Edge computing: Hype or ripe?

How telecom operators will create – or destroy – substantial value via edge computing

October 2021

Arthur D Little

Content

Authors:

Bela Virag
Managing Partner
TIME*, Vienna
virag.bela@adlittle.com

Ashish Sharma
Partner
TIME, New Delhi
sharma.ashish@adlittle.com

Andrea Faggiano
Partner
TIME, Dubai
faggiano.andrea@adlittle.com

Sean McDevitt
Partner
TIME, Boston
mcdevitt.sean@adlittle.com

*TIME: Telecommunications, Information Technology, Media & Electronics

Executive summary

With the expansion of technology worldwide, it is certain that more compute capacity closer to where the data is created and/or used will be needed. We can expect that need to grow by a factor of 10 in the next 5-10 years. But who will provide the needed compute infrastructure?

Great opportunities are on the horizon for technology developers (e.g., Intel, Microsoft, AWS), system integrators (e.g., Accenture, IBM), solution providers (e.g., Siemens, Schlumberger), IT providers (i.e., distributors, local integrators, experts), hyperscalers (e.g., Microsoft, Amazon, Google, Oracle), telecom operators, and likely many others. The battle for the *edge* has begun and lines are being drawn.[1] All these players are assessing which part of the value chain to focus on: hardware, orchestration solutions, system integration, ecosystem, or system operation. And they are asking some key questions: Where is the value? What to resell and whom to allow to resell own services? How to spread investment and risk?

In this Report, we consider these questions from the perspective of telecommunications network operators. A few strong believers – including players like Verizon, Vodafone, and Telefonica – have the rest of the industry wondering: Are they right? Wrong? Too early? Too slow? We believe for the front-runners to be right, a few things need to be believed:

- There is a need to have compute infrastructure outside the central data center or cloud – not in on-premises, self-built environments but somewhere in between.

- Operators can add and capture value from the new computing infrastructure.

- Today's hype is relevant because today is about securing options; it's not yet about building rollout investments.

Substantiating these beliefs and justifying any investment requires: (1) the right strategy, (2) the skill to structure commercial arrangements so that win-win partnerships emerge and coopetition becomes possible, and (3) a sound balance between asset-heavy and asset-light versus risk in a fast-aging tech race.

Beyond the market-facing case for telecom operators, there is one more consideration: eventually, we expect networking functions and customer workloads to run on a single appliance – despite current multi-access edge computing standard designs. This means that network operators, when assessing an edge computing engagement today, need to not only consider any potential incremental revenue upside, but also which options to keep open in regards to owning or sourcing compute capacity for their future network function workloads. Thus, today is not about significant rollouts, but it very much is about keeping options open.

---

1   For the purposes of this Report, edge computing refers to the provision of computing services delivered by servers neither located in data centers nor on customer premises or devices, but instead on servers distributed nationally at multiple locations – up to 50 km from one another. We leave out adjacent terms, such as fog computing (computing on individual devices), the deep edge (on client premises), core location data centers, and edge data centers (sometimes called second- or third-tier locations).

1. The promised good

There are new announcements almost weekly of deals and edge computing technology partnerships between telecom network operators and IT companies such as Microsoft, Amazon, Google, IBM, and others – the "hyperscalers." These companies all share a common vision: the sheer volume of data will explode, driven by, among other things:

- Digitalization across industries.

- Transformation of operations technology (OT), such as machine-control systems and so on, to information technology (IT).

- All types of Internet of Things (IoT) applications that enterprises and governments will deploy or that consumers will consume.

- New, more capable consumer gadgets and new forms of entertainment and experiences.

- An increased use in robotics and autonomous-guided vehicles, both on the ground and in the air.

All these will demand new locations for computing. The computing power will have to be physically closer to the place where the data is being generated (e.g., sensors, video cameras, game engines) and where the results of the computations are being consumed (e.g., consumers, actuators, robots). The argument is that only edge computing can satisfy the needed characteristics for such compute infrastructure, in terms of:

- Latency; that is, the time it takes to turn a sensor's signal into an actor's action.

- Ability to store, aggregate, and process large amounts of often unstructured data.

- Satisfying the need to keep data protected and potentially on-premises, or at least away from central data centers.

- Openness to an ecosystem of software applications and data pools.

In essence, there is a battle emerging for the winning IT platforms on the edge. As shown in Figure 1, the battle is being fought on two fronts:

1. **Location of compute** – on-premises, near-premises (operator's edge), or in a central data center.

Figure 1:   The battle for the edge
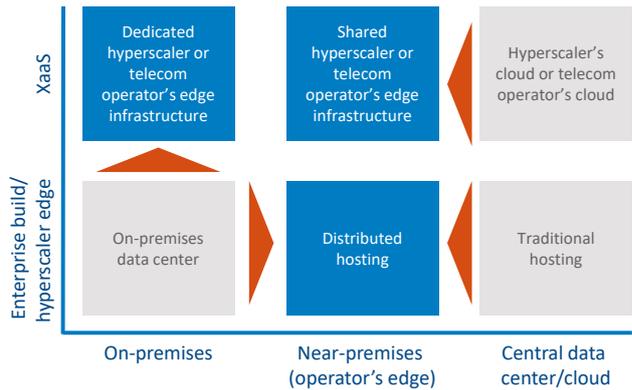
Source: Arthur D. Little analysis

2. **Make or buy** – whether it is an operated service, a self-built infrastructure (e.g., VMware, Red Hat, OpenShift), or a hyperscaler edge infrastructure (e.g., Microsoft Azure Edge Zone, Amazon Outpost).

To drive edge demand, telecom operators should motivate their clients to undertake the following:

- Deploy and shift workloads from central data centers or the cloud to near-premises distributed hosting or to on-premises dedicated edge compute cloud.

- Substitute on-premises data centers to near-premises distributed hosting or to on-premises dedicated edge compute cloud – often shifting CAPEX to OPEX.

There are two different rationales to convince clients to do so. While security and latency concerns seem to drive the former, market forces, driven by industrial digitalization, seem to support the latter.

The idea to encourage clients to process data locally to avoid transporting data over longer distances to a central data center or cloud seems especially at odds with network operators' purpose to transport data efficiently. Thus, the question may arise whether network operators are better off transporting the data or providing the local compute? We believe this depends on market ICT maturity and the specific client use case. Where data is best processed is not driven by networking cost, but rather by the use case itself.

2. Creating value: The edge of opportunity excellence

The opportunity – Hero use cases

To summarize opportunity for telecom operators, the "hero use cases" for edge computing include:

- **Artificial intelligence (AI), machine learning (ML), augmented reality (AR)/virtual reality (VR)/mixed reality (MR), and robotics/drones** – leveraging the advanced technology ecosystem while avoiding having to operate the infrastructure.

- **Video analytics/computer vision** – avoiding the cost, time, and effort of transporting video streams through networks.

- **Data aggregation** – avoiding the transport of large data volumes to achieve lower latency and to avoid public data centers.

- **Device offloading/gaming** – enabling a subscription business model and removing barriers of current model.

- **5G apps** – enabling customer experience with 5G networks.

AI, ML, AR/VR/MR, and robotics/drones

Enterprises that wish to employ AI, ML, or AR/VR/MR will require compute capacity nearby with high performance and low latency. Since these services often run on hyperscalers' platforms, such enterprises may be among the first to consume private edge compute services. They will use them, for example, for IoT use cases slowly finding their way into production realities, including robotics, sensor/actor-based control/automation, and other control systems (e.g., supervisory control and data acquisition systems – or SCADA – the control architecture that includes computers, networks, and control surfaces for high-level process supervisory management).

This need for local compute capacity, however, may expand beyond the enterprise campus into the public space, where similar applications will need to offload some compute from their devices to edge compute capacities. As such, we can expect to see two types of industrial edge computing: one dedicated to a particular location and another to supporting devices that move around a geography (whose compute capability must move with the device to remain in proximity).

Video analytics/computer vision

Video analytics (e.g., optical video as well as X-ray, lidar, or point clouds [ultrasonic, etc.]) and computer vision can either be processed on the individual device or on an edge compute infrastructure. Transporting hundreds or thousands of video signals through networks is less economical and may violate latency requirements (e.g., video overlays, real-time analytics, robotics). However, it is beneficial both economically and performance-wise to shift such workloads from the individual "intelligence on the device" domain to an edge compute infrastructure.

On top of baseline video analytics and computer vision, navigation support, traffic control, and mapping require significant on-site compute power and data aggregation capabilities. In particular, if data aggregation or data processing needs to happen in the public space, the edge is probably the nearest place to do so.

Data aggregation

Aggregating data for the purpose of analysis shares similar needs to video analytics: transportation is more expensive than placing compute infrastructure closer to the data generation. But there are two additional reasons for putting data aggregation into an edge environment:

1. **Latency.** If data needs to be aggregated in-line (while a production is running, a robot is moving, etc.) or if the volume of data to be aggregated is too voluminous to transport, the only solution is to do so nearby and then act upon the aggregated data.

   However, given that there are three places in which to aggregate data (i.e., on-premises, near-premises, or in a central data center/the cloud), we can expect only a subset of all use cases to be processed on the edge. The chief driver for edge processing is when "on-premises" is really in the public space. In these use cases, the edge may be the natural choice.

2. **Data sovereignty.** Some use cases demand data be kept away from any central infrastructure (e.g., data sets with limitations due to national security, community/municipality demands that should be kept local). If keeping data decentralized is essential, then it can either be on the sensors themselves (which, by their nature, are more difficult to manage across their lifecycle), in on-premises facilities, or in near-premises edge compute environments.

Device offloading/gaming

Most small smart devices (e.g., smart watches, glasses) are paired to mobile handsets and offload certain compute requirements to nearby phones. There are multiple reasons for this, including power provision, processing power, and business model. However, phones are not the final answer to nearby processing, as they, too, are run off of batteries. The engineering of ever more on-device computing will eventually become less economical than offloading – for AR glasses, phones, and other smart devices alike. This is certainly true for AR, which finds its first applications in B2B. But is no less true for cloud gaming, for example, where the business model shift is an interesting incremental aspect.

The case for "gaming in the cloud" rests on two pillars:

1. **Games will not continue to be installed onto devices.** Customers want to play more games than they will buy and install on their PCs, and they want to play PC-quality games on their mobile devices, essentially driving the cloud-gaming-as-a-service model. And, if networks perform well enough, there is no reason for consumers to invest in powerful gaming hardware, be it a PC, console, or mobile device. Cloud gaming shall instantly provide consumers with the most stunning audiovisual gaming experiences. For game publishers, it extends the market beyond those with deep pockets for a gaming PC to include occasional or ad hoc players, wherever they are, with a subscription or an ad-based business model to consume a plethora of games.

2. **The question to consider is: where will gaming content be computed?** There are not that many options. Games are more power-hungry than ever, including the electricity that sustains them. Data centers are designed to efficiently deliver increasingly green power, but the next data center may be located too far away from the network. One could argue that a less than 10 ms latency is needed for competitive gaming, but this is only a niche. If the gaming industry adds AR, VR, and MR to its gaming experiences, requiring the capture of head/eye/body movement, not being near the data center will create headaches and dizziness – literally. Thus, to provide computationally more sophisticated games on ever-smaller, battery-powered devices such as

glasses or goggles, these workloads must be close to the consumption. This is where the edge makes most sense.

But let's not forget, that cloud gaming is only one example, where device offloading may be sensible. While B2B examples include AR glasses in warehouses, assembly, training, and so forth, B2C examples include education, communication/entertainment, e-commerce, and more.

5G apps

Increasingly, telecom networks and the related functionality evolves into being mostly software-based and no longer appliance-based. This is true for cellular as well as for fixed-line services. And this enables 360-degrees integration between applications and the network to effectively enhance customer experience. For instance, if a network can anticipate congestion within the next 100 ms for a specific user consuming a video stream, it can signal to the video-encoding engine to lower the encoding rate and thus avoid the little spinning circle – in real time.

The technology for such real-time integration is not yet ready/available, but with multiple equipment providers claiming to offer software-based, cloud-native, real-time networking functions, we can expect this to change. Microsoft already has announced its intent to place radio access network (RAN) functionalities onto its Azure portfolio for communications service providers. Since RAN requires real-time processing capabilities, making use of network information to optimize application behavior will grow in relevance and importance.

Cloud players vs. telecom operators – Next best moves

The value chain to edge computing takes place between backhaul, facility, and RAN, as illustrated in Figure 2. If we assume that the top contestants to capture value from edge computing are hyperscalers, software vendors, system integrators, and telecom operators, it is clear that the approaches will differ greatly.

Each player has a stab at the edge computing market – with different chances of success. Here, we assess them in two groups:

■ **Type 1: cloud players** – hyperscalers, some software vendors (SAP, Oracle), and systems integrators (IBM, Accenture).

■ **Type 2: telecom operators** – telecoms and their offspring, such as TowerCos, fiber companies, and other telecom infrastructure companies.

Figure 2: The value chain to edge computing
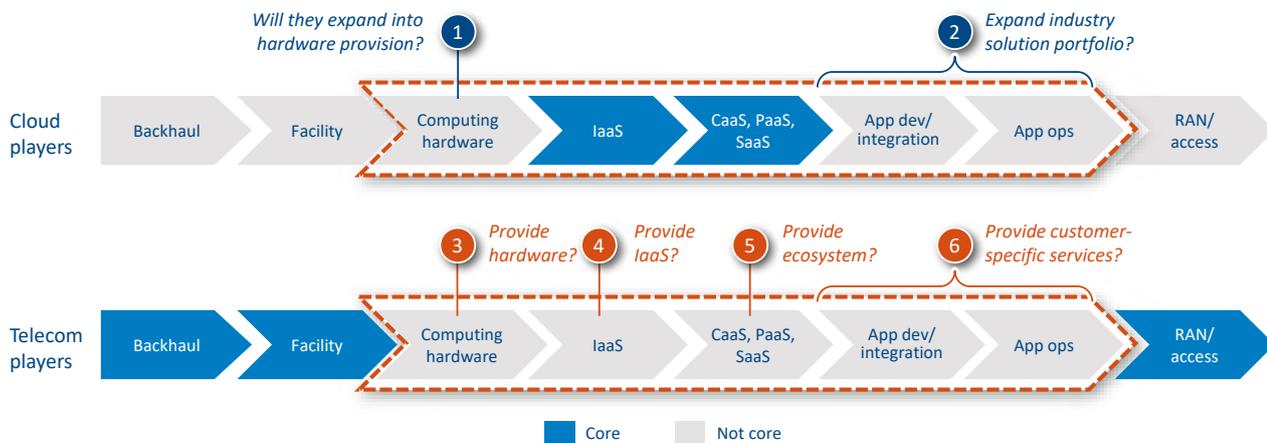
Source: Arthur D. Little analysis

In the evaluation of opportunities, this Report excludes the value of powered, secured, and connected real estate (backhaul and facility) as well as the value of RAN/access, as these are unchanged in any of the development opportunities and must form the base case. We can assume that approximately 15%-25% of the total value of edge computing is in these areas. This part can be captured by telecom operators, TowerCos, and/or other infrastructure investors with their different plays.

Beyond backhaul, facility, and RAN/access, players should consider the following moves:

1 **Computing hardware.** Cloud players could expand into the business of providing small-scale data centers across a country. This may seem simple, but it is a very different business to operate a few centralized data centers than to operate in hundreds of locations. Unlike cloud players, telecom operators are typically familiar with these challenges.

2 **App development, integration, and operations.** On the other side of the value chain, cloud players could expand into the application development and integration business as well as the application operation segment, if they are not already doing so. One such vertical is the telecom industry, with its network functions that run on standard cloud infrastructure – a possible anchor tenant. Although the networking software applications will not run on the same appliance as client workloads, the initial setup effort can be shared.

3 **Telecom players to deploy hardware.** Telecom operators could deploy computing hardware and offer that capacity to cloud service providers and customers. Given that telecom operators are recognized for their capability to manage distributed technical assets, this seems like a natural fit.

4 **Telecom players to provide IaaS services.** Telecom operators could go one step further and provide infrastructure as a service (IaaS). While utilizing a hyperscaler's technology to do so is one option, it is not the only one. There are reasons to do so utilizing other technologies, too, including license cost, data regulatory regime, differentiation, and so on. At the same time, hyperscalers provide a focal point for an ecosystem of software developers, which other technology solutions cannot provide at the same breadth.

5 **Telecom operators could move beyond IaaS** and provide containers as a service (CaaS), platform as a service (PaaS), and software as a service (SaaS). While their attempts to do so have not been successful on a broad basis, some segments within these fields do allow for successful entry of telecom operators.

6 **Telecom players to provide customer-specific use cases.** Telecom operators could also select a few verticals and provide application-level services specific to these verticals. Vertical candidates include automotive, public institutions, railways, gaming, street retail, drone space, and others. While this approach seems to be emerging, and a few telecom operators have already placed their bets and still more are actively thinking about how to embark on this journey, it is (1) not easy to do, and (2) many have failed in the past.

The quadrants shown in Figure 3 illustrate the likelihood of success and the expected value of the six options. From this matrix, we learn the following:

- **The largest value is in industry- or customer-specific solution provisioning.** While cloud players have a higher chance of success when it comes to developing industrial solutions, as these scale globally similarly to cloud players themselves, it remains to be seen whether telecom players can do so as well. There are examples in which telecom operators have shown a great ability to enter into customer-specific or even overarching use cases (e.g., surveillance and alarm services, in-car services). But there have also been many failed attempts. The question is therefore: why/how/what segment should telecom operators enter into?

- **Telecom operators can improve value capture if offering operating services of technical infrastructure.** However, these typically require CAPEX investment and may be too risky to engage early on, as server CAPEX quickly becomes dated.

- **Since the likelihood of success is relatively limited on IaaS, CaaS, PaaS, and SaaS plays related to ecosystems, we recommend that telecom operators not engage in this area.** Telecom operators have mostly dropped out of the battle for these ecosystems. However, their strategy should be to endorse and support the creation of such plays to stimulate the overall market and increase margin capture from backhaul, facility, and RAN, as well as potentially moves 3 and 6 (see figure below).

Figure 3:   Edge computing value matrix
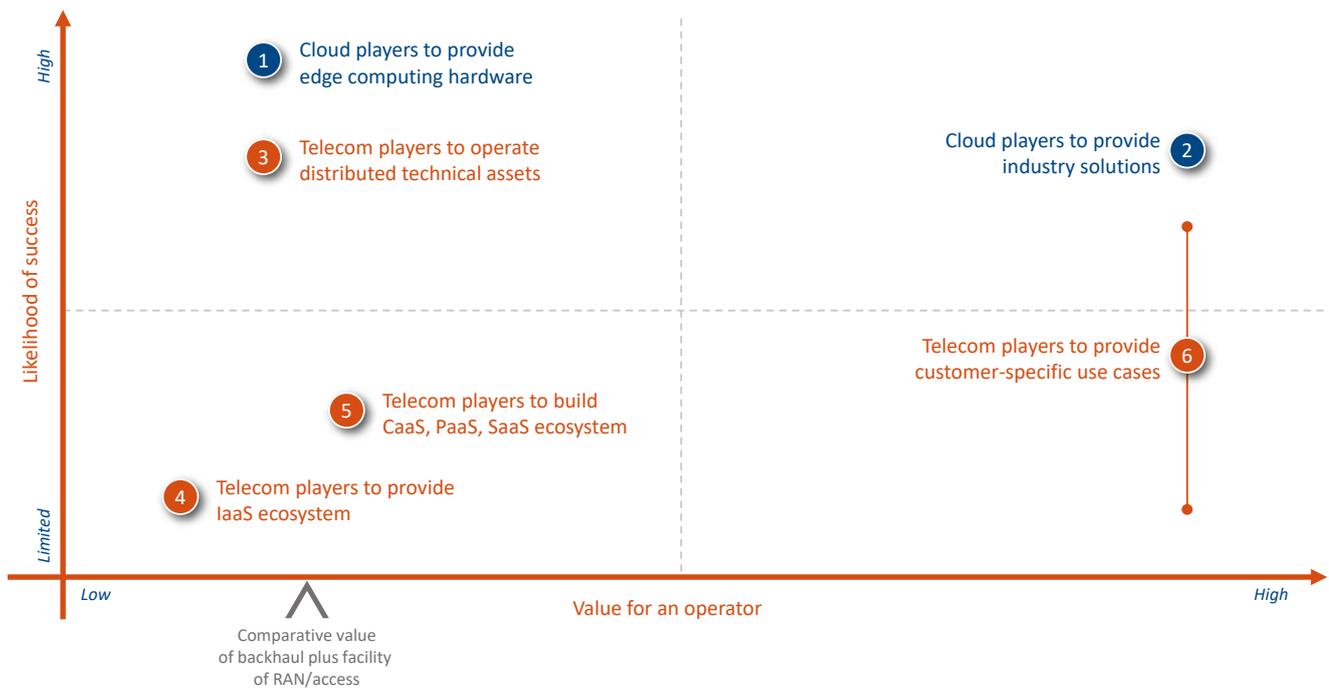
Source: Arthur D. Little analysis

3. Destroying value: Myths from the edge

There are several myths surrounding edge computing that can lead telecom operators to poor decisions and potential value destruction.

Myth 1 – The edge market is huge

Various analysts have forecasted significant market growth, with some estimating more than 25% CAGR and market sizes reaching over US $15 billion by 2025 or as much as more than $60 billion by 2028. In contrast, the total cloud computing market has been forecasted to reach more than $500 billion by 2025 to as much as over $800 billion by 2028. Whether or not these estimates are accurate, they suggest that the edge computing market merely reaches a maximum of a 10% share of the total cloud computing market. This is the near-premises segment, so excludes any on-premises spend. Since we can expect higher unit costs for edge computing than for classic public cloud services, the volume share for that segment is even less. As a result, we can assume that there will not be enough space to significantly overbuild an area with competing infrastructures.

Myth 2 – The market is growing fast

While we see "digital" being accelerated, particularly due to COVID-19, this does not mean that the edge will benefit to the same extent from this acceleration. Most common corporate workloads currently do not require edge computing. In the advent of new use cases or the creation of new devices, this demand could surface, and wider deployments could take place.

Given that the first one to meet demand is the winner, a "build it and they will come strategy" may seem appropriate. However, since technological evolution is still very fast, taking big, uncovered bets is exactly that: you need to be certain you can capture an infrastructure-backed position in this space and accept that it may take some years.

Myth 3 – CDN is a killer app

Akamai claims to be "the largest provider of edge computing services by far," with 300K servers deployed in 4,000 locations.[2] While this is truly impressive, Akamai CEO Tom Leighton also claims that this is equivalent to a $2 billion business, if reported separately. And, as he elaborates, putting all other content delivery network (CDN) players together will not come close to Akamai's footprint. (CDNs ensure content is stored and provided close to the content consumer.)

However, most of this is storage and less is compute. Thus, if market forecasts are correct, this would imply the market will find workloads and deploy infrastructure that are 10x CDN providers' current volumes in the coming four to five years. The growth therefore will come from the opportunities discussed earlier, and CDN will be a smaller share in the total edge computing market.

Myth 4 – Edge computing always reduces latency

If the services in the edge function fully autonomously and don't require any "call home" for any reason (e.g., for authentication, encryption keys, or even some logic), the edge will reduce response latency, sometimes even significantly.[3] The moment the application needs to call home, that latency advantage begins to disappear.

Some additional points about latency:

■ Despite current hype around low-latency requirements being the promised land for service providers, we have not yet been able to identify use cases or business cases that are sizeable and demand low latency in the near term. In many geographically smaller countries, latency requirements for most if not all applications are easily met when utilizing one or only very few data center locations if connected via fiber infrastructure.

---

2   Source: Dano, Mike. "Akamai CEO: We dominate edge computing." Light Reading, 30 July 2020.
3   Our calculations show that reducing the distance of the compute setup from 500 km cable length to 100 km fiber length (i.e., after the radio access) will reduce latency from 2.5 ms to 0.5 ms in one direction. Consequently, connection setup time is reduced by more than 50% – from 60 ms to less than 30 ms (including 5G) – and a 3 MB package will be transported in 90 ms instead of 135 ms.

- It can be assumed that regular applications perform significantly better if the latency is being reduced between the data used and the compute. Thus, central data centers or cloud environments will always need to have their data nearby to perform. If that is not desirable or possible, edge computing becomes an alternative – whether deployed on-premises or near-premises.

Myth 5 – Edge computing costs the same as data center computing

Deploying and operating edge computing infrastructures is more costly on a per-unit-cost basis than deploying and operating computing infrastructures in data centers. These disadvantages in per-unit-costs are incurred for service and maintenance, the casing/ruggedizing/physical protection per device, and so forth.

4. Cutting-edge perspective and predictions

Who should invest?

We have identified three possible categories of investors in edge compute infrastructure: enterprises, telecom service providers and their offshoots (e.g., TowerCos, neutral hosts), and hyperscalers (including emerging ones).

Most likely, all three will form partnerships to fund, invent, and drive edge computing and showcase the results. Examples include AWS's partnership with Verizon, Vodafone, SK Telecom, and KDDI based on the AWS Wavelength service or Microsoft's partnership with Vodafone, Rogers, AT&T, Telefonica, UAE's Etisalat, CenturyLink, Proximus, NTT, and other operators based on Microsoft Azure Edge Zones or Azure Stack Edge.

From the perspective of enterprises, investments into edge computing infrastructure enables capturing all use case value of whatever more advanced digitization means in their context: new products and business models, more productive and safer manufacturing and logistics processes, a safer and healthier public, and so on. The investment into the compute infrastructure is often the smallest part of the entire use case.

For hyperscalers, such investment allows them to get closer to their customers and expand their global ecosystem of application providers to their enterprise customers. This makes their platforms more attractive to their ecosystem of application developers, which is particularly relevant in the context of industrial digitalization, and especially as operations technology accelerates its transformation to IT.

On the one hand, telecom operators can definitely capture value from the foundational services, such as backhaul and rentals, access network, and potentially from provisioning of the IT infrastructure. Their own cloud computing services, on the other hand, have often not achieved the aspired successes in their respective markets. Success, however, varies among markets and positioning. One hindrance is that many telecom operators are limited to national boundaries. This limitation inhibits meaningful access to the often-global technology business models of solution providers. Therefore, telecom operators should focus on solutions that are valid in a local context if they are not software only but also require some physical involvement. This focus on local will greatly increase chances of success and opportunities for achieving defendable margins.

For telecom offshoots, such as neutral host providers and TowerCos, investment in edge compute infrastructure is likely a sound strategy. There are two reasons neutral host providers should provide edge compute infrastructure: (1) there is not enough money in the market for any significant overbuild, favoring "sharing business models"; and (2) it is their core business to provide infrastructure. However, they need to attend to the fact that they are used to the margins and financial structure of long-lived, nonperishable assets. This is very different in the IT world. The IT infrastructure business has shorter lifecycles than TowerCos generally do, causing new types of risks that require assessment.

Edge computing is more "big bet" than robust strategy

Edge computing requires a deep, multifaceted commitment. Electrical, networked, monitored, and managed real estate across countries is an increasingly valuable asset that telecom operators presently own or contract. This value has been proven for network services – both mobile and fixed – and for CDN and related services. The bet is: will edge computing be one more category following the same equation?

There is no clear answer to this question, yet. So far, most telecom operators have failed to capture value from computing services, and computing services providers have increased costs and CAPEX for telecom operators, without them benefiting equally (for multiple reasons). For telecom operators to succeed with edge computing, the scenarios described below would have to exist.

Enterprise workloads

- Enterprises will continue to drive their cloud migration programs.

- Enterprises will migrate to the cloud not only for servers located in data centers but also for the compute demand in other facilities, including factories, shop floor, office buildings, and so on, to enable AI, ML, AR/VR, and robotics/drone use cases.

■ Enterprises will not revert to automation of their own virtual or nonvirtual compute infrastructure but instead will utilize provided cloud environments. (In this case, it is likely that dedicated edge compute infrastructure close to or on-premises is a feasible option for local workloads.)

Consumer workloads

■ Immersive experiences and offloading:

– Devices such as watches, glasses, VR/AR/MR, and so on, gain scale.

– Less power is consumed to transmit a signal than to compute the experience (e.g., computing the image processing, facial recognition, position/rotation/rendering of VR content).

– Battery-efficient signal- and data-processing chips exist at lower cost and power budget, enabling offloading to an edge compute infrastructure. In this case, it is likely that there is a rational demand for an edge compute infrastructure. However, it is a chicken-and-egg problem, thus either a presumably disruptive device or application would need to meet investment appetite before launch, or such a move would grow organically beyond today's local networking link between the device and the smartphone.

■ Local sensory for mass markets:

– Services requiring local sensory are being created and deployed (e.g., local weather, traffic, parking). Such services either:

  – Generate more data than is transported.

  – Generate data that is too sensitive to transport over long distances to a central facility.

  – The data needs to be processed with lower latency than a central location would allow.

– The required data processing can be done in a locally and distributed infrastructure setup.

  – The result of the computation is to be transported upstream or to be used to effect local infrastructure/control systems. (Most likely, security and public safety are a design concern. Even though such use cases may be utilized by private institutions, we can expect public involvement/interest to be a key driver.)

Governmental/public space workloads

■ Public service or safety-related use cases gain ground in the public space and require local data processing – for any of multiple possible reasons.

■ It becomes evident that such services are more efficient if run on shared, general-purpose hardware and not on special-purpose hardware.

– Special-purpose hardware manufacturers unbundle their integrated setups. In this case, we can expect a lengthier process of standardization, changes in industrial behavior and competitive logic, as well as the emergence of publicly desirable technology (e.g., surveillance, self-driving vehicles).

What models can telecom operators consider following?

Telecom operators' models fall into three broad categories: asset-light, asset-heavy, and dedicated or shared:

1. **Asset-light.** An example is AWS Wavelength, which is a revenue share model for operators. This service is targeted at shared setups, so is less likely to be positioned on a customer premise but is close by, in the network. The deployment aspiration is to cover geographies rather than multiple singular or individual locations.

2. **Asset-heavy.** An example is Microsoft's Azure Edge Zones with carrier. The fundamental idea is to place Microsoft's compute infrastructure into the next 5G data center to allow for very low latency computing while having the full public cloud platform services available to the operator's customers.

3. **Dedicated or shared.** An example is Microsoft's Azure Edge Zones, which can run in both a connected and a standalone, or private, fashion. These services are similar in their business model to Amazon Outpost, Google Anthos on bare metal, or IBM Cloud Satellites. Some are pickier about hardware than others, and some are better integrated with the public cloud services of their creators than others.

Operators that want to get into the IaaS and PaaS game close to the client, or even on-premises, must make choices. These choices are determined by the willingness to invest CAPEX for a specific client, or in general for shared infrastructure and the fitting revenue model that comes with it.

Conclusion

There are many reasons why it is certain that more computing infrastructure will be needed closer to where the data is originated or utilized. This fuels the edge opportunity. A battle is emerging around the anticipated value creation; hyperscalers, system integrators, telecom operators, and so on, are all staking their positions and are mingling to shape all kinds of partnerships. This battleground is shaped by two dimensions:

1. Will workloads from centralized cloud environments start to become closer to premises and be more distributed?

2. Will enterprises transform or augment their on-premises data centers to cloud environments or to decentralized data centers? (This would be equivalent with a CAPEX to OPEX shift.)

Edge computing will evolve both on client premises as well as slightly remote. Any comprehensive edge computing portfolio must therefore include on-premises and near-premises solutions. The edge is a clear opportunity for hyperscalers to sell their cloud technology stack – and ecosystem of application developers.

For telecom operators, the opportunity to provide managed, powered, and connected space to technology providers is also clear. The fear that this move would invite competition is not founded. On the one hand, often, few locations are sufficient to cover any low-latency demand. Next, any competitor could easily find alternative locations for their setup – including the client's premises. And finally, telecom operators will need to partner with hyperscalers in all cases.

The key questions telecom operators must ask to determine the attractiveness of edge investments include:

■ **Strategy**

– What is our right to play (beyond backhaul, facility, and access network)?

– How do we balance CAPEX investment with monetization opportunity? To what extent can a small investment hold positions open for the future, before fully committing?

– How can we stimulate a CAPEX to OPEX shift for on-premises data centers of our clients?

– What clients/use cases/domains do we want to invest in our local market?

– Do we work with hyperscalers' technology and therefore comply with their business model, as depicted in their license conditions? Or do we deploy others' technology?

– How can we use the idea of application integration for the hottest applications to differentiate our network quality for users?

– Will such an infrastructure reduce our own CAPEX into network capacity?

– Will we need it and want to own it for future generations of mobile or fixed networks?

■ **Collaboration**

– How do we structure commercial arrangements that incentivize hyperscalers to collaborate in our market and incentivize workload deployment to the edge (i.e., to avoid competition between hyperscalers' central infrastructure or other technologies for enterprise or public customers and the newly erected edge infrastructure)?

– For which segments do we prefer a CAPEX-heavy partnership model over a CAPEX-light one?

■ **Competition**

– How do we keep hyperscalers from eating into our market with cellular networks as a managed service?

– How do we avoid depleting our early investments before recovery?

This is the time to secure options. It is the time to stimulate demand, drive the transformation of digital infrastructures, and forge partnerships – with technology suppliers, systems integrators, hyperscalers, and others. It is clearly not yet the time to invest in significant rollouts, but to gain clarity on strategy, collaboration, and competition.

Notes

Contacts

If you would like more information or to arrange an informal discussion on the issues raised here and how they affect your business, please contact:

Austria
Karim Taga
taga.karim@adlittle.com

Belgium
Gregory Pankert
pankert.gregory@adlittle.com

China
Yusuke Harada
harada.yusuke@adlittle.com

Czech Republic
Jiri Steif
steif.jiri@adlittle.com

France
Julien Duvaud-Schelnast
duvaud-schelnast.julien@adlittle.com

Germany
Michael Opitz
opitz.michael@adlittle.com

India
Ashish Sharma
sharma.ashish@adlittle.com

Italy
Giancarlo Agresti
agresti.giancarlo@adlittle.com

Japan
Shinichi Akayama
akayama.shinichi@adlittle.com

Korea
Kevin Lee
lee.kevin@adlittle.com

Latin America
Guillem Casahuga
casahuga.guillem@adlittle.com

Middle East
Andrea Faggiano
faggiano.andrea@adlittle.com

The Netherlands
Martijn Eikelenboom
eikelenboom.martijn@adlittle.com

Norway
Lars Thurmann-Moe
thurmann-moe.lars@adlittle.com

Poland
Piotr Baranowski
baranowski.piotr@adlittle.com

Russian Federation
Alexey Pankov
pankov.alexey@adlittle.com

Singapore
Tomasz Izydorczyk
Izydorczyk.Tomasz@adlittle.com

Spain
Jorge Abril-Martorell
abril.jorge@adlittle.com

Sweden
Agron Lasku
lasku.agron@adlittle.com

Switzerland
Michael Opitz
opitz.michael@adlittle.com

Turkey
Coskun Baban
baban.coskun@adlittle.com

UK
Nicholas Johnson
johnson.nicholas@adlittle.com

USA
Sean McDevitt
mcdevitt.sean@adlittle.com

Edge computing: Hype or ripe?

How telecom operators will create – or destroy – substantial value via edge computing

Arthur D. Little

Arthur D. Little has been at the forefront of innovation since 1886. We are an acknowledged thought leader in linking strategy, innovation and transformation in technology-intensive and converging industries. We navigate our clients through changing business ecosystems to uncover new growth opportunities. We enable our clients to build innovation capabilities and transform their organizations.

Our consultants have strong practical industry experience combined with excellent knowledge of key trends and dynamics. ADL is present in the most important business centers around the world. We are proud to serve most of the Fortune 1000 companies, in addition to other leading firms and public sector organizations.

For further information please visit **www.adlittle.com** or **www.adl.com**.

www.adl.com/EdgeComputingHypeOrRipe