

JANUARY 2023

MOBILE EDGE FORUM 2022

KEY FINDINGS

REPORT SPONSOR:

intel

Editor's note: The teams from *RCR Wireless News* and *Enterprise IoT Insights* recently convened a wide range of telecoms and enterprise stakeholders to examine all things mobile edge computing at the Mobile Edge Forum virtual event. Here we present key findings from the program with contributions from *Enterprise IoT Insights* Editor James Blackam, and *RCR Wireless News* editors Kelly Hill, Sean Kinney, Catherine Sbeglia and Juan Pedro Tomas.

With mobile edge computing, technology is the easy piece

By Sean Kinney

The current operator-led model of delivering mobile edge computing and, at a higher level the combination of MEC and 5G, is all about optionality—meeting customers where they are by being cloud agnostic to accommodate buyer preference, by selling point solutions or a turnkey stack, by working out on-prem, public carve-outs or hybrid connectivity/compute models. But, at the end of the day, is that really what enterprises want? Can the analysis paralysis that comes with optionality limit market adoption? And is the technology piece, regardless of the wrapper it's in, secondary to the services and applications that will drive meaningful digital transformation of the enterprise.

In pursuit of answers to the above questions, we asked a range of experts, including infrastructure and service providers, as well as the enterprises being asked to execute on new types of technology strategies. Through the course

of these conversations, a few truths became clear and potentially raised more questions than they answered.

First, as aptly stated by an Orange executive, delivering the value of 5G and mobile edge computing isn't a matter of infrastructure, rather it's a matter of services and applications. But to create a unified value proposition combining a network, distributed compute and an application layer, is very much an exercise in ecosystem development. And that ecosystem needs to coalesce in a right-on-time manner to address the business problems that cost companies time and money.

Second, it's important that sellers of mobile edge computing solutions put themselves in the end users' shoes, particularly as it relates to system level performance. If you've got a solution that combines an operator network, a hyperscaler compute node and a developer's application, the user doesn't care where or why latency has entered the mix. The solution provider(s) need to think holistically about quality of experience and quality of service, and confidently deliver a bullet-proof

service level agreement based on real world conditions, not based on aspirational KPIs and not limited to the particular piece of a solution one party directly controls.

Finally, and returning to the double-edged sword that is optionality, there's no one-size-fits-all deployment or delivery model that will win out in the market; the market is far too varied to accept a cookie-cutter approach. But, at the same time, the sell side value chain needs something replicable and scalable that comes with a reasonable time-to-return. This (re)emphasizes the need for studied and focused ecosystem development with an eye on building horizontal platforms with the inherent flexibility to address key vertical use cases. Computer vision is a perfect example—with some minor tweaks, cameras tied to distributed compute running analytics software can be used for access control in commercial buildings, quality assurance on a factory floor, parcel sorting in a logistics facility, monitoring inventory in a retail establishment, rinse and repeat, and repeat, and repeat.



Four market trends shaping mobile edge computing

By Sean Kinney

Some mobile edge computing trends are clear, but questions remain

Tantra Analyst Founder and Principal Prakash Sangam shared his perspective on how stakeholders should consider mobile edge computing investments, and looked at how enterprises and consumers stand to benefit from distributed cloud computing infrastructure. Sangam laid out four vectors, and an overarching capability, that will make or break the success of mobile edge computing.

Sangam's four vectors shaping mobile edge computing are highly-reliable connectivity, ultra-low latency, extreme capacity and power-efficient compute for artificial intelligence and machine learning. Readers of RCR Wireless News will no doubt recognize those first three vectors as staples of the 5G value proposition while the former is certainly coming of age as semiconductor firms look to power new

experiences requiring advancements in chip design and the latest-generation cellular air interface.

To his latency point, Sangam pointed out that this isn't just about application latency, compute latency or latency over the wireless link; rather it's a measure of the total latency needed to deliver a particular outcome. "Latency experienced by the end user," as he put it. For the network operators, compute providers, device OEMs and others involved, this means "making sure that you not only provide low latency but provide the exact latency" needed to deliver quality of service and quality of experience performance thresholds, and to adhere to enterprise service level agreements.

Looking at the need for power-efficient computing, Sangam called out that edge computing nodes, by their very nature and in comparison to centralized

cloud computing facilities, don't have the same luxury of space, power, thermal dissipation, and so on. "When you're looking at edge cloud deployments, they're highly constrained," he said. "When you consider all of that, the compute has to be extremely power efficient...AI/ML in the edge cloud is a big topic on its own."

He continued: "These are the four vectors I would say...that define edge, that allow you to provide the things that are expected from the edge. But it's not just offering all of them, but the ability to scale across these vectors... If you have to be at the peak of all of these vectors, you will end up at lots of money to implement and deploy those things which may not be cost-effective for the applications you're running... The whole idea is you scale the network across these vectors to provide cost-effective connectivity and compute that is needed for the specific type of application or service. That is the key to providing a good edge."

Questions to ponder as MEC evolves

With obvious opportunities for mobile edge computing—a [projection from Gartner](#) that more than half of enterprise-managed data being created and processed outside of data centers by 2025 for example—the picture still isn't crystal clear. Sangam



Prakash Sangam

sketched out the still-developing dynamic between mobile network operators and hyperscalers.

“Edge computing is where these two ecosystems either merge or collide or compete,” he said. “They have to collaborate.” He observed that right now it would appear operators are in the driver’s seat, although worth noting that major hyperscalers seem keen to capture the opportunity that sits at the confluence of private networks and edge computing.

Another question Sangam posed as worth considering is the need for multi-operator, multi-cloud deployments. One operator working with one cloud provider can be problematic for an enterprise with operations in multiple countries, for example. Complications can also arise when one operator works with one cloud provider, and that cloud provider is not the partner of choice for the end user. “It has to be multi-cloud and multi-operator. That is also complicated,” he acknowledged.

Bottomline, Sangam characterized mobile edge computing as bringing a “middle layer” that avoids the cost and time of hauling data back to the cloud for processing. Leveraging the intelligence of that middle layer, allows for faster action based on data intelligence. “No matter what you call it, it is very clear from the industry trends that the value and a lot of the intelligence is moving to the edge of the network... Immediate processing and intelligence will happen at the edge cloud.”



Grasping the opportunity – are telcos ready for the edge?

By Sean Kinney

How do 5G and edge complement one another?

The distribution of compute out of centralized data centers and closer to where data is generated is seen as a key complement to latency-sensitive 5G services. Drawing from that, operators are approaching edge computing as a way to leverage 5G network investments, deliver differentiated services (largely to enterprise customers), and, from there, generate new service revenues. It sounds easy on paper but has proved much more difficult in the real world. To understand the opportunity and associated challenges, and get an idea of where the industry sees itself today, we asked the experts.

Before diving into how telcos need

to be thinking about investing in and monetizing mobile edge computing, and network resources, we need alignment on the definition. Nitu Kaushal, managing director of intelligent edge for Europe at Accenture, said the “edge” could be a device, an on-prem compute node and, key for operators, in the network. “And this is really where we see MEC flourish and also define how we get to provide edge-based services which are coupled with the network. For us, where we see MEC really playing a role is through a compelling 5G network or a 4G network, a mobile private network, it’s about building an edge there.

Verizon has deployed both public mobile edge computing sites and on-prem systems designed for specific enterprise use cases; the operator also works with all three major hyperscalers. Babatunde Akindele, Verizon’s



Babatunde Akindele

director of technology strategy, said the operator puts “edge” in two buckets. “One, we announced something we call the public MEC, which is really having some of our cloud partners having that footprint within our network and making it available to be used for mass-scale applications...because you’re looking for latencies that are lower but not microsecond latencies...And then we get to what we call our private MEC which...is really on the premise coupled with a private 5G [or] private LTE network, being able to set that micro latency that is required for some really mission critical use cases.”

NTT Data’s Sharad Sharma, vice president and head of network and B2B business, expanded the aperture to “multi-access edge computing...When you talk to the customers more and more in this space, they always have a question about why mobile, why not Wi-Fi, right?...We talk more about multi-access edge; while definitely there are technologies, mobile technologies, playing an important role in bringing data and applications closer to the customer for different use cases.”

Kaushal described 5G and mobile edge computing as going “hand-in-hand together...For the carrier, it means enabling the enterprise or the business

to create services for real-time needs. So what needs to be driven with low latency as low as maybe even five milliseconds. And then what I’m also thinking about is for the consumer, it also means they could be experiencing brand new experiences that have not really been on the table today.” Specifically, she called out AR and VR as latency-sensitive services with positive outlooks for both enterprise and consumer adoption.

Sharma noted the importance of a 5G Standalone network deployment as a requisite to “actually enable the real 5G use cases in the industry. But is that available in the industry today? Are the operators ready with that? No...When we deploy the standalone, and we deploy it on premises with the edge, then only it makes a real use case...Without having an edge with the Standalone [5G] private network and infrastructure, we cannot actually get the benefit of all the 5G use cases that we are kind of looking for.”

Akindele pushed back on that assertion of Standalone 5G being a requirement for impactful mobile edge computing use cases...“From a 5G and MEC perspective, they go hand-in-hand...[but] it doesn’t mean you can’t do anything with LTE alone.” Further to that point, he said MEC is about latency but also about reliability. “Reliability is

key when you take an example of an industrial enterprise, right? Reliability is very important. Availability is super important for a lot of processes.”



Balancing tech and business at the edge

There are opportunities for both public and private mobile edge computing implementations, augmented by 5G, to deliver on consumer and enterprise use cases, and create new service revenues for operators and others in the value chain. But this opportunity is not without challenges, including the move to a Standalone 5G network and 5G as part of a larger “multi-access” edge continuum. Despite the issues that need to be sorted, the panel agreed that aligning technology investments with use cases and clear business outcomes is key to the development of a mobile edge computing ecosystem.

“For me, MEC is...about the use case,” Kaushal, managing director of intelligent said. She noted apparent use cases related to V2X and other



Sharad Sharma

automotive facing services.

“I see MEC playing really strong there...

The use case is really, really relevant.

The technology is there,

but I think the integration of it is still at the early stages.”

Akindele further examined the distinction between the technology and the business pieces of MEC. “Technology, in and of itself, is not a solution... This has to be a digital transformation story and we need to be able to help the enterprise understand how you take this journey, what you do first, and how you start to derive ROI versus just a use case approach that may not help you justify deployment of infrastructure that may be expensive.”

That means, he said, that large enterprise customers have to consider not just one use cases but multiple aligned use cases when mapping out ROI from mobile edge computing investments. “We’re not just talking one use case. So if I’m an enterprise and I’m going to spend millions on deploying an infrastructure, I’m not looking for one use case. I need to see a critical mass around the couple of processes that I can then transform to start to derive these values.”

In terms of use case development, Sharma cautioned that “the industry is dying with” proofs on concept, and it’s

time for operators to include not just customers’ business cases, but also the operators own ROI. “The caution comes out in the operators’ conversations is, ‘If I deploy this

MEC infrastructure...and create that availability of MEC, what does it bring?’ How much business does it bring to the operator? So that’s very important...It’s not just about the infrastructure..”

This echoes Akindele’s comment on technology not being a solution. As Sharma put it, “You have to transform, you have to address your problems, or you have to identify the opportunity enabled by MEC or 5G or whatever new infrastructure technology you’re talking about.”

Fragmentation, data sovereignty, and transitioning from communications service provider to digital service provider

More problems with scaling out MEC and delivering positive returns to sellers as identified by the panel. “In Europe, it’s still very fragmented,” Kaushal said. “We have different operators, we have different policies...around security, reliability and also with latency. And we see this also as not just a latency topic but just data management end-to-end for our customers. So when you start to think about that, then providing

something that’s a seamless solution feels really hard.”

A lack of industrial-grade devices—“They’re not there today,” Akindele said. “All these devices need to have the capability to utilize this. And then, in terms of technology itself, Release 16 of 3GPP, ultra reliable low latency, time-sensitive networking. These are all capabilities that need to be available...The device ecosystem and the network vendor ecosystem needs to also be able to deploy some things.”

Back to that whole how do operators make money off of mobile edge computing point. Akindele cited Verizon’s early investment. “You build it, they will come...You provide the infrastructure because until people can experiment, until they can start to use it, they don’t really know exactly how they will derive value. We can help them do that. And then it’s the same thing for 5G...Now what needs to happen is for each and every one of us, MNOs and stakeholders, to create the ecosystem that is required to help drive broader use cases.”



Nitu Kaushal



What key applications are helping Vodafone monetize MEC?

By Catherine Sbeglia

Vodafone on MEC: 'We believe in its potential'

Vodafone's Global Edge Computing and 5G Principal Joanna Newman and Senior Global Product Manager Elizabeth Rumsey recently talked about "the perfect storm" of tech advancements that's been brewing the past several years. And while 5G may be at the heart of this storm, Vodafone has also been focused on the coming opportunities around mobile edge computing (MEC), and in doing so, has identified five key areas to spur MEC monetization: Mixed reality, video analytics and streaming, live video production, V2X communication and artificial intelligence at the edge.

"We looked at the coming opportunity with MEC a few years ago ... because we recognized the potential for the 5G application space to be over a billion-euro market, just within

our European footprint over the next five to potentially 10 years," Rumsey shared. "We do believe in [its] potential, and we believe that one of the key factors to bring that potential to life is to make the experience easy for our customers." Therefore, she continued, Vodafone has prioritized bringing together 5G connectivity and the edge compute to create a "seamless infrastructure and application experience."

While Vodafone didn't question the potential of edge computing, it did initially have to ask the familiar question of "is there a killer use case?"

"Over the period ... there's been a question mark around, is there a killer app for edge? I would say, in my experience ... the answer to that is no," said Rumsey, adding that she believes this to be a good thing. "Because that shows the extent to which edge computing has use cases

that go across a wide range of industry verticals," she continued.

Mixed reality, video streaming and V2X at the edge

In partnership with Festo, Vodafone has been exploring AR and VR training applications. Newman provided more details: "There's a lot of really interesting use cases coming from the AR/VR training, especially for national utility type organizations. Whether it's an electricity grid or a gas turbine, or a solar farm, et cetera, there's a lot of mechanics and training that needs to be done to be a skilled worker that can go out in the field and be able to maintain things," she stated.

As part of its 5G pilots in Italy, Vodafone has conducted MEC trials in the live video streaming category, where it is enabling ambulances to connect directly to the hospital to provide live video streams of what is happening inside the ambulance. These streams allow first responders to share patient information and diagnostics with hospital staff before the patient even arrives on site.

"The hospital [is] receiving remote instructions, so that they [can] do their best to help the patient. And when the patient arrive[s] at the hospital, the hospital [is] best prepared to



Elizabeth Rumsey



Joanna Newman

support that patient,” Rumsey said.

In this particular example, a reliable, low-latency connection is needed on both sides:

The hospital where the video is being streamed and inside the ambulance where the stream is being produced live.

Newman acknowledged that for some use cases, like V2X, which are critical in the sense that they could be life affecting, edge computing is “a little ways away” from rollout. However, she did confirm that Vodafone is currently participating in some trials with V2X in Germany. Vodafone began developing an open V2X platform called Safer Transport for Europe Platform (STEP) in November 2021, and this September

announced that it is now available to developers in Germany.

While the speakers didn’t provide specific examples of AI at the edge — perhaps that is because it’s such a broad application — Newman did detail several additional scenarios in which intelligence at a local edge makes a material difference, from gaming to healthcare to government security.

“Don’t forget that [with] edge processing, because it’s done closest to the customer, you can also very easily comply with some local processing directives that apply to in-country data processing, be it for healthcare records or data sovereignty requirements,”

Newman explained. “There’s a lot of different elements where edge can be helpful [by] taking advantage of its low latency and its local processing to comply with local regulation that it can be quite useful.”

According to Rumsey, the majority of Vodafone’s customers using multiple cloud environments and the vast majority those that are, are working with more than one public cloud provider. “Therefore, our strategy is one of non-exclusivity,” she said, adding that Vodafone is committed to helping its customers manage this multi-vendor environment through its own multi-cloud platform. She also claimed that half of its customers are moving beyond the distributed cloud. “Which underpins our investment that we’ve been making in MEC,” she said.



Testing, low latency and the edge: Five things to know

By Kelly Hill

With latency becoming an increasingly important part of [enabling new services and monetizing 5G](#), just what does testing to assure low latency look like?

Rich McNally, senior director of mobile service strategy with Spirent Communications, discussed real-world observations and lessons on latency testing; here are five key takeaways from his observations.

1. Testing generally falls into three types, at different points in the life cycle. McNally said that Spirent breaks out testing in 1) design and launch assessment, 2) continuous network monitoring and 3) customer experience management.

Service design and launch might include something like an edge zone launch, he said, and would include testing that includes security assessment and also answers business questions like, is the service ready for launch in the given city or market, and how does it compare, performance-wise, to the competition?



Rich McNally

Continuous monitoring including typical service assurance and network change management testing, as well as service-level validation. But in changing networking, he said, Spirent is seeing that there is “an even more critical need for latency, segmentation and fault isolation, particularly due to the partnerships between operators and cloud providers, and the importance of quickly triaging where latency issues exist within hybrid networks.”

Additionally, he said, “Realizing that potential of edge services is going to require attention to making your enterprise customers happy. ... It’s really important to be able to model mission-critical enterprise app traffic ahead of launch, but also [ensure] that there’s a way to address customer support and performance issues that come up along the way, and to base that on quantified controlled results, and not just anecdotal complaints.”

2. Different tests have more or less usefulness in an edge scenario—but new approaches are being developed as well. Ping roundtrip time has long been the most common option and still has its usefulness, McNally said—“But it really has its shortcomings,” he added. Given routing differences in traffic types, pings can miss the priority given

to normal edge traffic because the ping itself is routed differently than application data. “It also doesn’t delineate between up and downlink latency,” McNally said. “This is important because different edge applications can have different dependencies on either direction. And if you’re not measuring it, you’re guessing, really.”

UDP (User Datagram Protocol) one-way measurements make up for some of these shortcomings, he added, since application data is routed via UDP and you can also get separate statistics on uplink and downlink latencies. McNally noted that packet sizes can be varied with UDP tests and emulation of different use cases and traffic profiles is possible. Using both ping roundtrip and UDP one-way measurements can cover a lot of bases for launch assessment, benchmarking, enterprise application emulation, and various service assurance use cases, he added.

But he also highlighted a new approach being worked on by the Broadband Forum, which Spirent has implemented, called DeltaQ or QED. It is UDP-based and separates uplink and downlink statistics, and “goes

another step further by allowing you to see latency contributions from geographic distance, from serialization/routing, and from network loading,” he explained. QED also ties in the concept of a quantitative time agreement, or QTA, in order to realize SLA commitments with customers, McNally said.

3. In real-world testing, edge zones successfully provide lower latency. Spirent conducted tests on publicly available edge services in various part of the world on a regular basis, and recently tested four cities and two networks in North America and Japan, using commercially available handsets. Latency ranged from less than 10 milliseconds to the low 20 ms for the downlink, and less than 10 ms to the high 20-ms range in the uplink for 5G traffic only.

But consistency of latency was problematic—and that’s important to a stable low-latency service. Downlink latency showed, generally, more consistency in cities tested in the U.S., while uplink latency was more consistent in Japan; however, the specifics varied by city.

4. There are multiple factors that impact latency, and some of them are going to persist for awhile. These



Stephen Douglas

include handoffs between different access networks to serve target customers in various geographies and challenging radio frequency environments for MEC applications with high mobility. McNally said that one big lesson that Spirent has learned is that service providers have to plan ahead to avoid delays in testing that can happen due to network security triggers in both edge and cloud zones. “You have to be careful not to trigger denial of service sensors and things of that nature” where normal testing procedures could accidentally set off a service shutdown.

5. There is still a lot to learn when it comes to enabling edge services—but good testing practices will help answer questions and reveal areas to work on. McNally said that active testing, both end-to-end and within a specific network, is “really critical to “capture issues before users do” and that testing capabilities need to be at-the-ready in order to quickly turn around solutions when issues arise.

McNally said that when it comes to latency, consistency is critical—and that in some cases, it may eventually overshadow throughput as the most important aspect of a service. With all these new edge services and applications that could come to be, there is no single latency standard that will suit everything. “When testing, it is important to model different data footprints of important applications. There’s not one test case that covers everything out there,” he said.



How does Open RAN intersect with edge computing?

By Kelly Hill

Mobile network operators, cable companies and hyperscalers are all investing in edge infrastructure, with market drivers including use cases from private cellular networking to edge robotics in industry—as well as disaggregated open networks, said Bejoy Pankajakshan, EVP and chief technology and strategy officer for Mavenir.

Pankajakshan said that a major shift in processing is coming. “Today, if we look at how the applications are being processed, less than 10% of the traffic is actually processed outside of a traditional data center, which comes from the operator,” he said. “If we look at 10 years out, the projections that I’ve seen, it talks about up to 75% of the processing would actually happen outside of the traditional data centers.” Expected verticals with anticipated growth in edge use include remote healthcare diagnostics and telehealth as well as industrial edge computing.

Pankajakshan sees Open RAN development intersecting with, and being an additional driver for, edge computing. In particular, he highlighted the role of the RAN Intelligent Controller, or RIC—which he compared to an app store for the RAN, where innovation can play out with more agility and speed than in traditional, single-vendor implementations.

The RIC, he noted, comes in two forms: The [near-real-time RIC](#) and the [non-real-time RIC](#). Both host applications: In the near-real-time RIC, those are xApps, while. The difference between then is the closeness of the control loops, Pankajakshan explains. “If you’re implementing an application where the control loop needs to be executed in less than a second, then you would run it on a near-real-time RIC. ... And when you have longer control cycles, then you would run it in the rApps,” he says.

That application-hosting ability makes the RIC in itself a “flavor” of edge compute, he says, bringing intelligence deeper into the network—and potentially hosting additional third-party apps that make use of the



Bejoy Pankajakshan

available Open RAN network APIs.

“Why this is critical for edge computing is because, by the fact that you’re moving some of the radio resource management and control functionality out into the centralized edge infrastructure or radio intelligent controller, you could now have third-party applications leverage the APIs that come out of these boxes to build applications that could be used for enterprise or consumer use cases,” Pankajakshan said. “You could now host applications that are much closer and have a per-user or per-device control mechanism in place.

He offered up use cases such as traffic steering or quality-of-experience optimization, in which edge computing could be applied to steer a particular user to get better throughput or latency. “If a user is in a

congested area, you could optimize the experience of the user even though the rest of the users in that locality may not get the best experience. ... You enhance the end experience that the enterprise or consumer could leverage.”

That means more opportunities for monetization as well as a variety of human or machine user experiences that could be enabled via edge intelligence—and an opportunity for operators in 5G that was missed in 4G, in which third-party players ultimately developed APIs that were more consumable and simpler than operator-exposed APIs, and drove adoption by developers. What does it take to actually implement such use cases? A lot of expertise in the RAN combined with machine learning, Pankajakshan says. “You would go through detailed

network modeling, you would build algorithms that can deliver better performance and then modify and retrain these models so that it gets better optimized over time,” he explains. “You build a library of network intelligent applications that can be leveraged by the edge applications, and that could be a third-party developer or the operator themselves utilizing this edge application to run their network better and deliver better experience for the edge applications.

“As you go higher up the value chain here in terms of intelligence, you use better algorithms, you use better machine learning and AI algorithms here to deliver a model which can deliver better results as opposed to a closed system,” Pankajakshan concluded. “So we see Open RAN with [the RIC] and the intelligence it brings into the network near the edge as one of the key building blocks for the operator to build intelligence at the edge of the network.”



How to create predictable services over unpredictable networks

By Juan Pedro Tomas

Service management can be “extraordinarily” cumbersome and expensive

Nancy Shemwell, chief operating officer at Trilogy Networks noted that some of the advanced services that will require low latency will be offered through a mix of networks, which can face a challenge from a technical perspective.

“If we look at the market today, there

is an outpour of advanced services, and they're going to start to require low latency. But if you look at the network, it's really a mishmash of networks. Because you're talking about fiber, you're talking about copper, you're talking about wireless. You have multiple generations, both public and private and satellite," Shemwell said, adding that management of all of these elements is "extraordinarily" cumbersome and expensive.

"The underlying services are just not reliable enough," according to Muriel Médard, NEC professor of Software Science and Engineering MIT and Chief Scientist at Steinwurf. "So what do you do? Do you just get multiple services and have them all delivered to you? Or do you actually try to make the service that you have better?"

Médard continued: "The right question to ask is: At a particular time, in a particular location, what are the resources that are available to a particular user for a particular application? ... Is this somebody who's just making a big massive download for something they're going to use later, or is somebody who's using something interactive?"

"If you look at another network that we use all the time, which is very reliable, which is part of our infrastructure, you look at the electricity, and so in the electricity domain, the network figures out basically what resources are available. And the electricity is, in effect,

something which is completely fluid. And what we can do with coding, which is a way of mixing together data to basically render it fluid and much more manageable, is that then we can manage resources, the way that we're used to managing resources in the grid, that provides us power," Médard said.

SUBHEAD: Who is motivated to solve for network latency and reliability?

When asked by who had the motivation and incentives to solve this issue, Médard responded that the stakeholder who decides to solve the challenges will emerge as the winner of the equation.

"But the question is who's motivated to do it right now, the hyperscalers are really coming from the application level. And of course, their motivation is to make the applications run. Well, if you have a very delay-sensitive application, and the delays are not being met, the customers are not going to use that application, the people who are selling the application are going to suffer from that. So you see, the hyperscalers are sort of coming from

the application and going down in order to make their applications run well, you then have the network operators. What happens to a large extent is the

Nancy Shemwell



network operators have been having to service the applications but not necessarily getting the full benefit of it, but bearing the brunt of a lot of the provisioning that they have to do," she said.

"And they have, to a large extent, relied on the equipment providers. Now, the equipment providers are of course, very motivated to provide high quality equipment, but they're not necessarily that incentivized to reduce the overall use of bandwidth by their clients. You know, they don't pay for that directly. That's really in the client's wheelhouse exactly how they manage that. So there's been a little bit of paralysis there, to some extent the operators are the most incentivized, but then not necessarily the most enabled."

"What we've been seeing sort of emerge in the nexus of those three large components of the market, which are all essential components, is SD-WAN. And if you look at SD-WAN, some of the software-defined networking is coming from the hyperscalers. Some of it is actually coming from the equipment providers, and some of them are standalone or very heavily connected in deep partnerships, or with some of the operators, I think, SD-WAN is sort of fulfilling that gap," Médard said. A software-defined Wide Area Network (SD-WAN) is a virtual WAN architecture that allows enterprises to leverage any combination of transport services — including MPLS, LTE and broadband internet services — to securely connect users to applications."



Muriel Médard



Uber on how AI and edge can improve customer service fulfillment

By Juan Pedro Tomas

Uber is “always trying to get computation close to data”

Madan Thangavelu, director of engineering at Uber, provided some insights on real world applications of AI and edge to improve customer service fulfillment. “When I think about edge, there’re just two principles that I’m always thinking about, which is in computer science and building large scale systems, we’re always trying to get computation close to data.”

He continued: “All the big data systems do that. And then the other is get your computation close to your users. And that’s where we do. When you’re talking about central cloud, we get the data centers out in so many different geographic regions, we get our CDNs, we get our points of presence. So in some sense, MEC is a new entrant into this but there are people who have been trying to achieve the various aspects of edge in production and through these other means. And it’s not an entirely

new concept when it comes to B2C systems,” Thangavelu said.

Thangavelu also provided insights on a particular use case in which Uber is implementing edge technology for the smooth provision of some of its services. “MEC, it’s not an entirely new concept when it comes to B2C systems. Today, we do things like the caching, and you probably heard about the way Netflix has built their open connect, where they add the ISPs, they cache their content in order to not have huge latency. So clearly, caching and those kinds of applications have come into a lot of the B2C and people have found innovative ways to push that around,” he said.

Mobile edge computing to stream GPS location between rider and driver


“One use case that’s interesting that I can call out is, especially in a company like Uber, you have the drivers and the riders on the street. We have this case where GPS locations are extremely important to get on time. And if you think about a round trip where we have to trust the GPS locations, and then extract—essentially, stream it to the rider to know where the driver is—seems like a very straightforward use case. Now, ideally, what you want is, by the time those two people are close, by within a few blocks, you really need to get a very high throughput exchange between the driver and rider.

And a lot of times that will fundamentally change the stress involved in getting a pickup happen. So at such very simple use case, we can’t push a lot of the tech down; we cannot do that at the device. Because there’s an exchange going on between the driver and rider, we do need to scrub clean up, fit into the map, and then essentially transfer. That’s where we would want to put some of that [edge] capabilities,” he said.

Thangavelu also talked about the emerging role of artificial intelligence (AI) and machine learning (ML) when it comes to managing data at the edge.

“Companies largely put their ML model training with all the data in the central cloud, they build their models and a lot of times the devices actually send the data to the backend. So that’s one model. And the other one is, clearly you have to put your ML model down into your app. And when you think about companies like Snap, when they do VR, you want to do object detection, so you can put a lot of that model down to the app,” he said.

“I feel like a lot of the really critical AI that is meaningful to a business is mostly on the devices. And those will be the ones that will have the first move out, because now you get the leverage of moving fast, you don’t have to wait for a device firmware update to get your AI out to the devices,” the Uber executive added.



Madan Thangavelu



Edge business models in practice—operator-owned and delivered as a service vs. enterprise-owned and operator managed

By Sean Kinney

The value of edge—whatever definition you use—is clear; but who should pay for it?

As 5G networks, both public and private, are scaled out by operators and integrated into enterprise technology workflows, the ambiguously titled “edge” has become a focal point for debate. The premise is clear: to fully leverage the throughput and latency of the 5G air interface to move data, apply some sort of intelligence and turn that into a decision, compute power has to be distributed out of centralized data centers to overcome the limitations of the speed of light. What’s less clear is who should own and operate what. And with many things 5G-related, the answer is, it depends.

Daniel Beazer, a senior analyst with Analysys Mason framed the issue: “The problem can actually be quite simply stated, which is that actually delivering services, be it compute, storage, streaming, CDN, what have you, from a...massive

centralized cloud node or massive data center campus from a interconnect point in Northern Virginia or London or Frankfurt... is something that makes very good sense economically, but that gets a bit shakier as we go out to something that’s really distributed and a long way away from the center. So I guess the question is, who’s going to pay for it and what’s that business model?”

‘It’s a smorgasbord right now in the early days of edge computing’

Equinix specializes in digital infrastructure, particularly data centers, colocation points and other internet connection services. Service providers use Equinix’s platform and infrastructure for a range of services like VPNs, SD-WAN, security applications and 5G network deployment. Orange for example, has used Equinix’s data centers for almost 20 years to host its backbone network, and also works with the firm on expanding its telco cloud environment.

As it relates to the metro edge, Zachary Smith, Equinix’s global head of edge

infrastructure services, explained the focus is “to provide an ecosystem of partners and capabilities at those locations, and make the effort and the friction less for getting the right thing in those places connected to the right network.”

He continued: “Where we think that we can get enough aggregation to make economics for what we’re going to consider to be more general purpose distributed use cases, really more economical.” With partners like Orange, Smith said, the goal is “bringing their networks and allowing local breakout and interchange of traffic to occur in those metro points instead of backhauling all to a few centralized locations.” Looking at highly latency-sensitive services in the sub-15 millisecond range, “I still think we’re a little ways off from...being able to help customers reach all of those.”

Looking at the business models currently in play in the market, Intel Senior Director of the Smart Edge Product Milan Djukic said, “There’s some situations where the enterprises are either, because of a strategic decision that they’re taking or their internal competency and skillset, where they’re more leaning towards working with

a partner to purchase some business outcome as a service.”

In other scenarios where the end customer “is a little bit more vertically integrated... they’re more open or willing to purchase the infrastructure



Daniel Beazer



Idir Fodil

themselves and deploy and manage their own various use cases.”

Calling out variances related to particular vertical segments and geographies, “It’s a smorgasbord right now in the early days of edge computing,” Djukic said. “To be successful in this type of environment, we really need to be flexible in how we package and how we deliver our solutions with our go-to-market partners.”

Back to Orange, the operator’s International Networks Transformation Director Idir Fodil said customers require both centralized and decentralized computing resources, and that having access to both architectures of telco cloud, “We can offer our customers with two...capabilities.” Edge capabilities allows Orange to adhere to business-critical SLAs regarding latency. “Being with latency less than 10 milliseconds, you cannot achieve this with centralized cloud.” This dual telco cloud operating model lets orange provide “customers with new services, onboarding them very quickly.”

Fodil said that, in his opinion, having a large footprint of smaller, distributed data centers gives “quick return on investment” compared with large, centralized cloud resources. He also referenced Orange’s collaboration with Equinix: “We are relying on the expansion of this footprint with Equinix

on this end-demand PoP in order to be in capacity to quickly roll out and deploy new services” including, 5G roaming, SASE, 5G cores, voice, and other services.

Developing an as-a-service mobile edge computing offering

In terms of how end users are communicating how they want to consume edge computing and related services, Smith said a consumption-based business model tied to specific business outcomes is pushing companies who build and operate the platforms to “figure out how we can bring that value equation to customers quicker.” He also called out the importance of extending equipment lifecycle, understanding carbon footprint and recycling components where possible in business model development.

“How can we bring more as-a-service consumption models to our ecosystem...which

does move the risk around,” is the first big question, he said. “And then number two, also making sure that we’re looking at all aspects of that infrastructure reuse, extend the life, find different use cases...[to] maximize on the ROI.

Djukic, discussing the challenge particularly around on-premises edge, is that “one size does not fit all...What works in industrial versus retail versus healthcare, these are extremely unique requirements. And integrating into these brownfield environments can be extremely complex and costly.” But, he said, Intel is racking up edge use cases with a demonstrable positive ROI, which allows them to easily quantify the upside on investment.

“I think there’s enough openness in the industry, whether it’s with the partners here on the panel or with Intel to come at our packaging differently to make the





Zachary Smith

numbers work so that value can be extracted by the end customer.”

Following Djukic’s commentary around on-prem deployments, Beazer

turned the discussion to multi-tenancy, shared edge platforms. When operator’s take this approach, he reckoned “use case stacking” is part of unlocking an attractive business case.

Continuing to use its partnership with Equinix, Fodil of Orange cataloged edge services that can be offered to multiple customers. “We have the SASE, which is very important for our B2B customers. We have the 5G—it may be private networks. We are seeing demand for the 5G and 4G local breakout...We have deployment of voice service” and the ability to legacy voice networks into an as-a-service layer that can be tailored to specific customers.

Back to the core question of who should own edge computing node, Fodil said partnering with Equinix gives Orange the ability to rapidly and flexibly deliver services to their customers. “We kind of focus our business on providing our customers with services...We partnered with Equinix to rollout all, let’s say, the down layer, which is the hardware, compute, storage, etc... We have business model and TCO that... at the same time benefit for both Equinix and Orange.”

The key question for mobile operators looking to monetize investments into

mobile edge computing is what role do they play in the value equation? The answer seems to be in serving as an aggregation point that’s also a service-delivery node.

To the aggregation point, Equinix’s Smith said, “We have to break out that network. And so there’s a great place for operators to come in and provide a virtualized substrate of connectivity, other services that are core fundamentals, and kind of unify that in multiple markets. I think it’s going to mean that operators have to get comfortable with working on a global basis. They can’t just be in one market. It’s going to be very, very hard to offer the kind of breadth that we see customers looking for at an enterprise basis.”

What’s the operator view on that? Fodil sees the opportunity at two levels. Using distributed telco cloud infrastructure to aggregate customer connections and services is a “great opportunity,” he said. Realizing the opportunity means operators “need to tackle the last mile...not the last mile in terms of telecommunication, but the last mile in terms of services.”

The second opportunity, he said, is providing enterprises with service level agreements against services. “The main important

point at this telco cloud is that it allows aggregating the customers and also aggregating services or [a] portfolio of services to our customers.”

The layer cake and the land grab

Back to the evolving mobile edge computing business model, Smith said the work Equinix is doing with Orange, deploying a virtualized telco infrastructure PoP, isn’t a quick “turn-it-on” type of project. “There was so much work that we had to add to our platform and...to meet the needs of a high-production network use case that Orange had. And Orange had to invest very heavily in its software stack in order to be able to consume that kind of heterogeneity and infrastructure and still deliver its service in operations.”

In addition to partnerships with operators, he’s also seeing movement around the serverless edge and front-end or low-code multi-tenant edges focused on enabling developers to build edge services. “I think we’re going to see this kind of layer cake and luckily, at least from my perspective over the last one year, it’s gone from be-

ing, “Can we think about moving into multiple locations?” to thinking and writing applications in a...kind of deployed everywhere model.

Intel’s Djukic characterized the MEC



Milan Djukic

business model as still in a “land grab” phase because things are still “in the early innings. But we believe there’s plenty of value to go around. And I think the mobile network operators are certainly going to be in a strong position to capitalize given their private mobility expertise and the innovative use cases that can be unlocked there.”

Given the nature of how the edge is being and built, deployed, operated and consumed by users, Djukic called out that there’s a role for systems integrators, cloud specialists and independent software vendors to play alongside operators. He said he expects there to be “multiple winners out there” and Intel trying to work across that ecosystem with an emphasis on developing use cases vertical by vertical.

Looking ahead, “What we’re probably going to see over the next few years out is then, once we have some pockets of edge computing deployments happening, the next phase will be how do these edges all play very nicely together?”

Attracting app developers to the developer-driven edge

Note the emphasis throughout on edge as a mechanism to deliver services and applications; beyond the telco realm, specific enterprise applications means specialist application developers need to be brought into the fold. This is primary piece of how edge will be put in front of buyers and also why the architecture is necessary in the first place.

As Djukic put it: “Developers are building their applications to run anywhere, and that’s really the jet fuel that’s going to drive this industry. And we’re keenly focused on developers, and we want to make sure that we’re putting the right edge platforms forward, the right hardware platforms forward, to enable those applications to run as efficiently as possible, whether it’s on perm, or on the network edge itself.”

To best attract developers, Fodil said the approach should focus on building APIs that can be quickly consumed by developers, including instantiation, testing and installation. Last word to Orange: “It’s not a matter of infrastructure, it’s a matter of services and applications.” ((«»))

Featured Companies



Intel

Intel’s mission is to shape the future of technology to help create a better future for the entire world. By pushing forward in fields like AI, analytics and cloud-to-edge technology, Intel’s work is at the heart of countless innovations.

Need guaranteed leads? Thought leadership? Incremental content marketing opportunities?

Sponsor an *RCR Wireless News* multi-platform, editorial program and receive the following benefits:

Editorial Webinar – sponsorship includes 250 guaranteed leads, participation as sponsored guest and recognition as sponsor in all promotional materials. Sponsors will receive webinar registration and attendee list, responses to pre and post surveys and polling responses gathered during webinar.

Editorial Feature Report – in addition to recognition as sponsor in program promotion, sponsorship includes 250 guaranteed leads, distinct from webinar leads, one-page ad spread or advertorial in feature report, and responses to lead capture survey questions.

For information contact sales@rcrwireless.com

Fast facts about the *RCR Wireless News* digital network

- 382,000 monthly page views
- 170,000 unique monthly visitors to websites
- 81,000+ opt in newsletter subscribers
- 220,526 monthly video minutes viewed on *RCR Wireless News* Youtube channel
- 45,000 leads generated annually
- Industry leading demand generation programs and results



UPCOMING 2023 EDITORIAL PROGRAMS INCLUDE:

JANUARY 2023

5G—is it making (enough) money yet?

The move to CI/CD requires the move to CT-continuous testing

What will the wireless workforce look like in 2030?

Public cloud vs. private cloud—key operator considerations

APRIL 2023

The evolution of T&M in the move to disaggregated, cloud-native networks

When will small cell densification happen at scale?