Driving New Business Opportunities with Multi-Access Edge Computing and 5G

Publication Date: 15 11 2019

Author: Julian Bright

Summary

In brief

Multi-access edge computing (MEC) can be expected to play a pivotal role in fulfilling a range of use cases projected for 5G. The prospects for service innovation and new market opportunities such as business to business (B2B), business to consumer (B2C), and business to home (B2H), driven by MEC in conjunction with 5G, are extremely positive, thanks to features such as ultra-low-latency connectivity, real-time processing and management of data, and localized content caching. Hence, the technical features, commercial prospects, and implementation strategies for MEC in 5G networks, including developing the 5G MEC architecture and key technologies associated with 5G and MEC, are the subject of intense discussion.

Ovum view

Ovum is optimistic about the prospects for edge computing as a technology capable of being widely deployed in communications service provider (CSP) networks, despite the lack of commercial MEC-based services thus far.

In order to achieve the service-related goals of 5G and deliver on the business case, computing and connectivity need to converge, and MEC provides the best available technical solution if network operators are to move beyond offering an enhanced mobile broadband service and begin to offer new services and penetrate new markets such as the enterprise.

As CSPs move toward a software-based cloud-network architecture in preparation for 5G, the case for the development of an edge-computing model will become stronger. The combination of network functions virtualization (NFV) and telco cloud is providing CSPs with the foundation of a platform for MEC, given additional support by the move from today's virtual machine (VM) model to a cloud-native, container-based reference architecture.

In future, as 5G technology and networks evolve, MEC has the potential to provide a high-performance, on-demand, and cost-effective platform capable of supporting a growing number of use cases.

Key messages

- The founding principle of MEC is that connectivity and computing can work together to support a growing number of applications and services at the edge.
- Coupled with 5G networks, MEC can deliver high levels of service performance based on an open platform with low-latency, high-throughput connectivity to enable new business scenarios.
- Because the 5G core network is the network engine that ensures a guaranteed service experience, extending it to the edge is key for MEC service enablement.
- Service provider networks are already evolving toward a 5G-oriented, virtualized, cloud-based model that can be the basis of future MEC deployments.

- Cost-efficient MEC networks will require heterogeneous, on-demand, high-performance application lifecycle management integrated into a network functions virtualization infrastructure (NFVi) platform that supports both VMs and containers.
- Edge computing is attracting interest from a broad ecosystem of potential participants, but CSPs are well placed to play a key role.

Why 5G MEC is important

Edge-computing progress update

Edge computing represents the convergence of connectivity and computing at the network edge. In contrast to today's centralized architectures, where the processing and storage of data are concentrated at the heart of the network and access is governed and controlled from the core, MEC employs a distributed cloud-computing model that pushes these functions out to the network edge, where applications are hosted in highly distributed, generally small-scale and localized edge data centers closer to the device and to the end user.

Edge computing can be deployed in wireless or wireline network environments and in a wide range of use cases wherever low latency or the localized processing, handling, and management of large volumes of data are critical to the delivery of high-quality services, and it can also be used as a platform for hosting third-party applications and services.

Taking computing and core network functions to the edge will build an optimized and cost-efficient MEC network architecture

A MEC network is designed to handle real-time data processing and management, ultra-low-latency connectivity, and localized content caching in a highly cost-effective and efficient manner. The edge-network nodes or micro data centers are equipped with their own compute, processing, storage, and management capabilities sufficient to service one or many applications on a localized basis (see Figure 1).

Figure 1: MEC framework

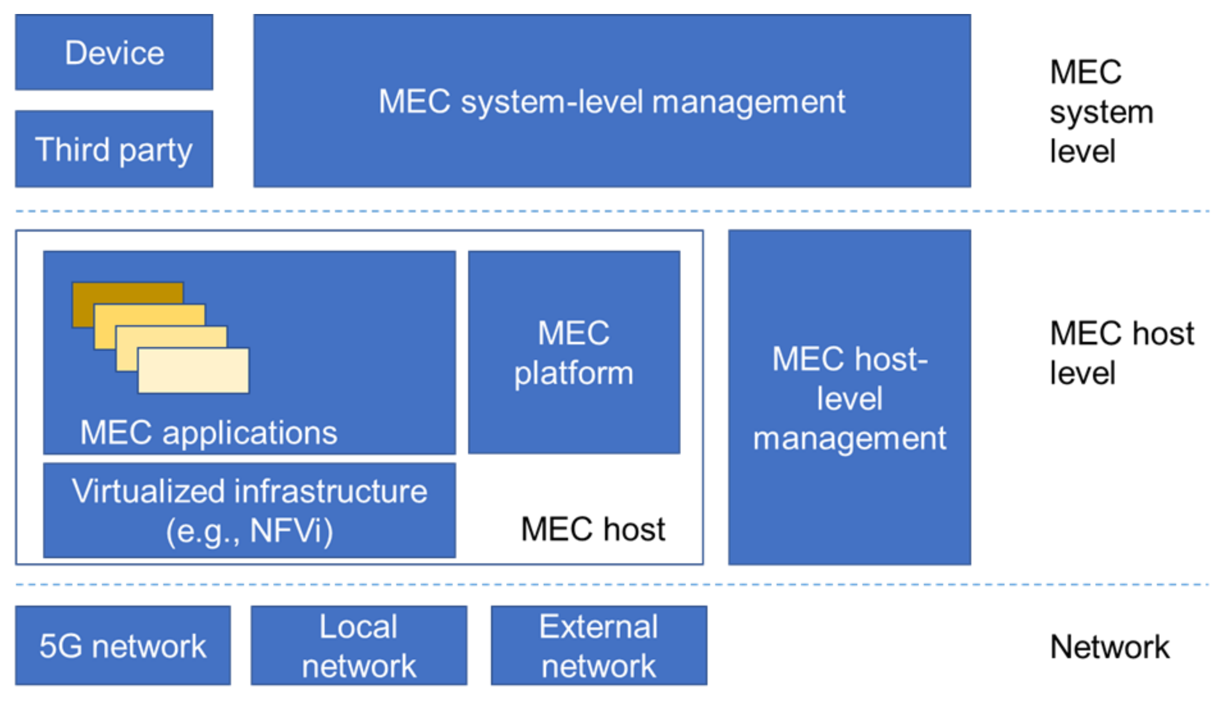

Source: Ovum/ETSI

Connection back to the central cloud is limited as far as possible to those management, orchestration, and application support functions necessary to service the wider edge network and to update the edge nodes. Thus the flow of traffic between the core and the edge network is minimized in order to reduce latency, cut loading on the core network, and save on transport costs such as backhauling.

CSPs can take the key role, while campus networks will lead the first wave of 5G MEC deployments

With the commercial rollout of 5G, CSPs are developing the capability to deliver unprecedented network speeds, greater capacity, and massively reduced latency. In order to better exploit and begin to monetize these new resources, many are already looking beyond 5G's early role as an enhanced mobile broadband technology and planning their 5G business cases around a range of B2B and B2C applications and services that include the control and monitoring of industrial machinery and processes (Industry 4.0), connected and autonomous vehicles, augmented and virtual reality (AR/VR), high-quality video, ultra-low-latency gaming, and a range of services for enterprise and industry verticals. Campus networks are expected to be among the first wave of 5G MEC deployments (see Figure 2).

Figure 2: Edge-computing use-case examples

Industrial IoT including Industry 4.0	Connected and autonomous vehicles	Augmented and virtual reality (AR/VR)	Ultra-low-latency gaming	Large-scale campus networks

Source: Ovum

The combination of network virtualization and telco cloud is already providing CSPs with the foundation for a new, more efficient and flexible network paradigm more akin to that created by the public cloud providers and internet players but with the added advantage of a high-performance connectivity layer connecting directly with users and devices. At the same time, the move to a cloud-native container-based reference architecture is building greater flexibility into the network and providing the capability to distribute resources closer to where they are needed.

By applying MEC's innovative network architecture and extending their newly acquired cloud-layer resources out to the network edge, CSPs can be well placed to exploit the opportunities for new services and revenue generation around these use cases. In a recent Heavy Reading survey, almost 85% of network operator professionals polled believed that edge computing would be critical or important to their network evolution strategy (see Figure 3).

Figure 3: How vital is edge computing to your network evolution strategy?

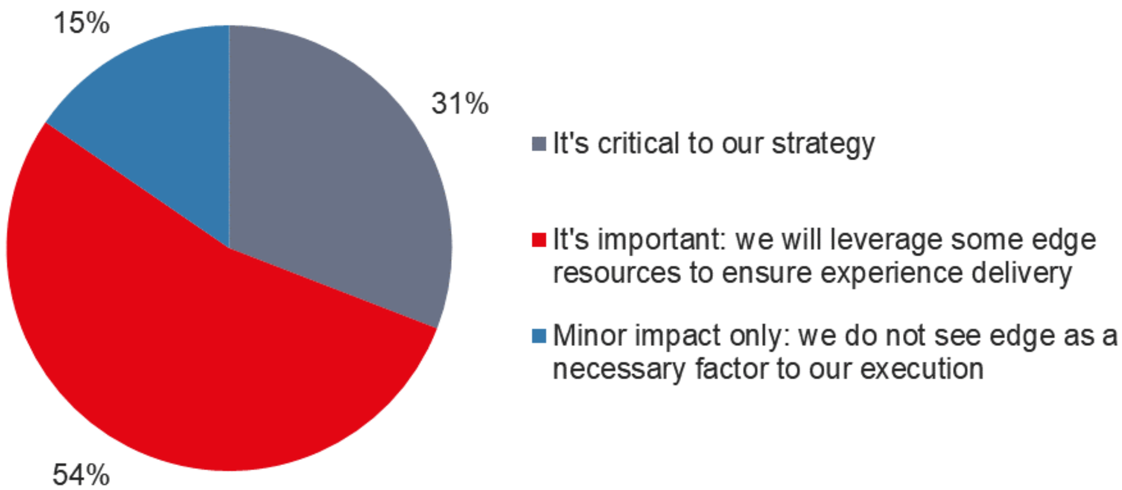

Source: Heavy Reading

Obstacles to MEC deployment

Above all, MEC needs to be as much about revenue generation as cost reduction if it is to be a commercial success. The challenges involved in implementing MEC are, therefore, not purely about the technology; in fact, many vendors argue that, technically, deploying edge computing is relatively straightforward. The problem rests more with concerns about the business drivers, which in turn depend on identifying the most promising use cases that MEC can support.

In this regard, there is still uncertainty about how and where operators should start to deploy their edge-cloud infrastructure. This is likely to be governed by a number of factors including which services and applications will require key MEC features such as ultra-low latency and how close to the end user or device the edge-computing resources need to be located.

In use cases such as VR or autonomous vehicle, where low latency can be critical, proximity to the point of service delivery might be the overriding consideration. In such instances, the base station could prove the optimum location, or an aggregation point could be provided for a network of more highly distributed, ultra-small-scale service nodes. In some instances, such as data capture for autonomous driving, the MEC node might even be located inside the vehicle.

In those use cases where a higher level of latency can be tolerated, but where the demands of local content caching or data processing require more capacity, a location further back in the network such as at a central office or aggregation point might be more suitable. In an enterprise deployment, the "edge" might be an Internet of Things (IoT) device or a gateway located at the customer premises.

Whatever the decision, service providers will need to have available or be able to free up the space required to locate, power, maintain, and service their edge-cloud resources.

Who will own the edge?

Though arguably best placed to build and run the edge cloud, CSPs may not have the field to themselves. Interest in the edge is growing among a number of potential stakeholders that are also likely to be in contention for the space. These include public cloud providers wanting to build their own edge clouds to either sell to CSPs for their networks or sell directly to end users, systems integrators that may feel they have the necessary combination of communications network and IT skills to build out the edge, and other ecosystem partners such as tower companies that see their own assets as being part of the cloud network.

In reality, many CSPs acknowledge that creating the edge cloud will depend on a broad ecosystem and require support from a number of players. Not least, the engagement with application developers will be crucial in driving the development of new applications for the edge, and it must include features such as service APIs in order to enable the exposure of underlying network information and capabilities.

ETSI MEC standards progress

Work in ETSI's Multi-access Edge Computing Industry Specification Group (MEC ISG) to define an edge framework for implementation in CSP networks is already well advanced: Phase 1 of its work program was completed in 2016, and the first set of Phase 2 specifications were released in March. These cover the architecture, framework, and general principles for service APIs and widen the scope of ETSI MEC to include any access technology and to take into account integration with NFV (see Table 1).

Table 1: ETSI MEC Phase 2 specifications

Specification	Focus
ETSI GS MEC 002	Defines the requirements for MEC with the aim of promoting interoperability and deployments. Focuses on NFV alignment since multi-access edge computing uses a virtualization platform for running applications at the mobile network edge. The specification also describes example use cases and their technical benefits for the purpose of deriving requirements.
ETSI GS MEC 003	Addresses the implementation of MEC applications as software-only entities that run on top of a virtualization infrastructure, which is located in or close to the network edge. A key new aspect of the Phase 2 version of this specification is the addition of MEC-in-NFV reference architecture, which defines how MEC-compliant edge deployments can be part of an overall NFV cloud architecture.
ETSI GR MEC 022	A report on MEC support for vehicle-to-infrastructure and vehicle-to-vehicle use cases. The report covers four use-case groups that are commonly known to the V2X communities, namely safety, convenience, advanced driving assistance, and vulnerable road user.
ETSI GS MEC 026	A specification supporting regulatory requirements in MEC.

Source: ETSI

ETSI MEC is recognized as a key contributing technology to 5G alongside developments such as SDN/NFV, cloud-based architectures, and network slicing. Support for edge computing is also addressed in 3GPP's ongoing work on the system architecture for 5G, and close cooperation between ETSI and 3GPP will ensure the coexistence of MEC and 5G.

Cross-industry cooperation will help create an open ecosystem for MEC

In an effort to establish cross-industry cooperation, various associated standards bodies and consortia are already bringing together network equipment providers (NEPs), chipset suppliers, providers of data center and computing hardware, software developers, and network test companies to help create a heterogeneous, on-demand, high-performance, and open MEC platform. Standards bodies and industry initiatives such as ETSI MEC and the Linux Foundation are deepening their collaboration in the drive to define a common vision for edge computing, while bodies such as the newly established Edge Computing Consortium Europe are focusing on the adoption of edge computing in areas such as smart manufacturing, industry, and enterprise IoT (see Table 2).

Table 2: Key organizations working on edge computing

Organization/initiative	Composition	Main area of activity
3GPP	Regional standards bodies, CSPs, network equipment and software vendors, chipset suppliers.	Incorporating support for MEC into the technical standards for 4G and 5G networks.
Edge Computing Consortium Europe (established January 2019)	Huawei, Intel, IBM, Arm, Analog Devices, Spirent, et al.	Encouraging the adoption of edge computing in smart manufacturing, industry, and enterprise IoT.
ETSI Multi-access Edge Computing Industry Specification Group (ETSI MEC ISG)	ICT companies, government and regulatory bodies, universities, and research bodies.	Producing specifications for a generic standardized framework for MEC including application environment and specific use cases and proofs of concept (POCs).
Linux Foundation (LF Edge Consortium)	IT vendors and NEPs (Dell, HPE, IBM, Ericsson), AT&T, NTT, Wind, Arm, Intel, Qualcomm, Baidu, Tencent.	Supporting open source projects relating to the development of an edge-computing software stack, including infrastructure.

IEEE	Professionals in the fields of technology and engineering.	Developing industry standards in a broad range of technologies.
Open Networking Foundation (ONF)	Mobile operators and web-scale providers.	Developing open source reference designs for operator edge-computing deployments.
Industrial Internet Consortium (incorporating OpenFog as of January 2019)	Large and small industrial companies, universities, and research institutions.	Helping define and promote fog computing, influencing standards, and building operational models and test beds.
Open Edge Computing Initiative	Carnegie Mellon University, Intel, Microsoft, Crown Castle, Vodafone, Deutsche Telekom, NTT, Nokia.	Creating a use-case-centric testbed for edge computing to provide application developers with experience of the benefits of 5G technology.
Open Edge and HPC [high-performance computing] Initiative	Atos, E4 Computer Engineering, Fraunhofer Fokus, Forschungszentrum Jülich, Huawei, Mellanox, Suse.	Recognizing the need for secure distributed data processing and low-latency communications in emerging 5G and Industry 4.0 environments.
OpenStack Foundation	Edge Computing Group.	Providing infrastructure software that can be deployed anywhere, including the edge of the network.
Telecoms Infra Project (TIP)	Co-founded by Facebook. Includes operators, network equipment and software vendors, application developers, and universities.	Focusing on lab and field trials for services/applications at the network edge. The scope includes 5G and IoT service deployments.

Source: Ovum

MEC and 5G: Connectivity and computing can work together to meet edge service requirements

The key feature of MEC is that it brings together the network, cloud, and devices in a converged ecosystem that will enable differentiated service innovation. While 5G provides the connectivity layer, MEC adds real-time processing and management of data with ultra-low-latency connectivity located at or driven from edge-network nodes or micro data centers connected back to the central cloud.

Network trials and newly launched commercial services are already demonstrating the potential of 5G to transform delivery of a range of enhanced mobile broadband and fixed wireless access services to consumer devices and homes. Support for high-quality voice and data, high-definition and interactive 4K/8K video, and new experiences such as AR/VR is possible thanks to 5G's network capabilities including high bandwidth, low latency, and support for features such as advanced video compression.

With access to the enhanced capabilities of 5G, network operators are already turning their attention toward new opportunities, not only to build on these existing services but also to expand their role in the market for B2B services, where they see a growing demand for new solutions capable of replacing or augmenting legacy technologies with high-performance, highly differentiated, and secure connectivity across a range of industrial, commercial, and public sector uses (see Figure 4).

Figure 4: Potential B2B applications enabled by 5G

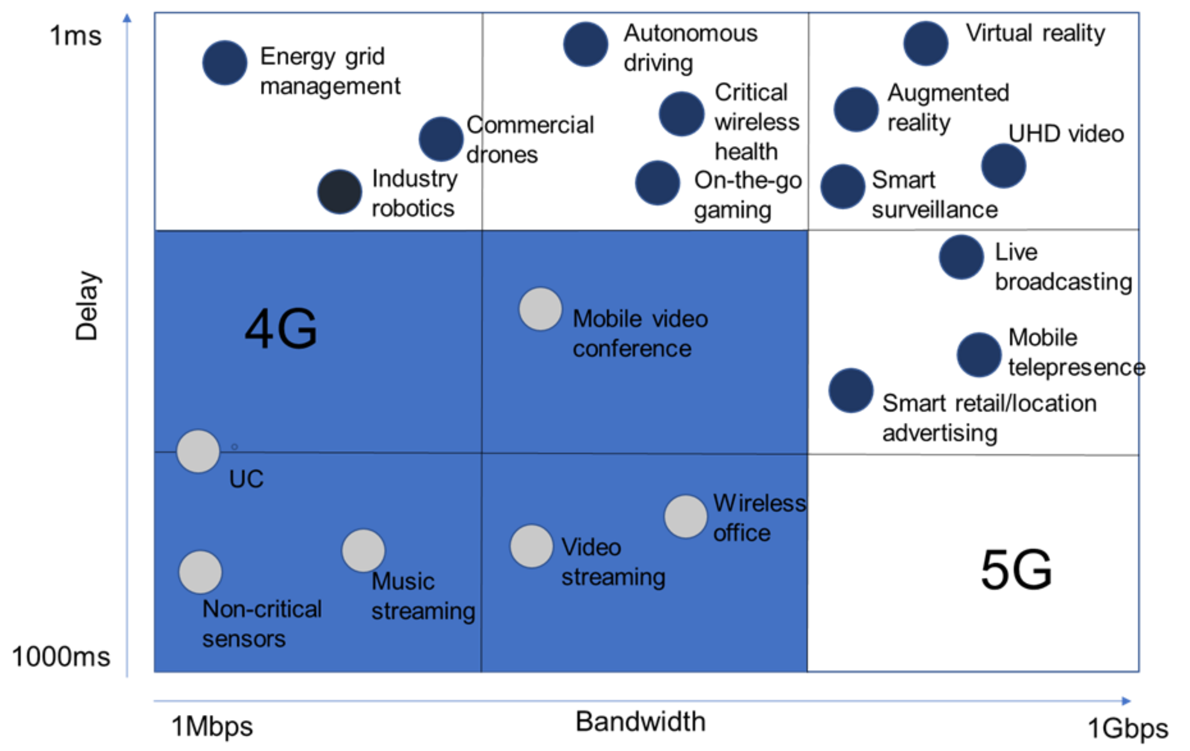

Source: Ovum

Key features of 5G MEC

Seamless mobility and extreme connectivity will ensure 5G's role as an enabler for MEC

The introduction of 5G is founded on CSPs' adoption of a virtualized, software-driven network model, both in the core network and at the edge, along with a cloud-based architecture. The 5G standards in their turn incorporate a number of key features and technologies that can be applied in order to support the decentralized architecture necessary for MEC implementation.

In the case of the 5G core, the evolution toward separation of the control and user planes (CUPS) in the 5G packet core is fundamental to the deployment of MEC, because it allows the user plane function (UPF) to be flexibly deployed at the network edge, while the control plane remains centrally located and can interface with other core network functions. This, in turn, allows for different traffic-steering options in both the uplink and downlink such as local policy-driven traffic forwarding, which can be applied by means of the insertion of an uplink classifier function (UL-CL) in the data path.

As well as offering low latency and high throughput for B2C services, 5G networks can provide additional features to support new B2B services via the edge cloud. These include network slicing for guaranteed end-to-end performance, reliability, and security and session and service continuity (SSC) modes for seamless, ultra-low-latency connectivity.

These features of 5G will be important in creating a deterministic network, capable of delivering differentiated and seamless services that can guarantee performance criteria such as latency according to the requirements of a particular use case but can also be capable of working within a range of latencies where the demands of a particular use case are less stringent.

Edge node integration and platform requirements

Operators' existing assets such as edge sites or central offices can potentially work as a starting point for edge-computing deployments so overcoming the limitations on computing and storage capacity close to the network edge. Applications that require the caching or processing of large amounts of data at the edge may be better sited at a central office or at metro aggregation points. In the longer term as 5G is rolled out, these sites can then also be used to host virtualized core network components and associated services (see Figure 5).

Figure 5: Examples of MEC deployment options

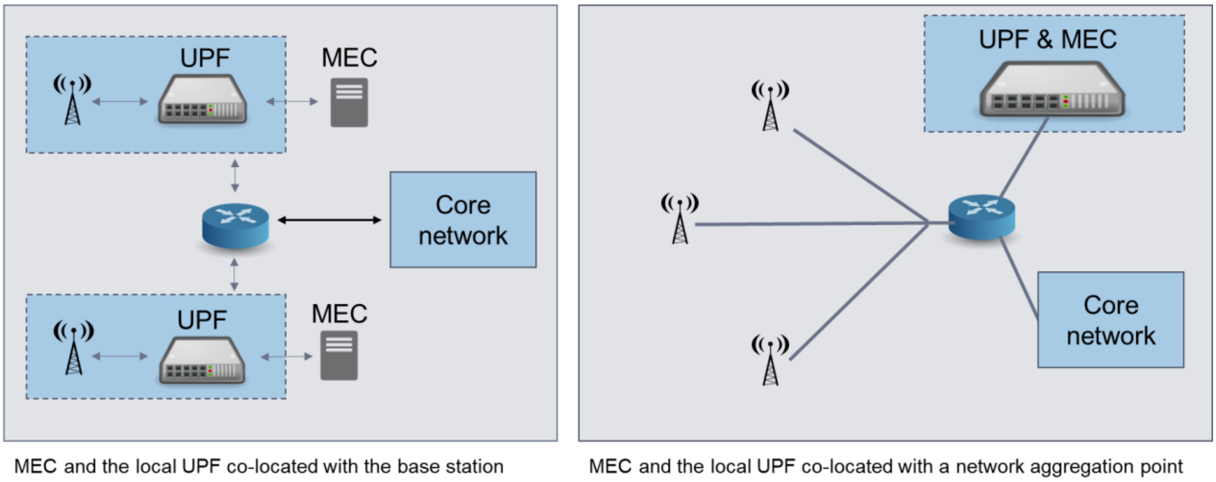

Source: ETSI/Ovum

Hardware platform and compute resources

Although likely to be space limited, edge installations must be able to support high traffic levels and provide the resources necessary to address a diverse range of use-case scenarios and performance demands at the network edge. Unlike in the network, where performance criteria such as ultra-low latency can be more readily guaranteed, the MEC platform presents significant challenges due to, for example, the delays incurred during the rendering and decoding of data traffic.

In order to achieve the lowest possible cost per bit and address the need for service acceleration, an open, heterogeneous computing platform is desirable, capable of addressing applications, content, core network service processing, and resource scheduling functions at the network edge.

Enhanced hardware, typically comprising compute and storage capacity and network interface controller (NIC), needs to be compact as well as flexible and scalable. A full-stack infrastructure approach with a one-box, plug-and-play solution featuring automatic remote maintenance may be appropriate in order to simplify deployment and operation.

In addition, a single MEC deployment ideally needs to support more than one use case in order to deliver maximum benefit and optimal use of resources. Thus, for example, a large-scale campus type MEC deployment for an enterprise might be used to share additional services such as B2C.

Centralized control and microservice support

Centralized orchestration and management (O&M) of the MEC platform is critical for seamless integration, resource management, control, and execution. Application server functionality will be required to give seamless support, via the network, to applications and services accessed via the edge nodes. This will ensure service continuity and content synchronization for applications such as V2X and VR where users are moving at the network edge.

Support for microservices is particularly appropriate in a MEC environment where services will be scaled down and modularized but where multiple software components are going into building solutions that will expand and evolve over time. In this context, adopting support for microservices can ensure support for an agile lifecycle management approach to the development, testing, and maintenance of individual applications and services at the network edge.

Convergent platform for multi-access

Because MEC will need to support multiple network access technologies including 2G/3G/4G and 5G as well as Wi-Fi and wireline, it will require a convergent platform in order to simplify these multiple accesses.

Evolution of 5G and future support for MEC

More technologies being introduced as part of 3GPP Release 16 will support MEC deployment for vertical industries, including industrial automation use cases. These include ultra-reliable low-latency communications (URLLC), 5G LAN, and Time-Sensitive Networking (TSN), the IEEE standard originally developed for Ethernet but which is being integrated with 5G to further support aspects such as reliability and synchronization and to guarantee low latency in industrial automation-type scenarios.

Trials and POCs are already demonstrating the value of MEC

The immense potential value of MEC is already amply demonstrated in numerous trials and POCs being carried out by CSPs across a range of use-case scenarios. These are targeting application areas such as AR/VR, cloud gaming, local content distribution, V2X and autonomous vehicles, intelligent transport systems, security and surveillance, manufacturing process automation and monitoring (Industry 4.0), IoT, fleet management and logistics, campus networks, smart stadia, and many more. Specific examples include

- control of automated guided vehicles (AGVs) in a factory environment in Switzerland
- a MEC-enabled augmented reality experience for visitors to a public light show in China
- parking management at American Dream, one of the biggest shopping malls in the US
- a 5G-ready digital pedestrian road-safety application in Germany
- automated quality inspection and monitoring of consumer products on a manufacturer's production line in China
- the interpretation and analysis of video data sourced from surveillance cameras in Japan.

At the same time, CSPs are beginning to deploy edge-cloud infrastructure in readiness for the introduction of commercial services.

Plans are particularly well advanced in the leading Asian markets of South Korea, Japan, and China. In South Korea, both SK Telecom and KT are deploying regional edge data centers with plans to target both the B2B and B2C markets. Meanwhile, all three leading Chinese operators, China Mobile, China Unicom, and China Telecom, have either begun large-scale, precommercial edge-cloud network construction or are engaged in pilot projects using MEC.

In Europe, BT is increasing the number of its metro exchanges tenfold in anticipation of providing an expanded edge-computing service offering for retail, enterprise, and wholesale customers, while Telefónica is exploring the future role of MEC as part of its global network virtualization program, UNICA.

In the US, AT&T is converting thousands of central-office locations into data centers and building an edge-computing test zone in Palo Alto, where it is working with third-party developers and startups to test use cases including autonomous car, AR/VR, and drones.

Conclusion

The edge business case can become a reality

A number of important questions about MEC remain to be answered, including which services and applications will require key MEC features such as ultra-low latency and when and where network operators will be ready to deploy MEC commercially. In order for the business case to stand up, CSPs may need to consider employing multiple use cases over a single MEC platform, since one MEC platform per application will be too costly. So, for example, larger-scale campus-type MEC deployments for enterprises might support many applications.

Nevertheless, the incentive for edge computing to be deployed in 5G networks appears to be growing, and it is becoming increasingly important to the 5G business case. The first commercial deployments appear likely to start in 2020 in markets such as South Korea and China with others following soon after. Above all, MEC can become a reality, particularly as more commercial 5G deployments begin to appear, services start to bed in, and yet more use cases start to emerge.

Appendix

Author

Julian Bright, Senior Analyst, Service Provider Network Infrastructure and Software

julian.bright@ovum.com

Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

Copyright notice and disclaimer

CONTACT US

ovum.informa.com
askananalyst@ovum.com

INTERNATIONAL OFFICES

Beijing
Boston
Chicago
Dubai
Hong Kong
Hyderabad
Johannesburg
London
Melbourne
New York
Paris
San Francisco
Sao Paulo
Shanghai
Singapore
Sydney
Tokyo