In partnership with:

Red Hat

STL PARTNERS

TELCO-HYPERSCALER PARTNERSHIPS: WHERE ARE WE IN BUILDING THE EDGE AND WHAT ROLE DO TELCOS PLAY?

Webinar: Questions and answers

How Private 5G is transforming enterprise business operations: lessons learned from real deployments

This document outlines the questions and answers received from the STL Partners and Red Hat webinar, '**Telco-hyperscaler partnerships: where are we in building the edge and what role do telcos play?'** which was hosted on Tuesday 29th November 2022.

You can watch the recording of the session, and also access the slides, using the link here.

Webinar questions and answers

Do you forecast a day when the edge may become disruptive to cloud in like 5 years?

At STL, we believe that the edge will be complementary to the cloud. Organisations do not need to make a binary choice between edge computing and cloud computing as the two models are compatible, and each is suited to different use cases. The key driver with edge applications for example will be for the growing need for real time data driven decision making such as with autonomous driving technologies which depend on both AI and machine learning systems. However, deploying an edge-to-cloud platform would enable organisations to employ each model where it makes the most business sense for them to do so while keeping information flowing throughout the enterprise network.

STL Partners

What was the primary underlying reason for your reduction in the forecast for edge deployment?

We have lowered our forecast from 2023 for edge deployment for three key reasons:

- Initial rate of deployment for China and India were higher than what we now expect. These two countries made up 70-75% of last year's forecast, therefore reducing the estimates has had a significant impact on total numbers.
- Longer than anticipated gaps between the 5G (NR) rollouts and edge deployments. For example, although Saudi Arabia was one of the earliest markets for 5G radio, the operators there have not yet launched network edge services.
- Slower rollout of subsequent edge sites after initial launch in leading markets, such as Germany, UK and South Korea. The operators there were early in deploying initial edges, but these have not increased in number significantly the last two years.

STL Partners

Do you believe that, at some point in the future, despite current partnerships made, hyperscalers might become an increasingly significant threat for telcos and eventually either get some operators out of business or merge with them? What would this all mean for the competition on the respective tech and telco markets?

As stated in our latest edge capacity forecast the two biggest drawbacks highlighted in the survey of working with hyperscalers were "reducing overall share of revenue from edge services" and the "threat of hyperscalers becoming telecoms companies". However, with regards to the question about hyperscalers potentially acquiring telcos, there are certain implications that we don't believe the hyperscalers want to get into if they were to acquire a telco, including pretty stringent national regulations. In many ways the hyperscalers are competitors, partners and vendors to the telecoms operators, and some are more willing to go head to head with operators in certain domains (e.g. AWS

in the private 5G space) whereas others have a stronger appetite to strictly work with operators (e.g. Microsoft with Azure for Operators and their acquisitions of Metaswitch and Affirmed).

In terms of the other point on telcos struggling to keep their customers while hyperscalers gain further traction, the related question is "how can telcos retain and enhance the level of value it's providing to the customer (developer or enterprise for example) in relation to the hyperscalers?". To this point, we've seen relatively successful business models whereby the telco has been co-developing use cases with a hyperscaler partner, where this involves a high degree of interdependence between the network and compute elements (inc. functionalities of the 5G network for example) and the requirements of the use case (greater network awareness, intelligence in applications and vice versa). This is actually critical because it positions the telco more as an equal, important, value-adding partner in developing and commercialising the application, versus a scenario where the telco is purely providing the connectivity.

STL Partners

Network-as-a-Service is a hot topic in telco. Are there many examples of telcos partnering with hyperscalers to deliver NaaS use cases, or is the industry still trying to properly define what NaaS looks like?

NaaS has certainly seen resurgent interest in 2022, with the GSMA, TM Forum, CAMARA and others leading the way in defining the standard models necessary for scale. There are many forms of NaaS, but in our upcoming Manifesto we are exploring it as a service enabling cloud-native networking over any cloud or physical network. We're exploring this as an enabler of what we're calling 'utility net compute' or 'edge-as-a-Service': end-to-end networks spun up, orchestrated and delivered with a high degree of automation in support of MEC use cases and private networking requirements. Keep an eye out for our updated Manifesto next week.

STL Partners

What is your definition of the Telco Edge? Packet Core location? CO/MSO location? CU location?, DU location?, Antenna site?

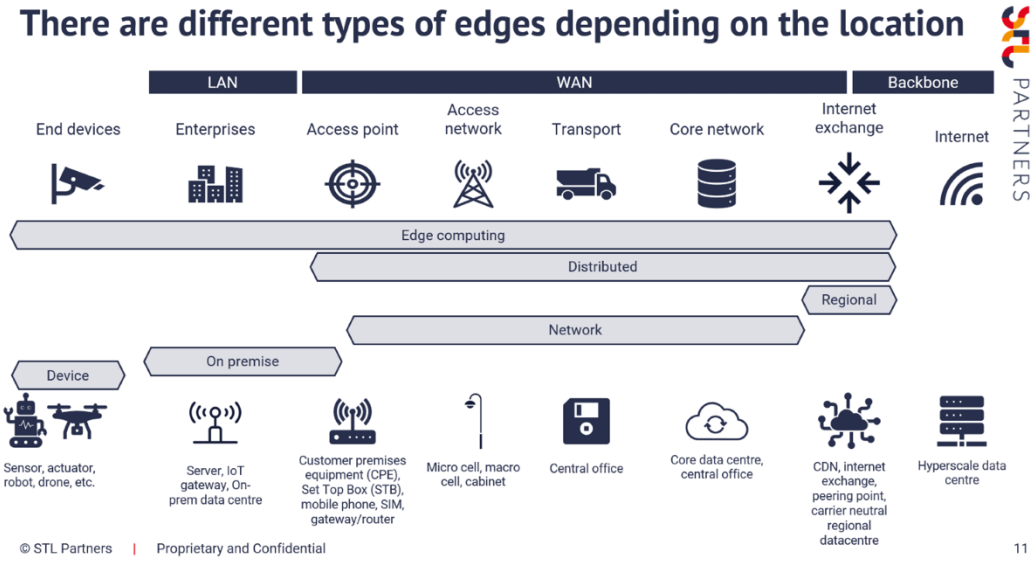

STL Partners

What are your thoughts on the smart cities edge computing opportunity?

Smart cities are a strong opportunity for edge computing, particularly as it relates to the network edge. The majority of use cases we are projecting to mature in the short-to-medium term involve on-prem edge infrastructure. As a result of its distributed, wide-area characteristics, smart city offers the potential for edge computing with the compute hosted in a more distributed manner. Video analytics solutions involving public transport and security are significant use cases which are already generating demand from a public safety and sustainability point-of-view.

STL Partners

How do you see the role of neutral hosts in those ecosystems?

Neutral hosts play an important role in edge computing ecosystems by providing a platform for multiple service providers and operators to access the same physical infrastructure. By providing a platform that is not owned by any single operator, multi-carrier, neutral host infrastructure will become particularly important for public, indoor use cases. As 5G SA deployment extends, consumer use cases for smart city, retail, cloud gaming, etc. will need this kind of connectivity to provide the premium connection that users are beginning to expect. Neutral host reduces the limitations of 5G indoors, whilst enabling multiple carriers to share/offload the infrastructure liability. Neutral host providers will need to integrate closely with the application ecosystem to create the scale necessary for their success, and we will see more and more of these players working with SIs, telcos, and ISVs in the coming 18 months.

STL Partners

How CSP's can handle the regulatory and data security issues if planning to move network functions or data into public cloud?

CSPs can handle the regulatory and data security issues of public cloud migration through the following:

- Ensure that all data is encrypted in transit to meet regulatory and data security requirements.
- Create control policies that limit access to cloud infrastructure and resources to authorized personnel only.
- Implement a security monitoring system to detect potential malicious activity in the cloud environment
- Conduct regular audits of their cloud to ensure that all regulatory and data security requirements are being met.
- Develop a detailed disaster recovery plan to ensure that any data loss or corruption is mitigated quickly and effectively.

STL Partners

Is the far edge (on-prem computing) key to the growth of the telco edge as the far edge is not a simple expansion of cloud technologies?

Yes, the far edge (on-prem computing) is key to the growth of the telco edge. The far edge helps organisations deploy use cases which need a reliable and secure way to access applications and services close to the end user on site. This is especially important for latency-sensitive and mission critical applications such as those related to the Internet of Things (IoT).

STL Partners

From an aspect of time synchronization, what is the latency thresholds being set from core to edge across the hierarchy of near edge, far edge, etc. How do you see this evolving?

Most operators, especially those in smaller countries with a higher population density, can meet the latency demands of key use cases with their limited edge deployments. Use cases like gaming, AR/VR, and video analytics generally get a round trip latency of 15-30ms which provides sufficient user experience to make the additional CAPEX of investment in edge resources unnecessary. Currently, there is very limited demand for uRLL (ultra-reliable low-latency) connectivity, though latency critical use cases will inevitably evolve and mature over time.

Round trip latency is also impacted by edge interconnect between operators. This can result in actual application latency being much larger than the network latency reported. On top of this, investments in 5G deployment are more impactful than the deployment of further distribution of edge assets (the latter often only provides latency improvements of 1-2ms).

STL Partners

What are your definitions of edge data centres (chart shown in intro slides)? Based on the numbers shown these seem to be sizeable structures, whereas there are models of micro edge DCs? How would this accelerate edge deployment?

Edge data centres are data centres that are located close to the end-user to allow for faster data access and processing. They are often used to reduce latency and improve the performance of applications, services, and networks. Edge data centres are typically placed at the "edge" of the network, such as in retail stores, hospitals, or other locations where data needs to be processed quickly and efficiently.

Models of micro edge data centres would accelerate edge deployment by providing a cost-effective, small-scale infrastructure that can be deployed quickly and with minimal disruption. These micro edge data centres can provide a low-latency, high-performance edge platform that can be quickly deployed to meet the needs of a variety of edge applications.

STL Partners

Do you see edge innovation as dependent on microchip supply chain?

Although microchips are important to various elements of the edge value chain (server, device, etc.), the scale of dependency could be debated.

As it relates to server capacity, there is already excess capacity in the network today available for edge computing applications. Telco PoPs and uCPEs are two examples of areas where operators are looking to leverage existing network capacity for edge functions. As server capacity increases over time, more microchips will be deployed, but the timing of this should allow supply to expand.

5G/edge devices are still very much in their infancy, an issue often noted in our conversations with ecosystem players as a key barrier to scaled adoption. Though many innovative devices will leverage microchips to maintain a minimal footprint, once again, the time for these devices to reach any kind of scale should give the microchip supply chain the necessary time to scale with demand.

In summary, though the edge ecosystem leverages microchips throughout the value chain, it is not wholly dependent upon it for innovation and development.

STL Partners

How CSP's can handle the increase in Opex when moving into cloud, when Opex is one of the biggest concerns for CSP's?

Telco cloudification in general does not necessarily mean an increase in OPEX. Granted, as telcos start sourcing their VNFs and CNFs as a service from the hyperscalers, there are increased OPEX costs which can cause issues as the deployments scale if the telco is charged per SIM connected or according to traffic volume. This being said, these increases are offset by reductions elsewhere, such as the decrease in direct OPEX from operating these functions in a telco's own facility, using its own staff, etc. Work must be done to develop the charging model in order to capture these savings, though a large portion of potential OPEX reduction comes from automation, which cloudification helps to drive.

STL Partners

So how do you get your workforce with the right skills for market penetration with the technology moving so fast?

Although the technology is moving fast, this market is still extremely nascent and the ecosystem is still forming. To better ready your workforce, CSPs should be in close discussion with the end customer to gain an understanding of their pain points and priorities, mapping these to the internal skills and capabilities of the CSP to find the gaps where investment should be made. This allows the CSP to understand the fundamentals of the ecosystem as it is being formed, rather than jumping to understand the dynamics of an infantile market.

Verizon has released a press release on Sept 12 this year where they have announced that "Verizon deploys more than 8,000 vRAN Cell Sites rapidly marching towards their goal of 20,000". How do I reconcile this with your research?

Our research on the growth of edge capacity does not include vRAN sites. Our research looked exclusively at third party enterprise applications and distributed core network functions. In order to get a better view on the state of vRAN and O-RAN deployments, our Telco Cloud Tracker tracks such deployments in detail.

STL Partners

A question for Joanna Newman - what partnerships has Vodafone made already with hyperscalers, how does Vodafone pick which company to partner with in this space, and is it planning to further expand its collaboration efforts with hyperscalers? Is Vodafone at all concerned about the role

hyperscalers play in their relationships with telcos in general, and whether Vodafone's business could be affected negatively because of these partnerships?

Vodafone has partnered with AWS to launch Multi-access Edge Compute (MEC) services delivered with AWS Wavelength in the UK and Germany. This is sometimes referred to as Public MEC. Vodafone has also partnered with Microsoft to provide Dedicated MEC, usually on premise. These partnerships were years in the making, undergoing multiple assessments for best fit for our respective customers, and are accompanied by our own stack in Italy. We use research from STL, and others, to monitor if these partnerships are meeting the needs of our customers.

Vodafone

If you have any questions intended for Red Hat, please contact: tim.otto@stlpartners.com