Growing Enterprise Revenues
Edge Insights Service

Kindly supported by:

Red Hat

STL PARTNERS

Executive Briefing

BUILDING TELCO EDGE: WHY MULTI-CLOUD WILL DOMINATE

Developing edge-enabled services is complicated; there are key decisions to be made at each level of the value chain, starting with infrastructure. Leveraging results from a survey conducted by STL Partners with 150 CSPs, this report outlines why edge infrastructure will be multi-cloud and what telcos should do about it.

Dalia Adib, Edge Computing Practice Lead | dalia.adib@stlpartners.com | October 2021

Preface

The document has been prepared by independent consulting and research firm STL Partners and commissioned by Red Hat®.

It is based on the output of an extensive survey interview programme conducted by STL Partners with 150 telecoms operators globally, as well as STL Partners' continuous research programme on edge computing.

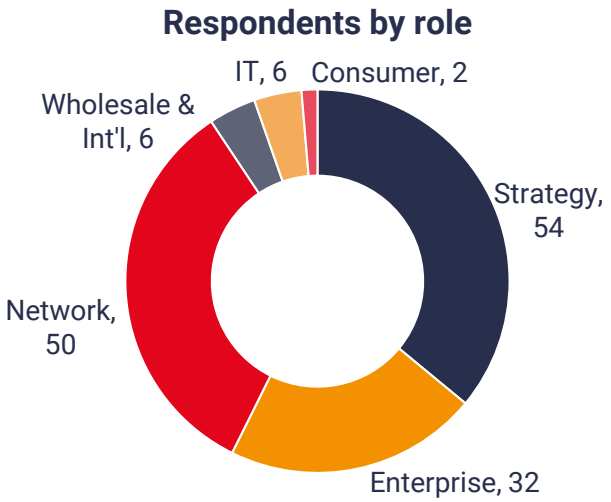

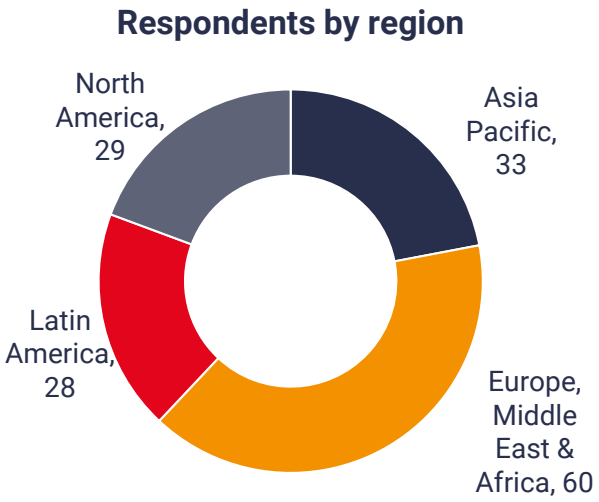

STL Partners maintains strict editorial independence. Mentions or allusions to companies or products in this document are intended as illustrations of market evolution and are not included as endorsements or product/service recommendations.

Executive Summary

There are three key factors to consider to build the CSP edge

Communication service providers (CSPs) are important stakeholders that can determine the growth of the edge market, namely because they are the only ones that can build a significant proportion of it – the network edge. However, even building the infrastructure is a challenging task, as there are three factors that CSPs need to consider:

1. **Convergence**: The extent of shared infrastructure between network functions and consumer/ enterprise applications varies between operators, depending on the maturity of the operator's edge strategy. Our survey found that CSPs in North America and Europe tend to favour a higher degree of separation in the hardware and software stacks, in contrast to CSPs in Asia Pacific.

2. **Organisation**: Within CSPs, different decision-makers and influencers are owning the edge opportunity, and they will build different edges to meet their needs. For instance, enterprise teams will prioritise taking edge applications to market, while network teams would seek to consolidate edge infrastructure to serve both telco and end-customer applications.

3. **Hyperscaler partnerships**: CSPs have been partnering with hyperscale cloud providers across various edge domains, but mainly to deliver network edge cloud services to third-parties (enterprises and developers).

Figure 1: Most hyperscaler-CSP activity is to support CSPs' ability to monetise edge for third parties

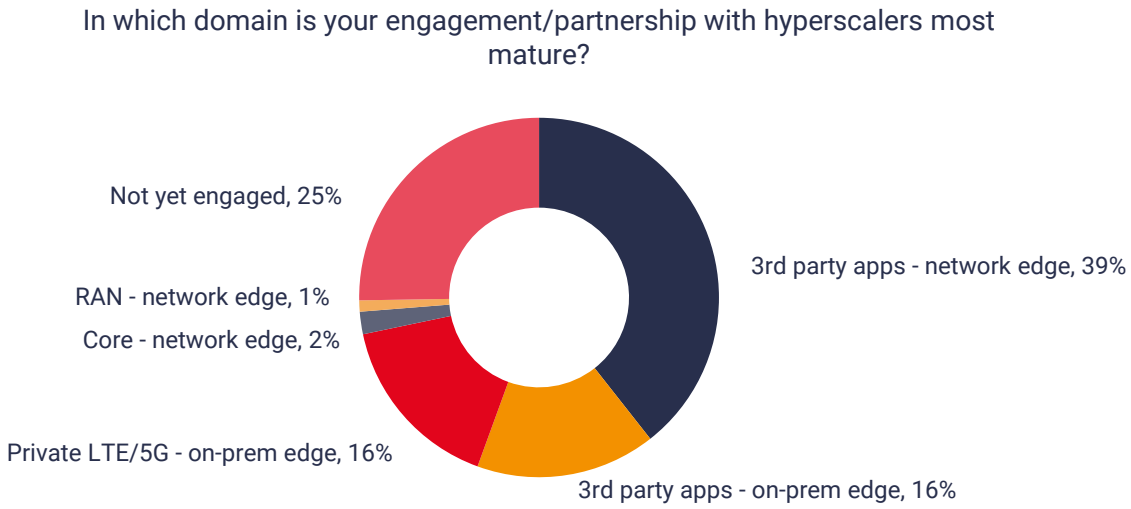

The edge will be multi-(edge) cloud

Every telecoms operator's edge strategy will be shaped differently by these three factors above: how converged their edge infrastructure will be, which part of the organisation owns their edge strategy, and how they choose to partner with hyperscalers. As a result, it is inevitable that the CSP edge will be made up of hybrid environments.

Although hyperscalers are viewed as credible partners to support CSPs to build their edge due to the perceived strength of their cloud platforms and developer ecosystems, most CSPs believe only a minor proportion of their edge will be hyperscaler-provided. For example, on average, 33% of network edge infrastructure built by CSPs to support third-party applications will be provided by hyperscalers.This means that most CSP edge infrastructure will be a multi-cloud environment.

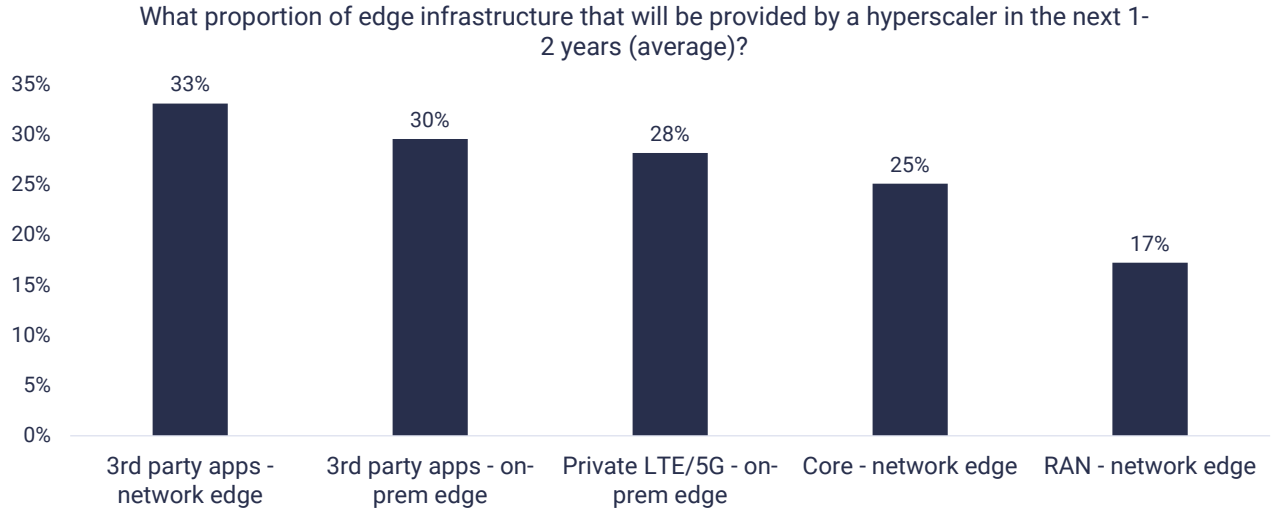

Source: STL Partners survey, n=136

CSPs must build capabilities and partnerships today to support their edge business

Going forward, CSPs must evaluate which partners they should work with and how depending on their use cases, go-to-market strategy and technical needs throughout the value chain. For example, a key area where CSPs can build IP and provide value is in the PaaS layer and orchestration; providing tools to help their internal operational teams and customers monitor services end-to-end using a single pane of glass and manage workloads across different infrastructures easily. Regardless of the commercial edge strategy, CSPs must be clear on organising their edge partnerships strategy. This is something STL Partners will be covering in the future in our Edge Insights Service.

Table of Contents

Table of Figures

Defining the edge

Edge computing will remain a focus for telecoms operators for the foreseeable future, both to optimise the network and enable new, third-party applications and services. In fact, 70% of survey respondents believe investment levels of edge computing for supporting third-party applications will increase over that for internal network infrastructure in the next five years.

This report explores how telecoms operators will build their edge computing business, infrastructure and services, and the role multi-cloud will take in this. Before diving into this, it is worth defining this confusing and complicated space. At a high level, edge computing refers to cloud-native computing (and storage) being brought closer to the end-device or source of the data, rather than centralised in a remote, hyperscale data centre.

The telecoms industry has been exploring the role of edge computing for over four years, starting when network functions virtualisation (NFV) began to make real strides. The initial interest was in mobile edge computing (MEC), but this has now evolved to multi-access edge computing to incorporate fixed networks and non-cellular networks too. Outside telecoms, there is edge compute capacity in regional data centres provided by third parties centres, e.g. data centre operators and cloud providers. These are often in untapped geographies, such as Tier 2 cities. In addition, there is edge compute at customer premises, e.g. business campuses or factories.

We outline the scope of edge computing below. There is a full spectrum of possible edges from devices to regional data centres. Some of these edge locations may be owned and/or operated by communications services providers (CSPs). The CSP edge contains the most relevant types of edge for CSPs: network edge and on-premises enterprise edge. They contain infrastructure either owned by a telecoms operator (e.g. a CSP data centre) or operated by one (e.g. network CPE at a customer site).

Figure 2: The spectrum of edge computing locations

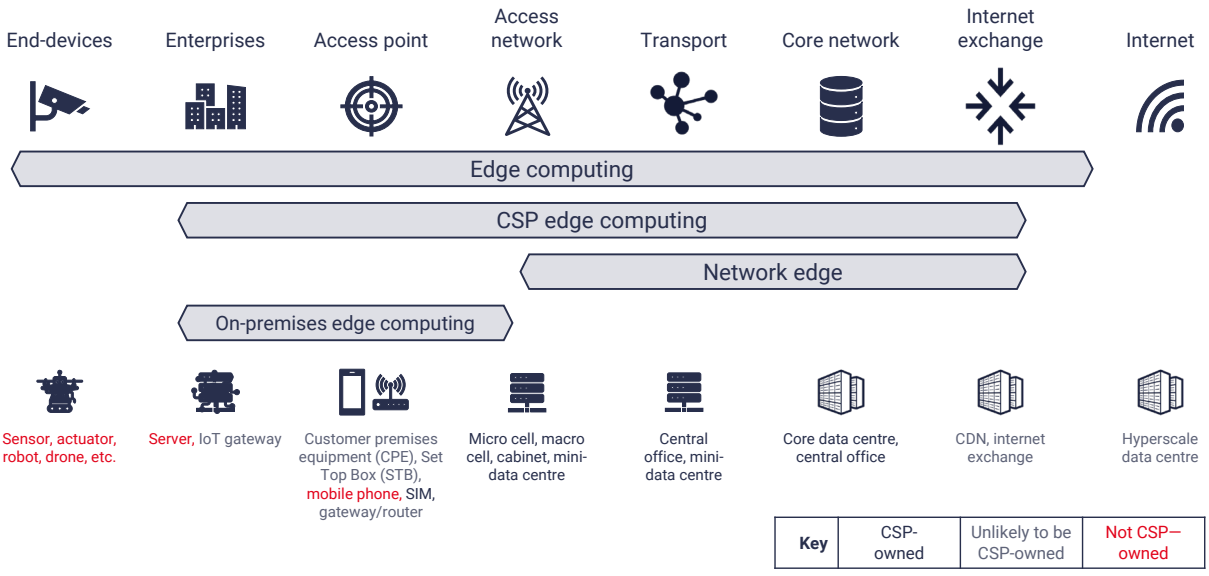

Source: STL Partners

There are two main types of applications that can be processed on CSP edge computing:

1. Telecoms applications that run, protect and monitor the network – i.e. CSP's own network functions;

2. Consumer/enterprise applications - which CSPs may provide for third-party customers.

STL Partners has been supporting the telecoms industry in exploring the opportunity to provide services and solutions to third parties by leveraging their edge computing infrastructure[1]. These could include enterprises deploying IT applications locally to comply with data sovereignty laws, developers using edge to optimise their applications, IoT solution vendors using edge to reduce latency for mission-critical applications, etc. Our survey highlighted the importance for CSPs in investing in the infrastructure for these applications. On average, CSPs believe that 40% of edge computing investments in the next 1-2 years will be used to support these applications, rather than be used for network functions infrastructure.

Figure 3: Defining edge computing within telecoms

	TYPE OF APPLICATION	
	External applications	Telecoms applications
TYPE OF EDGE	**Consumer/enterprise applications** e.g. SaaS, industrial applications, IoT, AI/ML	**Network functions** Virtualised or containerised network functions
On-prem edge e.g. on-prem mini data centre edge device in building/factory universal CPE	**Using on-premise edge compute infrastructure to run consumer/ enterprise applications** e.g. video analytics, industrial applications, healthcare applications, AI/ML, etc.	**Using on-premise edge compute infrastructure to run network functions and services** e.g. private LTE/5G network, SD WAN / universal CPE services, Wi-Fi gateways, etc.
Network edge e.g. telco data centre central office base station	**Using network edge compute infrastructure to run consumer/ enterprise applications** e.g. video analytics, cloud gaming, CDN, fleet management, AI/ML, etc.	**Using network edge compute infrastructure to run network functions and services** e.g. distributed core and RAN network functions, virtual CMTS etc.

Source: STL Partners

Although the edge computing market is nascent, there are emerging use cases that seek to take advantage of edge computing's main benefits. These include offering the flexibility that comes with the cloud more local to reduce latency, improving reliability, keeping data secure, and offloading processing from the end-device. However, use cases are at different stages of maturity; some will be deployed in the next two years in early adopter markets, others are more than five years away from commercial, wide scale deployments.

[1] See STL Partners' Edge Practice: www.stlpartners.com/edge-computing

	Exploration stage 5 years+ from commercial deployments	POC stage 2-5 years from commercial deployments	Production Already being deployed or will be deployed in next 2 years
Example use cases	• Precision monitoring & control (on-prem edge) • Real-time traffic warning (network edge)	• Predictive maintenance (on-prem edge) • Drone detection (network edge)	• Edge CDN (network edge) • Cloud gaming (network edge)
Examples from CSPs	Vodafone, HERE, Porsche (real-time traffic warning)	AT&T, Vorpal, Microsoft (drone detection)	Verizon & Bethesda (cloud gaming)

Source: STL Partners

Telecoms operators are keen to leverage edge computing to grow revenues, particularly in their enterprise business. There are different strategies emerging: one is to focus on enterprise connectivity and networking, another on developing a horizontal, cloud-like platform for developers, while a third focuses on building end-to-end solutions for specific verticals.

Figure 4: Types of edge services and business models

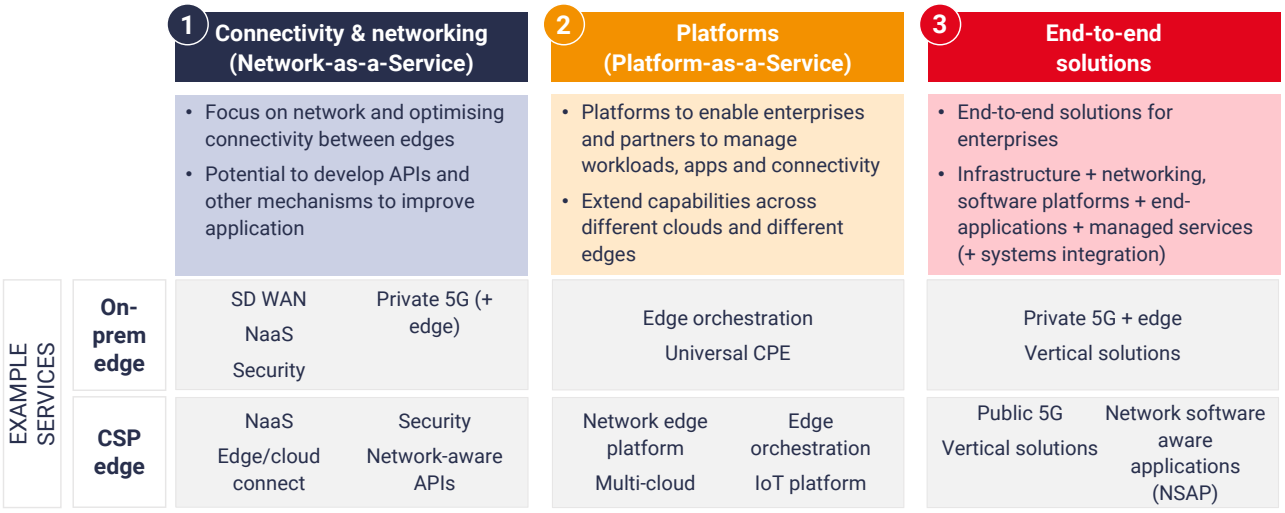

Source: STL Partners

The challenge with any new technology is that it takes time to educate the market and engage the innovators who will build the applications that will leverage its potential. Edge computing is complex, because it has a unique ecosystem that spans several industries: cloud, telecoms, industrial, traditional ICT, plus specific vertical sectors. In order to build an edge-based solution, there needs to be adequate infrastructure (facility, hardware, connectivity, edge cloud) *plus* the applications and services, *and* these need to be integrated so they work together seamlessly.

Figure 5: The edge value chain

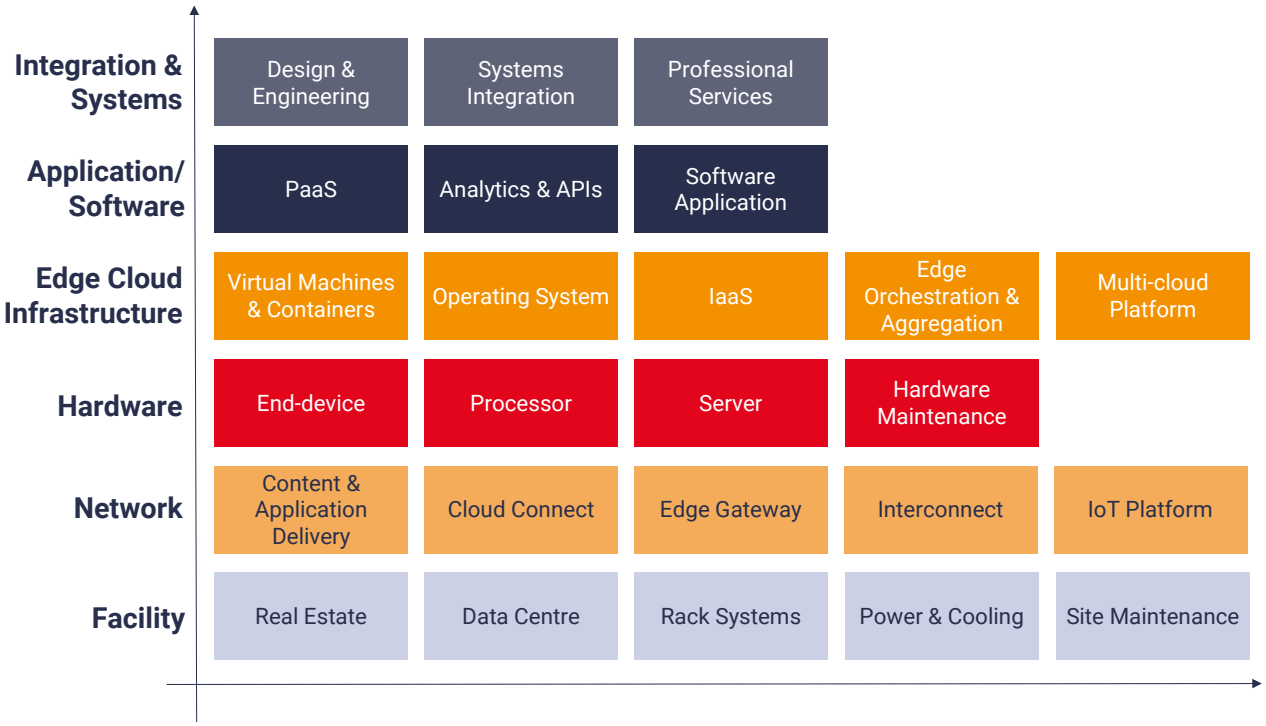

Source: STL Partners Edge Ecosystem Tool

Regardless of the business model and services strategy a telecoms operator chooses to pursue, it will need to first determine how best to build its edge infrastructure to optimise results. This report will dive into three key questions CSPs are still trying to evaluate:

1. How should telecoms operators build edge computing infrastructure that can support both enterprise applications and network functions?

2. To what extent should telecoms operators work with partners, particularly the hyperscalers, to build their edge and take services to market?

3. How can telecoms operators effectively work with the ecosystem?

Laying down the foundations: Options for building the CSP edge

Telecoms operators are in a good position to build the edge: they have highly distributed facilities, which are well-connected and well-powered. Nevertheless, some of these facilities require additional investment to be ready to host data centre-like infrastructure. This investment is non-trivial, however it is happening anyway as part of the transition to NFV and cloud-native, software-based networking, accelerated by 5G. This allows some costs to be amortised more effectively for edge computing.

Edge computing is in a chicken-and-egg situation: without infrastructure, a market cannot develop as there is insufficient capacity to build applications that would leverage edge computing. On the other hand, without a market, it is difficult to make a business case for infrastructure investments.

Most CSPs are building their edge with a piecemeal approach to minimise risk, yet capitalise on the opportunity in the next 2-3 years. Through our research and discussions with CSPs, we have found that there are three factors that are top-of-mind for managers making decisions for edge infrastructure:

1. **Convergence**: Should we build the same infrastructure platform to support both our internal, network functions and enterprise applications?

2. **Organisation**: How should we organise internally to make edge infrastructure decisions?

3. **Hyperscaler partnerships**: To what extent should we work with the hyperscaler cloud providers (HCPs) to build the edge computing stack?

Convergence

The original concept for edge computing in the telecoms sector stemmed from network virtualisation efforts. CSPs were building distributed data centre-like facilities to support their network function infrastructure and consolidating network functions onto server infrastructure in place of functional hardware appliances. This meant there was spare capacity in those facilities to support customer applications and maximise economies of scale by sharing as much of the infrastructure as possible.

However, reality has diverged from the theory. The requirements for network functions are different to that of IoT/IT/business applications that will be hosted at the edge. Network teams have arguably stricter security rules for controlling data and access and the type of hardware for network functions is different than that needed for enterprise applications. For example, radio access network functions need specialised hardware,

> *"When I think about converged edge, some things come to mind: privacy, separation, security. Telco edge applications on the edge contain a lot of customer data, need to be very cautious on how to approach this and want to understand the technology completely."*
> **VP Planning & Engineering, Fixed-Mobile CSP headquartered in North America**

such as field programmable gateways (FPGAs) and hardware accelerators. By contrast, video analytics, gaming or AI-intensive applications may need GPUs and other technologies for enhancing graphics processing.

Despite these limitations, there are different elements of the stack that can be shared between network functions and consumer/enterprise applications, beyond the site itself. Our survey demonstrated that there is a wide range of opinions in the telecoms industry on the level of convergence. Only 12% of respondents globally believe that edge infrastructure should be entirely converged (facility, hardware, application platforms and orchestration). On the other side of the spectrum, 13% of respondents feel that the infrastructure should be entirely separate and only the site itself should be shared. However, there are slight differences across regions.

Figure 6: CSPs in Asia Pacific are generally more favourable to having converged edge infrastructure

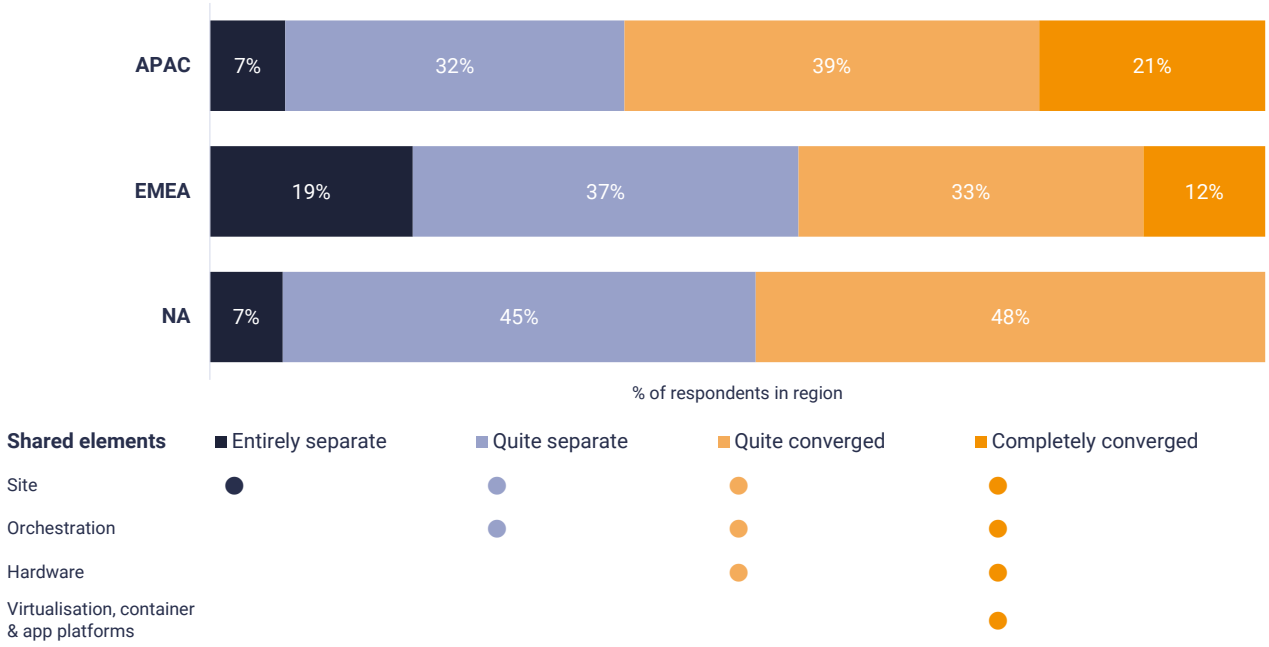

Source: STL Partners survey 2021, n=115
Note: LATAM has been removed as most respondents were not yet evaluating infrastructure for third-party applications

CSP respondents in North America were least supportive of having completely converged infrastructure, whereas Asia Pacific operators prefer convergence. One of the reasons for this is organisational dynamics, which will be explored in the next section. Another reason could be the level of maturity of the operator's edge strategy; CSPs in North America and Europe have gone through a process of initially attempting to share the infrastructure and operating models, but have since found it to be too challenging to manage each domain's different needs. As a result, most CSPs in those regions prefer a degree of separation in the hardware and software stacks.

Organisation

Organisation plays a part in edge strategies because different teams within the CSP's business influence decisions on infrastructure and partnerships. Edge touches different business functions: network teams (who tend to own and control edge sites today), IT (for platforms and supporting systems), wholesale (for CDN-like business opportunities), enterprise (for B2B services) and consumer (for B2C opportunities, such as cloud gaming). Some of these are budget holders/decision-makers, while others influence the edge strategy.

The edge infrastructure for NFV infrastructure (NFVi) is usually owned by the network teams, with some software decisions being controlled by IT, or a converged networks/IT business function. However, ownership of edge for enterprise applications is more fragmented. In some CSPs, the enterprise teams have a significant role to play in determining the strategy, including how to build the edge stack and with whom to partner.

Figure 7: Budget holders of edge infrastructure for NFVi differ from those for third-party applications

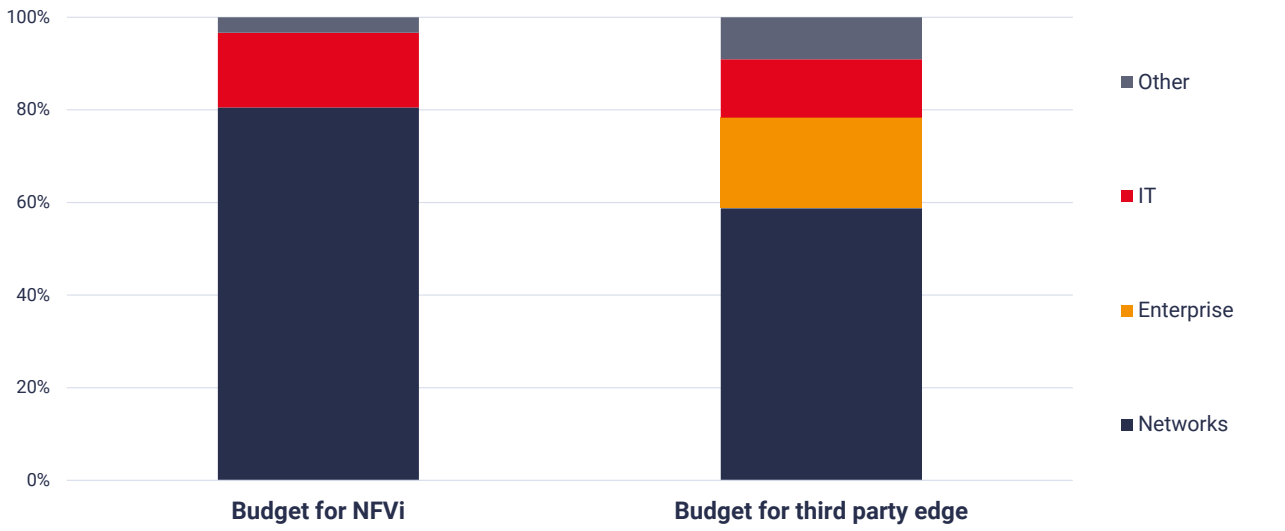

Source: STL Partners survey 2021, n=149

Decisions on edge infrastructure may differ across functions, as each has a different set of priorities. Whereas network teams are focused on building a high-quality network, improving customer experience, and optimising CapEx, enterprise teams prioritise being first-to-market and growing revenues from B2B services. As a result, enterprise teams focus on building the edge infrastructure quickly to accelerate go-to-market and be more responsive to their customers' demands. The network teams, on the other hand, prefer a consolidated approach, developing an edge stack that can support both the network and non-network applications' needs. The likely outcome is that there will be different edges to serve these different stakeholders' needs, at least in the short term.

Figure 8: Commercial and technical teams have different business priorities

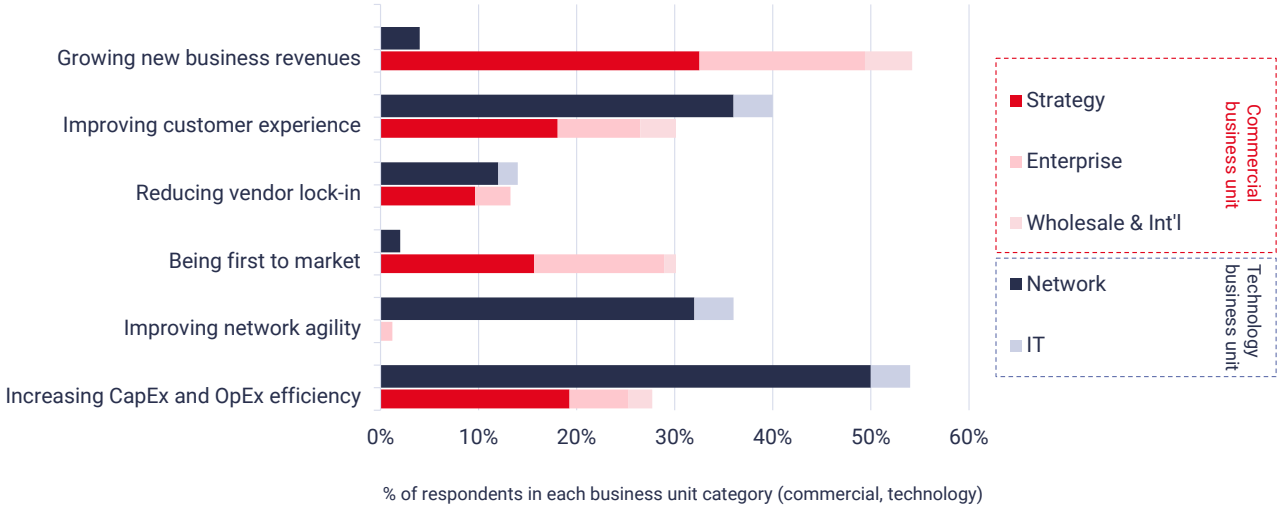

Source: STL Partners survey 2021, n=137

Hyperscaler partnerships

The hyperscale public cloud providers, predominantly AWS, Microsoft Azure and Google Cloud, have been expanding their cloud offering to the edge and developing new solutions in telecoms[2], in parallel. For example, Microsoft has acquired two telecoms vendors (Metaswitch and Affirmed) and AT&T's 5G network infrastructure orchestration capability[3] in the last two years. It also has a portfolio of edge services, building on its edge computing device offering Azure Stack Edge[4]. Together, these have allowed Microsoft to offer multiple partnership models for telecoms operators[5].

In edge computing specifically, many of the partnerships between telecoms operators and hyperscalers to date have focused on the 'third-party apps – network edge' domain. This is where the hyperscaler implements its edge stack (hardware and software) at a CSP's network edge site and the two partners offer network edge cloud services to third parties (enterprises and developers). The survey results highlight that this domain is where partnerships are most advanced today, followed by 'third-party apps – on-prem edge' domain. 55% of CSPs surveyed see these two domains as the most mature areas for hyperscaler-CSP edge partnerships.

[2] See previous STL Partners report: Telco edge computing: How to partner with hyperscalers
[3] Data Center Dynamics, 2021: AT&T: Microsoft cloud deal covers the Edge and network stack, not core network function
[4] Azure Stack Edge: Azure-managed equipment designed to run outside of Azure's hyperscale data centres
[5] See previous STL Partners report: Microsoft, Affirmed and Metaswitch: What does it mean for telecoms?

Figure 9: Most hyperscaler-CSP activity is to support CSPs' ability to monetise edge for third parties

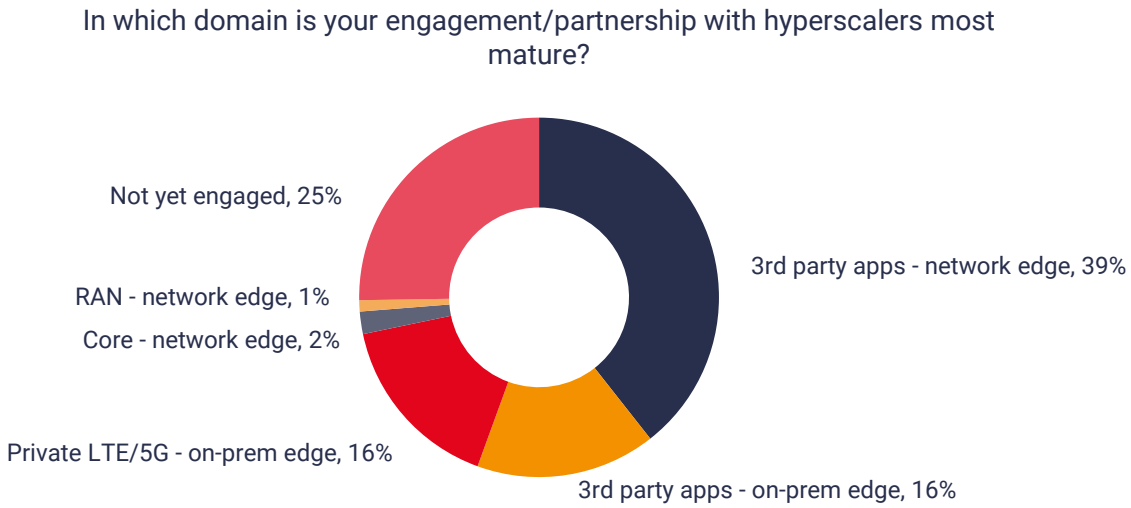

Source: STL Partners survey 2021, n=134

The nature of the partnerships differs across edge domains, in terms of who is responsible for the computing infrastructure, connectivity, applications, services, etc. The below diagram demonstrates examples of partnerships, highlighting who is responsible for each part of the stack and the edge domain it is associated with.

Figure 10: Example of hyperscaler-CSP partnerships in edge computing

Partners		Azure + AT&T		AWS + Verizon / Vodafone / Telstra / KDDI / SKT	Azure + Telefónica (Tech)	AWS + Telefónica (Germany)	
Edge domain		(5G core) NFs	Apps	Consumer/enterprise appslications	Private networks	Private network (5G core)	Apps
		Network edge	Network edge	Network edge	On-prem edge	Network edge	On-prem edge
Application & platforms	Application	5G core NF vendors	Application developers /enterprises	Application developers / enterprises	NF vendors (incl. Metaswitch, Affirmed)	Ericsson	Application developers /enterprises
	Application platform / PaaS	AT&T / vendors	Azure	AWS	Azure	Ericsson	AWS
Edge cloud infrastructure	Orchestration	Azure		AWS	Azure	AWS	AWS
	Containers / hypervisors	Azure		AWS	Azure	AWS	AWS
Hardware		Azure		AWS	Azure	AWS	AWS
Facility		AT&T		Verizon / Vodafone / Telstra / KDDI / SKT	Enterprise	Telefónica (Germany)	Enterprise

Key: Hyperscaler manages; CSP manages; 3rd party manages

Source: STL Partners

The hyperscalers are perceived to be strong partners in helping CSPs to monetise the edge, as they can provide elements across the value chain, to varying degrees:

- **Infrastructure**: Compute and storage hardware (and cloud software) to run platforms and third-party applications
- **Platforms**: Software platforms that allow developers and enterprises to onboard and manage edge applications across edge (and cloud) infrastructure
- **Go-to-market**: Taking services to market jointly and engaging potential (enterprise) customers
- **Delivery**: Implementing and operating the services in the field (not an area of focus for hyperscalers, as they often rely on partners to do this, but CSPs surveyed did see a role for hyperscalers here too)

50% of CSPs surveyed believe the strength of the hyperscalers' cloud platforms is the main benefit of these partnerships, with a further 43% favouring the ability to accelerate go-to-market. However, there are CSPs who are wary about relying too heavily on hyperscalers; 43% of those surveyed perceive them to be a threat to their business. Others are more concerned that partnering reduces overall share of revenue (40%) and limits the ability to differentiate (31%).

Given these varied opinions on the level of hyperscaler involvement, combined with the fact that there are different requirements demanded by each edge domain, edge infrastructure will be multi-cloud, with some of it built through hyperscaler partnerships and some by the CSP with other parties. Even though the '3rd-party apps – network edge' domain is the most mature domain for hyperscaler partnerships, the survey suggested that only a third of the edge infrastructure will be provided by a hyperscaler in the near term.

Figure 11: Edge infrastructure will be a hybrid of hyperscaler-provided and CSP-built

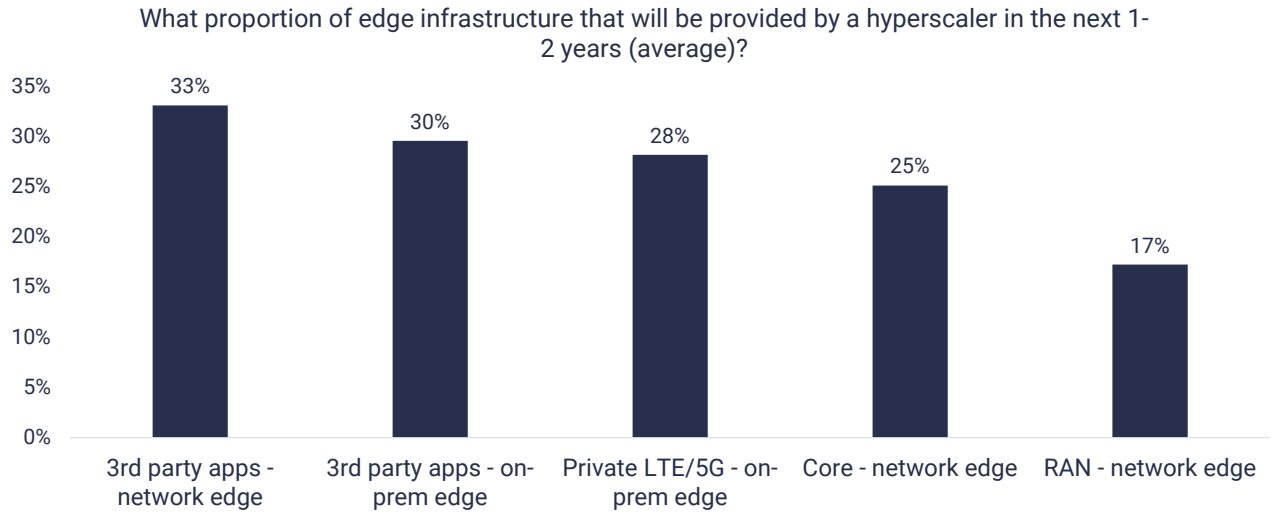

Source: STL Partners survey, n=136

There is no single edge – it is multi-cloud

Given the varying factors affecting telecoms operators' decision for how to build the edge, it is inevitable that the edge will be made up of hybrid environments. The reality is that different telecoms operators are partnering with different companies to build their edge solutions, hyperscalers being just one type. We are already seeing multi-partner scenarios in North America, such as:

- AT&T partnering with Microsoft, HPE and Google;
- Verizon with Nokia, Red Hat, AWS, Microsoft;
- Bell Canada with AWS and Google.

There are differences across regions in their views on which hyperscale cloud partner is the optimal partner to work with for edge computing. This is likely due to the perceived strength of the partner in their region, as well as the types of applications they feel the hyperscaler can support best. Generally, AWS is seen as the market leader (by market share), particularly in the ISV/application space; Azure in enterprise and, more recently, networking; whereas Google has strong capabilities in AI and machine learning. There are other large cloud providers, e.g. IBM and Alibaba Cloud, which are perceived as credible cloud (IaaS) partners too.

Figure 12: Differing opinions across regions on which cloud provider is the preferred partner for edge computing

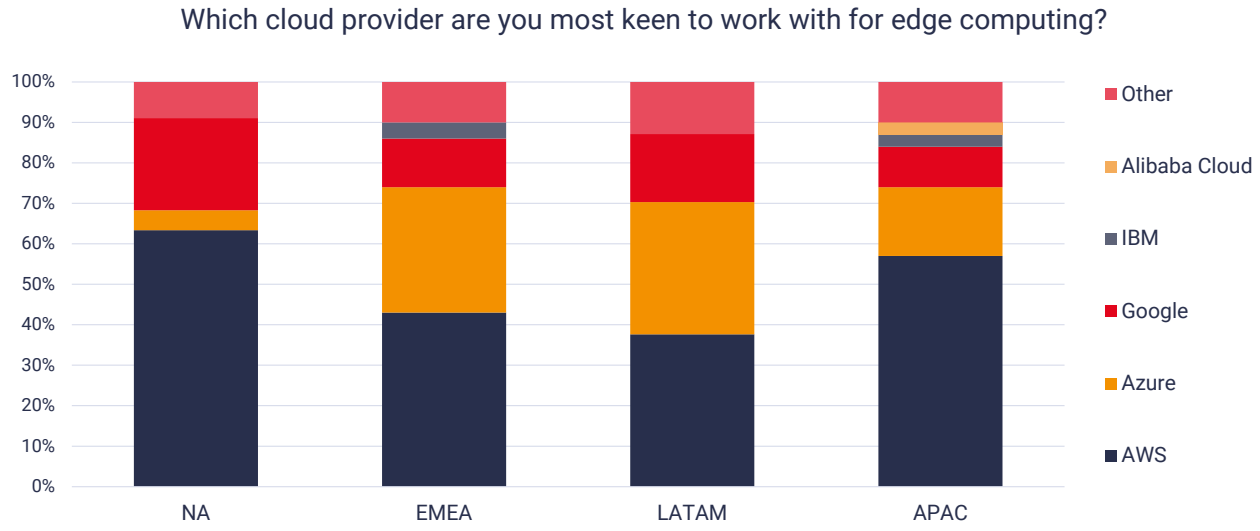

Source: STL Partners survey, n=128

This trend will continue as telecoms operators develop services and a go-to-market that ensures maximum economic benefits. Public edge cloud will be valuable for ISVs and applications that need global scale, for example cloud gaming, broadcasting, general enterprise SaaS applications, etc. Some

> *"For us, starting with the customer, they are also going hybrid cloud. They don't want to just go with AWS or Google, they will actually go to a number of them. So, we're not really worried and preoccupied about doing a deal with any one of them, we do a deal with all of them. We think where we provide a lot of the value is how you integrate them all so that it's a seamless solution for the customer, again removing that friction."*
> **SVP Enterprise, Fixed-Mobile CSP headquartered in Asia**

customers will need private edge cloud models, such as those in manufacturing, financial services and other data-sensitive segments, e.g. government institutions.

Although it is inevitable that edge will be multi-cloud, this fragmentation does pose risks for telecoms operators as they take edge services to market:

- **GTM and sales**: The different types of edges for different applications makes it more difficult for a customer to know what they need and more challenging for CSPs' marketing and sales teams to clearly communicate the proposition and help customers navigate the options. → CSPs should create internal and external tools that start with the customer's needs, then eventually lead the salesperson/customer to the appropriate solution.

- **Partner management**: Partnerships span multiple domains and different parts of the value chain. → CSPs need to be clear in how they categorise partners, what their roles are, which internal stakeholders are involved and must manage the relationship and the joint value proposition effectively.

- **Infrastructure management**: Both CSPs and their customers will need to reduce the complexity that comes with the proliferation of different vendors providing edge platforms and different parts of the edge stack. → A key area where CSPs can build IP and provide value is in the PaaS layer and orchestration; providing tools to help their internal operational teams and customers monitor services end-to-end using a single pane of glass and manage workloads across different infrastructures easily.

> *"Operators must find synergies between the layers of the stack - could potentially reach a stage where they are partnering with multiple vendors for NFVi and edge, leading to them having an overly complex system consisting of multiple parallel solutions."*
> **Strategy Director, Fixed-Mobile CSP headquartered in Europe**

Conclusions and recommendations: What CSPs should do next

Infrastructure is the foundation of any edge strategy, therefore it is paramount that CSPs have a clear edge infrastructure strategy to get started. This report has addressed the complexity of decision-making in this domain. Going forward, telecoms operators must acknowledge that infrastructure will be fragmented multi-cloud build capabilities in the following areas:

- Infrastructure management and orchestration
- Customer service platforms
- Operations and field management

Most CSPs agree, however, that they will not be able to monetise edge computing by solely focusing on edge infrastructure. There is a complex value chain that requires new ways of partnering with existing and new vendors. CSPs should evaluate partnerships along three dimensions:

1. **Use cases/verticals**: Partners vary in their ability to serve use cases in specific verticals. For example, at the hardware level, there are vendors who specialise in high-performance computing to serve AI-enabled applications, whereas others focus on general CPU. Targeting use cases becomes more prominent at the application layer; there will be specialist applications and it will be important for the CSP to be able to onboard these different applications easily to scale edge computing.

2. **Value chain**: An end-to-end solution (for an enterprise customer) brings together hardware, software platforms, applications and services. There will be partners who can offer the full stack, simplifying the role for the CSP, and others who provide best-of-breed components, needing more integration and orchestration.

3. **GTM/delivery**: CSPs do not only expect partners to help build the infrastructure and solution stack, but accelerate engagements with customers. This could be through edge cloud platforms that are easy to use by enterprises and developers, GTM channels and delivery/integration capabilities to help the customer adopt the solution. CSPs vary in their ability to take non-connectivity solutions to market; partnering with the right companies can close the gap as the telecoms operator develops these skills in-house.

STL PARTNERS

Research

Consulting

Events