# Cloud Versus MEC for Telcos and MSOs

## MEC Use Case Study

A Technical Paper prepared for SCTE by

**Yassar Abbas**
5G Innovation Lead Architect
Rogers Communications Inc.
8200 Dixie Road, Brampton ON
+14166252051
Yassar.abbas@rci.rogers.com

**Brad Tizard,** Rogers Communications Inc.

**Probir Khaskel,** Rogers Communications Inc.

 

# Table of Contents

**Title** **Page Number**



## List of Figures

**Title** **Page Number**

# 1. Introduction

The need for lower latency applications, decentralized processing power and increased adoption of IoT and connected devices are giving rise to the shift of data processing closer to the source commonly known as edge computing or MEC (Multi-access Edge Computing).

[1]Multi-access Edge Computing (MEC) offers application developers and content providers cloud-computing capabilities and an IT service environment at the edge of the network. This environment is characterized by ultra-low latency and high bandwidth as well as real-time access to network information that can be leveraged by applications.

This brings a unique opportunity for Telcos and cable operators to deploy compute resources leveraging their distributed facilities across their geographic footprint. Edge computing is particularly interesting in this era for MSO and Telcos with virtualization and decentralization of network functions moving user plane workloads closer to edge co-located with third-party applications to reduce latency.

The opportunities that edge computing presents are attractive; however, network operators need to identify the right use cases and applications that are candidates for edge computing.

In this paper, authors will share lessons learned from some of the MEC use cases like Cross-operator music jam session from different parts of the world, AR (Augmented Reality) based in-stadium fan experience during a live NHL game, last mile food delivery using autonomous robots and VR (Virtual Reality) banking experience.

These use cases have been tested and demonstrated in Roger's network leveraging MEC as part of the 5G Innovation and Partnership program.

This paper also provides a brief description of edge computing terminology and how network edge or MEC is related to a hyper-scaler public cloud. Please note that there is no standard terminology, and you may see other terms used to refer to the same thing in other material. However, the focus is really to highlight the difference between network edge, public cloud, and cloud edge.

[1]MEC initiative is an industry-standard specification group within ETSI to define standard and open environment which can allow interoperability across mobile operators, application developer, Independent Software Vendors (ISV), telecom equipment vendors and technology providers. However, the use of MEC terminology has been applied much more broadly than solely to ETSI standards and interoperability.

Although the focus of MEC use cases has mainly been on wireless access-based services but as the name implies, it is inclusive of all access network technologies including fixed and WLAN.

In some literature, MEC has also been called Mobile Edge computing due to its early focus on 5G and wireless use cases. The use of Edge computing terminology is very generic across the industry and subject to interpretation. For example, for user device manufacturers edge is on the device while for car OEMs edge is in the car. Our focus in this writing is the network edge use cases and not on standards and interoperability.

MEC infrastructure can be used to deploy operator's network functions (VNF) workloads or to host third-party applications. Operator own network workloads or VNFs are typically deployed on private MEC or private cloud however the use of public cloud has also been seen in the case of some greenfield deployments.

## 1.1. Edge Compute Benefits

There are three main benefits of edge computing when it comes to differentiating from the cloud.

### 1.1.1. Latency

Latency improvements are probably the most talked about attribute of MEC. Cloud computing needs processing of data in centralized data centres sent over the internet resulting in higher latency not suitable for applications like such as industrial applications or cloud gaming. MEC or edge computing in general is ideal for applications requiring low latency since data is processed closer to the source and edge of the network.

### 1.1.2. Bandwidth

Cloud computing requires transporting data over the internet far from data sources. This requires backhauling of the data traffic and increased bandwidth demands. In comparison to cloud computing, MEC can save bandwidth by computing at the edge and not requiring data to be transported to the cloud.

### 1.1.3. Security

With cloud computing data is processed and stored in remote data centres owned by cloud providers. MEC can alleviate those concerns by processing data locally and not transmitting it over the internet. With cloud computing data in transit can be protected using encryption and tunneling techniques at additional cost. Also, in some countries, there are data residency requirements that mandate local data processing and storage requirements.

# 2. Edge Compute Definitions

There are many different forums defining the specifications and standards for IT platforms enabling the use of MEC to locate the closest MEC and instantiate compute, storage, and other resources to host the application. It is not the purpose of this paper to delve into standards and MEC platform details but rather to describe MEC as a general concept. The use of MEC and edge computing has been used interchangeably in this document. However, for the sake of clarity to readers, we would define some key terms and nomenclature in the context of this paper.

## 2.1. Private MEC

Any edge compute infrastructure deployed dedicated to a single Enterprise is called private MEC or cloud. It is typically a single tenant serving only one organization or Enterprise. It can be deployed on customer premises aka on-premises (factory floor, warehouse, etc.) or on operator facilities very close to the network edge (e.g., cell tower location) to serve ultra-low latency use cases. Enterprises can combine the benefits of on-premises wireless private networks and private MEC to optimize latency and throughput while leveraging technologies like computer vision, AR/VR/XR, and machine learning.

For Telco and MSOs, private MEC or cloud is typically hosting the VNF (virtual network functions). With CUPS (control and user plane separation), user plane functions can be deployed closer to the edge resulting in latency improvement and bandwidth saving by not backhauling traffic to the central cloud data centre.

Azure Stack Edge and AWS outpost family are examples of Hyperscaler private MEC. Operators can also choose to deploy their own private MEC/cloud infrastructure and services.

### 2.2. Public MEC

Public MEC is an extension of the Public Cloud and supports multi-tenancy. AWS Wavelength zones are an example of Public MEC deployed at network operator facilities in core data centres. Public MEC is used to host third-party application providers or ISV (Independent Software Vendor) workloads closer to the network edge. Anyone can subscribe to Public MEC just like anyone can subscribe to Public Cloud.

AWS wavelength zones are an example of Public MEC offered in partnership with HCP and network operators.

### 2.3. Public Cloud

Public cloud is a cloud offering by hyper-scale cloud providers such as AWS, Azure, and GCP (Google Cloud Platform). Public cloud providers' data centres are centralized but geographically spread across regions over a wide area providing global coverage. When it comes to latency and bandwidth optimization, the physical placement of public cloud providers data centres plays an important role. Public cloud is ideal for applications needing large-scale compute and latency-tolerant applications.

### 2.4. Cloud Edge

Cloud edge is edge locations or POPs owned by HCP which are closer to users in comparison to public cloud providers' regions and availability zones. AWS local zones or edge locations POP is an example of cloud edge. A distinction should be made between HCP edge services versus cloud edge. HCP can provide edge services and infrastructure anywhere at the network edge or on-premises in partnership with network operators at POPs owned by network operators at the far-edge or near-edge.

### 2.5. Hybrid Cloud

A hybrid cloud is a mix of on-prem compute, private or public MEC, and public cloud environments. Workloads are deployed based on application latency, security, and data residency requirement.

### 2.6. Near-Edge and Far-Edge

Near-Edge and Far-Edge are the terms describing how far Edge is from the Cloud service provider. Network operators' central office or Cable Hub site can be classified under near-edge. An example of Public MEC like AWS wavelength zones is deployed at near edge sites like Telco central office or core data centres and cable Hub site.

Far-Edge is furthest away from the cloud and closer to the data source. Far-edge would include Network operator's sites deeper into the network, for example cell tower sites would fall into the far-edge category. Private MEC can be deployed at near-edge or far-edge.

### 2.7. Content Distribution Network

CDNs or content distribution networks are caches that are deployed closer to a user being served. CDNs are deployed at edge sites owned by caching providers like Akamai, HCP, or network operators.

 

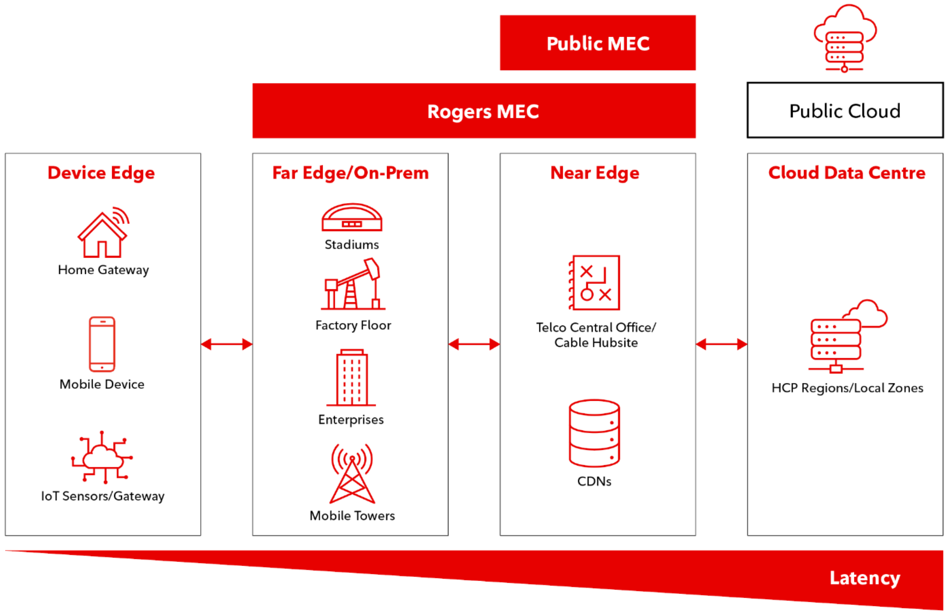


**Figure 1 – Edge Compute Locations**

The Figure 1 above shows possible edge compute locations and how they are categorized. MNO or Telco facilities fall under near-edge or far-edge categories and are good candidates for private or Public MEC. As applications are moved from the public cloud to the edge, latency improves. The actual edge cloud services or infrastructure can be built by the Network operator using a private cloud or in partnership with a Hyperscale cloud provider.

Roger's implementation of MEC for the tested use cases described is near-edge private MEC, and for the purpose of this paper, it will be referred to as Rogers MEC.

## 3. Tesbed Architecture

Figure 2 describes the generic architecture for Rogers MEC deployment common for all the use cases discussed in following sections.

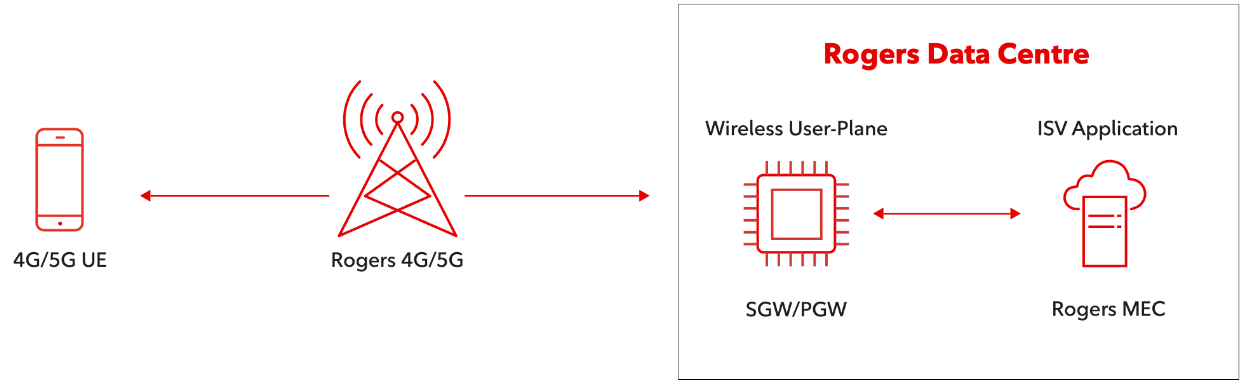


**Figure 2 – Rogers MEC Infra Deployment for POCs**

Rogers partnered with AWS to deploy Private MEC infrastructure collocated with wireless user plane functions (SGW/PGW) in one of Rogers's core data centres. AWS Private MEC called Rogers MEC hereafter was only hosting third-party applications and not wireless user plane network functions or VNFs. This resulted in optimizing the latency between 4G/5G UE (user equipment) and 3rd party applications hosted on the MEC server. For some use cases, a hybrid cloud model was used to deploy most latency-critical applications on MEC and more latency-tolerant workloads in the public cloud (AWS Canada central region).

5G Non-Stand Alone (NSA) mode was used in all Proof of Concepts (POCs).


**Figure 3 – MEC and Public Cloud Locations**

Figure 3 shows the location of Rogers data centre and public cloud for the purpose of use cases and benchmarking tests discussed in this whitepaper.

## 3.1. Performance Results

### 3.1.1. *Latency*

Figure 3 shows the location of Rogers's data centre and public cloud for the purpose of use cases and benchmarking tests discussed in this whitepaper.

When comparing latency between the public cloud (AWS Canada central region) and Rogers MEC, there is a reduction in network latency by one-fourth. The impact of latency have for particular use cases has been discussed in individual sections relating to the use case

 

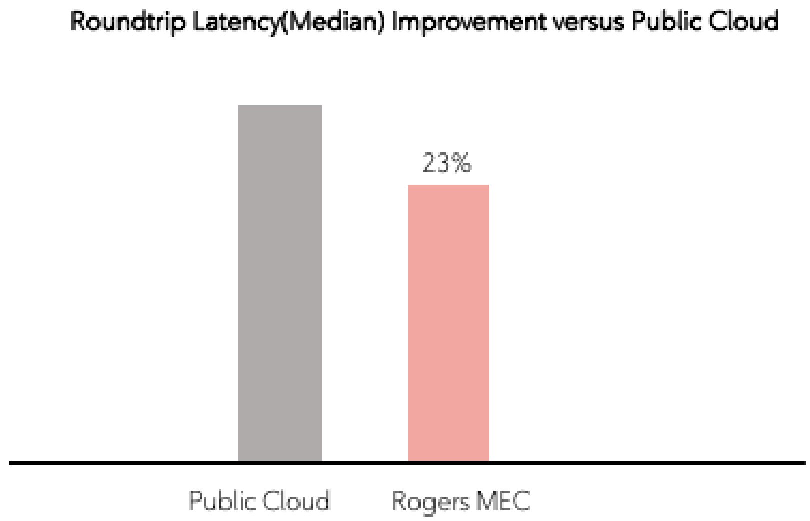


**Figure 4 – Latency Improvements Public Cloud versus Rogers MEC**

*Note: Roundtrip latency measurements were performed by sending continuous ICMP test packets from wireless UE to server deployed in Rogers MEC and Cloud. Results may vary depending on circumstances.*

### 3.1.2. Jitter

We saw around one tenth improvement in jitter while comparing Rogers MEC with public cloud.

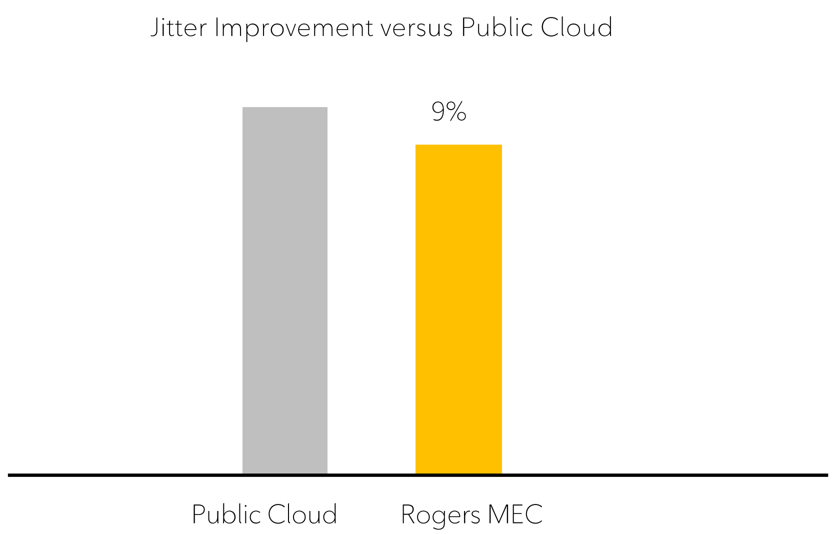


**Figure 5 – Jitter Improvements Public Cloud versus Rogers MEC**

*Note: Jitter measurements were performed by data collected using application logs*

 

# 4. AR based Sports Fan Experience

## 4.1. Overview

AR and VR technologies have been around for a while; however, their mass adoption has been somewhat limited due to various reasons including cost, customer experience, and battery life. Advancement of compute and cloud technologies along with more responsive transmission networks like 5G has brought a renewed interest in these technologies and made their adoption more realistic.

AR enhances a user's experience by superimposing digital assets onto what they are viewing in the physical world. With AR glasses or a smartphone, users can play games like Pokémon Go, envision how outfits looks like, or enhance fans sports viewing experience in a stadium or at home.

Rogers partnered with AWS, AR-based sports fan experience company called Immeriv.io, NHL, and MLSE to showcase the value of 5G and MEC technologies during a live NHL game at Scotiabank Arena in Toronto.

Users were able to watch live game in the arena on their smartphone phones or by wearing AR glasses with augmented features like player tracking, puck speed, and viewing player stats overlaid on the live game view. Puck and player tracking features have previously been demonstrated for in-home game-watching experience on TV but in arena fan experience during a live game is a more challenging environment.

This is because the user experience is negatively impacted if the AR view is not in sync with a real-life view of the stadium. This kind of application makes it a perfect candidate for AR content to be fetched over low latency network like 5G and hosted on MEC closer to the user being served.

## 4.2. Architecture

This use case was enabled by NHL puck and player tracking data collected through sensors deployed in the stadium, inside the puck, and in player jerseys.

Puck and player tracking data is generated, ingested, and stored on the premises server in the arena and also sent to cloud servers for analytics and other application consumption. To display game analytics in near real-time on the user's mobile phone or AR glass, puck and player tracking data needs to be sent from the in-stadium tracking application server to Immersiv.io tracking server on Rogers MEC and then to the AR application on the user device in low latency over Rogers 5G network.

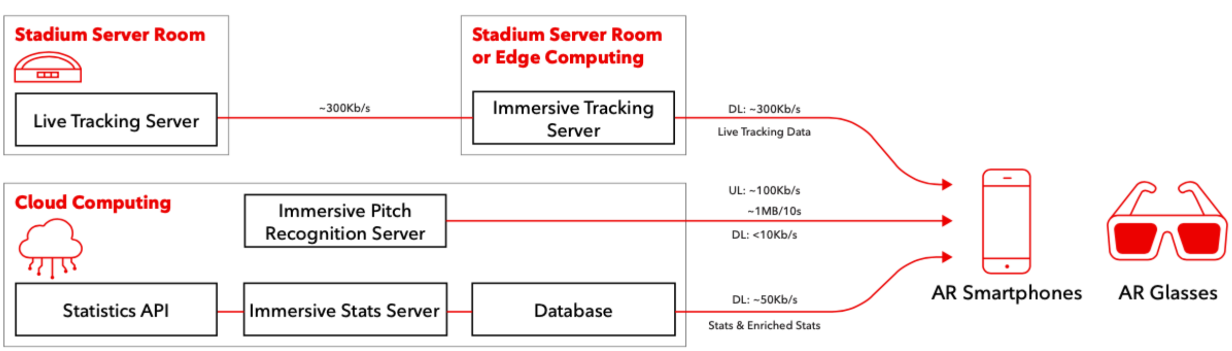


**Figure 6 – Immersiv.io Application Architecture**

Immersiv.io solution has the following main software components:

**Pitch Recognition Server:** Pictures of the ice rink or pitch in the stadium are uploaded periodically to the pitch recognition server by smartphone in a stadium.

**Stats Server:** Statistics server stores the league stats in the database by querying 3rd party providers and sends it to an in-stadium user.

**Tracking Server:** This server ingests puck and player tracking data from the in-stadium tracking server and sends live tracking data to smartphones in real-time. This data is most latency-sensitive and so the tracking server was the only candidate to be deployed on MEC infra.

To reduce the latency of Immersiv.io live tracking server data sent to UE, we considered three different options described below.

1. Host tracking application server on MEC co-located with wireless user plane functions (SGW&PGW) in the stadium.
2. Host tracking application server on MEC deployed in Rogers wireless data centre collocated with wireless user plane functions (SGW&PGW)
3. Host tracking Application server in Hyper-scaler closest regional cloud data centre with wireless user plane function (SGW&PGW) still deployed in Roger's data centre.

Option 1 was not chosen due to the cost and effort of deploying wireless user plane functions (SGW&PGW) in the stadium outweighing the benefits of latency improvement for this particular use case.

So, we deployed a tracking server application on MEC infra using Option 2. A tracking server was also deployed (Option 3) in AWS's geographically closest regional data centre (AWS Canada central) only to benchmark MEC performance versus the cloud.

In both cases, there were application workloads (like database, stats server etc.) which did not demand strict latency and were deployed in the public cloud.

The following diagram shows the high-level architecture of the demo with wireless packet core user plane functions co-located with MEC hosting immersiv.io tracking server while other application functions are deployed in the public cloud.

 

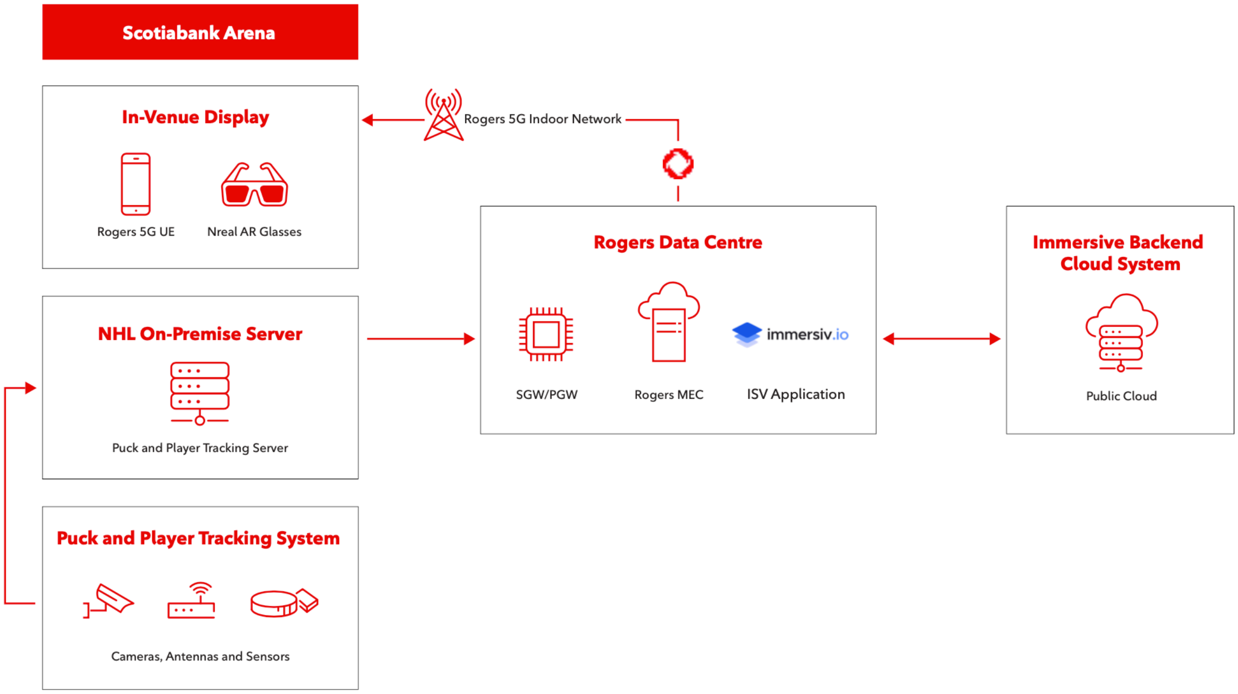


**Figure 7 – AR Based Sports Fan-Experience POC Architeture**

## 4.3. Conclusion

### *4.3.1. Latency*

When comparing latency between the tracking server and smartphone UE in the stadium, there was a reduction in network latency by one fourth as mentioned in section3.1.1

If the latency is too high for this use case, then AR overlays, the circles around the players, would be lagging the real positions of the players. The following Figure 8 shows an example of latency impact for AR-based application if latency is too high. The left-hand side of Figure 8 shows when the puck is tracked accurately while the right-hand side shows when the puck is inside the ice rink, but AR tracking shows it outside.

This kind of impact can be experienced if the public cloud region is geographically (or fibre distance is too long) too far from the network edge or network operator data centre.

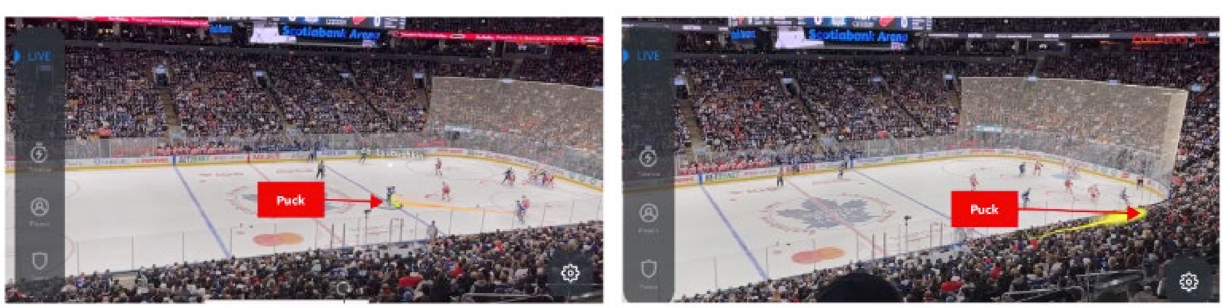


**Figure 8 – Impact of Latency on User Experience**

 

# 5. MEC based AR/VR Video Rendering

## 5.1. Introduction

Imagine being free to explore a virtual world, play games, attend events, or buy a new home as if you were actually there, without leaving the comfort of your home or office using a headset that covers your eyes. There has been a resurgence in the popularity of Augmented Reality (AR) and Virtual Reality (VR) in recent years. Thanks to improvements in processing power and cloud computing, these technologies have the potential to enable new, unique experiences for both consumers and businesses.

To begin bringing this imagination into reality, we searched for a best-in-class real-time 3D platform and collaborated with Unity to help us build a tangible VR prototype. Our next hurdle was meeting the demands of heavy Unity 3D graphics. As 3D visuals increase in quality to enhance AR and VR experiences, first, they require additional graphic processing power which, secondly, uses more battery power, and neither of which is abundantly available on headsets.

The solution was to completely move all major computational processing to the cloud. Cloud-hosted VR applications make use of the head-mounted displays in goggles for streaming ultra-realistic, high-fidelity, immersive virtual experiences instead of relying on ever-smaller headsets for rendering and processing. By doing this, only the pixels stream back to the headset, similar to a TV in traditional broadcasting. This will allow users to collaborate with others in much more immersive and extensive virtual workspaces than VR traditionally promises.

Streaming from the cloud, however, can create one hurdle: there is usually some physical and logical distance between the cloud and the headset which produces latency, which can stifle the experience. Latency is a key performance indicator in any modern viewing experience: from film actors being out of sync with their voices, news broadcasts waiting extensively before the interviewee responds, or Netflix reducing quality or buffering extensively. Everyone has knowingly or not experienced high latency, but in VR, high latency not only results in poor performance, low-fidelity graphics, and significant buffering like we traditionally expect – here, it is a recipe for nausea.

## 5.2. Architecture

In situations where Public Cloud is geographically very distant with high latency, a Private MEC can meet the low latency requirements of VR streaming. Rogers tested this use case, by offloading the graphic processing power from the VR headset onto a private telco MEC, using the architecture originally show in Figure 2. The entire solution added the components seen in Figure 9, and we were able to reduce the overall network roundtrip time and Jitter, as seen in Figure 4 and Figure 5 respectively, and in VR this improvement can be visually perceived by the average consumer.

3D Unity content for use inside the headset was developed separately offline and sideloaded onto the Rogers MEC server using out-of-band channels. Later, on demand, real-time rendering on the Rogers MEC streamed high-fidelity visuals over 5G back to the VR headset. Streaming with 5G minimizes RF latency while increasing bandwidth, and hosting MEC servers logically close to the client optimized the network latency. This allowed for more pixels and higher-definition graphics at a more stable rate, and thus a better VR experience overall.

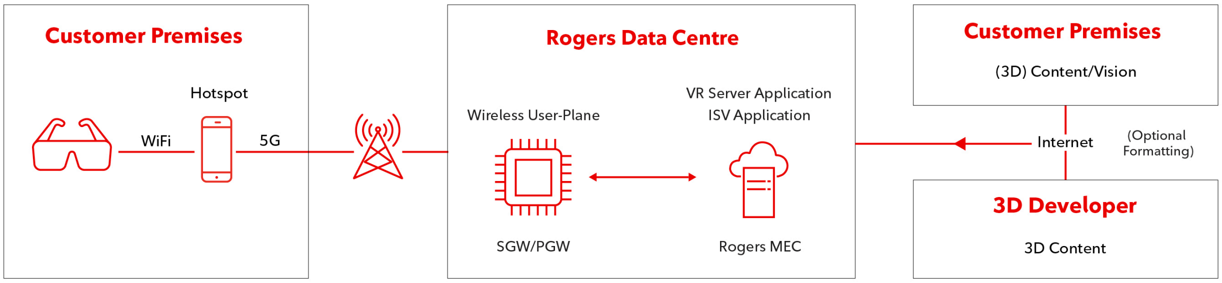


**Figure 9 – MEC based AR/VR Video Rendering Architecture**

## 5.3. Conclusion

Our architecture leveraged a VR headset connected to our Rogers MEC server via the internet, using Wi-Fi from a 5G cellular hotspot as last mile access technology. This end-to-end solution enabled us to leap over hardware limitations imposed by the headset and showcase a vivid and wireless VR streaming experience thanks to exceptionally high bandwidth, and low latency. It is reasonable conclude that the results, and VR experience overall, could be further improved by leveraging either 1) a headset with embedded SIM (eliminating the Wi-Fi hop), 2) a far-edge MEC deployment or 3) both of the aforementioned.

### *5.3.1. Battery*

As developers realize higher fidelity graphics from on-board processors there is an intrinsic trade-off: power – including heat, weight, and capacity. Manufacturers constantly consider ergonomics, and with a VR headset balancing battery capacity with the physical weight and strain on a person's head is critical. Ironically, larger batteries produce more heat by-product and as processors compute more complex calculations they too heat up. Uncomfortably warm components inches from your face encourage users to remove their VR headset. This problem quickly grows as the battery becomes more and more consumed.

These small factors altogether increase the chances you will take the headset off sooner, but more important decrease the immersion. As an application developer that is the opposite of what you want –you want to encourage retention by having your audience use the product as much as possible and 'get lost in the experience'. VR streaming significantly reduces the processing loads on the headset allowing processors to stay cool which reduces battery drain– rivalling battery life as if the device was just simply watching movies and encouraging users to finish enjoying themselves instead abruptly stopping to charge.

### *5.3.2. Bandwidth*

An on-premises solution and –in the right situation– Public Cloud can stream an identical VR solution over Wi-Fi, but each VR streams consumes a large bandwidth with preferred graphics. This requires a properly provisioned connection. Even still, any single access point on an average/properly provisioned Wi-Fi deployment is likely to be overloaded by multiple VR streams, and with real potential to overload the backbone IP network - especially in combination with everyday usage.

Here, 5G presents a unique solution with the opportunity to have dedicated high-capacity channels for each headset by pairing each headset with its own SIM card. For businesses with average capacity LAN networks even the slightest impact on the rest of the business can be a critical factor for adopting VR workloads.

 

#### 5.3.3. Remarks

It is much more enjoyable to work and collaborate remotely, play games, and explore environments with the incredible detail, lighting, scale, and lifelike physics that unrestricted GPU power provides. VR headsets have already evolved very quickly over the past decade, and as the industry matures and expands with new technology, we can expect a wireless and lifelike experience– unlocking and making practical use cases from: business to home entertainment; from indoors to outdoors; from single player to massively-multiplayer, emergencies, navigation, and more that we have yet to imagine.

## 6. Cross Operator Music Jam

### 6.1. Overview

Minimizing audio latency (or audio delay) between a group of musicians is critical for musicians to hear, play, and react to each other's performances. Musicians who play together in the same physical location, experience little to no audio latency between each other.

Musicians who have attempted to remotely perform online with widely used digital collaboration solutions such as video & audio meeting applications and audio-only communications applications have experienced challenges with synchronously performing together due to the significant audio latency within these solutions (+250 ms).

In this demonstration, three world-renowned guitarists performed a medley of rock 'n roll selections, each sitting in separate countries through the low-latency compute environment using 5GFF's Edge Discovery Service (EDS) API, with Open Sesame Inc. SyncStage ultra-low latency audio pipeline, with the 5G connectivity of Rogers, Verizon, and Vodafone.

SyncStage is an audio pipeline that is optimized for 5G and made available to application developers via Software Development Kit (SDK) to integrate with their application. Whenever a synchronized audio session is requested, OpenSesame's platform infrastructure powers the audio connections and communications between a group of synchronized audio users.

Edge discovery API lets the application select the closest MEC server in a country thus making sure there is a reduction in latency and jitter between UE and application. This helps the traffic path optimizations in case there are multiple MEC locations.

This demo not only achieved multi-country, glass-to-glass latency that enables the guitarists to jam together online, but also unlocked a new reference pattern for multi-edge, multi-region applications using an interoperable EDS API.

### 6.2. Architecture

The OpenSesame SyncStage platform consists of three major components:

**SyncStage Backend:** Built on Amazon API Gateway, the Backend enables session control across all the participating devices during a session. At the MWC demo, all the devices connected to the Backend were deployed in AWS us-east-1 region. The Backend checks with the 5GFF EDS APIs and provides optimal MEC service endpoint address and connection information to the devices in real-time.

**Studio Server:** Deployed on MEC infrastructure, the Studio Server streams content to and from the connected devices. An instance of Studio Server is deployed at each MEC location, and the devices connect to their nearest instance via 5G network.

**SyncStage SDK:** A software-only SDK for application developers to integrate into their application that provides ultra-low audio latency to a group of users to enable a vast range of synchronized audio experiences.

**Mobile Application using SyncStage SDK:** During the demo, the musicians plugged in their instruments and monitor to a 5G connected smartphone, running the application.

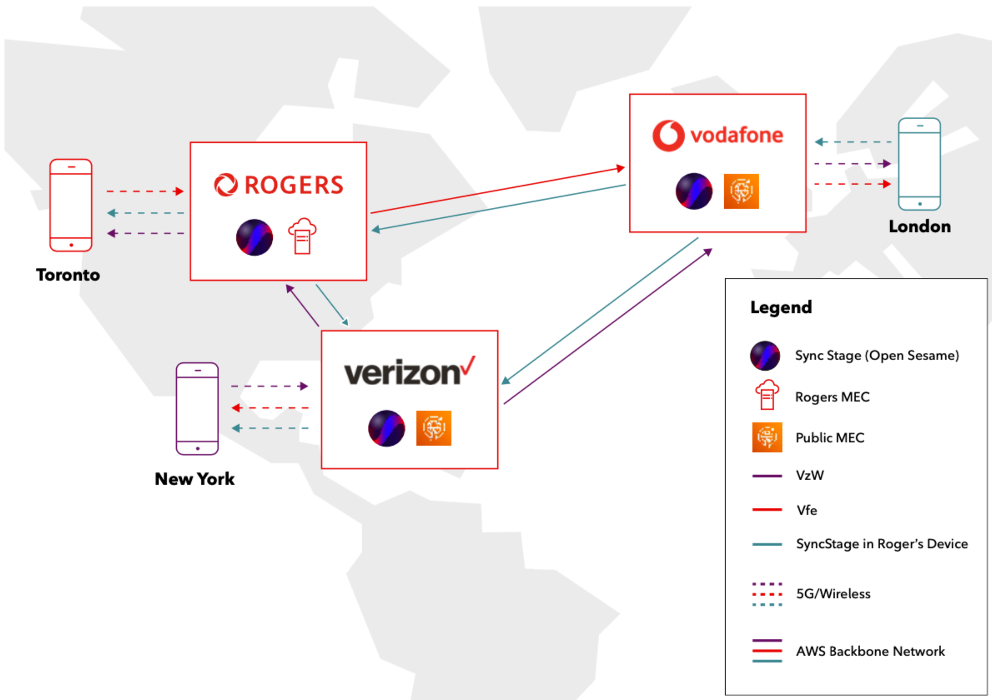


**Figure 10 – Technical Architecture for 5GFF Music Festival**

In New York and London, OpenSesame deployed the Studio Server instances at the Public MEC (local Wavelength Zones). In Toronto, the Studio Server was deployed on Rogers MEC, connected to the Rogers 5G network to resemble the setup in the other locations.

OpenSesame application running in 5G devices of the participating partner operators, got connected by IPv4 private network established among the three geographically distant MEC nodes through the AWS backbone network.

 

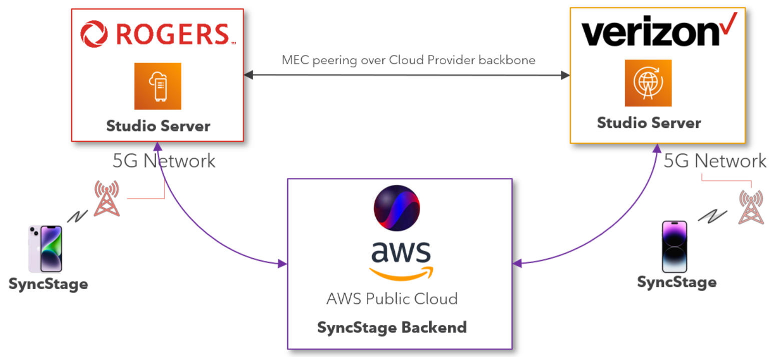


**Figure 11 – Edge Discovery Service API deployment**

When SyncStage mobile application initiates a call, on the control plane side, it queries to SyncStage Backend for optimal MEC service endpoint in a multi-MEC deployed network. The Backend checks with 5GFF Edge Discover Service (EDS) APIs and provides optimal MEC service endpoint address and connection information to the devices in real-time.

Once the optimal MEC endpoint is known, SyncStage application in the originating device requests a user-plane connection to the selected MEC endpoint in the home network that in turn establishes VPN tunnel with the partner MEC endpoint by using a private IP address and connects to SyncStage application in the target device.

While testing the demo setup, we observed musician-to-musician network latency. Some of this latency is attributed to the physical distance that the audio needs to travel to reach each performer. We experienced lower latency values between Toronto and New York, attributed to the shorter physical distance between these two cities. The longer distances between North America and the United Kingdom incurred greater latency. However, these latencies were neither noticeable to the live performers, nor was it evident in the session recording.

## 6.3. Conclusion

- MEC makes a good sense for the latency, jitter, and packet loss sensitive audio applications while keeping the traffic local in the MEC Zone where Local Zone and/or Region is geographically away. It measured close to one-fourth roundtrip latency reduction in Rogers' test environment setup by using Rogers MEC vs. public cloud for musician in Toronto. This test result is very much aligned with the findings presented in sub-section 3.1 Performance Results.

- Point-to-point physical direct connection between multi-operators MEC nodes would be beneficial to support latency, jitter, and packet loss sensitive use-cases, given there are use-cases to compensate for such cost. The estimated latency reduction in this specific application and use-case is 11-21% over the current MEC to MEC peering connection. This could be beneficial especially for intra-continental MEC deployment, favoured by distance and the law of physics.

 

# 7. Last Mile Food Delivery using Autonomous Robot

## 7.1. Overview

Rogers' objective has been to showcase the benefits of 5G network along with the capabilities of MEC. Rogers and OVIN (Ontario Vehicle Innovation Network) launched a challenge for SMEs (small to medium Enterprises) to come up with innovative 5G use cases in smart transportation and CAV (connected and autonomous) vehicle sector leveraging technologies like MEC.

LoopX, a start-up specializing in vehicle autonomy was selected to demonstrate the idea of last-mile food delivery using autonomous robot in Rogers' 5G powered smart campus testbed.

LoopX's last-mile food delivery solution uses electric-powered robots providing efficient delivery of food in an autonomous, contactless, and low carbon footprint manner. These delivery robots are designed to operate on sidewalks or bike lanes.

Delivery robots are equipped with LIDAR sensors, ultrasonic sensors, cameras for perception, GNSS receivers and IMU sensors for localization. Each robot is equipped with two onboard NVIDIA GPUs computers to host LoopX robot controller autonomy software.

Local sensors generate perception data on the robot to feed into its autonomy control software hosted on-board computer and GPU.

While there is a lot of buzz around autonomous or self-driving vehicles but, fully autonomous solutions would take time to mature given the regulatory and ethical aspects associated. In today's world, hybrid driving techniques involving a mix of autonomy and human driving (teleoperated and autonomous) are more practical and will continue to play an important role in the near future. we discuss MEC-based cooperative operation of multiple robots in this paper using Hybrid operation (autonomous and teleoperation with teleoperator driver intervention when required).

## 7.2. Architecture

For this use case, the MEC server in Rogers' data centre was hosting LoopX video session manager software and Mapping engine.

The mapping engine hosted on MEC plays a critical role in the collaborative operation of delivery robots sharing the same route. In case of an obstruction en route, robots can dynamically update the map on MEC enabling other robots to download an updated map and avoid routes blocked by an obstruction. Without a centralized mapping engine, Teleoperator intervention is needed to reroute the robot after detecting obstruction on the sidewalk or road. It is important to note that maps used for the autonomous operation of vehicles are not the same maps used for turn-by-turn navigation in use by human drivers. Autonomous vehicle operation needs HD maps which are more precise and specifically built for the safe operation of autonomous vehicles. Those maps need to be constantly updated.

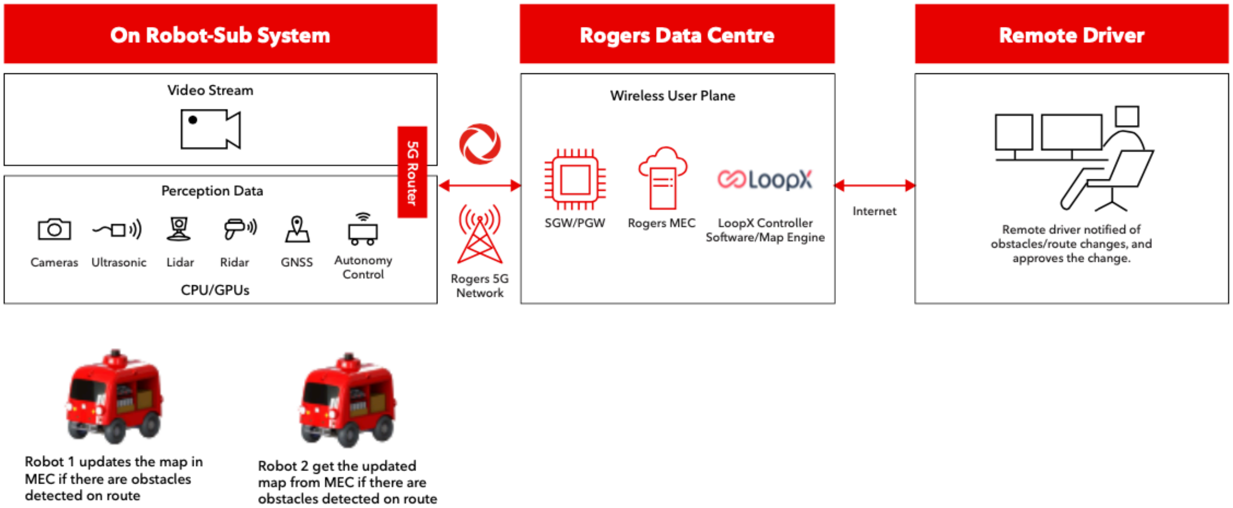


**Figure 12 – Last Mile Food Delivery Robot POC Architecture**

The other function which MEC software performed in this case was brokering the video session between Teleoperator and robots equipped with HD cameras before video streaming began. This improved the response time to start video streaming in comparison to controlling traffic using the public internet. Also, by hosting the video control on MEC, the overall security of the solution is improved.

Robots were communicating over Rogers's 5G network to application functions hosted on the MEC server in Roger's data centre. In this particular POC, the teleoperator driver was using wireline access over the internet to connect to MEC. The teleoperator driver ideally would connect to the MEC server over a dedicated wireline connection so that end-to-end control traffic does not cross the internet and the connection is private. This would enhance the overall security of the solution and provide more predictable service quality.

### 7.3. Conclusion

By hosting the controller software and mapping engine on the MEC server, the overall responsiveness and security of the solution was improved.

Due to project time constraints, the team was only able to offload only certain functions from onboard computers to MEC.

This particular use case can be further expanded to offload compute-intensive and more latency tolerant (comparing latency with on-prem versus MEC) workloads from Robots to MEC potentially extending battery life on robots and reducing cost for the compute and GPU resources.

It is worth noting that 5GAA has published various use cases related to autonomous vehicles area with corresponding SLR (service level requirements) which can be candidates for MEC. However, most of those use cases are uplink heavy and at scale, deployment of those use cases is a bit challenging at this point.

 

## 8. Overall Conclusion

As transport latency is a function of geographic and fibre distance between source and destination endpoints, latency benefits of operator MEC are largely defined by how close or far are public cloud provider regions and data centres are located to the user endpoint. However, traffic traversing over best-effort internet from Network operators to cloud provider data centre hosting applications results in higher jitter values not suitable for some applications although geographic distances are not that far.

Routing optimization and direct peering between Network operators and hyper-scale cloud providers also improve the overall latency for workload hosted in cloud data centres minimizing the benefits of MEC infra hosted relatively closer to public cloud provider data centres.

Regardless of the above factors, today's applications are built based on microservices architectural principles and can be deployed in a highly distributed way. Applications workloads which are latency or jitter-sensitive and bandwidth-heavy can be deployed at the network edge while large-scale workloads not so sensitive to latency can be deployed at hyper-scale cloud data centres.

The adoption of operator MEC would also largely depend upon the economic feasibility of moving latency-sensitive device edge workloads to the network edge.

There are a lot of opportunities to offload compute-intensive workloads like video rendering, video analytics and computer vision from the device edge or on-prem edge to the operator MEC or network edge. As most of these use cases are upstream-heavy, so the savings in compute cost should offset the network costs to build capacity and transport data to the network edge at increased volume and higher peak rates.

Uplink transport costs and peak capacity will become lesser of an issue with increased 5G spectrum availability and costs of data transmission going down.

It would also be important to note that low latency mission-critical use cases like remote surgeries and C-V2X need ultra-reliable network and latency reduction itself is not enough to host applications at the network edge.

Another area which can help with the wider adoption of MEC and make it commercially viable for Network operators is video caching at the edge using Network operator MEC infrastructure. Network operators are already hosting caching solutions provided by Akamai and other caching providers at their facilities.

Adoption of MEC is at a slow pace due to various reasons like the development of low latency use cases, commercial viability, lack of wider upstream spectrum, reliability of network (for mission-critical use case) and lack of 5G device ecosystem.

However, as some of those challenges ease with time. MEC has the potential to grow as a vertical for Network operators to be a source of additional revenue.

There are areas which need to be developed and need to be looked into further for MEC commercialization. For example, just-in-time workload orchestration, optimization of compute and GPU resources, scaling and supporting MEC for ISVs and facilities upgrades. Some of those are being addressed through ETSI, 5GFF and other forums and topics of further studies.

# Abbreviations

| | |
|---|---|
| AWS | Amazon Web Services |
| AR | Augmented Reality |
| C-V2X | Cellular vehicle to anything |
| CAV | Connected and Autonomous Vehicles |
| DOCSIS | Data over cable system Interface specification |
| EDS | Edge Discovery Service |
| ETSI | European Telecommunications Standard Institute |
| GNSS | Global Navigational Support System |
| GPU | Graphical processing unit |
| HCP | Hyper-scale cloud provider |
| ISV | Independent software vendor |
| MEC | Multi-access Edge computing |
| MWC | Mobile World Congress |
| OVIN | Ontario Vehicle Innovation Network |
| PGW | Packet Gateway |
| POC | Proof of concept |
| SME | Small to medium Enterprises |
| SGW | Serving Gateway |
| VNF | Virtual Network Function |
| VR | Virtual Reality |
| 5G | 5$^{th}$ Generation |
| 5GAA | 5G Automotive Association |
| 5GFF | 5G Future Forum |
| SCTE | Society of Cable Telecommunications Engineers |

# Bibliography & References

1. https://www.etsi.org/technologies/multi-access-edge-computing

2. https://www.smt.com/articles2019/puck-and-player-tracking-now-part-of-nhl-data-and-technology-offering

3. https://5gff.org/work-weve-done/?filter=whitepaper