analysys mason

Perspective

Assuring multi-access edge computing (MEC)

February 2022

Neil Kiritharan and Justin van der Lande

Contents

1. Executive summary 1

2. Landscape 2

2.1 Many new opportunities are arising from MEC 3

2.2 The Open RAN is changing network infrastructure and CSPs could use MEC to deploy it 3

2.3 CSPs could use public cloud provider infrastructure for enterprise MEC services 4

2.4 NaaS is opening up the ecosystem to new developers 4

3. Challenges 4

3.1 New architecture and standards always cause issues, and using MEC will only add to the problem 4

3.2 Infrastructure fragmentation and the multi-vendor aspect of MEC will add complexity 5

3.3 MEC's flexibility means that many enterprise service options must be managed 6

3.4 New enterprise services need self service and automated provisioning and monitoring 6

3.5 CSPs may have limited access to service metrics when using MEC 6

4. Solutions 7

4.1 Passive monitoring can provide detailed packet-based analysis 7

4.2 ML/AI can be used to dynamically understand the implications of applications 8

4.3 CSPs must understand the KPIs needed to monitor the QoE for specific application types 9

4.4 MEC assurance solutions must be dynamic and quick to deploy 9

4.5 MEC assurance solutions must support Open RAN, distributed Kubernetes and PCPs' stacks 9

5. About RADCOM 10

6. About the authors 12

List of figures

Figure 1.1: Overview of RADCOM ACE, an example passive monitoring solution 2

Figure 2.1: Overview of the networking stack in next-generation networks 3

Figure 3.1: CSPs' top challenges when deploying Open RAN 5

Figure 4.1: Overview of RADCOM ACE, an example passive monitoring solution 8

Figure 4.2: More-detailed breakdown of RADCOM ACE 8

This perspective was commissioned by RADCOM. Usage is subject to the terms and conditions in our copyright notice. Analysys Mason does not endorse any of the vendor's products or services.

1. Executive summary

5G networks are starting to affect both communications service providers (CSPs) and the support that they can offer other industries. The roll-out of 5G is acting as a catalyst for transformation in terms of the way in which networks are built, the types of services that are offered and the creation of new opportunities such as services based on multi-access edge computing (MEC). CSPs are gradually migrating from legacy network infrastructure to next-generation 5G networks with service-based architecture, backed by cloud-native, disaggregated and virtualised infrastructure. Edge cloud technologies are a significant enabler of this change; they bring computing closer to the point of service delivery to reduce latency and increase performance This is critical for the delivery of new applications such as network slicing, software-defined networking (SDN) and network function virtualisation (NFV).

MEC does not necessarily need to use 5G networks, but the two are closely linked due to the timing of the roll-outs of both technologies and the ability for 5G to support new services. MEC provides a general-purpose compute platform that can be used for both internal and external use cases. Internal use cases include supporting network functions or Open RAN solutions that require access to very-low-latency compute. More-advanced internal use cases support network slicing and the implementation of specific network services that require a flexible compute capability that is dependent on demand.

External use cases support enterprises that wish to take advantage of very low latency or local compute and storage capabilities for specific applications or requirements. The take-up of 5G enterprise services will be highly dependent on their reliability compared to that of current wireline services and their ability to support deterministic services and meet specific service-level agreements (SLAs).

Each MEC enterprise service will need to be assured in order to monitor a service's conformance to the KPIs outlined in these SLAs; metrics such as reliability, latency and throughput must be at the agreed levels. This means that 5G has a chance of commercial failure unless operational systems are changed. It is therefore imperative that a new approach is used to increase the level of automation for operational and assurance processes, to reduce costs, to improve efficiency and to enable 5G to be a success.

Assuring MEC-based workloads is complex due to the highly dispersed nature of the workloads and their geographical remoteness from what would have historically been hosted in larger data centres or provided through robust specialised networking equipment. In addition, MEC can be provided in collaboration with public cloud providers; CSPs may elect to deployed public-cloud-based solutions at the network edge.

Providing an updated assurance solution that can cope with the new infrastructure and service types that 5G entails is a key part of the required change in operational systems. Processing assurance at the edge can allow for automated optimisation at edge nodes, meaning that common control plane issues can be dealt with on local functions without having to send traffic across the network or clog up central management systems. In addition, MEC requires continuous monitoring across many nodes in different, unmanned locations; this calls for a monitoring solution that can provide detailed, packet-based analysis.

Passive monitoring helps to provide deeper insights into network performance. Real end-user data flows build up over a period of time, thereby allowing for trends to be identified for each element or network service. SLAs are monitored by assessing real user data to help to align service monitoring with actual customer usage. Moreover, using real data means that the issues that are identified are focused on the services that are being used by customers; simulated traffic may not accurately reflect real-world performance. Passive test data can also be used to pinpoint issues when fine-grain data is needed for fault finding and to help set up new SLAs that contain KPIs based on real end-user data.

Passive monitoring provides a highly detailed analysis after events occur; data can be collected on metrics such as bandwidth use, application use and signalling performance for deeper troubleshooting and root-cause analysis.

Figure 1.1 shows an overview of RADCOM ACE, a passive assurance system. It offers:

- network visibility that does not affect performance
- service assurance with real-time processing that also integrates inputs from other probes or assurance offerings
- network insights that are built upon a deeper analysis of service monitoring and that are automated, containerised and integrated into network functions.

Figure 1.1: Overview of RADCOM ACE, an example passive monitoring solution

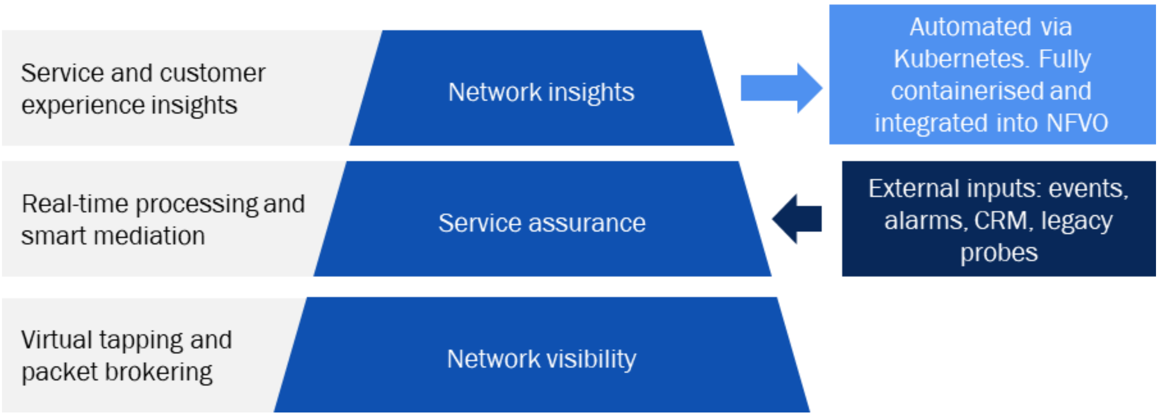
Service and customer experience insights
Network insights
Automated via Kubernetes. Fully containerised and integrated into NFVO
Real-time processing and smart mediation
Service assurance
External inputs: events, alarms, CRM, legacy probes
Virtual tapping and packet brokering
Network visibility

Source: RADCOM, 2022

2. Landscape

5G networks will affect the telecoms industry in several ways. For example, they will increase CSPs’ potential for digital transformation, add opportunities for new enterprise use cases and provide the ability to support MEC. CSPs will gradually migrate from legacy network infrastructure to next-generation 5G networks with service-based architecture, backed by cloud-native, disaggregated and virtualised infrastructure. Such infrastructure will become increasingly significant in the network stack, as depicted in Figure 2.1. Edge cloud technologies will be a significant enabler of this change; they will bring computing closer to the point of service delivery in order to decrease latency and increase performance. Network slicing, software-defined networking (SDN) and network function virtualisation (NFV) will also become crucial.

Figure 2.1: Overview of the networking stack in next-generation networks

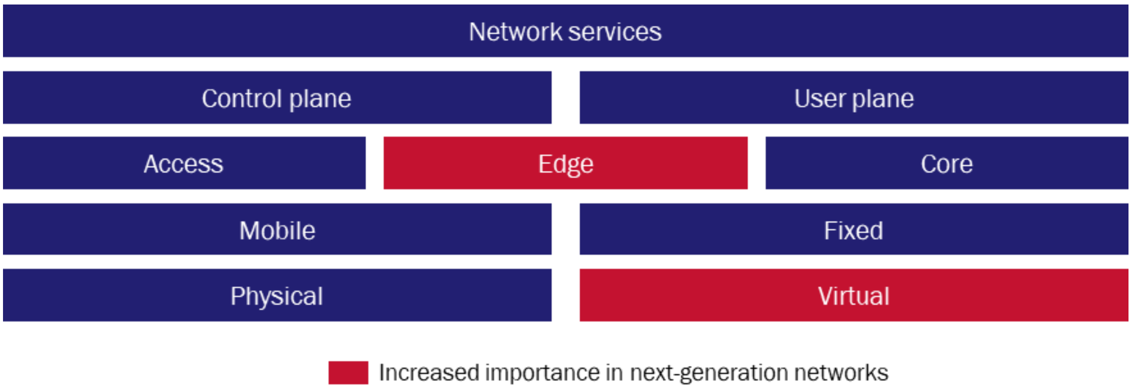

Source: Analysys Mason, 2022

2.1 Many new opportunities are arising from MEC

MEC will facilitate some of the most important new opportunities in the 5G era. Many digital use cases of the future will need a combination of ultra-low latency and network slicing, and this in turn will require edge computing and 5G capabilities. MEC offers a platform for both internal and external use cases (internal use cases are those related to CSPs' network services, while external use cases are those related to applications for consumers or enterprises).

External MEC use cases will benefit from low latencies due to the close proximity of connectivity and computing to the applications. Edge cloud will also provide storage capabilities for new enterprise applications. The benefits for internal MEC use cases are primarily related to the virtualisation and distribution of 5G core network functions to the edge; this will enable more-agile capacity planning and will improve the efficiency of capital allocation and dynamic CPU and memory utilisation.

2.2 The Open RAN is changing network infrastructure and CSPs could use MEC to deploy it

MEC will also help CSPs to deploy the Open RAN. Traditional RAN solutions use dedicated software and hardware from specialised vendors, while the Open RAN uses open standards and white-box solutions from different types of vendors. The results of a survey of 300 CSPs conducted by Analysys Mason in 3Q 2020 highlight the key benefits and drawbacks of the Open RAN.[1] The main benefits are a reduction in the total cost of ownership (TCO), an increase in supplier diversity, a reduction in the time to market for new services and a broadening of the innovation base. The concerns raised by the survey respondents included the cost and complexity of integrating and running a network with components from different suppliers. MEC addresses these concerns by providing a platform upon which cloud-native, multi-vendor solutions can be deployed to increase the flexibility and agility of network operations.

[1] For more information, see Analysys Mason's *Open RAN: ready for prime time?*

2.3 CSPs could use public cloud provider infrastructure for enterprise MEC services

CSPs own or have access to many locations; this should prove useful for edge computing infrastructure. CSPs have cell sites, central offices and metro data centres in areas with high population densities (often near to enterprises) in order to maintain good population coverage. If CSPs were to offer edge cloud capabilities in these locations, they could provide an attractive proposition for enterprises with new, edge cloud use cases.

Public cloud providers (PCPs) are already targeting enterprise verticals such as the healthcare and manufacturing sectors with potential edge computing use cases. They have developed versions of cloud technology stacks that can be used on enterprise platforms, and have built developer ecosystems around platform services that are particularly relevant to edge cloud applications, such as IoT device management and AI analytics. CSPs could partner with PCPs to take advantage of this infrastructure and increase their network capabilities at relatively low incremental cost. Such partnerships would also provide PCPs with access to revenue from enterprise sites without having to build out costly edge networks across geographically distributed locations.

2.4 NaaS is opening up the ecosystem to new developers

Network-as-a-service (NaaS) allows resellers of wholesale connectivity to sell networking flexibly and on-demand, and is a business model that is likely to become more common in the future. Resellers tend to have their own front-end systems for selling connectivity, so open APIs that can be integrated into these systems will increase flexibility and reduce integration work. Open APIs will also enable resellers' automations for provisioning, monitoring and bill flow (if available) to work directly with the network. Furthermore, open APIs for NaaS can increase innovation, both in front-end systems and future enterprise services, thereby helping to develop applications that call up certain slices or bits of the network in specific ways. With exposure to inventory management systems, open APIs could also be used for quoting. For example, they could be used to autonomously check whether the assets and capacity are in place to support a client without checking with the underlying wholesale network provider. NaaS is currently more common in the fixed market due to the popularity of wholesale fibre reselling, but it is likely to become more relevant in the mobile space during the 5G era as CSPs seek to share towerco infrastructure, either through NaaS directly or via similar open API practices that will enable better infrastructure sharing.

3. Challenges

3.1 New architecture and standards always cause issues, and using MEC will only add to the problem

5G has the potential to generate many new opportunities for CSPs, but it is possible that too much will change in the way that networks, infrastructure and services are delivered for CSPs to profit from the new technology. 5G is expensive to deploy; the RAN market alone will be worth USD62.7 billion worldwide by 2026 as a result of mobile network operators (MNOs) deploying 5G networks.[2] When the costs for the 5G core and 5G transport networks are added to this, it becomes evident that 5G represents a significant investment for CSPs in terms of software, hardware and new services for their mobile networks.

[2] For more information, see Analysys Mason's *RAN worldwide forecast 2021–2026*.

There are always issues with initial implementations when moving to new architecture and standards; CSPs have to deal with teething issues when changing interactions between components or within the software stack. However, this issue will be exacerbated when adopting 5G due to the significant move from legacy networks to a disaggregated, cloud-native, virtualised infrastructure with service-based architecture. The dramatic difference between next-generation and legacy technologies presents plenty of challenges before even considering the impacts of the additional requirements caused by the introduction of MEC.

3.2 Infrastructure fragmentation and the multi-vendor aspect of MEC will add complexity

MEC will add further complexity due to the fragmentation of the RAN ecosystem. As mentioned previously, MEC will provide a cloud-native, multi-vendor platform and will help CSPs to introduce Open RAN into next-generation networks. The multi-vendor aspect of MEC means that there are numerous possible combinations of vendor components, thereby increasing complexity because each potential cross-vendor interaction will need to be tested and optimised.

Similarly, the white-box nature of Open RAN will also support a large number of possible component permutations, thereby increasing complexity and management costs. Open RAN will present an opportunity for CSPs to make savings on the significant expenditures that 5G will require by reducing vendor lock-in and promoting competition, and so is likely to be widely adopted. However, it is possible that excessive fragmentation could mitigate or eliminate these cost savings. Figure 3.1 shows the concerns of 300 CSPs surveyed by Analysys Mason in 3Q 2020 regarding deploying Open RAN. Integration cost and complexity are considered to be the primary challenges by the majority of survey participants by a long way. Indeed, Rakuten Mobile and Dish Network are the only two major adopters of Open RAN for a large-scale commercial network so far.

Figure 3.1: CSPs' top challenges when deploying Open RAN[3]

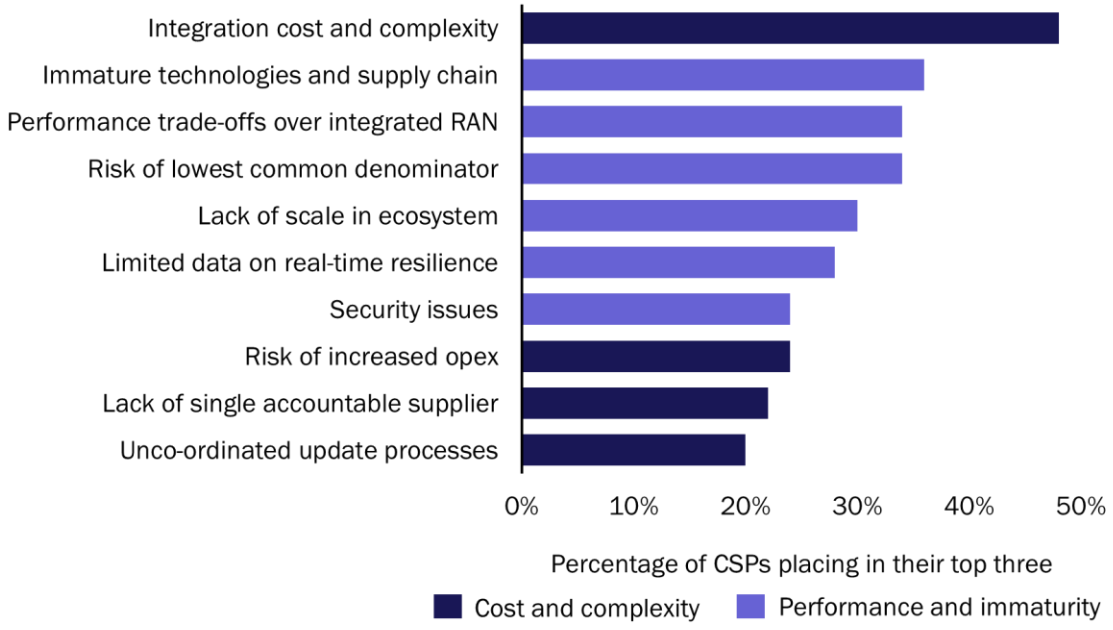

Source: Analysys Mason, 2022

[3] For more information, see Analysys Mason's *Open RAN: ready for prime time?*

3.3 MEC's flexibility means that many enterprise service options must be managed

The flexibility of MEC and the number of new enterprise service options that it enables will also increase the amount of testing and management required. On top of this, many CSPs currently use proprietary solutions to integrate OSS/BSS components and the underlying networks, with an overreliance on manual interventions for service design, activation and assurance activities.[4] Significant labour is required to manage an ever-increasing number of integrations; this increases the chance of manual errors and decreased the reliability of the service, thereby resulting in massive maintenance costs. Wholesale changes to the approach will be required for the cost-effective management of new MEC enterprise service options. Indeed, without automated service provisioning and assurance, new MEC-based infrastructure and services could hinder CSPs' 5G plans by excessively increasing operational costs.

3.4 New enterprise services need self service and automated provisioning and monitoring

Enterprises often manage cloud-based services such as infrastructure- and software-as-a-service (IaaS/SaaS) through online portals that enable them to order, provision, monitor and modify services on demand. As such, self-service experiences for new 5G MEC offerings may play a key role in encouraging adoption. In addition, new enterprise services enabled by 5G MEC are likely to be faster-moving and far more changeable than those that existed on previous networks; they will therefore need to have shorter set-up and deployment timescales to meet enterprise expectations. Both of these expectations mean that new 5G MEC enterprise services will require automated fulfilment processes that support self service and automated provisioning and monitoring.

3.5 CSPs may have limited access to service metrics when using MEC

The service assurance of MEC is challenging because it falls outside of the typical support structures. This is because MEC instances are potentially highly distributed over a large number of unmanned geographical locations and need to run automatically. For example, MEC nodes that are used as part of a telecoms network for delivering virtual network functions must have an availability of 99.999% if they are to support public safety services and new services such as autonomous driving. Moreover, MEC instances that are being used for 5G enterprise services will need to support new, more-specific SLAs in order to entice companies into shifting from their current fixed network services. Indeed, it is difficult to see MEC as a viable alternative to PCPs' or on-premises servers without guarantees of reliability and performance. Each MEC enterprise service will need to be assured to monitor the service's conformance to the KPIs outlined in SLAs.

In the future, CSPs may face additional problems when monitoring KPIs (for SLAs or otherwise). Currently, all network components are owned by CSPs, so these players have full access to all the performance information that is available. However, this may not always be the case when infrastructure sharing becomes more commonplace with the use of 5G and MEC. CSPs may find that they do not have access to all service metrics and parameters when they do not physically own the box running the virtualised network functions. Metrics may only be available through third-party interfaces (from the owner of the component), and there may be incompatibilities, delays or incomplete data. The control of the methodologies used to create third-party data may not be in the hands of the CSP.

[4] For more information, see Analysys Mason's *The TM Forum NaaS framework decouples the network and OSS to enable a two-pronged digital transformation*.

4. Solutions

Thus, the challenge is clear; 5G has a chance of commercial failure unless operational systems are changed. It is therefore imperative that a new approach is used to increase the level of automation for operational and assurance processes in order to reduce costs, improve efficiency and enable 5G to be a success.

4.1 Passive monitoring can provide detailed packet-based analysis

Operational systems of the future must support an updated assurance solution that is able to cope with the new infrastructure and service types of the 5G era. Processing assurance at the edge enables automated optimisation at edge nodes, which means that common control plane issues can be dealt with on local functions without having to send traffic across the network or clog up central management systems. In addition, MEC requires continuous monitoring across a large number of nodes in different, unmanned locations, which calls for a monitoring solution that can provide detailed, packet-based analysis.

Passive monitoring helps to provide deep insights into network performance using real end-user data (not synthetic data) and has a key role to play in the assurance of MEC. It may be preferable over other approaches for the following reasons.

- Some assurance solutions test network performance at an instant in time, which means that problems that do not occur at the specific moment of testing are not detected. Passive probes provide continuous data monitoring without impeding performance or adding traffic to the network and are therefore more likely to pick up trends of gradual degradation that active testing may not detect.

- Passive monitoring uses real performance data from end users, while active monitoring uses simulated data. Synthetic traffic injection has its benefits (notably that testing is ‘always on’ and can take place before networks are live), but real end-user data may differ from simulated traffic in unexpected ways, and the most important factor for enterprises is real-world performance. For this reason, new SLAs may contain requirements for the testing of real end-use data.

- Passive monitoring can provide detailed analysis after events occur, and collects data on metrics such as bandwidth use, application use and signalling performance for deeper troubleshooting and root-cause analysis.

Figure 4.1 shows an overview of RADCOM ACE, an example of a passive assurance system. It offers:

- network visibility that does not affect performance
- service assurance with real-time processing that also integrates inputs from other probes or assurance offerings
- network insights that are built upon a deeper analysis of service monitoring and are automated, containerised and integrated into network functions.

Figure 4.1: Overview of RADCOM ACE, an example passive monitoring solution

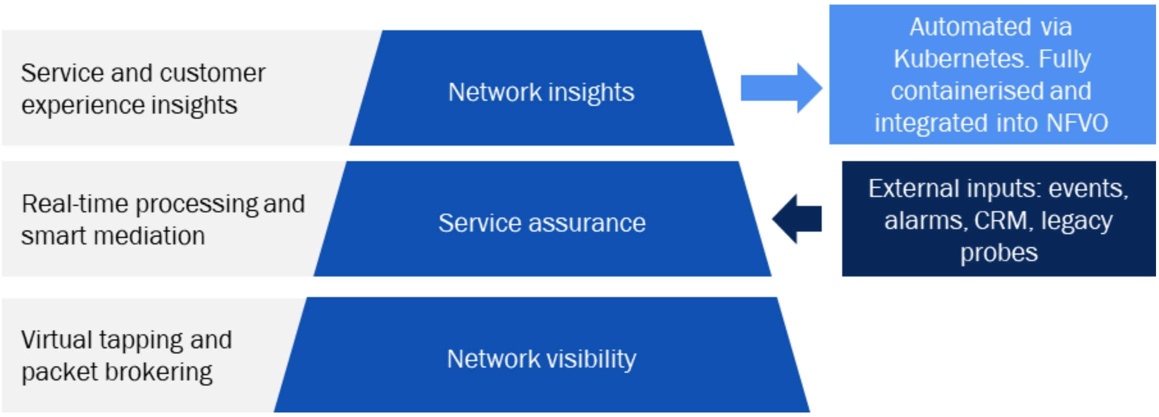
Service and customer experience insights
Network insights
Automated via Kubernetes. Fully containerised and integrated into NFVO
Real-time processing and smart mediation
Service assurance
External inputs: events, alarms, CRM, legacy probes
Virtual tapping and packet brokering
Network visibility

Source: RADCOM, 2022

4.2 ML/AI can be used to dynamically understand the implications of applications

Passive assurance will continuously generate significant amounts of data about network performance, and this can be used to create insights about the impact of new applications on the network itself. New applications can affect traffic engineering performance and security and can cause other performance issues. Machine learning (ML) and artificial intelligence (AI) will be required to dynamically understand the effects of applications on the 5G MEC network. They can be used to make sense of the detailed network data and to understand trends and anomalies. Figure 4.2 shows a more detailed diagram of the RADCOM ACE passive monitoring solution, and showcasing the use of AI/ML in the service assurance layer to generate deep, dynamic network insights.

Figure 4.2: More-detailed breakdown of RADCOM ACE

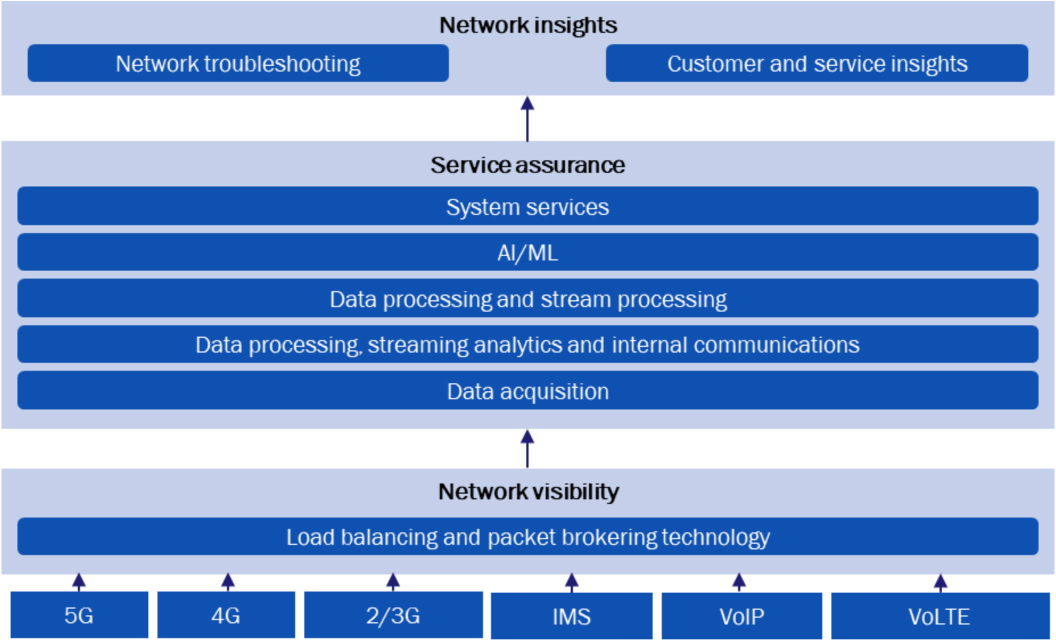
Network insights
Network troubleshooting
Customer and service insights
Service assurance
System services
AI/ML
Data processing and stream processing
Data processing, streaming analytics and internal communications
Data acquisition
Network visibility
Load balancing and packet brokering technology
5G
4G
2/3G
IMS
VoIP
VoLTE

Source: RADCOM, 2022

4.3 CSPs must understand the KPIs needed to monitor the QoE for specific application types

It is important that CSPs and assurance providers understand the KPIs needed to monitor the quality of experience (QoE) changes for different application types. Each class of application has specific characteristics that are important to functionality; for example, VR/AR applications that may be enabled by 5G will require ultra-low-latency of sub-1ms to avoid user nausea, but they will also require significant throughput to enable the streaming of high-quality assets or video. Autonomous vehicles provide another class of application with a different set of requirements; collision prevention and traffic management require extremely low latencies (if they are offloaded from on-board computers to MEC instances) and perfect reliability (>99.999%).

4.4 MEC assurance solutions must be dynamic and quick to deploy

There are many potential 5G enterprise applications, and these applications are likely to be updated on a much more frequent basis than those in previous network generations, so it is essential that MEC assurance solutions can be developed and deployed quickly. Assurance solutions that can be deployed alongside new edge services and dynamically adapt to updates or network topology changes will simplify application roll-out and allow CSPs to keep up with application development. For example, Amazon Elastic Kubernetes Service Anywhere (EKS-A) allows CSPs to run Kubernetes outside of AWS on their own infrastructure (including in CSPs' own data centres). RADCOM ACE works in conjunction with EKS-A to enable the quick deployment of dynamic edge services with instant optimisation.

4.5 MEC assurance solutions must support Open RAN, distributed Kubernetes and PCPs' stacks

Support for Open RAN and distributed Kubernetes is essential for any MEC assurance solution due to the significant benefits that these technologies can bring to next-generation networks. Indeed, Open RAN is attractive to CSPs due to the removal of vendor lock-in and the promotion of competition, and distributed Kubernetes allows for containerisation, the decoupling of applications from host infrastructure and easier cloud deployment. Rakuten and Dish already use Open RAN solutions in their commercial networks, so any assurance system that does not support such solutions will be incompatible. The number of such incompatible networks will only grow.

Maximising compatibility will also require agnostic support for PCPs' stacks. CSPs will partner with PCPs to reduce infrastructure capex by removing the need to develop aspects of the software stack. MEC assurance solutions that support PCPs' solutions agnostically therefore do not need to integrate with each PCP individually.

Standardisation such as this will become increasingly crucial as network complexity grows due to the increasing number of permutations of multi-vendor offerings and new enterprise services. In a similar way, any assurance solution for 5G MEC must be linked to 3GPP NWDAF data sources. NWDAF is a 3GPP standard that enables interactions between different vendors, thereby helping to limit potential problems related to the fragmentation of the 5G ecosystem. NWDAF retrieves event information from 5G core network functions and represents the data in a standard format, but also manages data and derived insights and handles data from anywhere within the network. NWDAF adoption would greatly help MEC to be successful, and so assurance systems for MEC must also support this standard.

5. About RADCOM

This perspective has been sponsored by RADCOM. Analysys Mason does not endorse any of the vendor's products or services.

RADCOM ACE is a passive monitoring solution that has been optimised for MEC and supports control/user plane separation (CUPS) correlation (using its patented solution), thereby providing CSPs with complete network visibility and troubleshooting capabilities for edge site deployments.

RADCOM ACE provides:

- network visibility that does not affect performance
- service assurance with real-time processing that also integrates inputs from other probes or assurance offerings
- network insights built upon deeper analysis of service monitoring, that are automated, containerised and integrated into network functions.

RADCOM ACE works in conjunction with EKS-A, which allows CSPs to run Kubernetes outside of AWS on their own infrastructure (including in CSPs' own data centres. This combination allows CSPs to quickly deploy dynamic edge services with instant optimisation.

The advantages of RADCOM ACE are as follows.

- **Fully cloud-native.** RADCOM ACE was developed from the ground up with a cloud-native design. It enables rapid deployment as containers or virtual machines on multiple cloud environments (public, private, and hybrid) such as Microsoft Azure and Amazon Web Services or on-premises. It uses Kubernetes or other leading orchestrators to automate the deployment, scaling and management of assurance components.

- **Advanced built-in AI and ML.** AI and ML capabilities are built into the solution from inception, thereby enabling a variety of AI-based use cases, such as anomaly detection, predictive and prescriptive customer experience analytics and predictive and dynamic slice quality KPIs.

- **Automated assurance.** RADCOM ACE automates solution deployment for the on-demand instantiation, scaling, healing and updating of a closed-loop approach to assurance. Kubernetes is used to control the containerised components lifecycle.

- **Provides real-time subscriber analytics.** RADCOM ACE enables CSPs to understand the customer experience from the RAN to core via the correlation of user plane and control plane traffic to identify customer-affecting problems in the network.

- **Built for 5G.** RADCOM ACE has full support for 5G standalone (auto-service detect, service-based architecture, service-based interface (SBI) deciphering and complex CUPS correlation, packet forwarding control protocol (PFCP), new aggregation and dimensioning).

- **Supports advanced 5G standalone use cases.** The solution smartly monitors private networks, edge computing and end-to-end network slicing.

- **Slicing and NWDAF support.** Individual slices can be monitored in real time to provide a plethora of KPIs and KQIs. Furthermore, RADCOM NWDAF enables closed-loop automation by processing real-time data from the assurance platform and triggering the relevant actions towards the PCF and NSSF, thereby resulting in the highest possible customer experience levels.

To find out more about RADCOM ACE, please visit the RADCOM website at:
https://radcom.com/solutions/5g-service-assurance/.

6. About the authors

Neil Kiritharan (Consultant) delivers strategy consulting projects in Analysys Mason's London office. He has worked with telecoms operators, technology vendors and national regulatory authorities, completing a diverse range of projects. Neil has significant desktop research and international benchmarking experience, using quantitative and qualitative analysis to develop market insights and recommendations. His recent project experience includes building market forecasts models, conducting primary research, leading interviews with C-suite executives, performing regulatory benchmarking and offering strategy recommendations on topics such as 5G, the Internet of Things (IoT), artificial intelligence (AI) and mobile edge computing (MEC).

Justin van der Lande (Research Director) leads the Applications practice, which is part of Analysys Mason’s Telecoms Software and Networks research stream. He specialises in business intelligence and analytics tools, which are used in all telecoms business processes and systems. In addition, Justin provides technical expertise for Analysys Mason in consultancy and bespoke large-scale custom research projects. He has more than 20 years’ experience in the communications industry in software development, marketing and research. He has held senior positions at NCR/AT&T, Micromuse (IBM), Granite Systems (Telcordia) and at the TM Forum. Justin holds a BSc in Management Science and Computer Studies from the University of Wales.

Analysys Mason Limited. Registered in England and Wales with company number 05177472. Registered office: North West Wing Bush House, Aldwych, London, England, WC2B 4PJ.

We have used reasonable care and skill to prepare this publication and are not responsible for any errors or omissions, or for the results obtained from the use of this publication. The opinions expressed are those of the authors only. All information is provided “as is”, with no guarantee of completeness or accuracy, and without warranty of any kind, express or implied, including, but not limited to warranties of performance, merchantability and fitness for a particular purpose. In no event will we be liable to you or any third party for any decision made or action taken in reliance on the information, including but not limited to investment decisions, or for any loss (including consequential, special or similar losses), even if advised of the possibility of such losses.