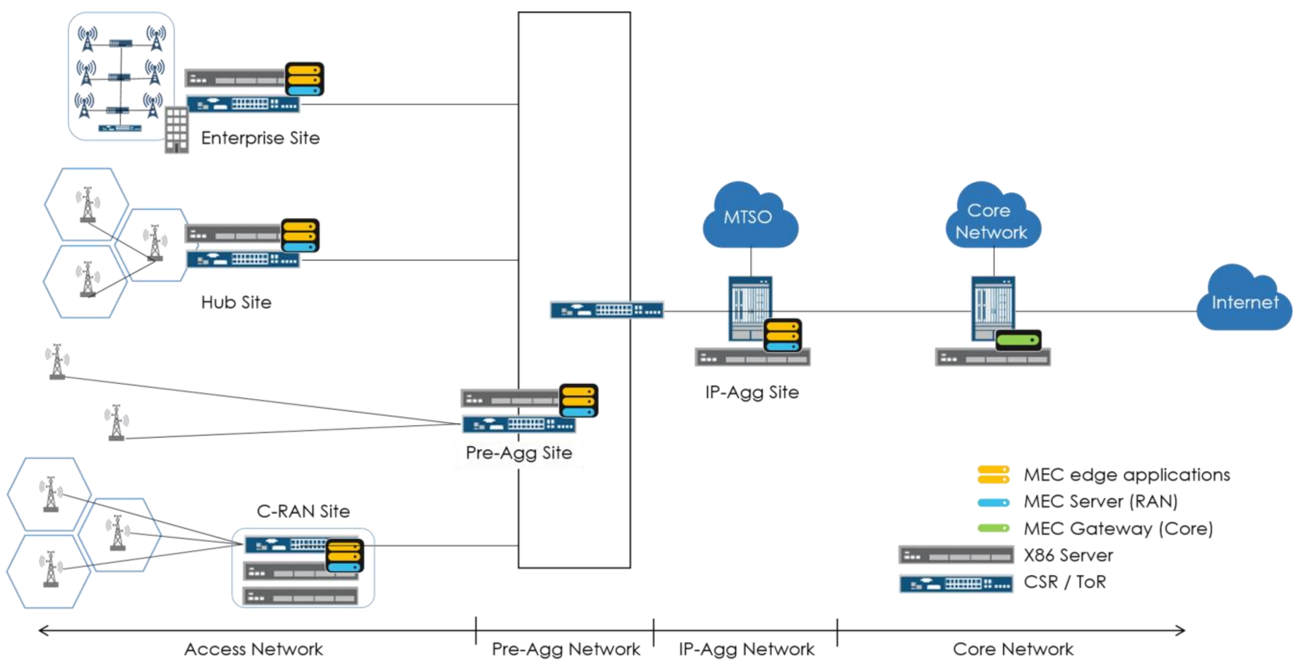

Source: Juniper Networks

In the macro cellular network, good candidates for MEC deployment are pre-aggregation or hub sites, which concentrate approximately 20-30 cell sites and enable a round trip latency to and from the base station of approximately 3-5 ms. Larger aggregation sites, which concentrate 100 or more cell sites, are also attractive, and even though this may result in higher latency, it would still be still lower than for services deployed "behind" the centralized core.

In the future, cloud RAN (C-RAN) hub sites that aggregate tens or hundreds of cell site equivalents may be good locations for MEC services. These C-RAN hubs are a form of distributed cloud and with virtual RAN (also deployed at the C-RAN site) there is an even greater opportunity to match application delivery to the prevailing radio conditions for a particular user or cell.

Mobile Edge or Simply Edge Computing?

Several operators interested in MEC have highlighted that the mobile network edge, is a subset of the converged network edge, and that a solution for edge computing across access networks is needed. There are important differences between the fixed and mobile edge. MEC is needed in mobile access, where GTP encapsulation is prevalent; however, standard IP-based edge compute is possible in other access networks that don't need mobility and GTP. The opportunity, therefore, is to combine these environments into a common, orchestrated SDN/NFV edge.

A simple example may be a user watching a video on cellular at one moment and over WiFi the next moment. This scenario illustrates the need for the edge solution to accommodate MEC for mobile access and generic IP edge for when the user "hands-off" the video service to WiFi. Naturally, converged operators wanting to offer a common service portfolio also want to pursue MEC within the context of generic "edge computing" across fixed and mobile access.

The concept is to build distributed cloud infrastructure to run VNFs to support end-user services (such as enabled by MEC), and to offer platform services for third parties. By turning distributed real estate into micro-data centers, operators are able to reduce end-to-end latencies, improve resiliency and optimize utilization. Access network operators have a physical footprint advantage at these types of locations and want to use this to their advantage.

This model fits well with emerging trends in mobile network architecture. Most operators today deploy highly centralized mobile packet core and SGi service LANs; however, to meet demand for greater data volumes and lower latency, many are looking to convert existing edge sites (such as used for 3G RNC or MSC technology or backhaul aggregation) into micro data centers, to enable them to deploy distributed, virtualized EPC networks, perhaps based on a control/user-plane separation ("CUPS") architecture. In essence, this moves a pure IP core closer to the customer.

In this future network, MEC can be deployed within one hop of a group of base stations with no impact on application layer mobility and remains transparent to the packet core. MEC and the SDN-based IP edge are, in this model, complementary and offer the operator great flexibility to scale and provide optimal service quality.

Summary

MEC is an emerging technology that enables mobile operators to host content and applications in the 4G radio access. The creation of the MEC Industry Specification Group by ETSI to specify a multi-vendor platform and service orchestration environment has given renewed impetus to edge computing for mobile operators. The lead deployments are likely to be where it is practical to deploy MEC in a manner transparent to the mobile core and where there are clear customer and operator benefits, which will likely be in enterprises, campuses, large public venues and mobile backhaul pre-aggregation sites, where the concentration of base station traffic and low latencies allow optimal MEC deployment.

About Juniper Networks

Juniper Networks (NYSE: JNPR) is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net or connect with Juniper on Twitter and Facebook.