Executive Briefing Service
Telco Cloud

Supported by

vmware®

STL PARTNERS

Executive Briefing

BUILDING TELCO EDGE INFRASTRUCTURE: MEC, PRIVATE LTE & VRAN

Edge infrastructure deployment is happening – but how big will it really be, and where will it happen? We quantify the five-year outlook for edge deployments across four key domains.

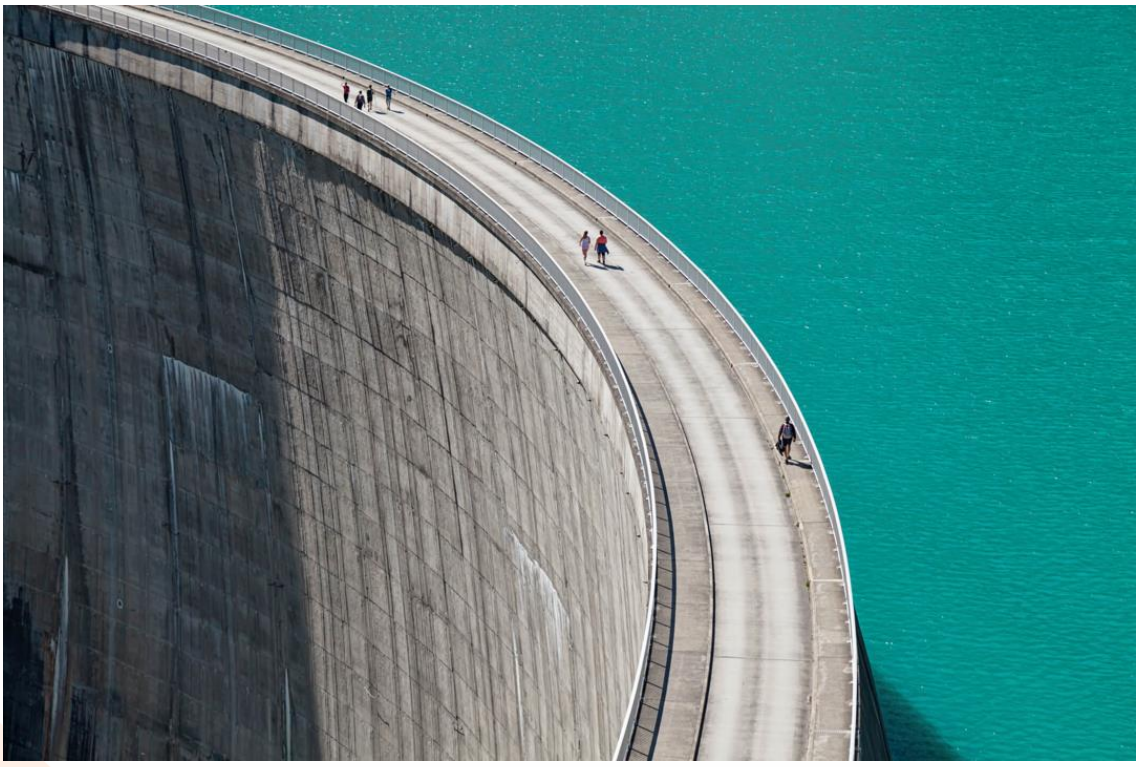

Tilly Gilbert, Senior Consultant | Matt Pooley, Telco Cloud Lead | August 2020

Executive Summary

Telco edge infrastructure build-out is coming

We expect investment in edge infrastructure in earnest to begin in 2020. It is a medium-long term opportunity, so operators should be not be expecting immediate returns, but those who move now should reap rewards later. In this report, we outline 5 key steps for operators looking to invest in edge infrastructure.

The edge computing industry has been catalysed by several key developments:

- High profile partnerships between hyperscale cloud providers (notably Amazon Web Services and Microsoft Azure) and telecoms operators focused on multi-access edge computing.
- Increased traction from operators looking for more open, virtualised and distributed network architecture, with high profile announcements such as Deutsche Telekom's vRAN platform.

Despite this, the domain is not yet mature. The industry lacks proof points which many will be looking for, to prove that the case for edge investment is more than just hype.

STL Partners has built an industry-first model to indicate the size of edge infrastructure deployments between 2020 and 2024 for different types of network operator, breaking this down across a set of use-case domains, edge locations and types of computing equipment. This allows us to:

- Identify where, when and how many edge deployments can be expected
- Highlight triggers, barriers and drivers that will influence deployment
- Adopt a holistic view across different domains

Figure 1: We predict a Tier-1 operator will have more than 8,000 edge servers by 2025, supporting edge workloads across multiple domains

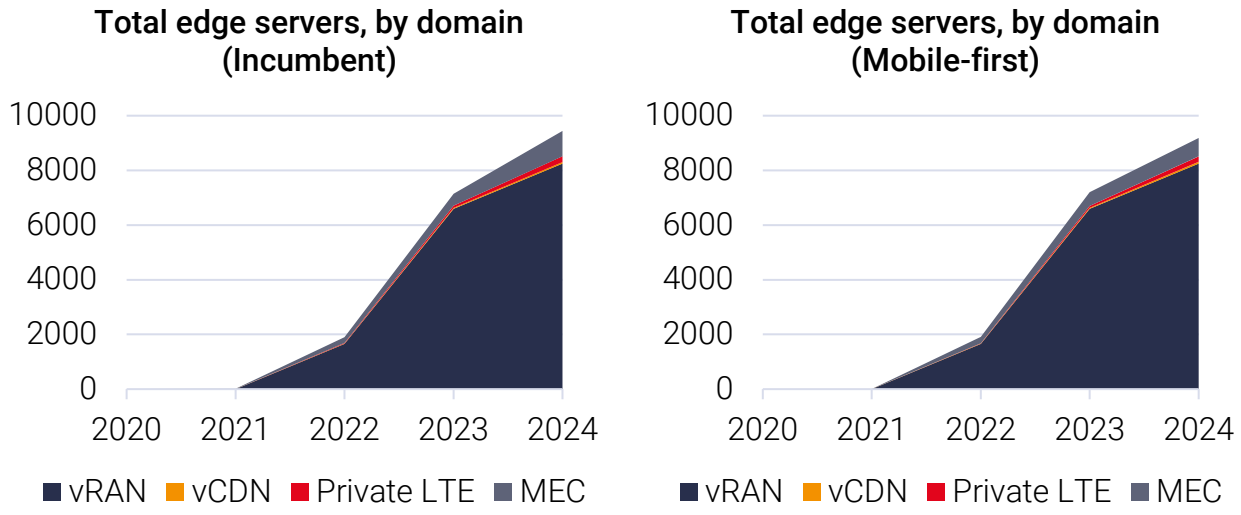

Source: STL Partners

Recommendation: learn operationally about edge while deployments are not too "edgy"

The model highlighted two clear messages for operators wanting to understand how edge infrastructure will develop over the next 5 years.

1. **The short-term edge opportunity does not require highly-distributed deployments – which is an opportunity to learn while keeping a lid on investments.** The most significant and earliest uses for edge sites will be to support vRAN, which can be done in sites close to the core of the network. This is an opportunity for operators to develop skills and expertise while the edge opportunity is both: primarily about delivering internal, network-centric use cases, and is not massively distributed. Operators will be able to identify benefits even with 2-5 more sites than you have currently.

2. **There is no short-term route to revenue – but now is the time to start exploring**. In our timeframe, MEC will not be a major revenue generator for most operators. Instead, it will be deployed in a tactical and opportunistic fashion. However, looking further out, 5-10 years and beyond, it is likely to grow into a significant opportunity. The crossover between MEC and a domain like vRAN should not be ignored: operators should organise to share learnings and ensure their physical edge sites are suitable to support all significant domains.

Based on these core takeaways, below are 5 practical steps operators should take today:

1. Underwrite / piggy back on vRAN to invest in edge sites that have foundational units of servers (likely to be able to serve demand for the next 2-3 years). These sites should have room to add more servers as demand increases and balance the network and third party application

requirements. This might mean, for example, keeping an eye on delivering end-to-end network latency below 30ms for as much of the population as possible.

2. Figure out security (physical and network) to enable third party applications to be able to run alongside your network functions. This means organising your technical spaces not just with your own network team in mind. Operators will want to be able to embrace wholesale / co-location opportunities where they occur and must plan accordingly.

3. Invest in delivering large, bespoke on-premise projects for trophy customers. This will be key to building the skills and partnerships needed to enable you to productise enterprise edge offerings in the long term.

4. Work with the hyperscalers, but don't put all your eggs in one basket. In the next 5 years, there will be MEC in telco's networks where the infrastructure has been built out by hyperscalers. However, this in and of itself is unlikely to be lucrative for the telecoms operator. With a similar model today, regional data centre operators barely break even from their business hosting the hyperscalers. Operators need to take an active role in exploring MEC use cases, developing the right relationships and engaging with end-customers to enable them to enable co-opetition.

5. Evaluate your role above and beyond providing infrastructure. Operators should consider innovative use cases and new business models on a domain by domain basis to ensure, where appropriate, they move up the value chain.

STL Partners – what's next?

STL Partners has several services which are designed to help operators evaluating "what's next?". For those wanting to understand MEC, the opportunity, use cases, ecosystem and potential partners they should be exploring, our edge ecosystem service and use case directory will be of interest. Those evaluating the business case and drivers for virtualisation should see our telco cloud hub and dedicated research stream.

Table of Contents

Table of Figures

Preface

The document has been prepared by independent research firm STL Partners, and commissioned by VMware. It is part of STL Partners' continuous research programme into the future telecoms operator and how to get there.

STL Partners maintains strict editorial independence. Mentions or allusions to companies or products in this document are intended as illustrations of market evolution and are not intended as endorsements or product/service recommendations.

Reality check: edge computing is not yet mature, and much is still to be decided

Edge computing[1] is still a maturing domain. STL Partners has written extensively on the topic over the last 4 years.[2] Within that timeframe, we have seen significant change in terminology, attitudes and approaches from telecoms and adjacent industries to the topic area.

Within the past twelve months, we've seen high profile partnerships between hyperscale cloud providers (Amazon Web Services, Microsoft and Google) and telecoms operators that are likely to catalyse the industry and accelerate route to market. We've also seen early movers within the industry (such as SK Telecom) developing MEC platforms[3] to enable access to their edge infrastructure.

In the course of this report, we will highlight which domains will drive early adoption for edge, and the potential roll out we could see over the next 5 years if operators move to capitalise on the opportunity. However, to start, it is important to evaluate the situation today.

Commercial deployments of edge computing are rare, and most operators are still in the exploration phase. For many, they have not and will not commit to the roll out of edge infrastructure until they have seen evidence from early movers that it is a genuine opportunity for the industry. For even more, the idea of additional capex investment on edge infrastructure, on top of their 5G rollout plans, is a difficult commitment to make.

This paper primarily draws on discussions with operators and others within the edge ecosystem conducted between February and March 2020. We interviewed a range of operators, and a range of job roles within them, to gain a snapshot of the existing attitudes and ambitions within the industry.

Below we highlight three situations that reflect the reality of edge computing for most operators today.

Reality #1: Organisationally, operators are still divided

For many operators, there is **no one clear team that has ownership over the edge opportunity and how they will address it**. Responsibility may be spread across several teams, including but not limited to:

- Innovation / R&D units focused on understanding future technology and its role within the operator.
- 5G teams who have been asked to explore the role of edge computing in conjunction with next-generation networks. The flexible and virtualised properties of 5G provide opportunity for new

[1] For definitions, see Where is "the edge"?
[2] See our edge hub.
[3] See SK Telecom: The road to the world's first 5G MEC platform

ways of delivering network functions and third-party applications that make edge an obvious synergistic technology.

- Network / technology units who may not have a direct 5G focus, but will tie edge computing into their remit. This means they can evaluate the use of networking technical spaces for edge infrastructure alongside other core technical drivers, like the separation of control and user plane data. Even within this, the mobile core, RAN, fixed core and fixed access teams themselves are often also significantly siloed in approaches and responsibilities.

- IoT teams who want to understand how edge computing could be used to optimise and enable at scale existing offerings and unlock new opportunities.

- Enterprise teams who see edge computing as a potential new revenue opportunity for their customers, often by combining connectivity with a MEC-enabled solution.

Depending on the role of an individual, attitudes and understanding of edge computing are likely to differ. Technical teams, for example, are likely to evaluate virtualised edge workloads along with their broader virtualisation efforts. Existing efforts, such as distributing EPC functions, will be the backdrop and the choice of how and where to develop edge sites will depend on, for example, the number of macro base stations it can serve.

In comparison, those teams with more commercial, customer-facing responsibilities typically view edge as an opportunity to move higher up the value chain, and offer more innovative solutions to customers. To illustrate, one individual evaluating edge computing within a Western European single country operator, felt that edge is one of the few areas where telcos should be able to differentiate their services meaningfully from OTT players.

Reality #2: The edge ecosystem is evolving fast

A common message from operators is that they have a high-level plan for edge computing, but are still figuring out who they would need to work with in order to implement it. As a new area, operators need to make decisions about what roles they will try to fill themselves, and where they will work with others to deliver a solution.

One North American operator described the challenge of “articulating to a new group of customers, application developers, why edge infrastructure is something they should consider paying a premium for”. Operators are trying to understand the new vendors or partners they may need to work with to address these new customers, be they application developers or industry verticals that telcos have meaningfully engaged with before. Even when operators themselves are the customer, like in the case of vRAN, there is still a challenge in deciding who to work with, with new challenger players like Mavenir and Rakuten establishing themselves as viable alternatives to the traditional vendors.

Reality #3: Operators are trying to predict, respond to and figure out what the "new normal" will be post COVID-19

Figure 2: There is uncertainty as to the impact of COVID-19 on operator's edge computing and telco cloud initiatives

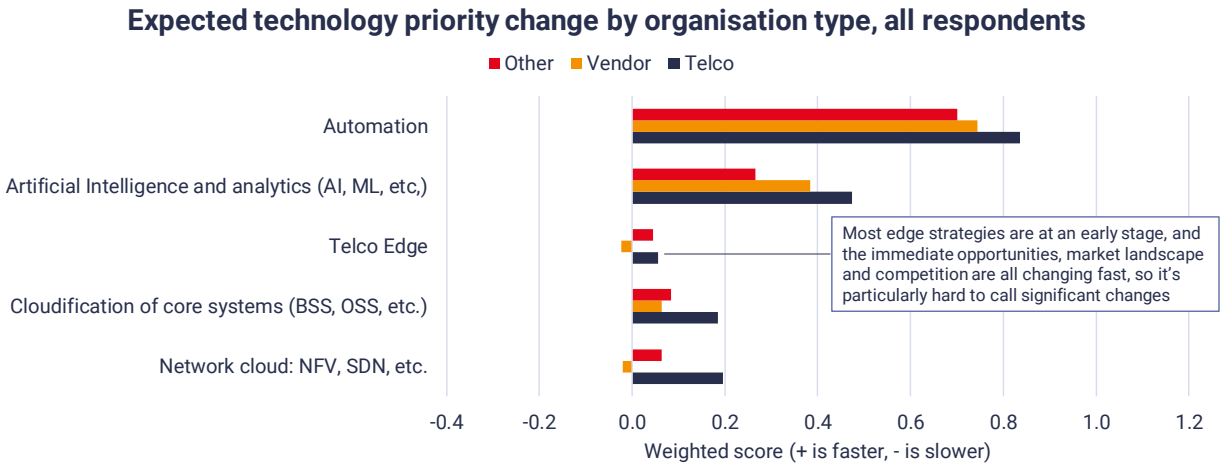

To get the weighted score, we weighted respondents choosing 'major uplift' = 2, 'minor uplift' = 1, 'minor cut' = -1, 'major cut' = -2.

These results include 202 responses people's input up to 8am UK time, 8th May 2020.

Source: STL Partners industry survey, STL Partners analysis

Figure 2 outlines the results from an industry survey led by STL Partners, that tracked the likely impact of COVID-19 on areas of technical innovation within operators. As operators navigate a time of economic uncertainty, we expect budgets in some areas to be squeezed. While more mature technical domains, like AI and automation, saw consensus across the board that investment would remain strong or even increase in light of recent events, areas like network cloud and telco edge were less strongly skewed. This creates a challenge for teams championing these areas, who may have to demonstrate faster or more concrete return on investment than is possible for a relatively greenfield opportunity and are likely to be battling these more established innovation areas for budget when investment cycles are reviewed.

One further reality in the edge ecosystem today, is there is a lack of reasoned quantification on how and why the industry will adopt it. This includes bridging the gap between the use of infrastructure for telco cloud (i.e. operators' own network functions) and MEC (i.e. third party IT applications). There is also a lack of clarity on exactly what is meant by "edge computing" and how it breaks down into physical real-world deployments. This paper will look to answer these questions.

Edge computing: key terms and definitions

Where is "the edge"?

There is no one clear definition of edge computing. Depending on the world you are coming from (Telco? Application developer? Data centre operator? Cloud provider? etc.), you are likely to define it differently. In practice, we know that even within these organisations there are differences between technical and commercial teams around the concept and terminology used to describe "the edge".

For the purposes on this paper, we will be discussing edge computing primarily from the perspective of a telecoms operator. As such, we'll be focusing on edge infrastructure that will be rolled out within their network infrastructure or that they will play a role in connecting. This may equate to adding additional servers into an existing technical space (such as a Central Office), or it may mean investing in new microdata centres. The servers may be bought, installed and managed by the telco themselves, or this could be done by a third party, but in all cases the real estate (e.g. the physical location as well as power and cooling) is owned either by the telecoms operator, or by the enterprise who is buying an edge-enabled solution.

Operators have choice and a range of options for where and how they might develop edge computing sites. Figure 3 starts to map some of the potential physical locations for an edge site.

Figure 3: There is a spectrum of edge infrastructure in which telcos may invest

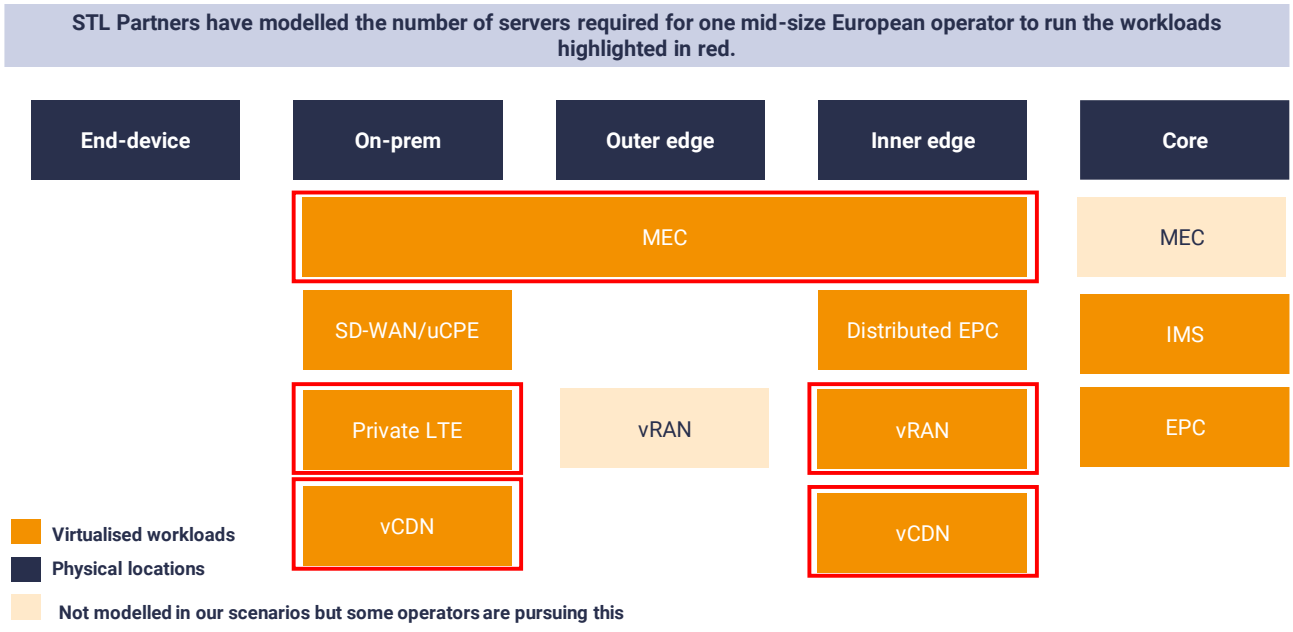

Source: STL Partners

- On the far-left hand side is the most disparate type of edge, at the level of an individual device such as compute on a smartphone. This is the realm of device manufacturers, rather than telecoms operators, and so won't be explored further in this paper.

- Moving to the right, however, is the opportunity for telcos to install and manage edge infrastructure on the premises of an enterprise customer. Telcos already manage infrastructure for customers today, for example to provide enterprise networking services like SD-WAN. In the future, we expect telcos to widen the use cases they will provide on-premises to customers (see **What applications & use cases will run at edge sites**).

- Moving again to the right, and to physical locations within the operator's network itself, we identify outer edge sites as those that will reside within an operator's access network, such as at base stations.

- Inner edge sites will likely be larger technical spaces that may, for example, be refurbished Central Offices. **We will refer to the inner and outer edge together as "network edge" throughout this report.**

- On the far right of Figure 3, we've highlighted operator's core network and some of the types of virtualised network workloads that are already running here. While our focus will be on edge sites (i.e. on-prem, outer edge and inner edge), it's important to highlight that operators may look to distribute their core functions too and so we expect some of these workloads to run alongside our edge use cases in the future.

Each operator managed edge site is likely to be built and designed differently, with different key use cases in mind. Figure 4 summarises how these spaces might differ in terms of:

- The data centre rating it is likely to have i.e. how much power is needed by the data centre and how much capacity it can support in terms of number of servers

- Number of base stations served by each edge site

- The indicative round trip latency the space could achieve

- An approximation of the number of sites within a small but populous country like the UK

Figure 4: There are three key types of telco edge site, ranging in size, location & type

Which location?	Type of technical space	Size / Rating	# macro base stations per edge site	Average distance to user	Max round-trip (4G)	Potential number per country (UK SP)
On-prem	Basic (Tier 0-2)	1-4 u / 3KW	1-5	<1km	<5ms	20,000+
Outer edge+	Micro (Tier 1-3)	6-12 u / 10KW	4-40	<5km urban, <20km suburban/rural	<10ms	1,000+
Inner edge*	Mini (Tier 3-4)	10+ racks / 0.5-2 MW	100-200	<200km	<20ms	10

+STL Partners envisages virtualised RAN workloads, at least initially, will run at inner edge sites.
***STL Partners has not modelled the number of servers needed to support core NFVI functions e.g. distributed EPC even though some of these may exist in the same physical location at the inner edge.**

Source: STL Partners

What applications & use cases will run at edge sites?

As Figure 3 depicts, there are four main types of workloads that we predict operators will run on edge sites by 2025.

1. **Virtual radio access networks (vRAN):** virtualisation of portions of the operator's own radio access network (primarily baseband functions) in order to run on common computing hardware.

2. **Private LTE networks:** using virtualised deployment of 3GPP-standard cellular network technology (4G/5G) to offer closed user group connectivity services to enterprise customers.

3. **Virtual content delivery networks (vCDN):** virtualisation of CDN applications to run CDN workloads on different types of hardware, adding flexibility by decoupling software and hardware layers.

4. **Multi-access edge computing (MEC):** providing a software environment on the edge to write and run edge-enabled software applications for customers.[4]

vRAN: Opex reduction and enabling next-generation network services

Rakuten Mobile,[5] which is building a national, end-to-end virtualised network (both core and RAN) in Japan, has come to be seen as a proof point for the benefits of virtualisation. Rakuten has reportedly

[4] Within MEC, there are a large number of potential use cases for both consumer and enterprise. Notable buckets of use cases, the drivers for moving them to the edge and the potential benefit have been explored in a recent report Telco edge computing: What's the operator strategy?
[5] For more information on Rakuten's RAN strategy, see Open RAN: What should telcos do?

claimed that its network will have 40% lower opex than legacy networks, although like-for-like comparison is problematic.

Reducing operating costs is the core business driver for most operators. Since renewal cycles for RAN equipment are typically relatively short, we expect to see early mover operators at least, making this commitment for their existing 4G RAN infrastructure, as well as 5G SA in the future. For planned renewals and expansions to the RAN, virtualised and more centralised infrastructure will require less capex, be cheaper to install, require cheaper civil works and cheaper power upgrades. This, in and of itself, drives the business case. To illustrate, 4G vRAN will likely be rolled out in circumstances such as:

1. In geographically-dispersed or currently under-served areas. In these situations, vRAN offers substantial potential cost savings and operational efficiencies in that it enables a dispersed estate of cell sites to be operated, managed and upgraded remotely from central locations (cloud data centres hosting the CUs), without having to dispatch engineers to each site to carry out work on physical BBUs.

2. In urban areas to support densification of 4G and 5G and to expand coverage in suburban and rural areas. It can be a more cost-effective, scalable solution vRAN technologies, for the purpose of supporting new locations and types of radio (e.g. small or micro-cells in dense urban or indoor environments).

Private LTE networks (i.e. those specific to individual enterprises, ecosystems, verticals, government bodies, etc.) and neutral host networks represent in our view the biggest short to medium term opportunity for telcos to create new business models from vRAN, as opposed to enabling them to do their existing core connectivity job more effectively – a category the above examples fall into.

Private LTE: Virtualisation for increased flexibility

Private LTE networks represent a relatively new network offering that operators can take to their enterprise customers. In many cases today, these deployments will not be virtualised, and will be provided by a proprietary equipment vendor, such as Ericsson or Nokia. By virtualising private LTE network functions, and running them on COTS hardware, operators will be able to deliver private LTE solutions with increased flexibility, scalability and customisation, and decreased cost.

As operators look to 5G, having experience managed virtualised private infrastructure will act as a glide-path to new opportunities, like providing private connectivity via a dedicated 5G network slice. One operator exemplified this end-goal, highlighting that operators do not want to have to handle a large amount of highly distributed network equipment, at separate sites for every enterprise. This will make monitoring, maintaining and upgrading these networks significantly more expensive – this is why they will look to slicing in the future, to meet enterprises' requirements with less distributed equipment.

vCDN: Monetisation from efficient networks and the potential of an edge PaaS play

Currently, operators are more likely to provide the physical space for a (v)CDN solution provider to put their equipment (co-location model), rather than to partner to provide a solution themselves. However, this could change as the benefits of new business models are realised. Operators are already virtualising their infrastructure, and would therefore be well placed to partner in a vCDN solution.

Figure 5: Most operators see vCDN as a means of increasing network efficiency, but there are revenue opportunities too

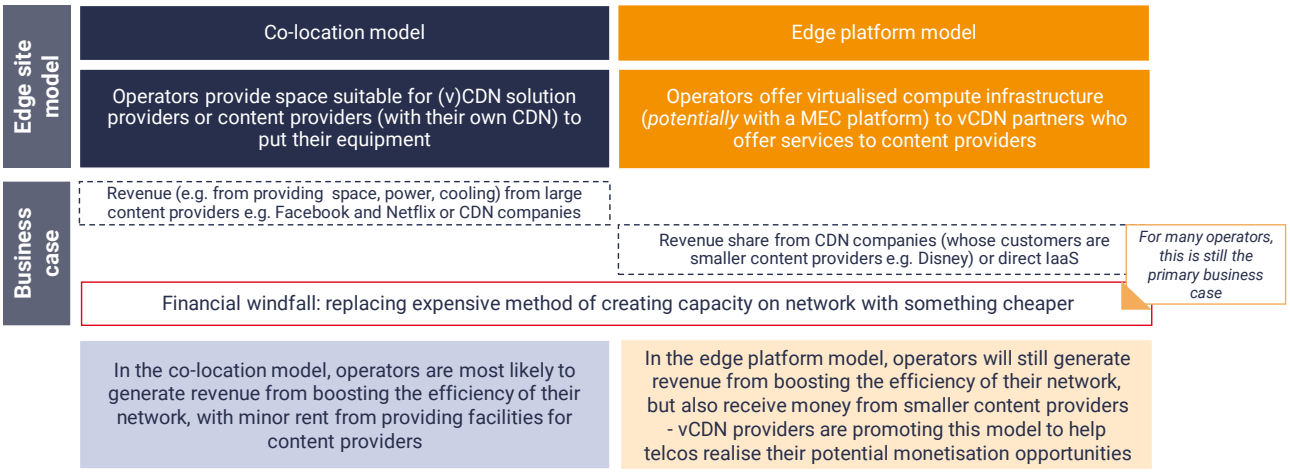

MEC: A new revenue opportunity to deliver new use cases

The opportunity for providing MEC lies in the sweet spot of applications that require capabilities that are both local compute-like and cloud-like. The core value proposition for edge is highlighted in Figure 6.

Figure 6: Use cases that fit in the edge sweet spot will have both local-compute like and cloud-like requirements

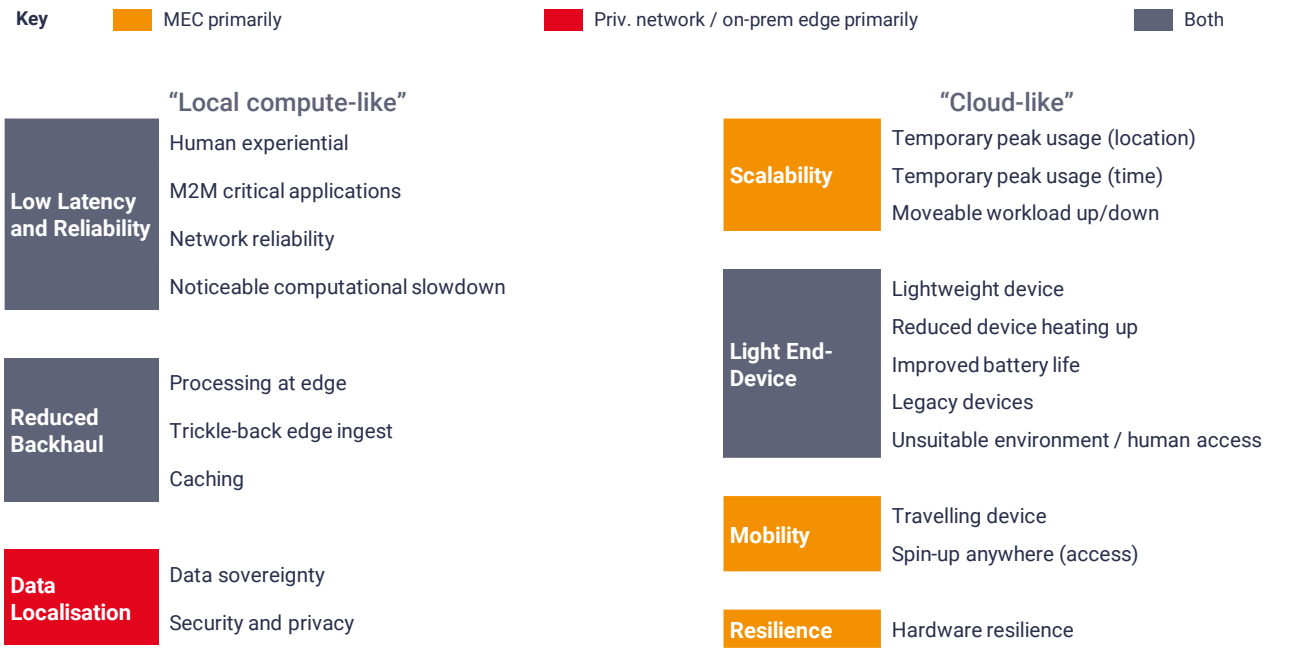

Edge infrastructure will support both internal and external use cases

These can be split into internal use cases, where the telecoms operator is essentially the end-customer or external use cases where the telecoms operator serves a third party like an enterprise or a media company (see Figure 7).

Internal use cases tend to focus on efficiency and cost saving for the operator and enables them to do their core connectivity role better. External use cases are opportunities for operators to generate new revenues via external customers. vCDN sits across these two categories. Commercialised CDN could see operators charging content provider customers, in a way that Akamai does today. Non-commercial CDN can be used by operators as an internal use case too, as caching content can reduce backhaul traffic and costs and improve customer experience.

Figure 7: Use cases leveraging telecoms edge infrastructure can be either internal or external facing

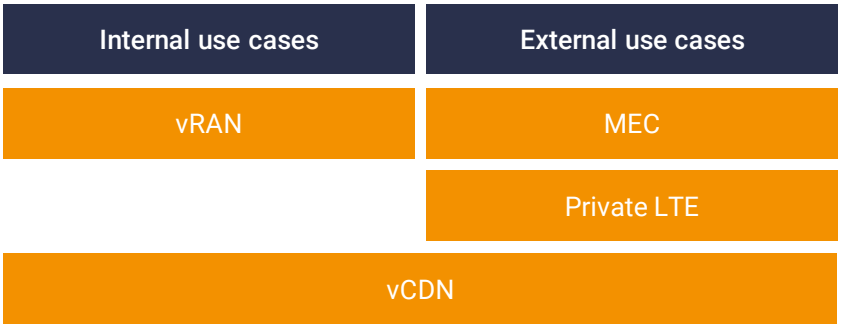

Source: STL Partners

What is inside a telco edge site?

There is no one clear picture of what a telecoms edge site will look like. As has already been highlighted, this will be different depending both on the type of edge (on-prem, outer, inner), but will also be specific to individual operators and their strategies. We discussed with 7 operators their plans over the next 5 years. Below are some of the key questions being weighed up by the industry at the moment, and the answers / attitudes that will have implications on what is inside telecoms edge sites:

Will I use the same physical location to run network function workloads (e.g. for vRAN) and enterprise IT applications?

For most operators, there are intuitive benefits to leveraging one site for both network functions and third party applications. One Western European operator described the "economy of scale" problem where operators could run the risk of having "three times the infrastructure they need" if they do not plan to support both within one physical space.

As most operators have limited capex to invest in building or refurbishing technical spaces, combining demand for such spaces can help to justify the business case, particularly if these spaces can evolve over time as demand for the edge changes and matures. Having the two collocated also makes it easier to securely and reliably connect to the edge, without the need to build out a separate connection. This can avoid added cost, delays and potential points of failure in the future.

However, there are some concerns. Mostly, these focus on how to ensure that the performance of the network is never impacted by third party applications. For our purposes, **we have assumed that while network functions and enterprise applications will likely run in the same physical locations, they are unlikely to share (compute and storage) infrastructure**. This means racks (which may be in separate rooms/cages) will either support network workloads or IT workloads for enterprises, rather than both at once. In some, but not all, cases, network functions and enterprise applications will share common architectures, toolsets and standard components.

How much non-virtualised infrastructure will there be at edge sites?

Edge sites, even at the inner edge, will be significantly more power and space constrained in comparison to the typical data centres operators use today e.g. for their core network functions. Non-virtualised infrastructure remains less flexible and less elastic than hardware which can run a range of virtualised workloads on top of it. All operators we spoke to, because of this, and because of the relatively greenfield build-out opportunity, will look to have almost exclusively virtualised infrastructure at their edge sites.

How much infrastructure installed and managed by third parties will there be within telco edge sites?

Operators are not the only party who will want to run their workloads at edge sites. Others, including but not limited to, CDN providers and hyperscale cloud providers, may also wish to put their workloads closer to their end customers. Whether and to what extent operators open up their technical spaces to third parties will have an impact on what these spaces look like, not in the least because of the potential security implications of third parties needing access to sites inside the operator's network.

With on-premise edge deployments this is a little different. Here, the enterprise, rather than the operator, has the final say on who has access to the physical hardware running the edge workloads. We predict in the next 5 years that there will be an almost 50:50 split between enterprises who look to their cloud provider (e.g. Amazon, Microsoft, Google etc.) to install and manage their edge infrastructure versus their connectivity provider or telecoms operator. In most cases, however, enterprise sites will be 100% one or the other, in comparison to network edge sites which will likely see racks of servers for different users within one physical site.

How much spare capacity will edge sites have (at least to begin with)?

For on-prem edge, this is a relatively simple equation, at least in the short term. Since most deployments are likely to be designed and implemented bespoke for a specific customer, the entity supplying the edge servers will need to gain a sense of what applications the end customer is expecting to deploy on those servers.

In fact, in most cases, we don't expect to see edge hardware being deployed without being part of a broader application or solution sale, since this is ultimately the service that the enterprise will be willing to pay for. In this case, there is not likely to be spare servers[6] and if additional compute or storage is required down the line, it will be up to the enterprise to ensure there is the space and facilities in order to do this.

For network edge, we predict that operators will look to invest in a foundational unit (typically at least 2 racks) first, and then may add racks as demand increases. This will often mean operators having a few initial use cases for their infrastructure lined up, some of which may be internal to the operator (such as supporting vRAN workloads) and some of which may be external (such as a few application developers who have committed to running their application at the edge). Once the space is established and delivering the performance required, they will look to broaden its usefulness, either by moving more workloads to the edge themselves, or opening the space up to a broader range of third parties.

[6] We have heard some demand for redundancy ensure that a reliability of service can be maintained, but when these servers are not in a true datacentre environment anyway, spare capacity is unlikely to the critical factor in delivering a reliable service.

How edge will play out: 5-year evolution

Modelling exercise: converting hype into numbers

Much research has been published about the edge computing *opportunity* – how and why it *could* impact on the telecoms business – but very little consensus has been reached what *will* happen in the coming years. How many edge deployments are likely to take place? Where will infrastructure be located? And how much computing resource will be required?

We undertook to build an industry-first model, with the objective to indicate the size of the edge infrastructure deployments for between 2020 and 2024 for different types of network operator, breaking this down across a set of use-case domains, edge locations and types of computing equipment. This allows us to:

- Identify where, when and how many edge deployments can be expected
- Highlight triggers, barriers and drivers that will influence deployment
- Adopt a holistic view across different domains

Our model is *not* a prediction or forecast of the amount of money telcos stand to make from edge computing, and we have not attempted to extrapolate the global size of the edge market. Instead, we use input data and assumption sets sourced from discussion with a variety of real telecoms operators and their partners across the world about *what* they are deploying, and *why*.

This means that we have been able to build an informed and grounded view of how edge will unfold in certain key markets in the next few years, without the risk of making inaccurate generalisations about the market as a whole.

Defining operator scenarios

Our discussions have shown that, while edge is on the radar for telecoms operators of all shapes and sizes around the world, currently most well-defined and resourced edge programmes are within larger, operators often referred to as "Tier 1". These include many household-name global operator groups, as well as certain single-country giants.

For the purposes of our exercise, we focussed on two hypothetical operators that would fit the Tier 1 moniker. Both are based in a hypothetical medium-sized European country, with around twenty million subscribers. These are:

1. **"Incumbent":** A former state-owned incumbent, offering converged connectivity, with significant fixed and mobile networking infrastructure.
2. **"Mobile-first:"** A non-incumbent mobile network operator (MNO), offering primarily mobile connectivity, with mainly mobile network infrastructure, and limited fixed network infrastructure of its own.

Defining key use case domains

To understand what edge infrastructure our operators will deploy and where it will be located, it is first necessary to understand what workloads the infrastructure is expected to *support*.

We outlined some of the likely workloads in Figure 3 – to reiterate briefly:

1. **Virtual radio access networks (vRAN):** virtualisation of portions of the operator's own radio access network (primarily baseband functions) in order to run on common computing hardware

2. **Virtual content delivery networks (vCDN):** virtualisation of CDN applications to run CDN workloads on different types of hardware, adding flexibility by decoupling software and hardware layers

3. **Private LTE networks:** using virtualised deployment of 3GPP-standard cellular network technology (mostly 4G/5G) to offer closed user group connectivity services to enterprise customers

4. **Multi-access edge computing (MEC):** providing a software environment on the edge to write and run edge-enabled software applications for customers

Defining edge locations

For the sake of simplicity, our model divides deployments across the locations outlined in Figure 3, with the exception of end-devices (as these sit outside the telco's network) and the core network (as this, by definition, is not the edge). This leaves:

1. **On-premise:** broadly referring to the final edge site in the access network (e.g. local site, in building, on a customer's premise)

2. **Outer edge:** broadly referring to an edge site within, or very close to, the radio access network (e.g. within a radio base station or similar or lower-tier backhaul aggregation points)

3. **Inner edge:** broadly referring to higher tier network aggregation points or regional/metro network nodes on a national backbone

Full definitions of the types and numbers of the technical spaces in these locations are outlined in Figure 4.

Defining types of edge computing equipment

Our final consideration is the computing hardware that will underpin edge applications at each respective location. Our model counts this hardware by quantity of servers required, divided across three different types of that we have seen deployed by Tier 1 operators. We use "server type" as shorthand for these:

1. **Common-off-the-shelf (COTS):** Generic, non-telco specific x86 computing hardware (server and storage) and could include standard graphical processing units (GPUs). Likely to come from a general computing vendor (e.g. HPE, Dell).

2. **Proprietary:** Hardware specifically marketed to support telco workloads, sold as a generic x86 device, but with adaptations designed to enhance compatibility with the vendor's other kit. Likely to come from a telco-specific vendor (e.g. Huawei, Ericsson, Juniper, Cisco) or a hyperscaler (e.g. Amazon Web Services).

3. **Specialised:** Bespoke, application-specific hardware designed to support specific use cases (e.g. a video processing application), based on programmable silicon or similar (e.g. FPGAs, smart NICs, app-optimised GPUs).

Our findings: edge deployments won't be very "edgy" in 2024

Our modelling shows that, for both of our hypothetical Tier 1 operators, the vast majority of edge servers in 2024 – more than 80% - will be at *inner edge* locations: that is to say, a relatively small number of large, fairly centralised technical spaces, closer to the core than the edge of operators' networks. In our examples, this meant around ten locations spread around a mid-sized European country to serve the same footprint as the operators' existing networks.

This stands in stark contrast to predictions that edge sites will number in the hundreds or thousands, spread across the *outer edge* and *on-premise edge*. So, at least in the short term, the vast majority of edge compute will be relatively centralised – not in the core of operators' networks, but close to it.

Figure 8: Most edge servers, regardless of operator type, will be on the *inner edge*

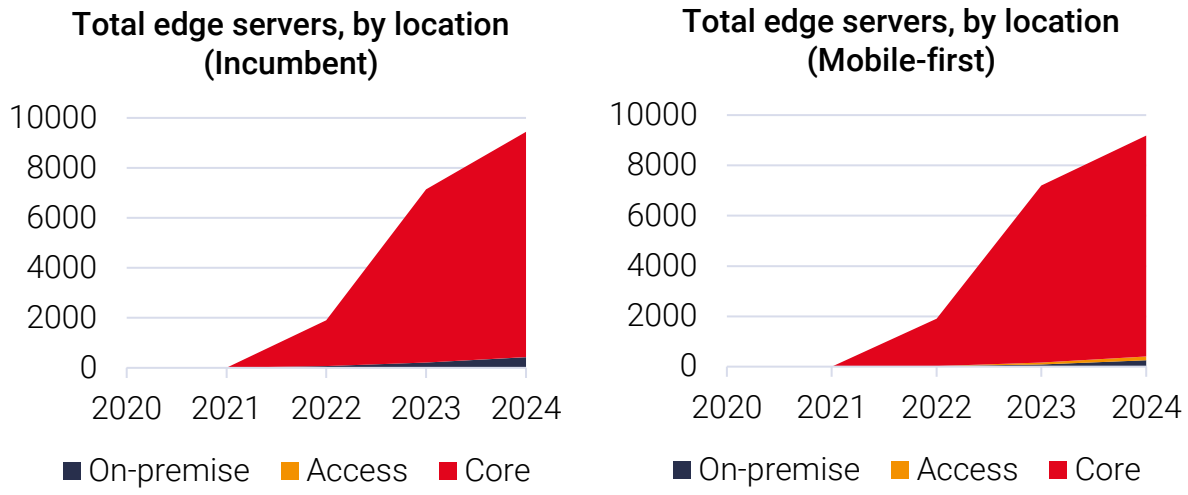

Source: STL Partners

This is not to say that on-premise edge computing will not be significant. As private LTE network functions are virtualised they will likely run on servers on-premises, alongside MEC functions and, in certain cases, vCDN.[7] However, since each on-premise edge site will likely require only 1-2 severs,

[7] In most cases, vCDN will reside at the network edge, but in certain cases it may also reside on-premises, for example at a sports venue or university campus.

whereas inner edge sites could be made up of around 100 servers each, in pure amount of compute and storage required, inner edge will drive these numbers in the short-term.

Short-term adoption of vRAN is the driving factor

The primary driver for edge compute deployment in the *inner edge* is operators' adoption of virtual RAN technology. Our model shows RAN workloads will be deployed in this location over any other, and the majority of the inner edge deployments shown in Figure 8 can be attributed to this:

Figure 9: vRAN is an inner edge (fairly centralised) play

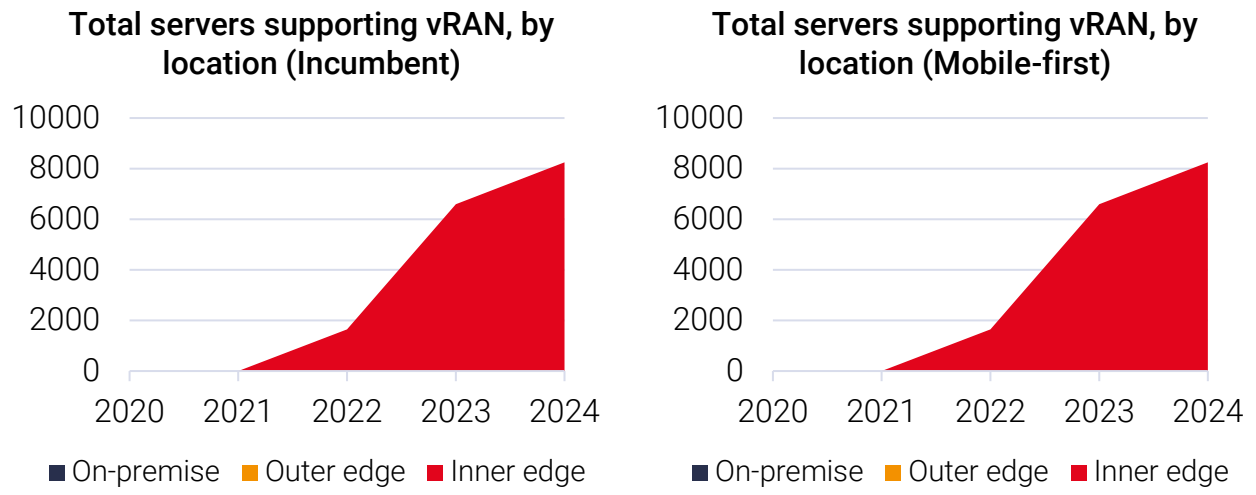

Source: STL Partners

To understand this, it is helpful to outline the wider trend in radio access networks, as described in our recent report[8]:

- **Virtual RAN** itself involves virtualising the radio baseband unit so that it is run as software on generic hardware platforms. This enables the baseband software and hardware, and even different components of them, to be supplied by different vendors…
- … but most operators are pursuing it alongside a complementary technology: **Centralised RAN (C-RAN)**, also known as cloud RAN. C-RAN involves distributing and centralising the baseband functionality across different telco edge, aggregation and core locations, and in the telco cloud, so that baseband processing for multiple sites can be carried out in different locations, closer to or further from the core of the network, depending on requirements.

The operators we modelled are following a typical deployment pathway for mobile operators in Europe, in which RAN functions will be gradually centralised over the coming years. To begin with, multiple physical baseband units will be brought together in locations close to the *outer edge*. Over time, these

[8] See Open RAN: What should telcos do?

locations will be further centralised – and eventually, the physical functions will be replaced by virtual functions running in a small number of regional datacentres.

What this means is that radio functions are being disaggregated from radio antennae, brought together in centralised locations, and becoming software. Or, in other words, vRAN is driving centralisation of workloads that were traditionally distributed at the extreme edge of telco networks.

vRAN is unique among the use-cases we have modelled because it is an evolution of a core service. Private LTE, vCDN and MEC all involve connectivity, but represent new service offerings, for which the operator must attract new, primarily enterprise customers (or convert existing customers), with the hope of extracting new revenues. Building a customer base will take time and investment. vRAN, however, is more of an efficiency upgrade to mobile networks which already exist and support millions of subscribers. As soon as it goes live, it will have users.

For this reason, the amount of infrastructure required to support vRAN within our five-year timeframe – for both our *converged* and *mobile-first* operators – is several orders of magnitude larger than the other use-cases. Approximately 80% of edge servers in use by operators in 2024 will be supporting vRAN workloads:

Figure 10: vRAN will dominate edge infrastructure deployments...

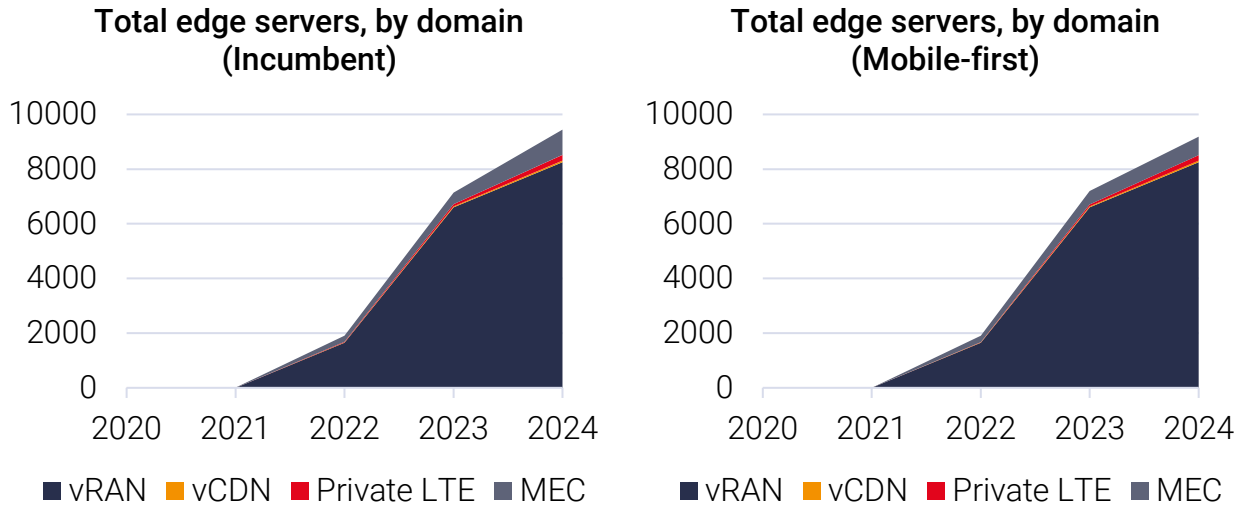

Source: STL Partners

The curve for adoption of vRAN shown in these charts is fairly steep – suggesting that operators will move fast to adopt it over the coming five years. In fact, our model assumes that by 2024, both of our operators will have virtualised 100% of their existing LTE (4G) baseband functions – corresponding to around 16,500 mobile base stations. This may seem ambitious, but it is very much in line with what the roadmap outlined by Western European operators. The reason we believe this roadmap to be realistic is two-fold:

1. vRAN is already being deployed in test environments – and based on approaches we have seen for other VNF deployments, we believe operators will scale as soon as it is proven

2. The handful of physical sites required to host the centralised vRAN workloads actually already exist (or at least are in the process of being built). These same sites are being used to support operators' ongoing programme to distribute user plane functions that currently sit in the core network closer to end-user – so the case for investment and process of building new infrastructure has already been made.[9]

It should be noted that our estimates do *not* include vRAN workloads related to 5G New Radio. Our research is clear that, for most operators, vRAN and C-RAN will be deployed solely to support efficient 4G networks, at least in the short term[10]. Over the next five years and beyond, 5G vRAN will also be adopted – in which case vRAN our model may understate the vRAN opportunity.

New revenues from MEC remain a longer-term opportunity

The other use-cases that we covered *will* be deployed in the coming five years, but on a much smaller scale. If we exclude vRAN figures, we can see a slow uptake, mostly driven by deployment of MEC servers:

Figure 11: MEC is the primary short-term revenue-generating edge opportunity

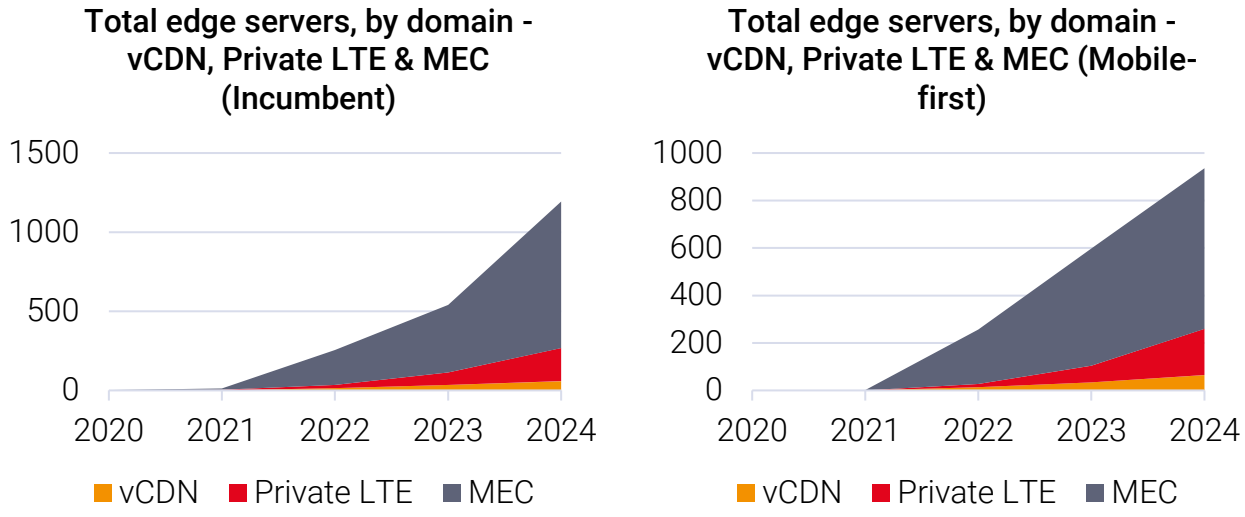

Source: STL Partners

Most MEC workloads are actually hosted in what we define as *inner edge* locations – not on enterprise premises or in base stations. This reflects a trend among European operators to start with fairly centralised edge deployments – what some are calling "edge hotels" – and only push deeper into the

[9] While the same may be true for the other domains, changes to the technical spaces (e.g. to allow third party access securely) are likely to be need to be more extensive.
[10] A full discussion of the rationale for this can be found in Open RAN: What should telcos do?

network when the business opportunity becomes clearer: These locations represent a fairly big sunk cost for the operator, and they need to maximise their use:

Figure 12: Incumbent MEC dominance driven by fibre to enterprise premises

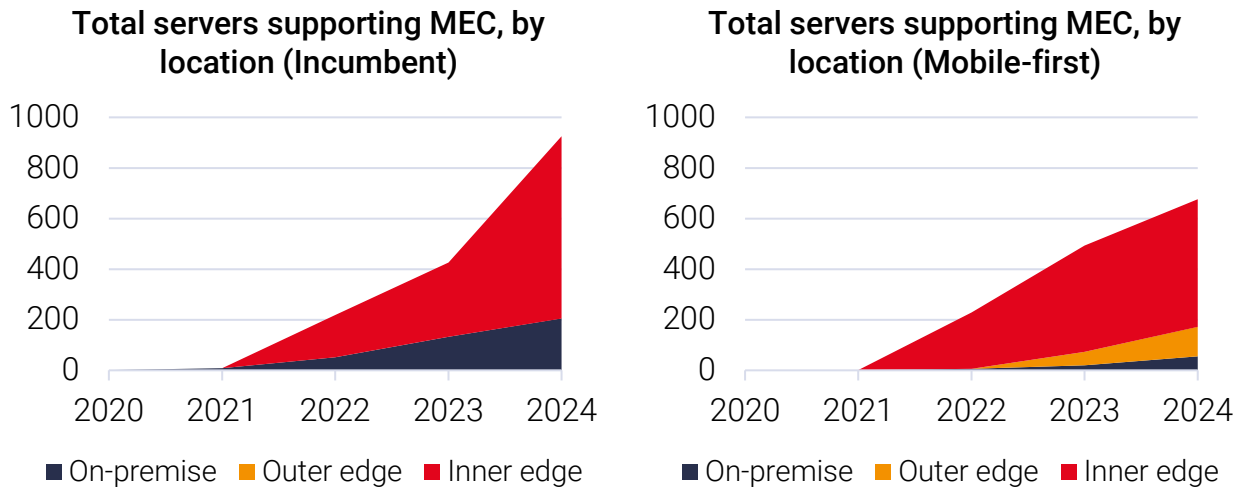

Source: STL Partners

Both Figure 11 and Figure 12 show that overall MEC deployments will be slightly larger for our *incumbent* European operator. Much of this is due to more aggressive use of the *inner edge* hotels. Concerns about latency due to hosting MEC workloads deep in the networks are mitigated by the *incumbent's* significant fibre infrastructure.

The *mobile*-first operator will have less access to fibre infrastructure of its own, which makes reliance on the *inner* edge locations more difficult. Instead, there will be more appetite from operators like the *mobile-first* player to make tactical use of its *outer edge* infrastructure, such as radio base stations, to host MEC workloads in "micro datacentre" environments. In practice, this could be to host workloads closer to particularly important public places with high footfall, such as train stations. However, the demand to do so this is not yet strong, and we expect this to happen only opportunistically in the short-term.

This difference is partly because some short-term MEC monetisation opportunities are thought to be from servers hosted *on-premise* – that is to say, as close to the enterprise customer as possible. An incumbent operator which offers fibre connectivity will tend to already have a broader presence on enterprise customer premises, and therefore will be likely to deploy a much larger proportion of *on-premise* edge computing equipment than a *mobile-first* player.

Short-term adoption is focussed on efficient operations, but revenue opportunity has not been dismissed

Our model shows significantly **more aggressive adoption of use cases with internal efficiency benefits, rather than to enable new external-facing services.** This can be exemplified with MEC, which, not being a cost-saving opportunity, is less aggressively pursued than vRAN.

This does not mean that operators will dismiss MEC, private LTE networks and vCDN as non-important opportunities over the next few years. Looking to 2025 and beyond we think these offer some of the strongest opportunities for telcos looking to find new routes to revenue beyond their traditional offering. Most leading European operators are working hard on this.

However, it is generally accepted that these are not clear-cut, tactical or short-term wins. They will require exploration, investment and education over the next 5 years, from building new partnerships and ecosystems to developing closer relationships with enterprises who will be making decisions about who they want to offer them these services. With new competitors (e.g. the hyperscalers with MEC or SIs exploring private LTE network opportunities), the playing ground for these more nascent opportunities is very much being formulated now.

Addressing the edge opportunity: operators can be more than infrastructure providers

The focus of this paper is primarily on edge sites and edge infrastructure, and the way that telcos will provide this over the next 5 years. Of course, this is not the only role that a telecoms operator could pursue within the edge value chain as a whole. Figure 13 breaks down the edge play into "edge sites" "edge infrastructure" and "edge platforms".

For our purposes, we have focused on scenarios 2 and 3, where telecoms operators provide the edge sites and, in many cases, the edge infrastructure as well. There are third parties, like Equinix, for example who will look to follow scenario 1, and provide edge hosting via their neutral data centres.

Figure 13: It is not clear which parties will be the primary player in edge hardware and platforms, though there is a role for operators in all scenarios

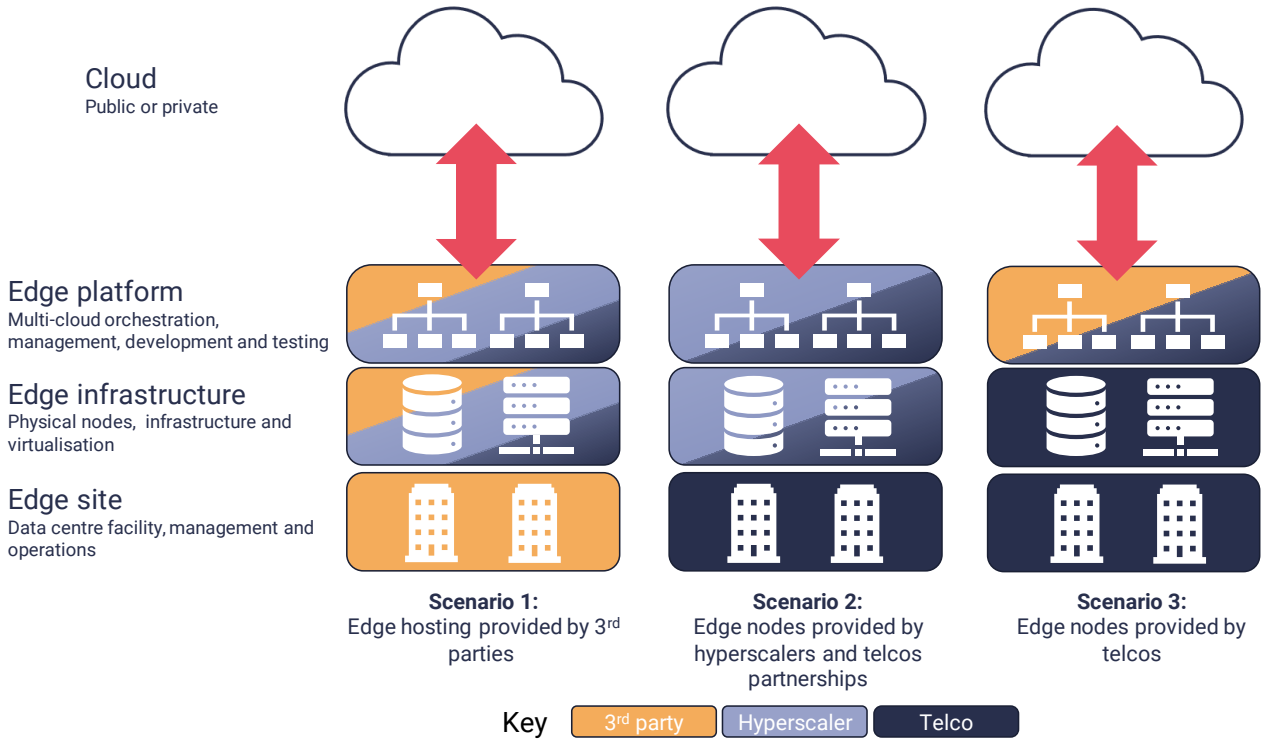

Source: STL Partners

When evaluating business models, operators may look to do more. STL Partners believes there are 5 main business models operators could look to pursue when it comes to edge computing. Operators are likely to need to evaluate at a domain level the opportunity to move beyond dedicated hosting (providing the physical site) or providing edge IaaS (providing the edge infrastructure).

Figure 14: There are a range of viable business models for operators with edge computing

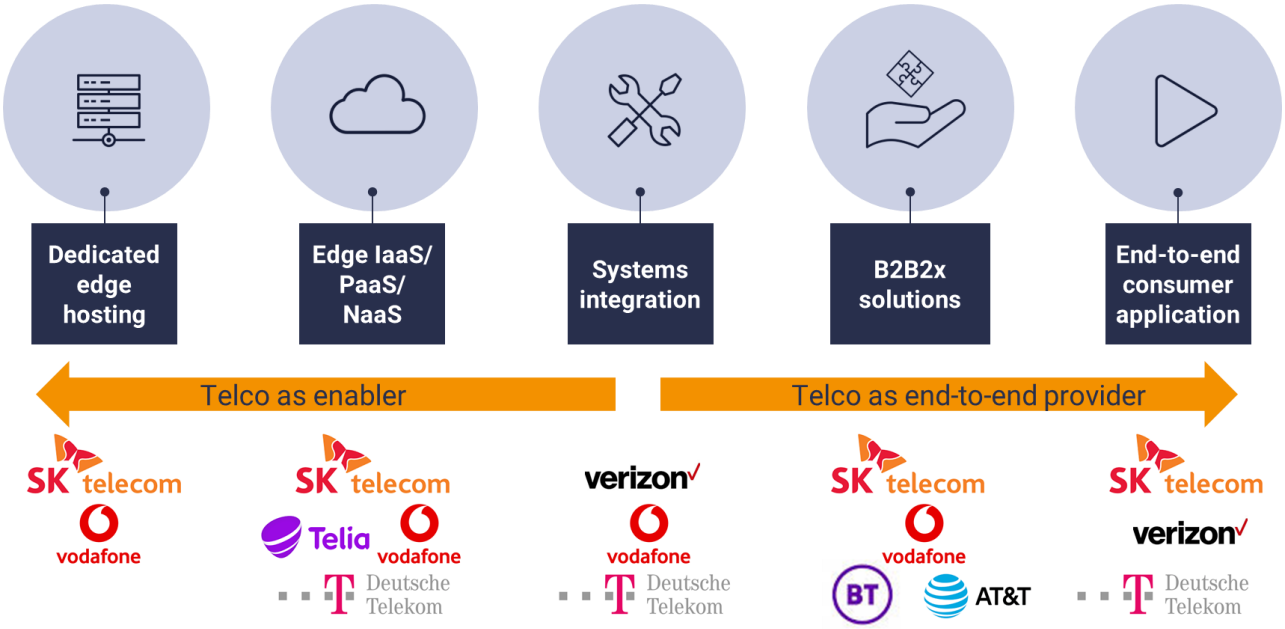

Source: STL Partners

We have seen evidence of operators moving from "enabler" to "end-to-end provider". As an example, within MEC, operators are well positioned to do more than just IaaS. The opportunity can be split broadly into two different types:

1. **Horizontal platform-as-a-service opportunity.** Where operators own the infrastructure, they could also offer a PaaS solution to ensure that access is given to that infrastructure in a useful way. That could mean, for example, making sure that edge workloads are always running on the nearest possible piece of edge infrastructure. An operator edge platform would need to provide developers with the APIs needed to run their solution at the edge, as well as services like a testing environment and analytics on performance at the edge.

2. **Vertical / use case specific opportunity higher up the stack.** Making up the business models on the righthand side of Figure 14 are roles for operators higher up the value chain. This includes being able to provide the systems integration that enterprises will often need, through to delivering the full end-to-end solution including ongoing customer support. In these cases, operators will need to invest in new partnerships and skills to make this possible and will need significant knowledge of the customer's industry. This means that, while the value capture is likely to be high, operators are unlikely to be able to do this across the board and will need to prioritise certain industries or use cases they will deliver this for.

Conclusions: practical recommendations for operators

Our modelling exercise highlighted two clear messages for operators wanting to understand how edge infrastructure will develop over the next 5 years.

1. **The short-term edge opportunity does not require highly-distributed deployments – which is an opportunity to learn while keeping a lid on investments.** The most significant and earliest uses for edge sites will be to support vRAN, which can be done in sites close to the core of the network. This is an opportunity for operators to develop skills and expertise while the edge opportunity is both: primarily about delivering internal, network-centric use cases, and is not massively distributed. Operators will be able to identify benefits even with 2-5 more sites than you have currently.

2. **There is no short-term route to revenue – but now is the time to start exploring**. In our timeframe, MEC will not be a major revenue generator for most operators. Instead, it will be deployed in a tactical and opportunistic fashion. However, looking further out, 5-10 years and beyond, it is likely to grow into a significant opportunity. The crossover between MEC and a domain like vRAN should not be ignored: operators should organise to share learnings and ensure their physical edge sites are suitable to support all significant domains.

Based on these core takeaways, below are 5 practical steps operators should take today:

1. Underwrite / piggyback on vRAN to invest in edge sites that have foundational units of servers (likely to be able to serve demand for the next 2-3 years). These sites should have room to add more servers as demand increases and balance the network and third-party application requirements. This might mean, for example, keeping an eye on delivering end-to-end network latency below 30ms for as much of the population as possible.

2. Figure out security (physical and network) to enable third party applications to be able to run alongside your network functions. This means organising your technical spaces not just with your own network team in mind. Operators will want to be able to embrace wholesale / co-location opportunities where they occur and must plan accordingly.

3. Invest in delivering large, bespoke on-premise projects for trophy customers. This will be key to building the skills and partnerships needed to enable you to productise enterprise edge offerings in the long term.

4. Work with the hyperscalers, but don't put all your eggs in one basket. In the next 5 years, there will be MEC in telco's networks where the infrastructure has been built out by hyperscalers. However, this in and of itself is unlikely to be lucrative for the telecoms operator. With a similar model today, regional data centre operators barely break even from their business hosting the

hyperscalers. Operators need to take an active role in exploring MEC use cases, developing the right relationships and engaging with end-customers to enable them to enable co-opetition.

5. Evaluate your role above and beyond providing infrastructure. Operators should consider innovative use cases and new business models on a domain by domain basis to ensure, where appropriate, they move up the value chain.

STL PARTNERS

Research

Consulting

Events

www.stlpartners.com