tmforum

HOW TO BUILD AND OPERATE AT THE EDGE

A TM Forum benchmark report

Authors:

Mark Newman, *Chief Analyst*
Tim McElligott, *Senior Analyst*

Editor:

Dawn Bushaus, *Managing Editor*

November 2020

Sponsors:

Netcracker An NEC Company

Red Hat

vmware

tmforum

Contents

The big picture

Edge computing is an evolution of public cloud that is happening regardless of how, where, when or whether communications service providers (CSPs) find a place in the value chain. Public cloud providers like Amazon Web Services (AWS), Google Cloud Platform and Microsoft Azure are driving the change by extending computing services to the edge of the telecoms network to manage sustained growth in traffic and improve customer experience. They are also making speculative investments in edge computing capabilities to meet the anticipated demand for IoT services.

The telecoms industry has an important role to play in edge computing because CSPs provide the connectivity that edge computing requires. The powerful combination of 5G connectivity and edge computing can unlock new latency-sensitive consumer and B2B services, such as gaming and telemedicine. Operators also own passive infrastructure such as telephone exchanges that could be suitable for housing edge computing equipment in some cases.

CSPs would like to bring together edge and network strategies. However, they are still in the early stages of deploying 5G, and the market for computing services is not mature enough for them to do so. This makes it difficult for them to determine where to collocate edge network functions.

Research methodology

Our research for this report includes interviews with more than 20 CSPs, public cloud providers, technology vendors and startups in the edge computing ecosystem plus an online survey of 55 CSPs from North America, Latin America, Europe, the Middle East and Africa, and the Asia-Pacific region, about 65% of them mobile operators. Nearly half of CSP respondents were from Europe and North America where edge activity is greatest.

About 40% of CSP respondents said they work in IT/operations support roles with job titles including CIO, SVP, VP, director and senior architect. Some responses also came from corporate strategy leaders and people working in CSPs' enterprise lines of business or network infrastructure groups.

Separately, we surveyed more than 40 technology companies, including hardware and software suppliers and systems integrators.

> “
>
> The survey results show that a few first-mover CSPs are deploying edge computing commercially, and more than a quarter are conducting trials.

But many operators are unsure whether their companies will pursue an edge computing strategy at all, and respondents in every group cite lack of solid use cases as the most significant barrier to moving forward.

CSPs' plans for edge computing

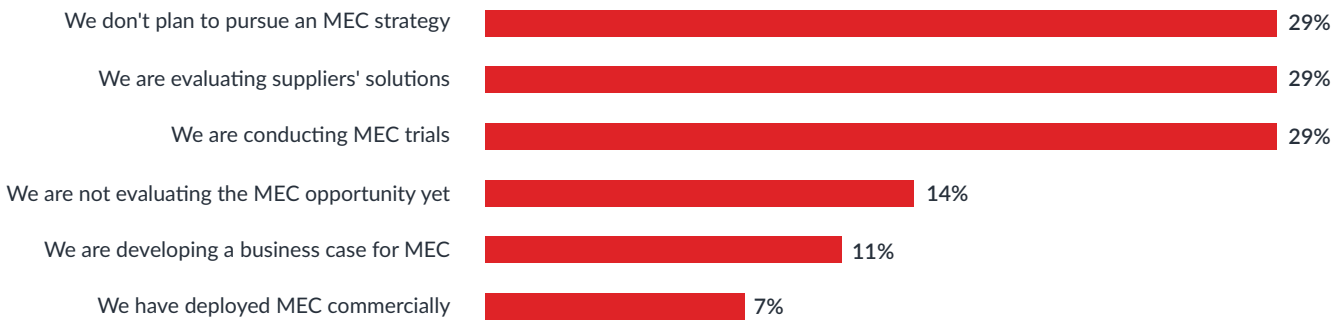

Common themes

Briefings with executives reveal several important findings that are consistent across companies and regions. For example:

Getting started is difficult – operators have put a lot of thought into their strategies around edge computing but are struggling to commit to and start executing a clear plan of action. Most activity comprises trials or proofs of concept, usually involving partnerships with public cloud providers and sometimes including third-party content providers and applications developers. Only a handful of Tier 1 CSPs have launched their own edge cloud services in select locations.

Public cloud providers are in the driver's seat – they are setting the timetable for CSPs' investments in and commitments to edge computing. Operators are turning to cloud providers as partners for their digital and network transformations, and cloud providers are urging them to deploy edge computing capabilities as part of these partnerships. CSPs are being forced to decide whether to proceed with their own edge computing plans and determine how they could affect partnerships. They are also trying to figure out the right balance between forging preferred relationships with cloud providers and adopting a multi-cloud strategy.

Deciding where to deploy edge computing capabilities is key – CSPs are conducting detailed assessments of where to deploy edge computing sites (either alone or with a partner) by mapping out the locations of their own B2B customers, factoring in where key verticals need connectivity.

5G timing is everything – CSPs want to combine deployment of edge computing and 5G core network capabilities. However, most operators are still in the very early stages of deploying the 5G core. Furthermore, many of the technologies and capabilities that could justify deploying a highly disaggregated core network, such as open radio access networks and network slicing, are still being developed and may not be ready for widespread deployment for two to three years.

CSPs are developing local and national strategies – in the B2B market, many requirements are to serve local campuses or factories, so operators likely will target these customers with a cloud provider partner. Nationwide capabilities are more likely to be part of a CSP's 5G core network strategy.

Private LTE and 5G networks will impact edge computing – in most cases organizations operating private mobile networks (for example, a port or manufacturing company) will take a hybrid on-premises and public cloud approach to managing IT workloads. This could remove CSPs from the revenue chain if industries build and operate their own networks and connect directly to the public cloud.

Latency is a major driver for the deployment of edge computing, though not the only one – the consensus is that 20-30 millisecond "in-air" latency will suffice for most applications. For small countries this means that edge computing may be required in only a handful of locations. Other drivers include delivering a better experience, capturing more of the enterprise market and reducing the cost of transport.

Using the local exchange is not a significant revenue opportunity – none of the CSPs interviewed for this report believe that renting space in their local or regional exchanges to hyperscale cloud providers will be lucrative. Furthermore, it is by no means certain that these facilities can be used to collocate edge computing gear because they lack space and power.

These conclusions from interviews are broadly consistent with the results of the online survey. However, it is worth noting that while our briefings were largely with Tier 1 CSPs that have already revealed their edge computing ambitions, the majority of operators deploying 5G are smaller players in their respective markets and do not have a strong enterprise business to drive edge computing.

Read this report to understand:

- Why CSPs are pursuing edge computing strategies – and why some may not
- How CSPs intend to deploy edge computing capabilities and what the best business models may be
- What the biggest challenges are in developing an edge strategy
- The role for partnerships with cloud providers in deploying edge capabilities and why many CSPs are concerned about building equitable business arrangements with them
- Why the access network is a source of leverage for CSPs and could put them in a position of strength with cloud providers
- How 5G and edge computing are closely linked
- What the impact of edge computing is on operations
- Why edge will not be a monolithic solution but will grow organically in hot spots of opportunity

Section 1

What is edge computing?

Asking what edge computing is may seem simple enough, but answers to the question can vary greatly depending on whom you ask. Indeed, companies spend a lot of time at events (remember them?) debating what and where the edge is.

Consider these varied definitions:

- Amazon Web Services defines edge computing services as providing "infrastructure and software that move data processing and analysis as close to the end-point as necessary. This includes deploying AWS managed hardware and software to locations outside AWS data centers, and even onto customer-owned devices themselves."
- Gartner explains that edge computing "represents an emerging topology-based computing model that enables and optimizes extreme decentralization, placing nodes as close as possible to the sources and sinks of data and content."
- Red Hat defines it as "computing that takes place at or near the physical location of either the user or the source of the data."

Ultimately, "the edge" can exist anywhere outside of centralized computing systems and closer to the point of consumption. For hyperscale cloud providers this means at end points beyond the reach of their global data center networks, while for communications service providers (CSPs), edge computing can be most relevant at the network edge. This could be at the base station, the local or regional exchange, or any part of their network where they choose to deploy or collocate edge computing servers. The edge can also, however, refer to new computing capabilities that are embedded within devices such as smartphones or IoT devices.

From the end user's perspective, edge computing might be something invisible, in that they don't know that a public cloud provider or CSP is using edge computing capabilities to improve efficiency by processing data closer to the user rather than in a data center hundreds or thousands of miles away. It could also be something that enhances an existing service – for example, delivering lower latency or a new way of leveraging data because it can be processed locally. End users may also experience it as an entirely new service that would not be possible in a traditional network – for example, IoT services that require low latency.

Edge & 5G

Edge computing and 5G are important to each other because edge computing facilities are necessary to reduce latency. Very low latency (1-10 milliseconds) is a key requirement for one category of 5G use case.

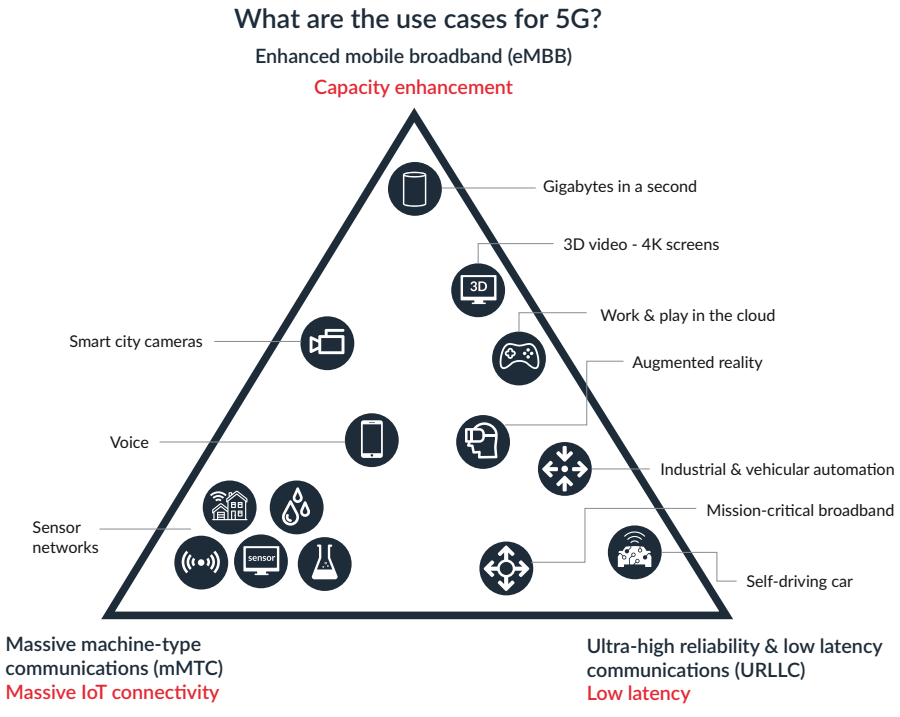

TM Forum, 2020 (based on ITU graphic)

URLLC services are the leading contender to deliver new revenue streams for telecoms operators. eMBB services are unlikely to provide an opportunity for premium pricing over LTE, and mMTC will only be required if and when more spectrum is needed for services that can be delivered today using variants of LTE or dedicated IoT network technologies such as LoRA or Sigfox.

Some observers contend that 5G needs edge computing more than edge computing needs 5G, because without edge computing 5G offers little advantage over existing network technologies. Edge computing, on the other hand, will gain ground with or without 5G.

It is still early days for edge computing, however, and it is not clear which use cases will be most successful. Both CSPs and hyperscale cloud providers talk about "vertical" use cases such as retail or manufacturing, but in practice there may be tens or even hundreds of use cases within each vertical, some of which may be unique to that vertical and others that could be relevant across multiple sectors.

Furthermore, cloud providers and CSPs have different assets and capabilities that they want to leverage for edge computing. As such, they have different ideas about which use cases to develop and support.

Emerging use cases

Most CSPs view edge computing and 5G as linked because of the low latency that 5G promises. They have gone through the process of determining latency required for specific services and which type of edge computing may be needed to achieve it. The graphic below, redrawn from one provided by Vodafone, shows the latency demands anticipated for several 5G services.

While a large number of use cases – mostly in IoT – require or benefit from the high speeds offered by 5G, relatively few require the ultra-low latency capabilities that can be delivered with a combination of 5G and edge computing. Many of these use cases also require processing of huge volumes of data, often in real time.

For example:

Cloud robotics – this is a specific field of robotics that requires vast amounts of processing power and storage resources to deliver real-time processing with zero latency so that, for example, a robotic arm is in the right place at precisely the right time.

Immersive augmented and virtual reality (AR/VR) – in an immersive AR/VR experience, the user loses awareness of the fact that they are in an artificial world. In the consumer market, immersive AR/VR experiences are likely to be enjoyed at home and so will not need 5G. Immersive AR/VR via 5G is more likely to be a B2B application, for example, for field workers. These types of 5G services are exciting for their potential, but it is unclear whether they can succeed, particularly in the case of VR for the consumer market where mass-market adoption is proving elusive.

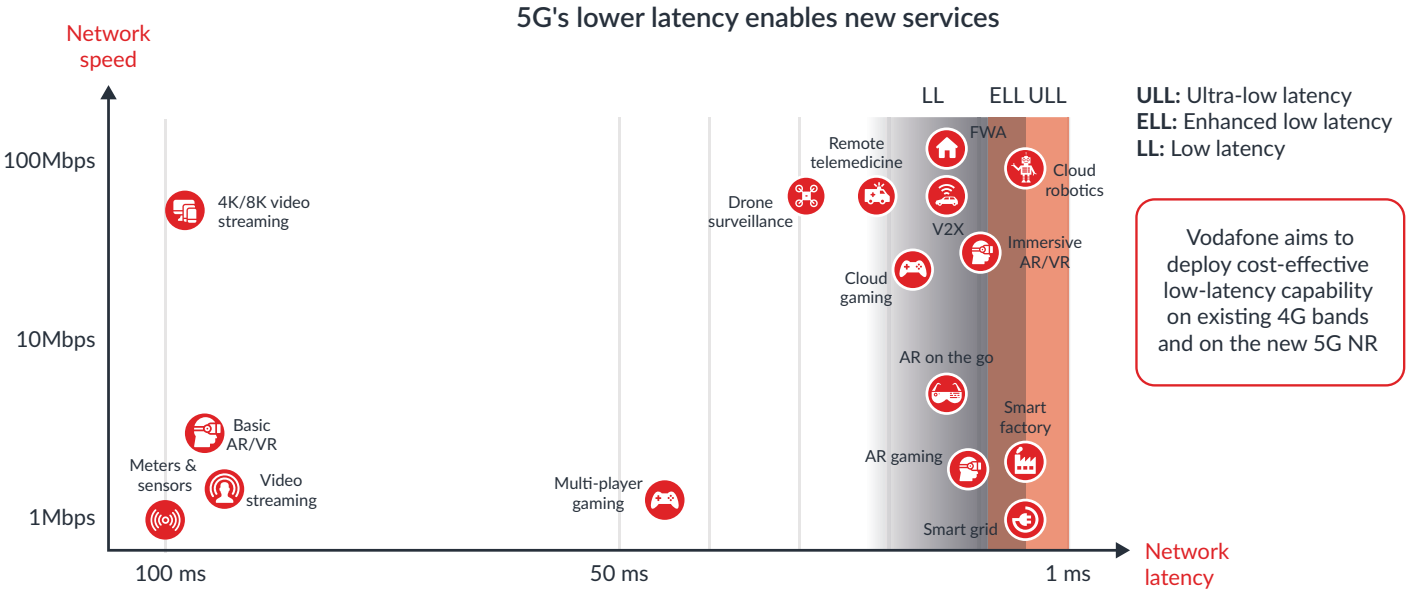

TM Forum, 2020 (based on Vodafone graphic)

Smart factory – a smart factory is one that uses computer-integrated manufacturing and relies on high levels of flexibility and automation enabled by wireless technologies. Real-time data processing either on or close to the premises is required.

Smart grid – a smart grid is an electricity network enabling a two-way flow of electricity and data using digital communications technology to detect and react to changes in usage. As smart grids become automated they will rely increasingly on real-time connectivity.

Edge computing is not only about 5G, however. US CSP Lumen (formerly known as CenturyLink) is pushing aggressively into edge computing even though it does not have a mobile business.

> "Edge computing is not limited to 5G use cases," says Felipe Castro, Director of Product Management at Lumen. "Edge computing can be done without 5G."

Castro notes that Lumen has a strong competitive position in edge computing because of its extensive fiber network, "widely distributed" real estate, and experience managing enterprises' cloud computing systems. The use cases he cites – manufacturing control systems, microgrid control, retail robotics, AR/VR, real-time video analytics, and tactile Internet – are similar to those touted by 5G operators, but Lumen's vision is supporting them using fiber and Wi-Fi connectivity instead.

All these use cases can be considered emerging applications which are unlikely to become mainstream in B2C or B2B markets in the short-to-medium term.

Cloud use cases

Many Tier 1 CSPs are targeting more generic edge computing services in the near term, and low latency is less of a driver in these cases. For example, Telefónica has deployed edge computing with VMware in the Spanish cities of Madrid and Barcelona and is offering infrastructure-as-a-service.

Juan Carlos Garcia, VP of Technology and Ecosystems at Telefónica, reckons that there is a good opportunity to address the needs of medium-sized companies that may have decided to stop investing in premises IT or that see an opportunity to separate data storage requirements, keeping some data on premises and some in an edge cloud.

In this scenario, edge computing represents just one of the computing capabilities alongside on-premises computing, private cloud and public cloud. As data volumes increase, companies will take increasingly sophisticated approaches towards managing their data, and edge computing will emerge as one element of a hybrid cloud approach.

CSPs and public cloud providers are looking at these opportunities for more generic computing services through a vertical lens. Google, for example, is prioritizing the retail market (given that Amazon is the world's largest retailer, other retailers are unlikely to want to partner with AWS). As part of Google's partnership with French operator Orange, the retail market will be a priority along with industrial and manufacturing verticals, which Orange already serves through its core legacy communications and connectivity business.

Consumer use cases

While the primary focus for edge computing services and 5G use cases is the B2B market, some of those companies are targeting consumers. Two use cases stand out: cloud gaming and AR/VR services.

In our briefings with CSPs, cloud gaming was mentioned regularly as a key application for edge computing. Online gaming requires low-latency connectivity, and cloud gaming could depend on the ability of operators to offer guaranteed minimum latency across their entire 5G footprints.

However, this does not mean that the 5G operators themselves will be the providers of these services. A more likely scenario is that existing publishers and providers of games, or new entrants, will launch services on the Apple and Google app stores and that CSPs will merely carry the traffic.

In the next section, we'll look at the need for standards and collaboration to advance edge computing.

Section 2

What's the role for standards and collaboration?

The terms multi-access edge computing (MEC) and mobile edge computing are sometimes used interchangeably, but MEC is actually an evolution of the earlier, more narrowly defined mobile edge computing. While some communications service providers (CSPs) believe that MEC requires 5G, others are not convinced. They agree, however, that standards and collaboration are necessary to advance the technology – and that coordination and consolidation of these efforts will be important.

The European Telecommunications Standards Institute (ETSI) has led the way on developing architectural standards for edge computing. In its initial 2015 white paper, ETSI described mobile edge computing as providing an IT service environment and cloud-computing capabilities at the edge of the mobile network, within the radio access network (RAN) and in close proximity to mobile subscribers. The aim was to reduce latency, ensure highly efficient network operation and service delivery, and improve customer experience.

MEC today is broader and more ambitious. ETSI has revised its description to explain that MEC is an evolution of cloud computing that moves application hosting from centralized data centers out to the network edge, closer not only to consumers but just as importantly to the data generated by applications. Now, companies view MEC as a key way to meet the demanding performance of 5G, especially with regards to low latency and bandwidth efficiency.

Does MEC require 5G?

Some CSPs with fixed-line assets such as US-based Lumen (formerly known as CenturyLink) and Singapore's Singtel, two differently sized operators in very different markets, say they are not convinced that MEC requires 5G because they believe their fiber optic networks can meet requirements for latency less than 10 milliseconds. The exception in Singtel's case is mobile gaming, which the company will deliver using 5G.

Lumen believes the reach of its fiber network will support most demand for MEC from enterprises. The company's recent rebranding emphasizes its fiber capacity – "lumen" is the measure of the brightness of light. Lumen has acquired several fiber networks including Level 3 Communications in 2017 and now operates a 450,000-route mile network in the US.

On the other hand, HKT, a company with a similar geographic and demographic profiles to SingTel (they both operate fiber networks serving densely populated areas) has stated that MEC is part of its 5G strategy. Paul Berriman, HKT's Group CTO, explains that a pure latency play is not enough of an enticement to HKT's customers and that 5G will be required to create services that have requirements other than or in addition to low latency.

Many operators believe they will need 5G to process and route large volumes of data, support real-time, high-bandwidth applications, and meet the demand for on-premises private networks. Enterprises see advantages in 5G over other technologies such as Wi-Fi, which may not provide sufficient reach, quality-of-service (QoS) guarantees or security. In addition, 5G is quicker and cheaper to deploy than fiber because it is wireless, and better able to be positioned where needed to be close to the customer.

It takes a village

To guarantee QoS in 5G networks, many operators and suppliers envision that MEC will leverage advances in network functions virtualization (NFV) and network slicing. Standards organizations and open source groups are collaborating to make this happen. For example, ETSI is working with 3GPP, NGMN, the Open Networking Foundation, the Open Network Automaton Platform and TM Forum to solve issues related to controlling traffic redirection to MEC applications, availability of network resources and end-to-end orchestration of slices.

Collaboration is also occurring elsewhere. CSPs and their suppliers are working together in various projects, initiatives, forums and task forces to pave the way for edge. Simultaneously, cloud providers and developer communities are working on MEC in open source projects that CSPs may or may not engage in. The table below highlights some of the most important edge computing initiatives.

To learn more about the work TM Forum members are doing to increase automation at the edge, see page 10.

Edge computing initiatives and projects

Initiative	Proponents	Hosted by	Primary goal
LF Edge	Network infrastructure and software suppliers, CSPs and enterprises	Linux Foundation	New umbrella organization that aims to establish open, interoperable frameworks for edge computing independent of hardware, silicon, cloud or operating system; currently operates 9 projects including Akraino Edge Stack, Beatyl and EdgeX Foundry which are described below
Akraino Edge Stack	AT&T Labs	Linux Foundation	Develop a set of open infrastructures and application blueprints for the edge, spanning a broad variety of use cases
Beatyl	Baidu, a Chinese technology and search engine company	Linux Foundation	Extends cloud computing, data and services to edge devices, enabling developers to build light, secure and scalable edge applications
EdgeX Foundry	Infrastructure and platform suppliers, embedded software companies and developers	Linux Foundation	Building a common open, vendor-neutral framework for IoT edge computing
5G Future Forum	América Móvil, KT, Rogers, Telstra, Verizon and Vodafone	América Móvil, KT, Rogers, Telstra, Verizon and Vodafone	To collaborate on global 5G specifications and interoperability to accelerate 5G and accelerate the delivery of 5G and mobile-edge computing-enabled solutions adoption
Cloud Native Computing Foundation	CSPs, major cloud providers and open source suppliers	Linux Foundation	Bring cloud-native technology including Kubernetes to the edge
Global MEC Task Force	Globe Telecom, HKT, PCCW Global, SK Telecom, Singtel and Taiwan Mobile	Bridge Alliance	Develop Asian MEC ecosystem
Open Edge Computing Initiative	Deutsche Telekom, Intel, NTT, Verizon, Vodafone and VMware	Carnegie Mellon University	Drive adoption of open edge computing with edge application providers, telecom operators and cloud service providers
Open Network Automation Platform	Originally led by AT&T and China Mobile, now includes many CSPs and major software suppliers	Linux Foundation	Develop policy-driven management, orchestration and automation of the edge
Open Networking Foundation	AT&T, China Unicom, Ciena, Comcast, Edge Core Networks, Google, Intel, NTT Group, Radisys, T-Mobile, Tech Mahindra and Turk Telekom	Deutsche Telekom, Facebook, Google, Microsoft, Verizon, and Yahoo! in 2011	Foster open source software development and provide standards liaisons and collaboration
StarlingX	Asian operators and cloud native technology providers	The OpenStack Foundation	Provide a container-based infrastructure for edge implementations
Telecom Infra Project	BT, Deutsche Telekom, Facebook, Intel, Telefónica, Vodafone	Facebook	Drive infrastructure solutions to advance global connectivity
Telco Edge Cloud	GSMA members, in particular China Unicom, Deutsche Telekom, EE, KDDI, Orange, Singtel, SK Telecom, Telefónica and TIM	GSMA	Federate multiple operators' edge computing infrastructures to give application providers access to the global edge
TM Forum's Open Digital Architecture and Autonomous Networks project	TM Forum members, particularly Orange and Vodafone	TM Forum	Build the Open Digital Architecture to support next-generation networks and increase automation at the edge

TM Forum, 2020

Collaboration is key in helping CSPs go to the edge

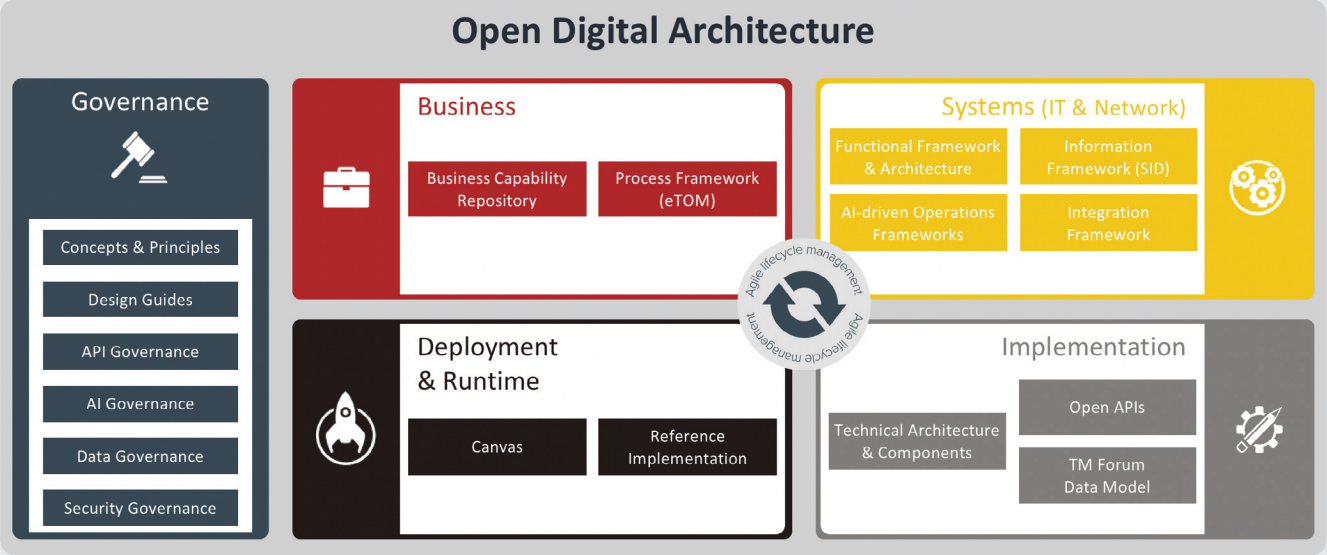

TM Forum members are working together in the Open Digital Architecture (ODA) and Autonomous Networks projects to automate management of complex new services and architectures at the edge. Importantly, the teams have begun collaborating with other standards organizations and open source groups to further define edge operations and support.

Part of the Open Digital Framework (see page 39), the ODA defines standardized, interoperable software components which are organized into loosely coupled domains. These components expose business services through Open APIs that are built on a common data model. The ODA team is developing machine-readable assets and software code, including a reference implementation and test environment to help CSPs and their suppliers increase agility and time-to-market. The Autonomous Networks project is using these assets to create an industry-standard framework for autonomous networks, one driven by use cases and based on autonomous domains and intent-driven interactions.

When it comes to edge computing specifically, TM Forum is mapping ETSI's MEC architecture to the TM Forum Information Framework (SID) and Open APIs. Forum members are creating a technical report that will include the results of these efforts. The goal is to incorporate the learnings into the ODA Edge Compute Implementation Framework.

Perhaps even more importantly, members are working on ways to help CSPs monetize 5G and edge computing services (see the discussion of Catalyst proofs of concept below). Common challenges CSPs and their partners face at the edge include:

- **Scaling edge applications** – inconsistent management of standalone solutions and lifecycle models makes it hard to scale applications. More automation in the management of network and IT components from the design and planning phase through decommissioning is needed, as is a more holistic approach to orchestration.
- **Lack of application development platforms for the edge** – software vendors are still developing apps in isolation, which leads to problems with time-to-market and porting edge applications to different platforms or edge cloud platforms.
- **Too much variability** – while open source communities make it easier to develop, deploy and scale edge apps, the projects often differ in scope. The communities need to come together to consolidate efforts and collaborate to drive critical mass in adoption.

Proving the concept

Members of the ODA and Autonomous Networks teams are participating in TM Forum Catalyst projects to address these issues.

The EDGE in automation project, which is championed by BT, Orange, TIM, Telus, Verizon, Vidéotron and Vodafone and includes AWS, Hansen and Microsoft as special advisors, showcases edge computing as a service (ECaaS) in a real-world scenario. The team demonstrates how ECaaS can support first responders in emergencies and provide the best performance for all applications during public and private events, rendering them smart venues. The use cases are very different, but both demand rapid, dynamic reconfiguring of network resources and instantiation of applications at the edge to deliver and monetize critical services in a timely way.

Watch this video to learn more:

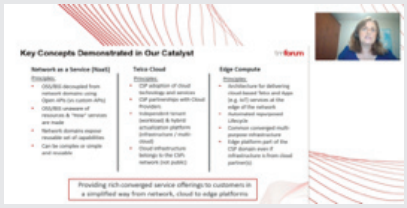

The *Developer is king!* project, which is championed by Globe Telecom and KDDI Research, seeks to help CSPs monetize 5G and edge by looking at two sets of consuming stakeholders: enterprises or other organizations in verticals like education, automotive, healthcare and manufacturing, plus application developers. Globe and KDDI also sponsored a precursor to this Catalyst, *Becoming edgy*, which provides a blueprint for an edge cloud lifecycle manager, demonstrating how CSPs can implement 5G use cases for enterprise vertical markets quickly and at scale.

The focus on developers in the latest project aims to level the edge services playing field for CSPs. Most developers focus on hyperscale platforms because it isn't feasible for them to use different tools, processes, platforms and technologies for every CSP and only reach a small audience. But this could change with use cases that require 5G and edge computing.

In the Catalyst demonstration, the team looked at a timely remote learning use case where a CSP provides the service enabled by 5G and edge to include live streaming and chats. In the demonstration, video quality is low for students but could be improved by a sponsor paying for reduced latency.

"This team is doing some groundbreaking work," says Dave Milham, Chief Architect, Service Provider Engagement at TM Forum. "They're really starting to focus on where the money is."

Watch the team discuss the project:

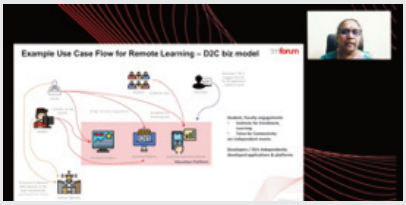

To learn more about these projects or to get involved in TM Forum's collaborative work at the edge, please contact TM Forum CTO George Glass.

Consolidation is needed

These initiatives represent only a small portion of the collaboration occurring across industries to make edge computing a reality. The large number of projects is good for accelerating adoption and not doing so myopically, but some observers believe it may be time to consolidate efforts on edge computing or take a new tact altogether.

> "At the moment there are too many groups working their own technology angles, and what they are writing hasn't really come to anything," says Volkmar Hammer, Head of Network Standards, Vodafone. "The industry needs to be more solutions-driven and cautious about who they think they are building this technology for. Too often they like to build stuff because it is a good idea – MEC is a good example."

He adds: "You can build it quite easily, but it may not meet the customer's needs. So, I think we have to go back to customers and find out [what their needs are] and think bigger about the endgame we're trying to pursue."

In the next section we'll look at the kinds of companies that are developing edge computing capabilities.

Netcracker Digital BSS/OSS Brings Automation, Agility and Monetization to the Edge

With the rollout of 5G networks, increased power at the edge and partner-enabled ecosystems bringing new services to market, CSPs need modernized BSS/OSS and orchestration that meet the new network, IT and business demands these services create.

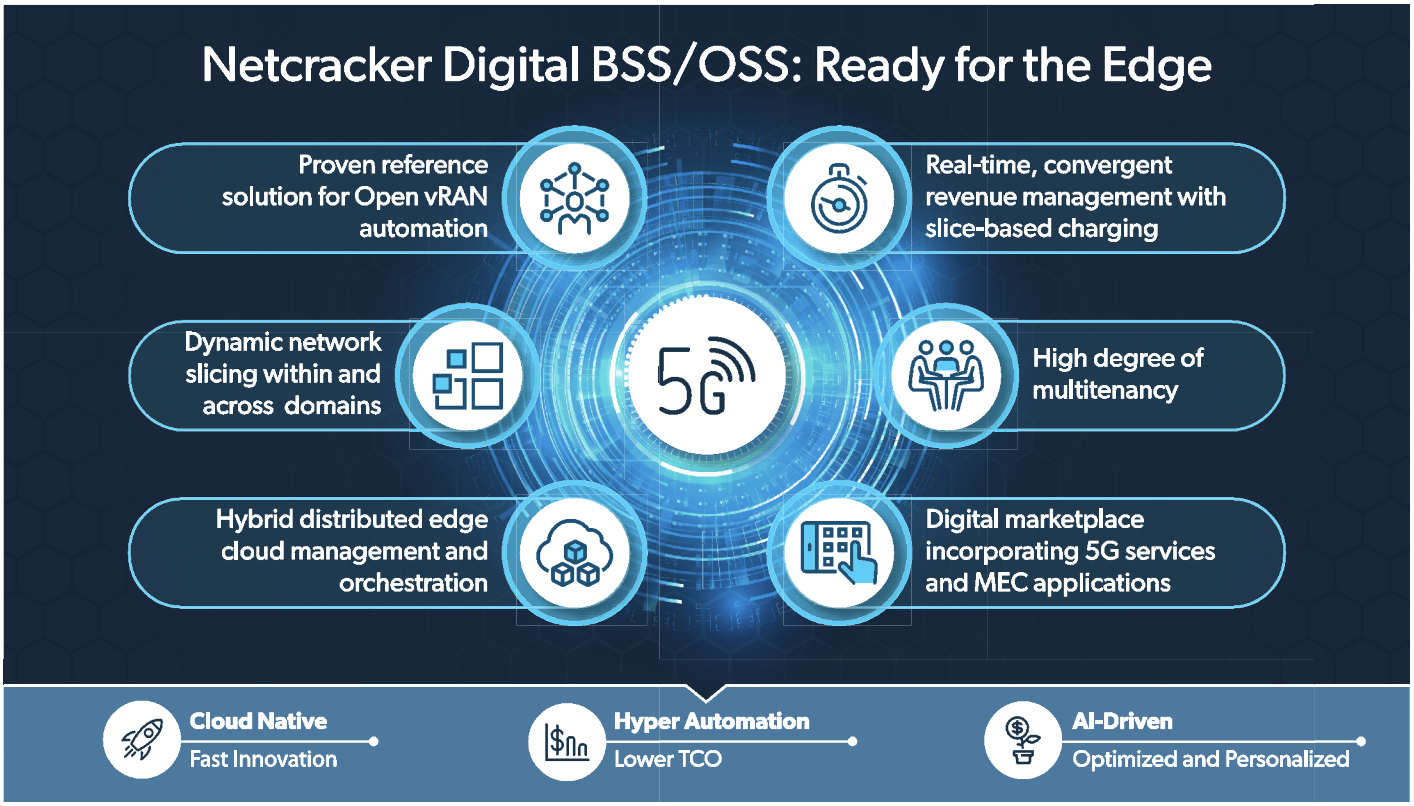

A unified approach at the edge

Edge cloud is a critical component of digital transformation as enterprises digitize and automate as many parts of their business as possible. In light of the COVID-19 crisis, businesses need to speed up their efforts to engage with their customers in a more digital way while accommodating remote workers and autonomous operations.

Edge cloud can take on many forms depending on the service and applications it is hosting, the edge cloud provider as well as the location of the edge cloud. While some applications with stringent requirements will be on premise or network edge-based, others will be better suited to regional or core data centers to maximize cost efficiency.

The success of edge cloud services is dependent on the ability to intelligently unify edge cloud orchestration and monetization across multiple edge cloud domains and multiple providers to meet the demanding requirements these services need for optimal performance.

Netcracker Digital OSS: A new operational model

Netcracker Digital OSS, based on the Netcracker 2020 Digital BSS/OSS portfolio, has been modernized and enhanced to automate and optimize edge cloud services, including Open vRAN, MEC applications and 5G network slicing, across multiple cloud platforms.

Netcracker Digital OSS uses a new operational model to remove manual tasks and drive hyper automation and agility across network edge and core domains, services that cross domains, business processes and channels that are quickly moving towards open 5G networks with cloud-native functions and highly demanding edge cloud applications. Our Digital OSS functions are cloud native, enabling CSPs to deploy only the functions they need for their specific edge strategy. Netcracker is working closely with a strong ecosystem of CSPs, cloud providers and application vendors to accelerate industry adoption.

Automating Open vRAN

Open vRAN provides CSPs with the ability to choose multiple best of breed vendors for different parts of the solution. With the base band unit further disaggregated and the majority of the solution virtualized compared with previous architectures, orchestration becomes essential to manage the dynamic nature of Open vRAN resources and services.

Netcracker Open vRAN Domain Orchestration, based on Netcracker Digital OSS, comprises a suite of orchestration and OSS systems, together with advanced analytics, to automate the end-to-end service lifecycle from planning and design to activation, configuration and assurance across the entire domain. Orchestration systems include Network Orchestration and VNF Management for vCU/vDU resources and Service Orchestration for the entire physical and virtual Open vRAN domain.

Evolution towards MEC Applications

MEC (Multi-access edge computing) creates new opportunities for CSPs to deploy high-performance services at the network edge with optimal quality and cost efficiency. Netcracker MEC Orchestration is compliant with the evolving ETSI ISG standard (MEAO) and supports placement and lifecycle management of MEC applications at the optimal edge location. It also includes an edge host for MEC application lifecycle management, catalog management of MEC applications and descriptors, MEC application instances inventory handling as well as integration with NFV Orchestration, Resource Inventory and Service Orchestration for end-to-end orchestration. MEC Orchestration supports VNFs and containers and can be deployed on any cloud platform including public cloud and telco cloud.

Slicing the network

CSPs can benefit from new business opportunities by allocating network slices to serve the specific needs of a service or vertical market. However, network slices must be provisioned and scaled on demand to meet the dynamic needs of enterprise verticals, which can only be achieved though end-to-end automation.

Netcracker Service Orchestration supports the 3GPP Customer Service Management Function and Network Service Management Function for dynamic orchestration of network slices that cross multiple domains. Within the edge domain, Netcracker Network Orchestration supports the Network Slice Subnet Management Function for dynamic instantiation and lifecycle management of slices within the edge domain.

Building the intelligent edge

Netcracker's suite of Digital OSS components is fully modular and cloud native, enabling CSPs to choose the functions they need to build their intelligent edge platforms. For example, CSPs can start with orchestration functions for Open vRAN and expand at a later date with MEC orchestration and network slice management functions. Alternatively, they can start with SD-WAN, add VNF/CNF applications and evolve with MEC applications and network slicing. Netcracker enables CSPs to build the intelligent edge that continuously evolves to meet their business goals.

Monetizing new edge services with Netcracker Digital BSS

The success of edge services will depend on the ability of CSPs to efficiently and rapidly monetize these new digital services. Netcracker Digital BSS has been modernized and enhanced to provide distributed rating and charging, including 5G 3GPP-compliant Converged Charging System (CCS), to many thousands of edge sites. Support of slice-based and cross-slice charging scenarios, cost-efficient support of multi-tenant charging requirements and the ability to scale on-demand to support growth in the number of subscriptions/devices are key characteristics that are supported by Netcracker Digital BSS.

In addition, these edge sites, collectively and individually, will support a wide range of applications and services, which can be dynamically changed at different times of day, for example to relieve congestion by relying on the dynamism of cloud-native network functions and infrastructure. Netcracker Digital BSS easily adapts to allow operators to get the most out of their resources for internal efficiencies as well as to give their customers the best quality of service.

Section 3

Which kinds of companies are developing edge capabilities?

While revenue projections for edge computing services and infrastructure vary, observers generally agree that they are potentially lucrative markets. This promise is attracting several types of players, some hoping to target businesses and/or consumers with edge-based services and others hoping to enable service providers. They include hyperscale public cloud providers, communications service providers (CSPs), other infrastructure providers such as content delivery networks (CDNs) and application developers, and companies that supply operational and business support systems (OSS/BSS) to telcos.

Many IT and telecoms research firms have published forecasts for edge computing services. While the projections look great in presentation slides, it would be extremely dangerous for a company to base their edge strategy on them, given that edge computing is still such a broad and undefined category.

Edge computing use cases can be divided into two categories: emerging applications such as IoT and applications that are complementary to public cloud services and indeed are viewed as the evolution of public cloud. Market forecasts are often similarly divided into two areas of focus: edge computing and cloud computing. The graphic below shows a wide span of revenue forecasts.

Cloud & edge forecasts vary

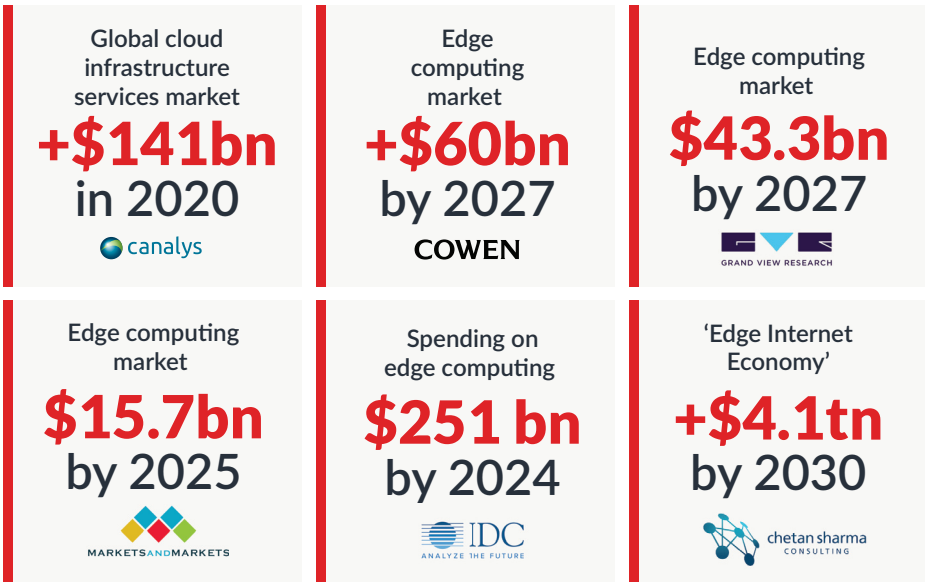

TM Forum, 2020

Predicting edge's future

Research firm Canalys suggests that the market for global cloud infrastructure services, which includes nascent edge computing services (i.e. public cloud), will be worth more than $141 billion this year. The company believes that the "big four" cloud providers – Alibaba Cloud, Amazon Web Services (AWS), Google Cloud Platform and Microsoft Azure – will account for roughly two thirds of this.

Analysts at US investment bank Cowen forecast that the edge computing market will be worth more than $20 billion by 2024 and more than $60 billion by 2027. Their forecasts separate revenue from edge computing from revenue from public cloud services, and they expect edge computing revenue to be roughly two-thirds as much as public cloud revenues annually in the medium to long term. Meanwhile, Grand View Research and MarketsandMarkets are more conservative in their forecasts for the edge computing market.

Some forecasts include a wider range of products and services. IDC, for example, forecasts that spending on edge computing will reach a staggering $251 billion by 2024. The firm breaks this down further suggesting that edge services will earn 46.2% of the revenue, hardware 32.2% and edge-related software 21.6%.

The most bullish projection comes from Chetan Sharma Consulting which predicts that the "Edge Internet Economy" will be worth $4.1 trillion by 2030. This forecast includes revenue from services that will be enabled by edge computing such as IoT, and Chetan Sharma divides revenue into existing use cases and emerging use cases. Interestingly, the company predicts that starting in 2019 revenue for both types will grow at a similar rate, but by 2030 revenue from new use cases will be almost twice as high as from existing ones.

Public cloud providers

Public cloud providers are investing heavily in their platforms, and they are expected to dedicate a significant proportion of their investment to edge computing. They see edge computing as a natural evolution of their public cloud business in that it will improve customer experience by putting computing closer to the point of consumption, and as an enabler of new services and capabilities including IoT and AI. While CSP revenue continues to stagnate, revenue from cloud infrastructure services is growing by 25% to 50% per year.

Canalys estimates, for example, that in the second quarter of 2020, spending on cloud services rose by 31% year on year to $34.6 billion (see below). This is equal to roughly 10% of telecoms operators' B2B revenues. However, cloud providers are catching up with telcos' B2B revenues globally.

We estimate that enterprise services generate about $100 billion per quarter for CSPs. However this revenue is flat, with gains in new services and broadband barely compensating for the decline in voice revenue. As such, we can size the telco B2B business at roughly three times as large as the cloud infrastructure business, but based on current growth trends, it will only be twice the size by 2022.

For more about hyperscale cloud providers' edge computing strategies, see Section 4.

CSPs

As noted, mobile operators are interested in edge computing because of 5G, but it is an opportunity for fixed line operators as well. Multi-access edge computing (MEC) is a network architecture concept that is relevant because it involves running applications closer to the connected end user and leverages existing assets such as base stations, data centers or telephone exchanges.

Worldwide cloud infrastructure services spending

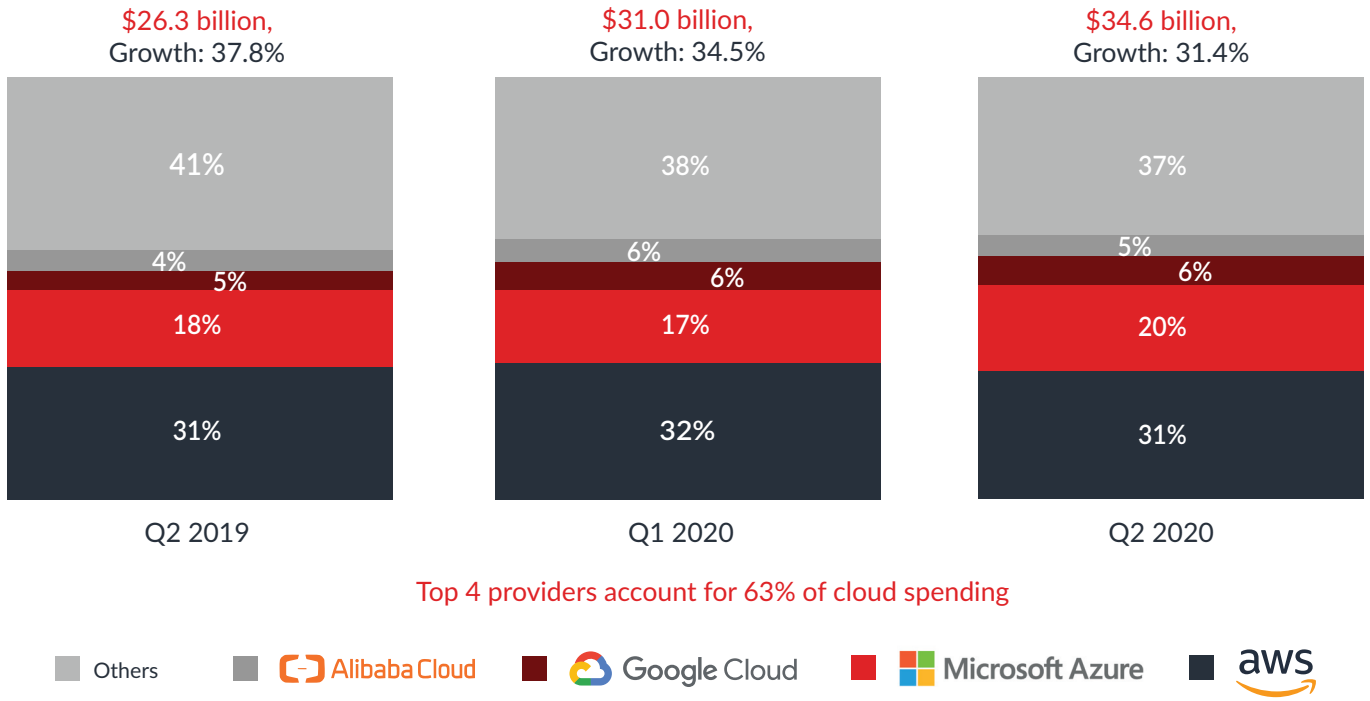

TM Forum, 2020 (source: Canalys estimates, July 2020)

CSPs leading the effort to roll out MEC typically meet at least one of the following criteria:

- They enjoy a strong or dominant position in mobile markets.
- They operate robust B2B lines of business.
- They own extensive physical assets such as local and regional telephone exchanges that can house edge computing gear.
- They have good existing relationships or partnerships with public cloud providers.

Most of the operators we spoke to meet all these criteria, or at the very least, three of the four. One exception is US telecoms operator Lumen which plans to deploy edge computing nodes in close to 100 locations in North America in order to provide latency of 5 milliseconds or less.

Most mobile operators are not dominant in their local markets, however. They largely focus on consumers, and in terms of physical assets they operate base stations and a handful of core network locations. Furthermore, they may have sold off their cellular towers and passive infrastructure to a third-party tower company.

It is much less clear how these operators can play a role in edge computing, at least in the short term. In the medium to long term they could deploy new 5G core networks in such a way as to provide consistently low latency across their geographical footprints. This may give them an opportunity to partner with public cloud providers, assuming the cloud providers don't already have partnerships in the region.

For more on CSPs' strategies, see Section 5.

Other players

Hyperscale cloud providers and telecoms operators are not the only companies that own the infrastructure needed to provide edge computing services. Other companies such as cellular tower operators, tech companies, CDNs, data center collocation providers and edge startups are vying for a piece of the pie.

Tower companies – cell tower owners are "among the fastest-growing companies in the real estate investment trust sector, as consumers use more data through mobile phones and tablets than ever before," according to *The Wall Street Journal*. Such companies are exploring whether to launch edge computing businesses. For example, Spanish tower company Cellnex sealed a €250 million ($295 million) deal in January 2020 to build 88 edge data centers across France for French mobile operator Bouygues Telecom by 2024, and Cellnex will build enough capacity to allow further tenants in the future.

The US technology giant also has edge ambition, as it has built 15,000 edge computing facilities globally which are used by all types of companies, from enterprises to telecoms operators, tower companies, and technology and network vendors. Customers include Vodafone, Nokia, China Telecom, Ericsson, Lenovo and RedHat. It also has a strong developer play. By building a strong ecosystem and encouraging its customers to work together, Intel may give CSPs an opportunity to carve out a strategy that is independent of complementary to partnering with public cloud providers.

CDNs – companies like Akamai are well positioned as edge providers because they already store video content in data centers that are physically close to viewers, thus speeding up delivery. In a blog published in August Akamai's VP of Product Marketing Ari Weil reflects on 20 years in edge computing, noting that "the Akamai Intelligent Edge is broadly recognized as the largest distributed network platform and has been delivering edge computing solutions to our customers virtually from day one." Akamai's Intelligent Edge Platform covers 300,000 servers in 4,000 locations and boasts more than 1,500 network partners. Akamai does not break out revenue from edge computing, but Weil suggests it is a $2 billion business. Akamai also has been diversifying recently into IoT and cybersecurity.

Data center operators – companies such as Equinix host private and public clouds and provide a service interconnecting the two. Executive Chairman Peter Van Camp compares Equinix's data centers to "international airports where passengers from many different airlines make connections to get to their final destinations." Equinix has 205 data centers in 25 countries and is planning expansion into edge computing and wants to play a role in IoT. In September Nokia selected Equinix as a strategic supplier for Nokia's Worldwide IoT Network Grid (WING) managed service. Nokia will leverage Equinix IBX data centers to deploy core and edge nodes to support mobile, IoT and cloud connectivity at global scale.

Startups – several start-up companies are building edge data centers or federating existing ones, aiming to serve edge cloud providers, CSPs and enterprises. For example, EdgeInfra provides data center and collocation services, while Mutable describes itself as "Airbnb for servers" and Stackpath focuses on building a secure edge platform. Initiatives to federate data centers and data center capabilities include MobilEdgeX and the ill-fated Edge Gravity which was set up by Ericsson with the lofty aim of building partnerships between CDNs, telecoms operators and data center operators to create a global edge computing capability. Ericsson revealed in June that it was ending the initiative.

In the next section, we'll dive deeper into the edge strategies of AWS, Google Cloud and Microsoft Azure.

“The edge is open: Why scale-out computing doesn’t exist without open hybrid cloud”

What is edge computing?

Edge computing brings computing, storage, and networking services closer to users and data sources. Several terms are used to describe the idea of moving computing resources to edge locations, and industry definitions are still being standardized. To add to the confusion, there is a continuum of edge tiers with different properties (e.g. distance to users, architecture, number of sites, size of sites, ownership, etc). The terminology for these different edge locations varies across and within industries. Terms like the Internet of Things (IoT) or 5G are often conflated with edge computing. IoT is an important use case for edge computing. 5G is, in part, a use case and it also is an enabling technology for other use cases.

At Red Hat, we understand computing to be distributed across the following tiers:

- **Provider/enterprise core:** Traditional “non-edge” tiers, owned and operated by public cloud providers, telecommunications service providers, or large enterprises.
- **Service provider edge:** Edge tiers located between the core or regional datacenters and the last mile access, commonly owned and operated by a telecommunications or internet service provider to serve multiple customers.
- **End-user premises edge:** Edge tiers on the end-user side of the last mile access, which can include the enterprise edge (e.g., a retail store, a factory, a train) or the consumer edge (e.g., a residential household, a car).

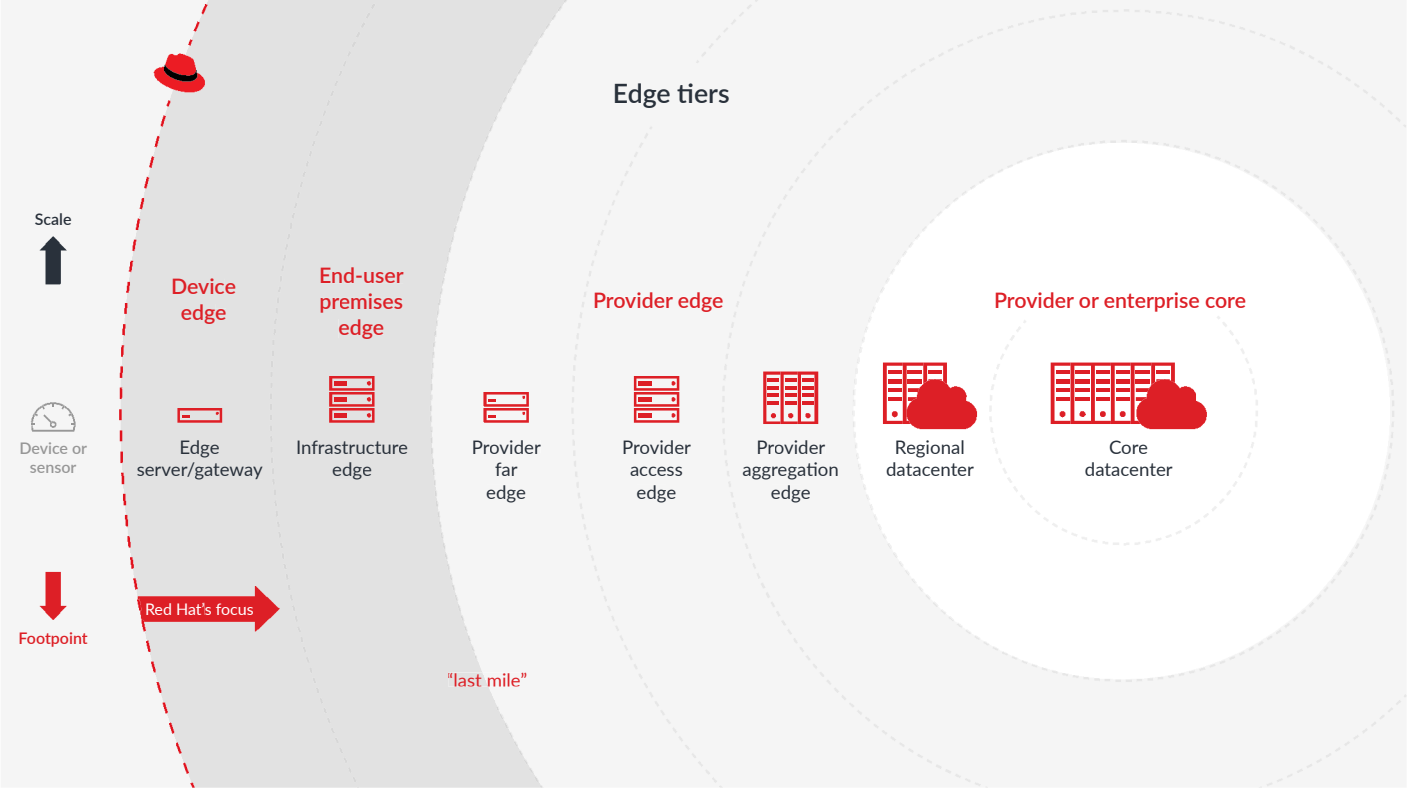

Red Hat's approach to edge computing

Red Hat helped drive the early standards around enterprise Linux®, fuel the adoption of Kernel-based Virtual Machine (KVM) as the open virtualization platform of choice, and build out Kubernetes into the juggernaut of container orchestration that it is. Red Hat, along with the open source communities and our ecosystem of partners, intends to help the business world capitalize on edge computing without fear of fragmentation or lock-in.

We are uniquely positioned to address the needs of organizations looking to extend their infrastructure and services to the edge through our platform, portfolio, openness, ecosystem, and record of continuous innovation.

Platform

Red Hat® OpenShift® and Red Hat Enterprise Linux are at the core of our open hybrid cloud and edge strategy. Kubernetes and Linux work together—the former orchestrates container workloads and controls policy, and the latter runs the workloads and enforces policies. Co-engineering our offers together is the secret behind our enterprise-grade security, performance, stability, hardware enablement, and commercial support critical for the enterprise. This combination is designed to run everywhere, from public clouds to private clouds, on bare metal or virtual infrastructure, at the core or at the edge. Customers benefit from a "write once, run everywhere" development experience and a homogeneous operations and management experience.

We have carefully redesigned Red Hat OpenShift and Red Hat Enterprise Linux CoreOS around industry best practices. This includes an immutable infrastructure and the Kubernetes Operator pattern, which is perfect for edge computing. These practices extend the highly robust, reproducible, and scalable container orchestration that Kubernetes is known for to the clusters themselves, the applications running on them, and (in combination with technologies that Red Hat has been pioneering) to virtual and physical infrastructures. As a result, a whole edge deployment can be consistently managed from top to bottom using Kubernetes primitives. And if you can do one deployment, you can do thousands of them.

We understand that edge computing use cases can come with harsh environmental conditions and legacy infrastructure—unlike clean, regular, carefully planned datacenters. Our platform is therefore designed to be modular ("batteries included but swappable"). It can adapt to running on a wide range of footprints --from customers' existing virtualization or cloud platforms; to modern, cloud-native infrastructures; to small form factor bare-metal clusters; to single-edge gateways and servers.

Open

Because edge computing solutions are made up of different technologies that spread across multiple hardware and software platforms, Red Hat's open source approach emphasizes interoperability and works against lock-in. Our customers can build their edge solutions with the most flexibility possible. We collaborate with a vast technology and community ecosystem to meet the needs of our customers' heterogeneous environments. The result is that our customers are allowed to quickly adapt their strategy as the market evolves without having to redo the work previously done.

Other large vendors are offering edge computing solutions that work only on their stack or platform. Customers using these solutions don't have freedom to mix and match needed components from third parties to build a better or more customized solution, leaving them trapped with their choices. Red Hat works on top of all relevant public cloud vendors and hardware vendors, making it easier for customers to change their minds, and their infrastructure, later.

Innovation can be crippled by fragmentation and, worse, proprietary models. When innovation is not driven by common, industry-wide open standards—and "edge edition" technologies are created instead—fragmentation results.

Red Hat and our industry partners are taking an active leadership role in creating these standards in upstream communities and industry working groups. Our strong track record of using the open source development model to deliver enterprise products allows our customers to make choices today that will remain flexible and valid for years to come. Red Hat is one of the most active participants in the core projects of the communities that we participate in. The experience, knowledge and participation of our teams allow us to deliver value to our customers. In parallel to our open source development model, we deliver products that we maintain with a life cycle that complies with the most demanding customer needs. Our product stability and reliability is the reason why we are preferred across business-critical use cases in stock exchanges, financial services, defense, and the public sector - in addition to those in telecommunications.

100% of Telecommunications Companies in the Fortune Global 500 rely on Red Hat! [1]

For further information, please visit: Red Hat's Approach to Edge Computing.

https://www.redhat.com/en/about/privacy-policy#process

[1] Citation: Red Hat client data and Fortune Global 500 list for 2020

Section 4

Edge strategies of hyperscale cloud providers

As noted in Section 3, revenue from cloud infrastructure services is growing by 25% to 50% per year. Cloud platform providers are well positioned to become important players in delivering edge capabilities to enterprises, and while some are potentially competitors to communications service providers (CSPs), more often they are viewed as partners. In this section, we examine strategies for three cloud platforms: Amazon Web Services (AWS), Google Cloud Platform and Microsoft Azure.

In the survey of CSPs conducted for this report, we asked which cloud provider is the best fit for operators' visions of edge computing. While more than a quarter of respondents expressed no real preference, Google Cloud scored slightly higher than Microsoft Azure and AWS.

Which cloud provider do CSPs prefer for edge?

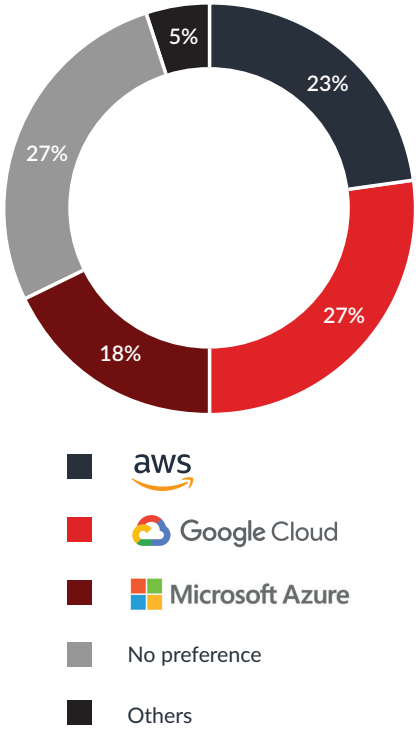

TM Forum, 2020

Comparing the Big Three cloud platforms Q2 2020

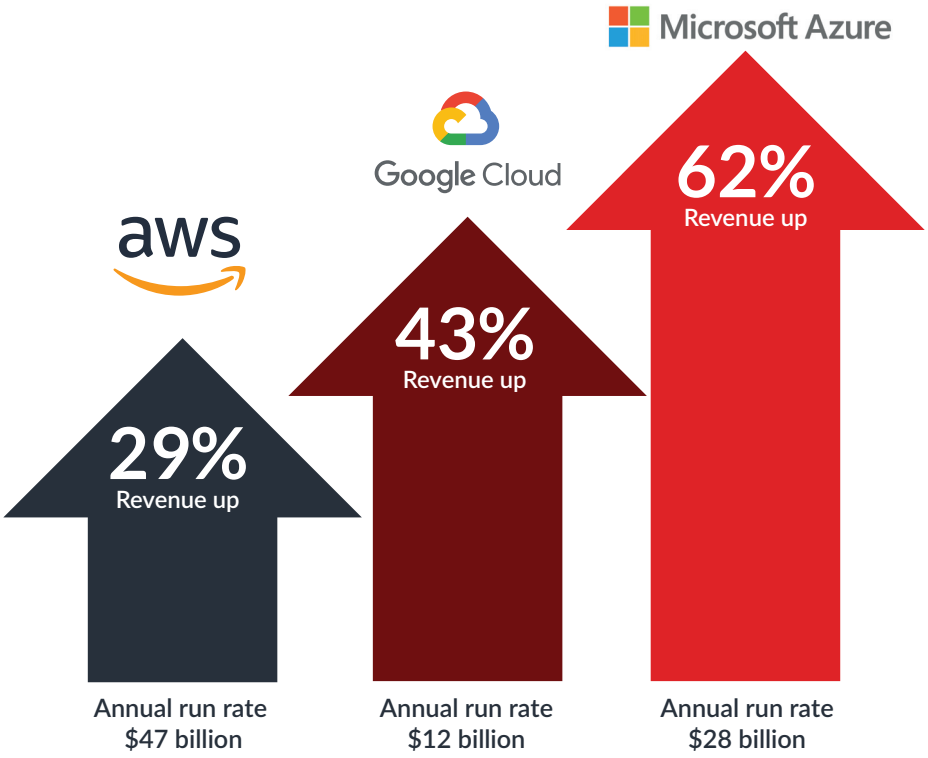

TM Forum, 2020 (source: Canalys)

Market shares of AWS, Google Cloud and Microsoft Azure are relatively stable based on data from research firm Canalys, but there are signs that both Microsoft and Google are gaining on AWS. In the second quarter of 2020, AWS' revenue grew by 29%, whereas in the previous three years it grew between 30% and 50%. The graphic above compares the companies' revenues, and the pages that follow offer additional analysis of their edge strategies, particularly as they relate to partnerships with telcos.

AWS

More telco workloads run on AWS than any other public cloud, and it was the first cloud provider to set up a dedicated business unit targeting CSPs in 2016, thanks to early customers including Comcast, Verizon and Vodafone.

The AWS cloud edge has 216 points of presence in 84 cities and 42 countries. AWS Outposts enables CSPs to run AWS infrastructure and services on their own premises, and AWS Wavelength supports ultra-low latency applications for 5G devices at the edge of CSPs' networks. Other tools include AWS Snowball, an edge computing, data migration and edge storage device that provides local storage and large scale-data transfer (for example, data that is continuously generated by sensors or machines).

At its AWS Telcom Symposium in July, company executives described four ways the platform can help CSPs:

Digital transformation and data center modernization – this principally refers to the ongoing migration of CSPs' IT workloads to the public cloud. AWS specifically refers to its partnership with T-Mobile in the US whose performance has improved in recent years.

Network evolution and business automation – AWS is enabling the shift to 5G packet core networks running in the cloud. Until very recently most CSPs were committed to building networks in their own private clouds. However, attitudes are starting to change as operators consider partnering with new network vendors, embrace concepts such as open radio access networks (RANs), and think more holistically about cloud computing across their technology and enterprise divisions. In August Telefónica Deutschland announced it is working with AWS and Ericsson to build a new 5G core network in the cloud and as part of an industrial proof of concept.

Applied AI – CSPs are using several AWS AI products and services, including the SageMaker platform, Amazon Connect, Fraud Detector and Personalize, in areas such as contact center transformation, reducing fraud and churn, network automation, and preventative maintenance.

Monetization of 5G, edge and IoT – AWS is positioning itself as provider of building blocks to help developers create applications across cloud, edge, IoT and 5G. US cable operator Comcast is expanding into home security using the AWS platform and adding capabilities such as voice and AI assistance to customer premises equipment. In addition, Chilean telco Entel worked with AWS to build an IoT platform called Reef to connect its network and OSS to onboard IoT devices.

In December 2019 AWS announced a major partnership with Verizon that involves bringing AWS products and services to Verizon's 5G network edge. The idea is to use AWS Wavelength Zones to embed compute and storage services within data centers at the edge of the 5G network.

At the time of the announcement Verizon claimed that the deal made it the first CSP in the world to deliver MEC services. Early trial customers included the National Football League and games publisher Bethesda Softworks.

Watch AWS CEO Andy Jassy discuss the Verizon deal:

Since then AWS has announced edge partnerships with Vodafone Business in Europe (the deployments initially will be in the UK and Germany), India's Bharti Airtel, KDDI in Japan and Korea's SK Telecom. AWS' long-term goal is to take the Wavelength service global.

The bottom line

AWS is the largest cloud platform, with nearly 40% market share, and it is viewed as the dominant platform. Some CSPs might see this as a strength because of AWS' reach and tools, but others may believe the company's dominance is a threat to their edge strategies and opt to partner instead with a player like Google Cloud, which is arriving later to the cloud platform party.

In addition to its early edge deals with CSPs, AWS is partnering with several suppliers of operational and business support systems (OSS/BSS) and network components including Amdocs, Nokia, Tata Communications and Mavenir. These partnerships can reassure CSPs that integration has already been validated and reduce the level of cross-domain communications when accessing systems or moving data between them. For more about OSS/BSS vendors' edge strategies, see Section 6.

Google Cloud

Google is the smallest of the big three public cloud service providers, but it is investing aggressively in the Google Cloud Platform. The company already employs 2,000 people in the cloud division and plans to take the workforce up to 5,000.

Google Cloud focuses in three areas: distributed infrastructure as a service, digital transformation platform and industry-specific digital transformation solutions (targets include retail, healthcare, financial services, media and entertainment, and manufacturing). The company unveiled a three-pronged strategy for partnering with telcos in March 2020:

Monetizing 5G as a business services platform – with its Global Mobile Edge Cloud (GMEC), Google aims to "deliver a portfolio and marketplace of 5G solutions built jointly with telecommunications companies; an open cloud platform for developing these network-centric applications; and a global distributed edge for optimally deploying these solutions."

Engaging customers with data-driven experiences – for example, Vodafone is migrating its data analytics to Google Cloud.

Improving operational efficiency across core telecoms systems – this is evident within the context of CSPs' ongoing efforts to transform their IT by moving applications to the cloud. Google is also partnering with Amdocs and Netcracker, both of which are in the process of making their existing products and services available in the public cloud and creating new cloud native services. In March, Amdcos and Google Cloud formed a strategic partnership to deliver Amdocs' OSS and BSS systems on Google Cloud and to create additional solutions and services. At the same time, Netcracker announced the deployment of its entire Digital BSS/OSS and Orchestration stack on Google Cloud.

In developing its cloud platform, Google aims to replicate its success with the Android platform but this time at the network edge. As part of its telecoms strategy announcement, the company unveiled Anthos for Telecom, which will take its existing Anthos cloud application platform to the network edge. Google wants to build an edge marketplace on Anthos which telcos could leverage by creating their own private label marketplaces for enterprise customers. Anthos for Telecom will also allow CSPs to run their applications wherever it makes the most sense.

Watch George Nazi, Global VP leading the Telecommunications, Media and Entertainment Industry solutions for Google Cloud, discuss the company's telecoms strategy:

Google has forged major deals with AT&T and Orange. The agreement signed with AT&T in March focuses specifically on edge and 5G. "We're testing a portfolio of 5G edge computing solutions for industries like retail, manufacturing, transportation that bring together AT&T's network, Google Cloud's leading technologies, including AI/ML and Kubernetes, and edge computing to help enterprises address real business challenges," Google said in its strategy announcement.

The strategic partnership with Orange, announced in July, focuses on transforming Orange's internal IT and on developing future cloud services, particularly in edge computing. Google is also seeking to partner with other companies that are building their own edge facilities including tower companies.

The bottom line

CSPs mostly see benefits to partnering with Google on an edge strategy:

Google has already amassed valuable experience building and operating edge computing facilities via its YouTube business which runs more than 2,000 edge locations deep inside telecoms networks as part of a global content delivery network. Google's edge computing architecture uses the same automation stack based on container orchestration as it has developed for YouTube, and through YouTube, Google has been entering into relationships with CSPs paying them relatively modest fees to allocate rack space in their data centers.

Google developed Kubernetes and is a leader in AI, machine learning and data analytics delivered across a global footprint, which is important as CSPs become more committed to cloud native approaches across their technology functions.

Importantly, CSPs do not view Google as a threat to their core connectivity business. Google is arriving late as a cloud platform provider and is positioning itself as a partner that can share the CapEx required to construct edge computing facilities. Such a partnership is particularly compelling for telcos that are not sure how to approach edge computing. In a recent *Light Reading* article, a Google spokesperson rebutted a claim that the company would one day be competing with CSPs saying, "We want to work side by side with CSPs to transform industries like retail, healthcare, and media and entertainment, ultimately delivering faster, more immersive experiences for the end consumer. We're committed to partnering with telecommunication companies to accelerate their digital transformation and unlock value from emerging technologies such as 5G, edge computing, and machine learning and artificial intelligence." For its part, Microsoft has given similar assurances.

Microsoft Azure

CSPs have traditionally viewed Microsoft as a provider of enterprise cloud services, but its public cloud platform and acquisitions of telecoms network vendors Affirmed Networks and Metaswitch make the company a major player capable of helping CSPs move network and IT workloads to the cloud.

Microsoft aims to help CSPs in the following ways:

- Modernize the telco workplace with products such as Microsoft Teams and Microsoft 365
- Streamline business support systems (solutions that address CSPs' requirements for moving applications to the public cloud)
- Deploy and optimize next-generation networks such as Azure Edge Zones
- Transform customer experience (for example, using AI and analytics)
- Accelerate growth and innovation through platforms that enable operators to unlock new revenue streams via partnerships with over-the-top content providers and by rolling out edge services in sectors such as healthcare, manufacturing, gaming and mixed reality

Microsoft uses the terms "intelligent cloud" and "intelligent edge" to describe the evolution of its cloud computing business. Intelligent cloud is ubiquitous computing, enabled by the public cloud and AI, while intelligent edge is an expanding set of connected systems and devices that gather and analyze data close to the end user.

Microsoft operates more than 160 edge sites in 50 countries which connect to Microsoft's global backbone network. In addition to these sites Microsoft has launched Azure Edge Zones with CSPs that connect Azure services directly to operators' 5G networks

The company's first significant edge computing deal with a telecoms operator was with AT&T in March 2020. As part of this deal the two companies plan to launch the first new Edge Zone in Los Angeles in spring 2021. Microsoft has struck similar agreements with Etisalat, NTT, Proximus, Rogers, SK Telecom, Telefónica, Telstra and Vodafone Business. In addition to launching Edge Zones, Microsoft is launching Azure Private Edge Zones for large enterprises to install on their own premises.

Learn more about how AT&T envisions partnering at the edge with Microsoft:

The bottom line

Microsoft's acquisitions of Affirmed Networks and Metaswitch in April and May give the company unique ability to target CSPs because both network suppliers already provide cloud-based network solutions to operators. In addition, the network solutions are highly complementary to Microsoft's Edge Zone strategy.

Microsoft describes Edge Zones as "an evolving platform built with customers, carriers, and industry partners to allow seamless integration and operation of a wide selection of Virtual Network Functions, including 5G software and SD-WAN and firewalls from technology partners such as Affirmed, Mavenir, Nuage Networks from Nokia, Metaswitch, Palo Alto Networks and VeloCloud by VMware."

Shawn Hakl, a Microsoft Partner and former senior executive in Verizon's enterprise line of business, leads Microsoft's edge computing strategy which includes Affirmed Networks and Metaswitch. He characterizes Microsoft's approach like this: "You have to have a telco-grade cloud; telco workloads are unique. I can say to my telco, 'Work with me and we will figure it out together,' or I bring in 1,000 engineers and I can make sure that their platforms are ready. I don't think you can be credible in telco space without engineers."

We'll discuss cloud providers more in the next two sections, where we evaluate CSPs' edge strategies and the potential for partnerships between CSPs and cloud providers.

Preparing for innovation at the Edge

Fully immersed in digital network and service transformations, service providers around the world are working with VMware to build, run, manage, connect and protect any app on any cloud on any device – from the core and access, to edge clouds.

The three tenants of a 5G network — mobile broadband, massive machine-to-machine communications and ultra-reliable, low-latency connectivity — all require performance at the edge. The edge alone, however, won't be enough. Delivering 5G service is going to require resources from throughout the network. Unifying edge resources with the end-points, access, and core clouds through a holistic management platform, will enable the agility, resiliency and performance service providers will need to thrive in 5G and beyond.

It Starts at the Edge

Driven by shifting market expectations, most enterprise application data will be created and processed outside the core clouds by 2023. The market behaviors which represent this application data have been changing since the growth of digital platforms. Not simply the natural bi-products of technical advances, the recent pandemic has also already contributed seismic shifts in demand, and 5G will have considerable impact in the near future. Soon, edge networks will need to anticipate application consumption patterns and respond accordingly in order to fulfill service expectations.

Luckily, advances in compute technology and network connectivity enable new, distributed architectures. Advanced capabilities at the edge compliment core clouds for applications with capacity and response considerations. This level of responsiveness is especially powerful for real-time and mission critical applications. The advanced edge, when unified with the unique capabilities within the public, private, and telco clouds, helps to define a horizontal network fabric. This fabric provides smarter end-points access to an intelligent network with ubiquitous control and consistent infrastructure operations management. With a horizontal platform in place, the intelligent edge can enable applications based on its needs. Delivering better quality of service with differentiated characteristics – latency, throughput, jitter - is possible with capabilities like network slicing and dynamic provisioning.

> "Edge, as a little bit of a mythical description, is how everything is going to get delivered from some cloud but it's going to magically become low latency and high bandwidth with the ability to get services rendered and delivered much faster."
>
> Shekar Ayyar, EVP and GM, Telco and Edge Cloud, VMware

Edges Aren't One-Size-Fits-All

Significant functionality will be distributed to the edge. But as the service provider and enterprise network distributes, a balancing act of capacity scaling and distributing workloads ensues. Physical constraints like space and power and operational constraints like cost will limit some edge clouds, so not all edges will have the same capabilities.

Core Cloud - With greater that 50 hosts, the core cloud will continue to manage IT, traditional applications, data archives and batch analytics processes. The core cloud will host the control and management planes.

Thick Edge - With fewer than 30 hosts, the Thick Edge will enable Mobile Edge Compute (MEC) applications, Infrastructure as a Service (IaaS), Containers as a Service (CaaS), cloud services and video delivery.

Medium Edge - With less than 10 hosts Medium Edge will manage mission critical, automated guided vehicles, augmented reality and virtual reality applications. It will include data planes and manage real-time inferencing.

Thin Edge - With 1 to 2 hosts, Thin Edges will enable real-time applications, Layer 1 processing, kiosks and IoT gateways.

Any Cloud

A robust edge network with the capability to engage any cloud, requires alignment with existing core transport capabilities. SD-WAN combined with central processing will provide the interconnectivity and traffic engineering capabilities service providers will need to support demanding edge use cases.

To enable the clouds themselves - public, private, core, access or edge – a converged compute environment with ubiquitous resource control and unified service automation will manage virtual network functions as well as cloud-native network functions. With horizontal computing capabilities in place and holistic orchestration to manage service lifecycle throughout an end-to-end network, third-party application providers will deliver 5G applications including AR/VR, gaming, security and beyond.

> "This is a world where you are not bound by the source of the capability, but freed for the limitless potential of your business."
>
> Pat Gelsinger, CEO, VMware

Any Application

Moving into world of on-demand, customizable services, there needs to be enough agility at the edge to allow service providers to quickly create and deploy new services. To enable microservice-based, cloud-native applications from the edge, modern API developer toolkits will leverage a horizontal automation layer. Application lifecycle management and workload placement capabilities are managed from a single interface to engage any cloud.

> "…move from development with static infrastructure and ticket-driven provisioning, into a much more progressive cloud-like, API-driven provisioning and consumption model."
>
> Craig McLuckie, VP R&D, Cloud Native Apps, VMware

Application developers will code to native cloud APIs, like VMware Tanzu, an Azure or AWS, to leverage best of breed capabilities. On the same single platform, developers will also be able to create composable services to be stitched together across the multi-cloud ecosystem.

Any Device

> "…work at the speed of life, so you can pick up any app on any device."
>
> Sanjay Poonen, COO, Customer Operations, VMware

With the growth of IoT, the trajectory of remote work, and the inevitable advent of 5G-enabled enterprise use cases on the horizon, unified endpoint management is a "must." Controlling device security and enabling remote triage aren't just conveniences for a pandemic-altered operation. With a portfolio of endpoints and the high-network-demands of some very prescriptive IoT devices, the edge network is going to be sensor-rich.

Intelligence will be paramount to assist in improving experience and operations. Telemetry, location and identity attributes will have to be captured from all end-points with secure, authorized access control.

With unified endpoint management in place, there will be a new crop of solutions to help realize the value of "any device". Content acceleration, is a good example, as it will require both composable infrastructure and private 5G connectivity.

VMware at the Edge

VMware creates a consistent network, data and application fabric for multi-cloud infrastructure operations. This intelligent edge solution centralizes workload and infrastructure automation which isolates network issues and speeds up service delivery. A horizontal platform, VMware solutions enable service providers to build solutions once, and deploy them end-to-end without the need of custom hardware. With operations management integrated.

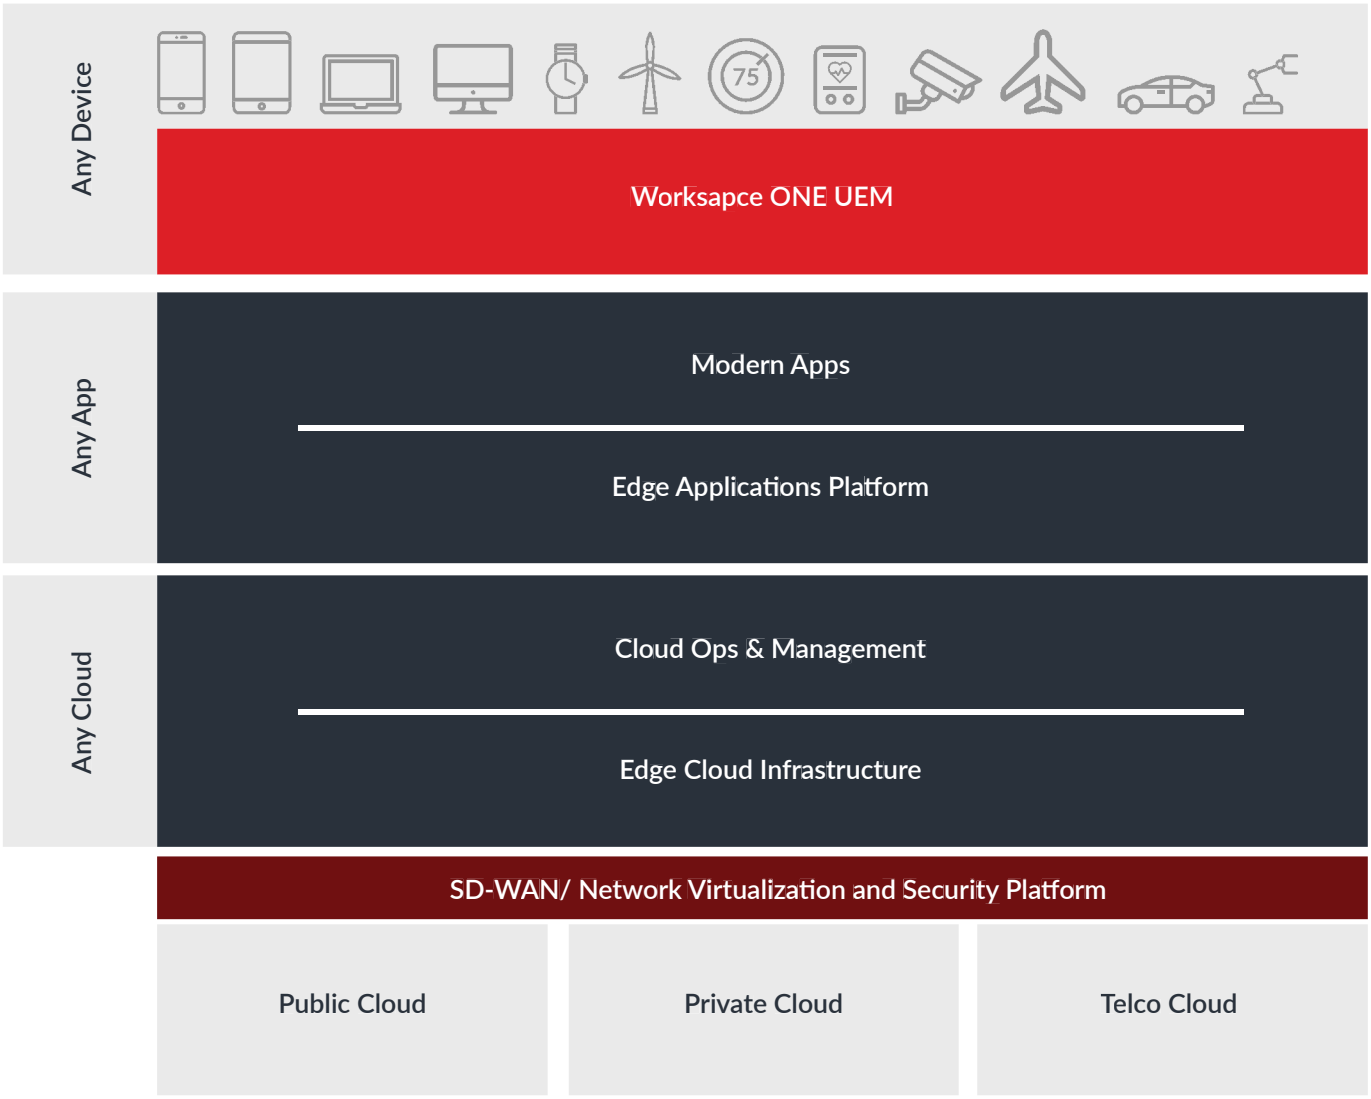

service providers can close the automation loop and leverage application flexibility and portability while balancing stringent response timing and edge footprints – thick, medium, thin.

VMware has intrinsic security including end-point policies to deliver data protection and secure service edges. When paired with embedded AI/ML engines to provide contextual intelligence for application and infrastructure awareness, the solution enables just-in-time decisioning.

Reducing carbon footprint at the edge and breaking down network silos, VMware solutions accelerate application innovation to create the agility needed to succeed at the edge.

Section 5

CSPs' edge computing strategies

The communications service providers (CSPs) surveyed and interviewed for this report agree that edge computing is full of promise. Indeed, more than a third of survey respondents said their companies have already deployed multi-access edge computing (MEC) in some locations. However, there is a risk that these results could create the extremely misleading impression that many operators around the world are busy signing up large numbers of edge customers. In practice, a small number of CSPs are making relatively modest investments in small-scale distributed edge deployments and in private network deployments in edge locations such as transport hubs (ports, for example) and manufacturing facilities. Furthermore, many of these deployments are trials or proofs of concept, paid for in part by CSPs themselves and their partners.

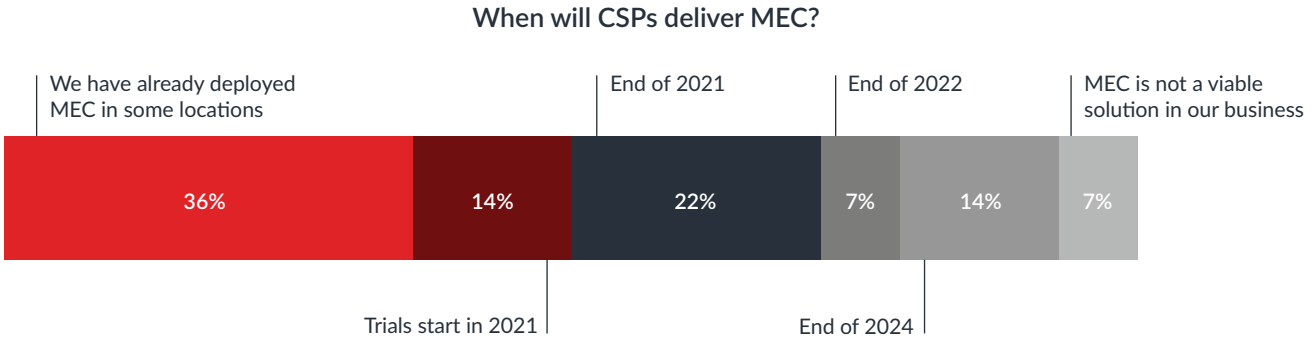

For example, Spanish operator group Telefónica has deployed its own edge facilities in Madrid and Barcelona offering infrastructure-as-as-service and network-as-a-service edge storage based on VMware. Telefónica will launch edge computing in two more cities this year and an additional six over the next two to three years. The company is targeting mid-sized companies that do not have established relationships with public cloud providers.

US CSP Verizon, meanwhile, has launched 5G edge for IoT in Atlanta, Boston, New York, San Francisco and Washington. The service is offered in partnership with Microsoft Azure as an on-premises, private multi-cloud solution.

What are the drivers?

As noted in the introduction to this report, our research includes interviews with more than 20 CSPs, public cloud providers, technology vendors and startups in the edge computing ecosystem, and an online survey of 55 CSPs from North America, Latin America, Europe, the Middle East and Africa, and the Asia-Pacific region, about 65% of them mobile operators. Nearly half of CSP respondents were from Europe and North America where edge activity is greatest.

We believe our research provides a clear picture of what CSPs think about edge computing and what their plans are. Any vagueness comes down to operators' own uncertainty about their role in delivering edge services. While our survey covers all types of CSPs, our interviews were largely with early-mover Tier 1 operators that have an edge strategy.

Among survey respondents, the biggest driver for multi-access edge computing (MEC) is improving customer experience, followed by meeting latency requirements of new applications and increasing revenue through new enterprise services. However, not all CSPs are convinced that they can make a business case for edge computing. Indeed, more than half of respondents said they are not pursuing an edge strategy or are undecided (see graphics on the next page).

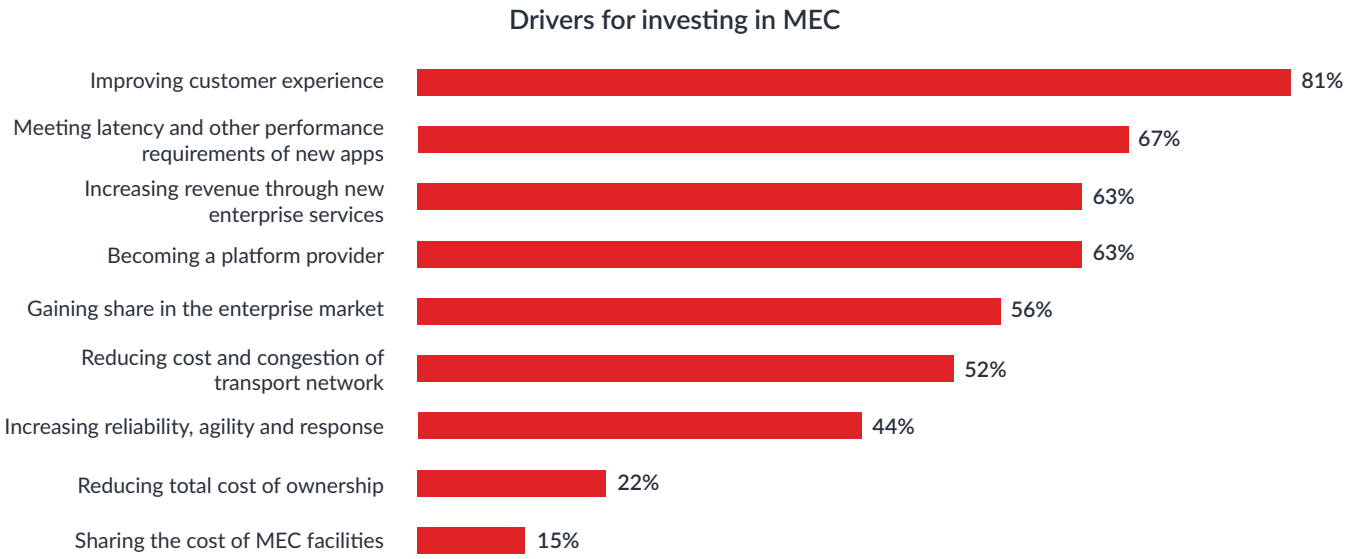

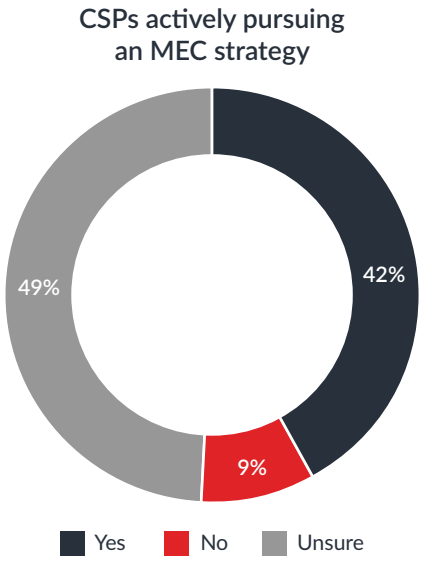

Many of the companies that are not sure about MEC or are not planning to deploy edge capabilities do not think there are enough viable use cases. Or, in some cases, they recognize viable uses but are waiting for the business case to be proved – that is, they want to know how much they will need to invest and what the potential return on investment is.

An unspoken truth, however, is that not all CSPs have a strong enough foundation in delivering B2B capabilities on which to build an edge strategy, at least not without the power of a partnership with a cloud provider (see page 32)."

Where to deploy?

The number of distributed edge computing sites that a CSP needs depends on the size of the country in which it operates and the latency requirements of specific applications or use cases. For example, with its plans to deploy edge computing facilities in up to 10 cities over the next two to three years, Telefónica can offer 3 millisecond latency within a city, 6 milliseconds in provinces and 11 milliseconds in bigger regions. The company believes that with guaranteed latency of up to 30 milliseconds it can meet requirements for up to 95% of applications.

There are two main locations for CSPs to deploy edge capabilities: in their existing private data centers and points of presence (local/regional exchanges), or on the customer's premises (for example, in private mobile networks). As operators expand into edge computing they will mainly use these existing locations and facilities, but some will also seek out new locations.

Existing facilities

As mobile operators build out 5G networks they will have the opportunity to collocate new edge computing facilities, but there is no blueprint for this. In the first phase of 5G most operators are leveraging existing passive infrastructure in their radio access networks (RANs) and the existing mobile packet core network. 5G radios are being collocated with LTE radios.

However, in the second phase of 5G deployment, operators will build new virtualized core networks that allow them to create new services and capabilities for their customers. This will involve building a network edge made up of small data centers that could also house third-party edge computing equipment.

UK mobile operator Three, for example, opened its core network comprising 20 data centers across the UK in July 2019. Rival operator BT currently houses its core network functions in eight major exchanges in the UK but has plans to extend these to nearly 100 metro exchanges. BT has an advantage because it can deploy its core network in the local and regional exchanges used for its legacy copper network, while Three, which has no legacy wireline facilities, must build its edge computing facilities from scratch.

French telecoms group Orange plans to roll out between 100 and 200 edge computing sites across Europe. Even though France is Orange's biggest market, the company operates across all of Europe through its B2B division Orange Business Services.

It is unlikely that incumbent fixed network operators would ever want – or need – to convert all their local exchanges into edge data centers. For example, Australian telco Telstra has 8,000 physical locations across the country but has identified between 300 and 500 sites that it would like to use for edge computing. Even so, the company recognizes that it is likely to gain access to most of the potential market for edge computing services with fewer than 100 edge computing sites.

In some cases, CSPs' edge strategies may be influenced by the evolution of the RAN, particularly open RANs that use cloud native architectures. An open RAN approach involves moving some functions to a generic (centralized) compute platform instead of a purpose-build hardware platform.

CSPs could house centralized cloud functions in the same central offices that house edge computing facilities. However, in practice it will be some time before open RANs become commercially viable for a large proportion of networks, so it can't be figured into any edge computing strategy for the next two to three years.

Even if operators do own central offices, there is no guarantee that they will be able to use them for edge computing. As a senior executive at a German CSP comments:

> "We're talking to hyperscalers about housing their equipment, but it's not that easy. Amazon equipment is just too bulky – it doesn't fit physically. They need a more compact figuration, and the servers need to use less power."

Vodafone's MPN-as-a-service offer

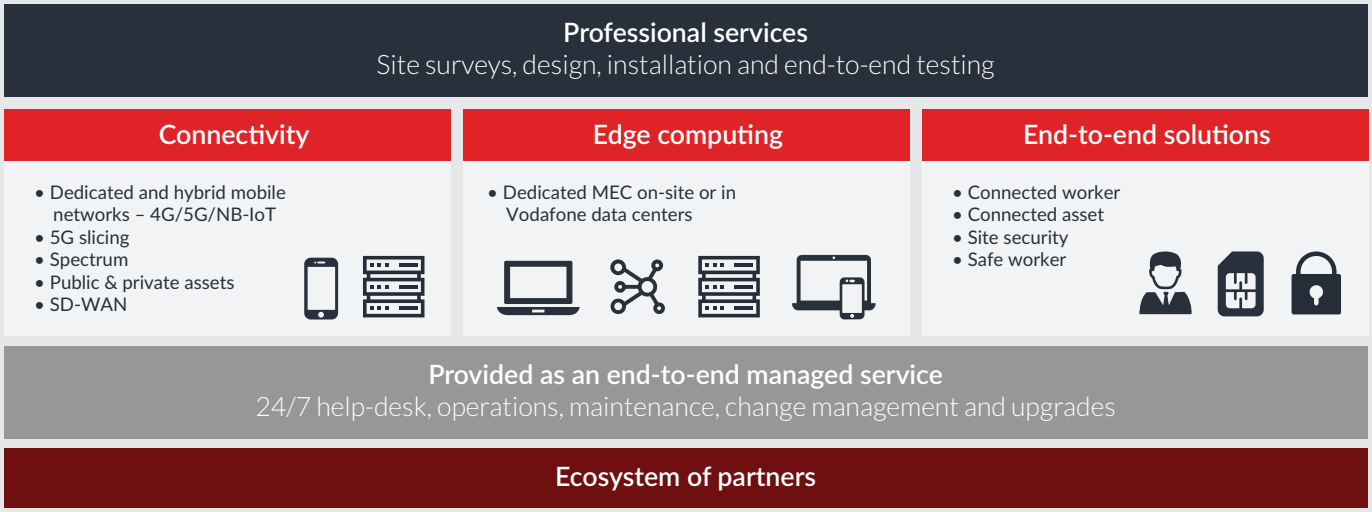

TM Forum, 2020 (based on Vodafone graphic)

Mobile private networks

Another option is to deploy edge computing facilities on the customer's premises in mobile private networks (MPNs). Vodafone, for example, offers MPN-as-a-service, which includes connectivity, edge computing and end-to-end solutions (see graphic). The company targets sectors including manufacturing/Industry 4.0; transportation and logistics (airports, ports and warehouses); and oil, gas and mining. These are all verticals likely to suffer from poor or non-existent mobile network coverage.

While Vodafone is confident about its MPN-as-a-service offer, this is not a clear edge win for CSPs because many enterprises that own their own MPNs may opt to use unlicensed or shared spectrum and work directly with network vendors and public cloud providers.

Business models

The business models CSPs are adopting for edge computing depend on many variables, ranging from the size of the operator to their plans for offering platform-based services and the specific use case they want to address. Ericsson has identified four edge deployment categories or scenarios to cover the full range of use cases:

- **Full Edge** – the CSP owns the relationship with the customer and provides its own edge computing services. The operator is responsible for the go-to-market strategy, managing the relationship, integrating the solution and working with third-party developers that bring new services to the customer. To do this the CSPs will either need to develop their own platform businesses or leverage other developer ecosystems.

  This model is the most ambitious of the four from a CSPs' perspective, and given that CSPs have tried, and in most cases failed, to build and operate their own public clouds, it is questionable whether they have the skills to succeed. CSPs pursuing a Full Edge strategy typically will partner with hyperscale cloud providers.

- **Partner Edge** – the CSP owns the relationship with the customer (with whom it likely already has an existing relationship for telecoms services) but uses APIs to connect to third-party edge, private clouds and public clouds. The CSP is responsible for systems integration and must engage directly with developers, although these developers may be part of another edge or public ecosystem. There are a wide range of companies that operators may need to partner with in this role including cloud providers, systems integrators and specialist (vertical) platform providers.

  While many CSPs already bundle cloud with core connectivity services for small and medium-sized businesses, this represents an altogether different level of technical engagement. Still, it is much less risky than Full Edge because the operator is not responsible for building and operating the edge computing facilities. This scenario effectively means that the CSP is delivering a managed service, which is a role many are already comfortable playing.

- **Aggregator Edge** – this is a new concept but one that may be required to enable new applications in areas such as robotics, drones and virtual reality, where data-heavy workloads need to move between a distributed range of edge locations. The edge aggregator provides management and control for these applications and is a trusted partner of the telecoms operator which is integrated at the back-end. The aggregator pays the operator for edge connectivity and compute and then passes the cost on to the developer of the application.

  The aggregator role requires the operator to invest heavily in edge hardware, and if the operator is to provide regional or national services, it will need a large edge computing footprint. However, this scenario requires the emergence of new players to fulfill the aggregator role and partner with CSPs. Examples of this are MobilEdgeX, a Silicon Valley startup founded by Deutsche Telekom in 2018, and UK-based Ori.

- **Limited Edge** – in this scenario the edge computing service is delivered by a hyperscale cloud provider, a systems integrator, the enterprise itself, or a vendor specializing in a vertical such as manufacturing or automotive. The operator's only role is to provide connectivity, although it may have revenue-sharing deals in place (for example, with public cloud providers that have installed their edge computing components inside a CSP's point of presence.)

What's it worth?

While each of the four potential business models requires CSPs to play different roles in the value chain, operators are always responsible for providing connectivity and access to the site where edge equipment is collocated. The only exception is when the edge is at the customer's premises and the enterprise is using its own spectrum and network.

Ericsson estimates that only about 15% of the revenue from an edge computing service delivered to an enterprise comes from spending on site access and connectivity. The largest share of revenue, up to two-thirds of the total, will come from the application itself and delivering the service to the enterprise.

Predicted share of revenue for an edge computing service

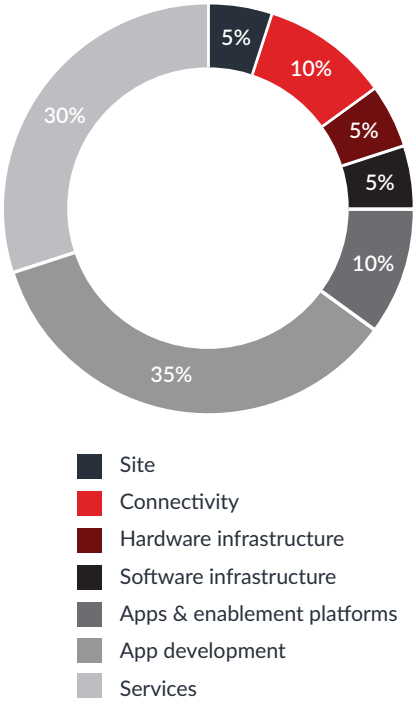

TM Forum, 2020 (source: Ericsson)

Which to choose?

In deciding where they want to play in the edge computing value chain, CSPs must consider the following:

- Whether the company should diversify into new lines of business
- Availability of internal skills and the capabilities of the organization more broadly
- How much they will need to invest in edge facilities
- The size of the market opportunity

It's easy to understand why CSPs may be reluctant to adopt a business model that guarantees only connectivity and site revenue when the real prize clearly lies in a relationship with the enterprise customer. But without connectivity and access, there are no edge services. So, CSPs have some leverage if they want to push for more lucrative deals with partners. Hyperscale cloud providers realize this and are developing commercial contracts with network operators that involve sharing revenue from the sale of edge computing services rather than merely paying them for site access and connectivity (see page 32).

The graphic below shows our assessment of the four edge scenarios, which finds that enterprises are most likely to adopt the Limited Edge scenario. We also evaluate the level of expertise CSPs must have for each business model and the level of investment required. Finally, we look at the potential for CSPs to collect a large share of enterprise spending on edge computing services.

If, as we expect, enterprises adopt the Limited Edge scenario most widely, CSPs will be relegated to providing connectivity and will not have an opportunity to earn a significant share of total revenue. However, if there is market demand for bandwidth-hungry, wide-area IoT services, the Aggregator Edge could become a powerful opportunity for operators.

The Partner Edge approach represents an opportunity for CSPs that already have a strong ICT services businesses, but it will still make up a relatively small number of total deployments. Orange, Telefónica and Verizon are pursuing Partner Edge opportunities.

Telefónica, Telstra and Verizon are also planning Full Edge deployments. CSPs are most likely to be successful with this business model during the short window of opportunity that exists before hyperscale cloud providers deploy their edge platforms, so in the next one to two years. Some CSPs may also be successful with a Full Edge approach in emerging markets where they have a strong background in B2B sales, where rules around privacy favor local players or where hyperscale cloud providers have no national presence. Again, we expect this to account for a relatively small number of deployments.

Given that CSPs continually invest a hefty amount in their networks, it might be logical to assume that they would be willing to invest in deploying edge computing systems. But our conversations with operators reveal that this is not the case at all and that their ability to make speculative, long-term investments is limited. However, unless CSPs make these investments, they will only ever be able to capture a small proportion of total enterprise spending on edge computing services.

Assessing edge business models

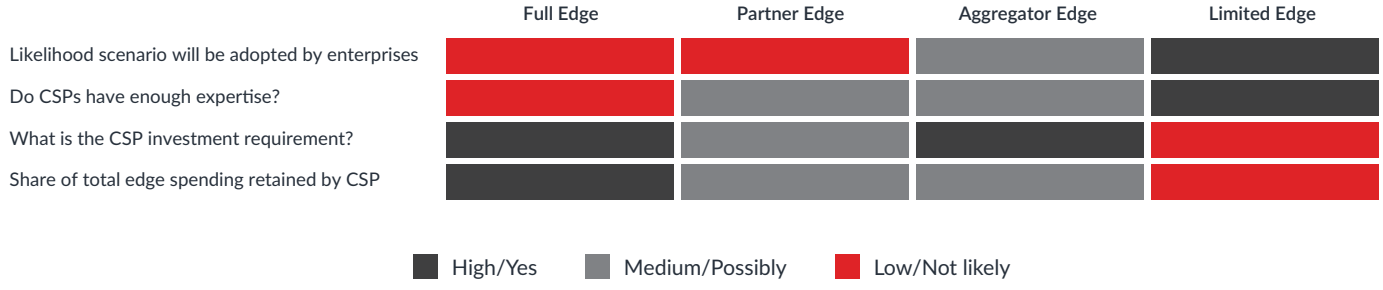

TM Forum, 2020

It won't be easy

It is difficult to make a case for CSPs taking a leading or dominant role in providing edge computing services to enterprise customers. Despite operators' aspirations to diversify their businesses, most enterprises do not consider them as providers of services beyond traditional telephony. This is certainly the case in Europe and the US but perhaps not in less developed markets where hyperscale cloud providers and IT service providers have a limited presence.

When it comes to delivering applications to enterprises, CSPs are not generally playing a role in B2B markets, although there are some opportunities for them to do so, such as in delivering emergency services that benefit from low latency and guaranteed quality of service.

If CSPs had their own developers, they may be able to create new applications that use edge computing. Large CSPs like AT&T and Orange have attempted to build developer communities. AT&T's effort was perhaps most successful, with the company as early as 2010 boasting thousands of developers. Some of those projects have waned or been handed off to open source communities (for example, ECOMP going to the Linux Foundation). Orange, meanwhile, continues to focus on building a developer community around its African payment services business. Most CSP developer community initiatives, however, largely have fizzled out since SMS messaging has migrated to over-the-top applications.

Many edge computing use cases involve delivering enhancements and improvements to existing services rather than creating new ones. As such, CSPs that already deliver ICT and cloud-based services to enterprises will invest in edge computing to deliver better customer experience and new services. For example, Verizon is a world leader in fleet telematics, and it is easy to imagine how edge computing could help the company deliver new low-latency services that benefit from lower-cost and faster data processing closer to the point of consumption.

In the next section we'll continue the discussion of CSPs' edge computing strategies, looking at the necessity of partnering with cloud providers.

Section 6

The power of partnerships

As noted in this report and others, the relationship between CSPs and hyperscale cloud providers is complicated because they are potential competitors and partners. CSPs largely have failed in their own bids to become cloud providers, selling off most of their data center assets. This is relevant to the discussion around edge computing because many enterprises see edge as an extension of or enhancement to their use of public cloud. Indeed, enterprises are likely to have the same opinion about buying an edge cloud service from a CSP than they do about buying a public cloud service from them, which means that cloud providers are going to be necessary partners for CSPs that want to have a bigger role in delivering edge capabilities to enterprises.

In general, operators believe there are two types of customers that they may be able to target without cloud providers as partners: medium-sized businesses and public sector organizations that own or operate vital infrastructure facilities, such as utility companies. Many medium-sized enterprises have relatively unsophisticated requirements for computing infrastructure, so they may be happy to buy services from CSPs rather than public cloud providers. CSPs often already provide connectivity services to such companies, so they may be able to bundle edge computing services with connectivity.

Operators may also have an advantage with government or public sector companies that require stringent security. In some countries, for example, government bodies may stipulate that edge computing must conform to specific security conditions at a national level, and CSPs are more trusted to deliver secure and reliable services.

> However, these are exceptions to the rule. Most CSPs believe they must partner with hyperscale cloud providers on edge computing. "As a general principle, we strongly believe that it is impossible to do it without these guys," says a senior executive working for Orange Group."

Partnership arrangements

In most cases the deals that are being struck between CSPs and cloud providers are based on more than just edge computing. There are three main opportunities:

- **Moving IT workloads to the public cloud** – most CSPs have only a small percentage of their network and operations workloads in the public cloud, but there is strong momentum to move much more over time, with many CSP support system suppliers embracing public cloud (see Section 7).
- **Moving the 5G network core to the cloud** – most CSPs have steadfastly refused to consider the possibility of adopting public cloud for their 5G core, but there are signs this could happen. In September, Telefónica Deutschland announced an "industrial use case proof-of-concept" which involves using Amazon Web Services (AWS) and its cloud infrastructure to virtualize the telco's 5G core.
- **Integrating CSP connectivity with public cloud services** – this is mainly about going to market jointly with 5G and edge computing services, although many of the target customers are likely to be pre-existing users of public cloud services.

As noted in Section 4, AT&T's agreement with Microsoft and Orange's partnership with Google are good examples of telco-cloud partnerships, but the focus and emphasis of these deals are slightly different. The AT&T-Microsoft partnership focuses mainly on migrating IT workloads into the Azure cloud, with a commitment from both parties to jointly bring to market "integrated industry solutions leveraging AT&T's 5G network and Microsoft's edge computing products and capabilities." Specific reference is made to voice, collaboration and conferencing, intelligent edge and networking, IoT, public safety, and cyber security, but there is no suggestion from either company that the deal is exclusive or privileged.

Orange and Google, on the other hand, describe their relationship as "strategic", with an emphasis on Orange using Google's expertise in data analytics and AI to build a next-generation data analytics and machine learning platform rather than on moving IT workloads to the cloud (although this is also part of the partnership). And when it comes to partnering on edge services, Google is a privileged if not exclusive partner.

Terms of the deals

The timing and agenda for CSPs' role in edge computing is being pushed by the public cloud providers which are putting pressure on CSPs to agree to revenue-sharing deals that allow them to collocate edge computing servers at the CSPs' point of presence. Their message to telecoms operators is clear: "If you don't sign up with us now, we will look for other partners because our customers are already knocking on our doors asking us for edge computing services."

While CSPs understand the potential for edge partnerships with cloud providers, they must answer three questions to determine whether to go forward with contracts:

1. **How much of the share of revenue will we collect?** By negotiating an agreement for a specified period of time (for example, one that lasts only two to three years) both sides have the opportunity to go into the partnership with full knowledge that they can renegotiate if the deal works out poorly for them.

2. **How will this impact our own plans for developing edge computing capabilities and services?** Most CSPs that are developing their own cloud computing capabilities are also happy to partner with hyperscale cloud providers; it is not an "either/or" scenario. The industry is still in a phase of experimentation, and by partnering with cloud providers CSPs are expanding the overall market opportunity and can make a stronger business case for preparing their sites to deploy their own edge computing gear and/or that of cloud providers.

3. **Do our central offices have enough space and power for edge computing gear?** This is an issue for CSPs and their partners. However, only a relatively small number of edge locations are needed in a medium-sized country. Furthermore, operators may be able to justify upgrading existing facilities or even building new ones as part of their 5G core deployments.

4. **How exclusive should the deal be given that our enterprise customers likely will use different cloud providers?** CSPs and cloud providers recognize the enterprise requirement for multi-cloud services. This does not, however, mean that all partnerships should be the same. CSPs may prefer the tools or capabilities offered by one cloud provider over another, and they may also see different market opportunities depending on the partner.

Looking ahead

Whether a CSP is entering into edge partnerships, going it alone or doing both, a starting point for targeting the enterprise market should include a detailed assessment of existing customers' locations and the potential to deliver services to specific verticals. CSPs will also need to find the right balance between providing ultra-low latency services in some locations versus guaranteed minimum latency for an entire region or country.

The focus for edge computing until now has been the B2B market, but CSPs don't have a strong track record selling to enterprises, which means they may have difficulty turning the focus to specific verticals. Going forward, operators should use their edge strategies to shape their overall B2B strategies.

With the focus on building out 5G networks, it may be tempting for CSPs to take a wait-and-see approach, allowing their network deployment strategies to determine whether they can play a role in edge computing. But this is risky because it could be two years or more before operators have a fully distributed network edge.

It is unlikely that any CSP will fully integrate its access network strategy with its edge computing strategy. Nevertheless, many of the executives we interviewed said that the ability to convert revenue from services that use network infrastructure to revenue from edge computing services would ultimately determine their success.

CSPs should develop an edge strategy that includes a holistic, cross-departmental approach. This requires coordination across the business, and it means taking a longer-term term view of the market opportunity. As a senior executive from Telstra explains:

"I don't think anyone has any misguided view that [edge] is going to be a major short-term (next two years) revenue stream. This is really about the next five to ten years."

In the next section, we'll look at how suppliers of IT systems and services intend to help CSPs target the edge.

Section 7

The role for enabling suppliers

Suppliers of telco IT support systems and application developers will play an important role in the success of edge computing. Indeed, nearly 70% of suppliers surveyed for this report said they believe edge computing is a significant opportunity for their businesses.

Support system suppliers typically provide centralized applications to manage and orchestrate networks, so the biggest challenge lies in extending these capabilities to the edge, wherever it may be. As noted, there is no widespread agreement about where the edge lies. In many trials, however, it is at the enterprise's premises, which could be an end point in a private mobile network or local area network, or in a public or private cloud.

Most suppliers have not deployed edge solutions yet, but several are participating in trials or anticipate deploying solutions soon (see graphic). In on-premises deployments, support system suppliers can potentially deploy their solutions as-is in enterprise data centers. This includes mediation, performance monitoring, traffic and performance management, onboarding, reporting, security, etc. Solutions do not necessarily have to be cloud native. However, this model likely will not work beyond the early phase of edge deployment as enterprises and operators themselves continue their transition to the public cloud.

Status of suppliers' edge portfolios

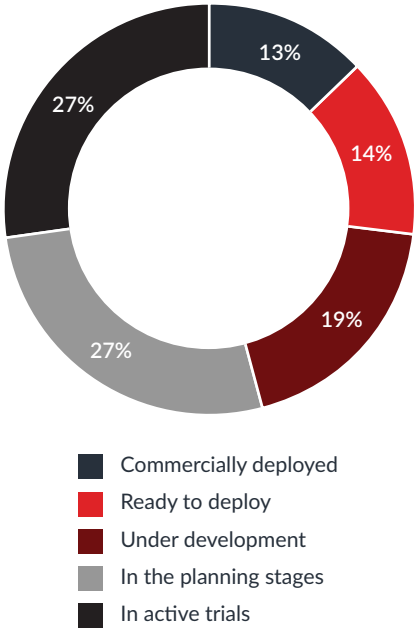

TM Forum, 2020

Depending on the environment (end offices turned edge data centers or colocation centers, 5G aggregation points, compact telco edge locations or distributed cloud centers), suppliers will also need the ability to deploy solutions as microservices in containers. In a 2019 TM Forum report on public cloud adoption, nearly as many CSPs said they would not put their operations support systems into the cloud as said they would (26% to 34%.) However, this was prior to CSPs and hyperscale cloud providers announcing partnerships to serve the edge computing market. Now CSPs would be wise to push their suppliers to accelerate cloud native development.

Read the public cloud report:

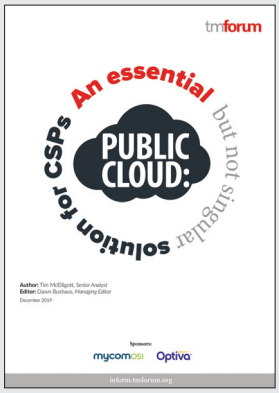

Vendors expect to be involved in private network deployments at various levels, with nearly 40% saying that they think up to half of edge deployments will be on premises in private networks. More than a quarter think it could be as high as 75%.

Percentage of edge deployments suppliers expect to be private networks on enterprises' premises

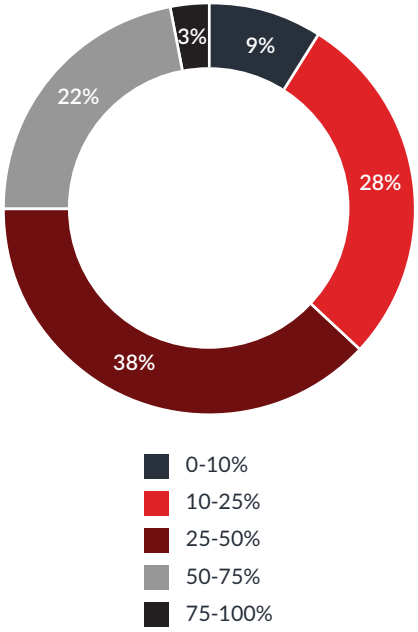

TM Forum, 2020

Early deployments like AT&T's with Rush Medical Center and others bear this out. This leaves out consumer-based edge networks for gaming and other high-bandwidth, low-latency applications which will not require private networks.

Suppliers' strategies

Following are brief assessments of several IT system suppliers' strategies for supporting edge computing services:

affirmed®

Affirmed Networks, a Microsoft company, has deployed its orchestration and automation platform in Etisalat's operations to accelerate automation of the Middle Eastern CSP's existing and upcoming digital services and infrastructure, and to pave the way for 5G and multi-access edge computing (MEC) services. Etisalat has also deployed the Affirmed Cloud Edge solution in a successful proof of concept as part of the company's MEC implementation.

amdocs

In March Amdocs announced a partnership with Google to deliver Amdocs' operational and business support systems (OSS/BSS) on the Google Cloud Platform. Part of this deal includes developing new enterprise-focused 5G edge computing solutions for various industries in effort to help CSPs monetize 5G networks at the edge.

blueplanet
a division of ciena

Blue Planet has delivered its multi-domain service orchestration and network functions virtualization (NFV) orchestration solutions to orchestrate virtualized network functions (VNFs) for several of its partners' solutions, such as Akamai's virtualized content delivery network and Affirmed Networks' cloud edge and virtual evolved packet core solutions. The Blue Planet Virtual Edge Automation solution automates and accelerates the initial launch and lifecycle management of NFV-based network services.

ERICSSON

Ericsson envisions CSPs playing two possible roles in edge computing: an aggregator, where a third party such as a systems integrator, application developer or cloud partner is responsible for delivering the enterprise application; or provider of a "limited edge", where the CSP provides connectivity, an enablement platform and some applications. To serve these scenarios Ericsson has developed the Edge NFVI platform optimized for distributed workloads and edge use cases. It moves traffic through a distributed network meeting the required latency and high throughput, and includes unified management of cloud native applications and VNFs.

Hewlett Packard Enterprise

Together with Schneider Electric and StorMagic, HPE launched a micro data center solution that delivers IT systems for edge computing environments. The solution provides resilient power, ultra-low latency connectivity and secure, on-premises servers and data storage. HPE also extended its GreenLake cloud services to the network edge with distributed environments for application containers, virtual machines and machine learning workloads. This summer, HPE also launched the Edge Orchestrator, software-as-a-service that enables telcos to monetize 5G networks and edge infrastructure by deploying new services at the edge using a catalog of edge computing applications.

IBM

In July IBM teamed up with Verizon Business to deliver 5G and AI solutions at the enterprise edge with an eye toward enabling Industry 4.0. Initial solutions will use Verizon's 5G Ultra Wideband network, MEC capabilities, ThingSpace IoT Platform and Critical Asset Sensor solution as well as IBM's Maximo Monitor with IBM Watson and advanced analytics. AT&T also deployed a MEC and 5G network in IBM's lab in Yorktown Heights, New York, to develop new capabilities for enterprises seeking a private, on-premises cellular service. IBM will also add its edge application manager and the open frameworks provided by Red Hat's OpenShift Kubernetes platform.

Netcracker launched a new suite of OSS/BSS products this summer with a focus on 5G/MEC acceleration, adopting multi-cloud platforms, managing complex partner relationships, optimizing the customer's experience and attaining end-to-end automation. The Netcracker 2020 Digital OSS/BSS and Customer Engagement products and professional services provide standards-compliant MEC orchestration functions to support placement and lifecycle management of edge applications at near and far locations.

Turkcell began rolling out of its edge computing strategy this year with virtual radio access network as its first use case, with Red Hat supplying the federated edge management though its Advanced Kubernetes Cluster Management platform. The operator plans to begin addressing low-latency use cases and near-premises data processing to industries next.

VMware is taking an IoT and hybrid cloud strategy toward edge computing. The company recently began collaborating with Lumen, which will deliver edge services using integrated VMware technologies. In September, VMware began operating its VeloCloud software defined-wide area network in Equinix's data centers and will offer it as a VNF from Equinix's Network Edge platform. Lumen will integrate VMware Workspace ONE and VMware Carbon Black to deliver a "Work from Anywhere" solution on their global edge infrastructure.

Role of developers

Edge computing represents the next frontier for developers of applications, but it is unlikely that a standalone developer ecosystem distinct and separate from centralized computing will develop, even though an explosion of activity in vertical sectors is expected with 5G deployment.

Most edge computing will use a combination of core and edge computing applications and services. Centralized computing will continue to support compute-intensive workloads, data aggregation and storage, AI and machine learning, coordination of operations across geographies, and traditional back-end processing, while edge computing will handle workloads that require near real-time processing.

Edge application development likely will happen, at least initially, as an evolution of development on public cloud networks. Today developers can build their applications on a single cloud, or they can take a multi-cloud approach. Large public cloud providers have their own developer communities which they are bringing with them to the edge, but it is unclear whether CSPs will do the same.

As noted in Section 5, large CSPs like AT&T and Orange have attempted to build developer communities but largely have not been successful. Public cloud providers are making their developer communities available to CSPs, and unless CSPs reach critical mass with their own cloud computing propositions, it is difficult to see how they will successfully build their own developer communities. It is more likely that they will roll out network edge capabilities in a way that maintains the possibility of opening them to third-party developers. This could result in hybrid application development across edge computing platforms and edge networks.

In the next section we offer some advice for companies that are getting started with an edge computing strategy.

Section 8

Make it happen – Strategies for realizing the edge opportunity

Edge computing is a large and strategically important opportunity for communications service providers (CSPs) – one at which they would have to willfully swing and miss not to capitalize on at least in some way. CSPs currently have a near monopoly on the access network connecting almost every "opportunity" in this space to the applications they need now and in the future when 5G is widely deployed. So, the questions really are: To what degree will CSPs be able to capitalize on edge? And how willing and capable are they to make the operational and cultural changes required to do so in a big way.

As our research shows, some large operators are already deploying edge capabilities as private networks or via partnerships with hyperscale cloud providers to deliver on-premises applications. In neither case, however, are CSPs fully leveraging cloud computing and getting all the associated revenue. To do so they must address a glaring problem: They have never been, nor are they now, well positioned to deliver B2B or B2B2X computing services. They do not focus on selling into verticals, and they will continue to struggle so long as their IT support systems and culture prevent them from changing.

Following are steps CSPs can take now to realize the edge opportunity:

Prioritize a strategy

As CSPs build out their new 5G core networks, they will have the opportunity to create a network edge that takes computing platforms closer to customers. In doing so they can support new services such as cloud gaming, which will require consistently low latency across the entire network. But this alone will not allow operators to deliver new revenue streams. They must either build out their own edge cloud computing capabilities to serve specific customers or groups of them, or partner with other companies such as public cloud providers. The approach that they take towards building out their own network edge should be part of an overall strategy for edge computing, one that involves partnering and building end-to-end capabilities, or in many cases acting as wholesalers or enablers of new services.

Focus on B2B

Most mobile operators target the consumer market, but the opportunity for monetizing edge computing lies with enterprises, particularly as 5G is deployed. Building out edge computing facilities or capabilities alone will not be enough. Edge should be part of an overall strategy and vision for delivering ICT solutions and services. Without a strong presence and reputation for delivering traditional connectivity services to the B2B market, it will be extremely difficult for mobile operators offering 5G to succeed in edge computing. This does not mean they cannot play a valuable role, however – for example, they could become enablers for cloud providers or other ICT service providers.

Create a cross-functional team

The survey conducted for this report reveals that no clear pattern has emerged in terms of responsibility for creating, owning and executing an edge computing strategy, but nearly 40% of respondents said it resides in a cross-functional team, which is encouraging. Whatever approach CSPs adopt, it is essential that there is representation and support at the board level. How operators exploit the edge opportunity has huge implications for their future network strategies. Their success in expanding B2B revenue or entering the market for ICT services for the first time hinges on getting buy-in across the company.

Find cloud partners

In the early stage of edge computing, CSPs and cloud providers are necessarily being opportunistic and doing what it takes to get edge services in place without being strict about support and business arrangements. In other words, they are learning as they go, collaborating and applying the technologies they know best while managing their own contributions independently. They are also bringing the enterprise into the process through co-creation.

Long term, however, concrete roles related to ownership of responsibilities and customers will be necessary, which means that CSPs and their partners will soon need to get beyond the "make-it-work" mindset and start defining roles around end-to-end orchestration; onboarding of customers and experience management; network and service management; application management; billing and compensation; building out and managing edge facilities; and more. In other words, they need workable and equitable partnership management.

Consider multi-cloud

Few companies are entering exclusive cloud partnering agreements. Indeed, to do so would be antithetical to the open rules of engagement adopted and promoted by cloud providers, platform providers and developers. Our survey finds that CSPs show little preference for one cloud partner over another, selecting them based on the enterprise opportunity rather than their particular technology approaches. In some cases, however, operators are forging tight alliances with specific cloud providers to offer the fullest range of edge services. Orange's deal with Google Cloud Platform, for example, involves making the cloud provider a "privileged" partner.

Each opportunity may call for being open to different ideal partners. However, being open also must apply to CSPs' networks and IT systems. They should consider transforming to an open framework, such as the TM Forum Open Digital Architecture (ODA), which defines standardized, interoperable software components organized into loosely coupled domains. These components expose business services through Open APIs built on a common data model. Adopting this approach allows CSPs to work with partners without a great deal of integration. Operators do face a challenge, however, in knowing whether to promote or plan for all hyperscale cloud providers to collocate their edge servers in the same locations, because it is not at all clear that CSPs' existing local exchanges will have enough space or power to do this.

Explore use cases

Part of the reason CSPs are experimenting at the edge and not moving faster could be because there is no killer use case. But there may never be one. Operators must work with partners to develop edge use cases. There is no monolithic edge service that operators can deliver market-wide like digital subscriber line or even 4G and 5G. Developing edge computing services requires a vertical mindset depending on which role a CSP wants to play. For example, some operators may take an ecosystem approach that allows partners who have a viable edge computing proposition and existing user base to determine which verticals to target. Other operators may choose to develop vertical use cases one on one with enterprise customers or through direct involvement with industry-specific associations, such as a manufacturing consortium.

Leverage developers

CSPs may be too late in the game to attract developers, but they can leverage other platform companies' developer communities. This includes not only the major cloud providers, but also companies like Intel which is a huge builder of edge data centers. However, operators must be crystal clear about their motivation for attracting developers and how they can monetize the use of edge applications. This will most likely require them to develop platform business models. CSPs should build 5G and edge lab environments dedicated to application development, and these should be open to other CSPs, hyperscale cloud providers and open-source platform providers, providing access to real-world telecom operations.

Increase automation

A full 82% of CSP survey respondents said that lack of closed-loop automation is a barrier to realizing their visions for edge computing. Closed loop processes involve collecting and analyzing data to figure out how the network can be optimized and then implementing those changes automatically. This capability has been improving incrementally in CSPs' operations since the turn of the millennium, but it needs to accelerate.

A more comprehensive, enterprise-wide approach to coordinating automation across processes, departments and business units is essential to the overall ability to deliver real-time services. Automation is also essential to streamlining the operations of the business. Automation will be necessary for edge computing in order to work seamlessly end to end across the ecosystem of partners that will be involved in providing the services.

TM Forum Open Digital Framework

A blueprint for intelligent operations fit for the 5G era

The TM Forum Open Digital Framework (ODF) provides a migration path from legacy IT systems and processes to modular, cloud native software orchestrated using AI.

The framework comprises tools, code, knowledge and standards (machine-readable assets, not just documents). It is delivering business value for TM Forum members today, accelerating concept-to-cash, eliminating IT & network costs, and enhancing digital customer experience.

Developed by TM Forum member organizations through our Collaboration Community and Catalyst proofs of concept, building on TM Forum's established standards, the Open Digital Framework is being used by leading service providers and software companies worldwide.

Core elements of the Open Digital Framework

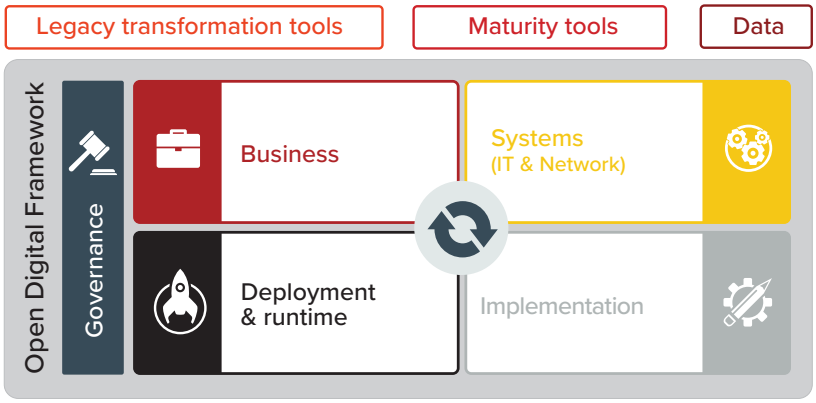

The framework comprises TM Forum's Open Digital Architecture (ODA), together with tools, models and data that guide the transformation to ODA from legacy IT systems and operations.

Open Digital Architecture

- Architecture framework, common language and design principles
- Open APIs exposing business services
- Standardized software components
- Reference implementation and test environment

Transformation Tools

- Guides to navigate digital transformation
- Tools to support the migration from legacy architecture to ODA

Maturity Tools & Data

- Maturity models and readiness checks to baseline digital capabilities
- Data for benchmarking progress and training AI

Goals of the Open Digital Framework

The aim is to transform business agility (accelerating concept-to-cash from 18 months to 18 days), enable simpler IT solutions that are easier and cheaper to deploy, integrate and upgrade, and to establish a standardized software model and market which benefits all parties (service providers, their suppliers and systems integrators).

Learn more about member collaboration

If you would like to learn more about the Open Digital Framework, or how to get involved in the TM Forum Collaboration Community, please contact George Glass.

TM Forum research reports

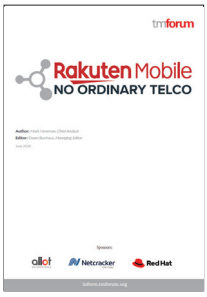

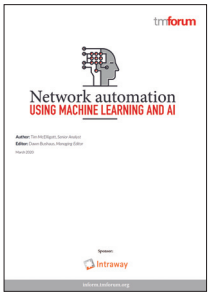

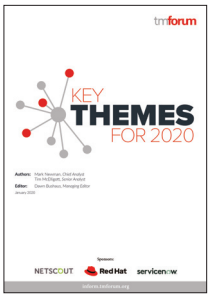

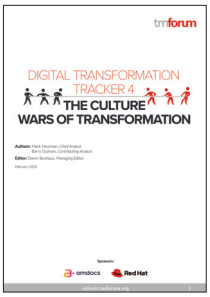

Meet the Research & Media team

Report Author:
Mark Newman
Chief Analyst
mnewman@tmforum.org

Report Author:
Tim McElligott
Senior Analyst
tmcelligott@tmforum.org

Report Editor:
Dawn Bushaus
Managing Editor
dbushaus@tmforum.org

Customer Success & Operations Manager:
Ali Groves
agroves@tmforum.org

Editor, Digital Content:
Arti Mehta
amehta@tmforum.org

Global Account Director:
Carine Vandevelde
cvandevelde@tmforum.org

Commercial Manager, Research & Media:
Tim Edwards
Senior Analyst
tedwards@tmforum.org

Digital Marketing Manager:
Anna Kurmanbaeva
akurmanbaeva@tmforum.org

Report Design:
Intuitive Design UK Ltd
info@intuitive-design.co.uk

Published by:
TM Forum
4 Century Drive,
Parsippany,
NJ 07054
USA
www.tmforum.org
Phone: +1 973-944-5100
Fax: +1 973-944-5110
ISBN: 978-1-945220-85-2

For more about TM Forum's Collaboration Community,
please contact **George Glass**