

In partnership with:



HOW TELCOS CAN LEVERAGE HYPERSCALERS FOR PRIVATE NETWORKS AND EDGE

Webinar: Questions and Answers

How telcos can leverage hyperscalers for private networks and edge: Q&A

*This document outlines the questions and answers received from the STL Partners and Red Hat webinar on Wednesday 19th April 2023. **You can watch the recording of the session, and also access the slides, using the link [here](#).** In this document, we seek to address the questions raised in the webinar that we were unable to address in the time available.*

If you have any questions not addressed in the webinar or this Q&A document, or want to hear more about our latest research or from our panellists, please contact:

- Miran Gilmore, Senior Consultant – miran.gilmore@stlpartners.com
 - Duncan Curry – dcurry@redhat.com
-

1. **Do you see value in having a ready to deploy solution to make it easier for customers to start a PoC and then providing the flexibility if customer want to scale? It just seems that customer do not know what they want at this point and a ready to deploy solution would help them figure it out.**

STL Partners: Yes, a ready to deploy solution would help speed up adoption, with the possibility of scaling the solution afterwards. The difficulty is that the range of deployment options relies in large part on the specific use case requirements and defining this broad-appeal 'killer' use case is (as has come out of the webinar discussion) perhaps one of the key challenges of scaling the on-prem edge/private networks market to the long-tail.

There is certainly an element of chicken-and-egg when it comes to educating enterprise customers on the value of the technology to find these use cases: while solutions need to come from a real customer business need and be co-developed with them, the lack of enterprise awareness of the capabilities of edge or private networks (not helped by the fragmented and complex ecosystem of options available) can stunt progress.

The key could reside with the developers of edge applications. Engaging with developers to specifically design around the unique capabilities of edge and private networks will foster the innovation needed to find those 'killer' use cases, as developers drive towards real business needs of the verticals they serve. From there, a ready-to-deploy solution could emerge with the potential to scale. You can [read more](#) on how telcos can effectively engage developers in edge, though it must be said that this is one of the key strengths of hyperscalers today and is a major driver of telco-hyperscaler partnerships.

2. **What can be the use cases of deploying kubernetes (k3) on Edge?**

STL Partners: K3s is a lightweight version of Kubernetes created by Rancher (acquired by SUSE) that is designed to work in resource-constrained environments and to run on very small devices. A recent announcement with [Home Depot](#) shows the speed at which a rollout of such a system can happen. In general, lightweight, single-node distributions of kubernetes like this have the advantage of being fast to deploy as pre-packaged units with minimal configuration and setup, simplified management once installed and ruggedised reliability, making them ideal for edge use cases with many thousands of locations and without extensive specialised IT teams within the enterprise, such as distribution and logistics (as in the Home Depot example) or in the retail and transport verticals.

3. **On the On-Premise / Far Network Edge, what are the needs of the Telcos and Application Vendors on Application Acceleration and SmartNICs and Embedded Network Access Platforms on DPUs/ GPUs etc., and how are they merging their focus on this area?**

Jake Hughes (Verizon): On the cusp of the fourth industrial revolution, the requirements for application acceleration are ever-changing; that being sad, Telcos need to focus on two things, or at least one of two things: finding partners that write best in class code and/or partners that build OT devices that embed said code. Leveraging those two guiding principles will guide the merging of focus.

Deep Mukherjee (Telstra): Customers are evaluating far network edge options for multiple options like efficient data processing and analytics. GPUs at embedded network edge can provide hardware acceleration for data processing tasks like real time analytics, AI inference and machine learning. There is an increased focus on providing developer friendly services. We are also evaluating partnerships to provide best in class service alongside our intelligent networking solutions to solve some of these industry challenges. There is also an opportunity for Telcos to partner with the application vendors to improve application and network performances. We are actively looking at how we can offload network related tasks from the CPU, reducing processing overhead and improving application response times.

Also see response to question 5 below

4. **Back to the performance optimization of the edge that Telcos need to optimize the edge in collaboration with the cloud and with edge user proximity aka "edge-native" way apps be it for manufacturing or logistics, what are the role of smartNICs for Telcos to include at the edge?**

See response to question 5 below

5. **SmartNICs and DPUs are being addressed by AMD, ARM and Intel in the form of RAN Boost and FlexRAN and Ericsson and Nokia re adopting or developing native APIs. What is the impact of SmartNICs on the far edge and enterprise?**

STL Partners: As you correctly point out, SmartNICs, or specialised processors and silicon (like FPGAs and ASICs that are programmable), are already being integrated by a number of vendors and OEMs into access networks to handle core networking tasks at sites near the edge of the network. This interacts with the enterprise on-premise edge (or "far edge") when the edge of telcos' networks are on an enterprise premise, and especially when on-premise edge computing is deployed with a private network in what we call the **Private 5G edge**. Here, shared hardware and MEC platform can bring benefits to both RAN/core VNFs (for the private network) and the enterprise MEC applications, some of which will require the increased performance, latency and power consumption of optimised silicon.

However, there are trade-offs (crucially with regards to cost) in choosing to use these smartNICs over CPUs or even GPUs that make decisions for enterprises to go down this route for their on-premise apps very use-case dependant. The table below gives an overview of the various capabilities and trade-offs. Perhaps the largest opportunity is in OT use cases in industries where these perform mission-critical tasks that run on specialist hardware today – e.g. robotics or computer vision in manufacturing industries. Capturing this OT opportunity has its own challenges but may well be the key to kicking off the on-prem edge/private network opportunity (see question 8 below).

| | CPU | GPU | FPGA | ASICs |
|-----------------------|--------|--------|-------|-------|
| Computing performance | Low | Medium | High | High |
| Power consumption | Medium | High | Low | Low |
| Cost | Low | Medium | High | High |
| Size | Small | Small | Large | Small |
| Latency | High | Medium | Low | Low |

Source: STL Partners

6. **Would AWS consider a similar model to MSFT (i.e. HCI) where their cloud services SW would run independent of the HW? For example, could AWS run on HP, IBM, Levovo, etc on customer prem?**

Zarn Kucel (AWS): AWS continues to remain customer focused as 90%+ is based on direct feedback from our customers. We are unable to comment on our competitor's approach to the market as we remain customer obsessed and not competitor focused. Regarding using AWS software on 3rd party hardware, we're not in a position to share any further detail at this time. However, where/when our customers do need to run workloads on 3P hardware like COTS and servers from the likes of Dell/IBM/Lenovo/etc. AWS provides services like our Enterprise Kubernetes Service - Anywhere (EKS-A) (<https://anywhere.eks.amazonaws.com/>) services like FreeRTOS a real-time OS for microcontrollers and IoT devices, and other services – [learn more here](#).

7. **The benefit of customer edge compute is to run cloud services and data analytics/inference locally. The benefit of edge telco services is to run network/security functions (sometimes Priv 5G) at breakout and orchestrated with edge apps. This provides the platform for industrial/OT automation. What is the best example of a telco + hyperscaler partnership solution delivering this capability?**

Jake Hughes (Verizon): In the manufacturing and transportation space, Ports have been an early adopter of edge functionality. International Trading Agreements and Regulation call for enhanced network security while concurrently delivering a seamless converged network as a service experience. On the flip side of that coin, we need to power the applications that run porting logistics over the hyperscaler's computing infrastructure. Telcos and Hyperscalers together power the converged network service functionality through data sovereignty and edge compute application performance.

8. **What are hyperscalers/Red Hat doing to enable the OT ecosystem? Cloud has captured Enterprise IT. Edge will fire when OT starts to run on Edge.**

Zarn Kucel (AWS): AWS extends the AWS cloud with new services and infrastructure options to bring applications closer to where they are needed. AWS Hybrid Edge services (such as AWS Wavelength, AWS Local Zones, AWS Outposts, and AWS Snow Family) empower builders to build solutions to serve both IT and OT needs. For example, AWS, partners, and CSPs are working together to build solutions that transform and digitize industries. These include: 5G-enabled smart fulfillment center with partners RiverPlus and Quicktron, continuous ML optimization for smart manufacturing with partners Intel and Megh, and Private 5G-enabled digital twin for construction with partner TCTS.

Some customers will have the need to run OT workloads over Private Networks. To help service this need, AWS announced (02'23) partnerships with Leading CSPs (including Deutsche Telekom, KDDI, Orange, T-Mobile, and Telefonica Tech) to Launch the **'Integrated Private Wireless on AWS'** Program, making it easy for enterprise customers to discover private wireless offerings. Using the Integrated Private Wireless on AWS portal, customers can explore private wireless offerings from participating CSPs, browsing by industry or use case. From the portal, customers can contact the telco of choice, who will design, deliver, operate, and support the private wireless solution. The program combines private 4G and 5G wireless technologies from leading CSPs with AWS services across AWS Regions, AWS Local Zones, AWS Outposts, and AWS Snow Family. The program will also leverage AWS's dynamic community of more than 100,000 partners from over 150 countries, offering additional validated solutions that run on AWS.

STL Partners: Enterprise IT teams have indeed in large part already been swayed over to public and private cloud, but OT has typically remained on-premises, with OT teams generally more sceptical. A variety of factors contribute to this: larger teams than IT are more reluctant to do what can be seen as 'outsourcing' and large IP value is attributed to operational technology that is often considered a crucial differentiating factor of these organisations.

The winds of change may already be underway in overcoming these distinctions, with increasing **IT/OT convergence**, but telcos and their partners wishing to have credibility with OT teams need to directly address specific challenges for OT to adopt edge computing: vertical specialisation can either be bought or built internally and partnering with **industry specialist SIs and ISVs** will be more successful than horizontal product-led strategies at winning over OT.

9. **What kind of solutions are Telcos using to manage a hybrid environment? For instance Azure stack edge/outposts/gdce and their own "edge" custom infra, in terms of application lifecycle, network integration , security etc.**

Jake Hughes (Verizon): At Verizon, we use our Edge Service Architecture to effectively create, read, update and delete policy tied to the packet core. Effectively paving the way for enterprise-grade network slicing. Edge Service Architecture components complement many of the compute and storage elements our Hyperscaler partners provide and manage for end clients in a jGTM motion.

Deep Mukherjee (Telstra): This is an interesting question and there is no one size fits all answer to this. Here at Telstra, we want to provide what our customers demand. Most of these requirements are driven by the current deployments, skillset and the type of application to be deployed. We provide

customers with a host of options that fits their requirements. This means providing access to best in class solutions with ASE/outpost tightly integrated with our professional and managed services, adaptive network solutions and security services. At the same time we also aim to provide customized edge infrastructure solutions which are highly scalable, developer friendly with out of box monitoring and management which are cost effective.

STL Partners: The variety of deployment models between telcos and partners (including hyperscalers) make a hybrid cloud environment likely. The benefits and downsides of each are outlined in the table below and more nuance for these various scenarios and the strategic approaches of the hyperscalers and a number of Tier 1 operators is explored in our report – [Telco edge computing: how to partner with the hyperscalers](#)

| Types of edge computing | Company own on-premises compute | Colocation | Private cloud | Public cloud |
|-----------------------------|---|---|---|--|
| What is it? | Company hosts its own compute and storage infrastructure on its own premises | Enterprise's own compute & storage infrastructure hosted in third-party data centre | Third-party data centre provides dedicated compute & storage infrastructure for the enterprise | Third-party (hyperscaler) provides shared compute and storage infrastructure for the enterprise |
| Benefits | <ul style="list-style-type: none"> + Company retains full control over policies and security + Ultra low latency + Data protection & sovereignty | <ul style="list-style-type: none"> + Reduced premises capital and operating costs vs on-prem + Data protection & sovereignty | <ul style="list-style-type: none"> + Lower compute infrastructure capital costs vs Colocation + Data protection & sovereignty + Scalability – dial up and down | <ul style="list-style-type: none"> + Lower compute infrastructure capital costs vs Private cloud + Access to hyperscaler tech + Scalability – dial up and down |
| Downsides | <ul style="list-style-type: none"> - High capital & operating costs - Difficult to future-proof: easy to be out of date - Tough to build in-house skills | <ul style="list-style-type: none"> - Compute and storage infrastructure costs remain high - Future-proofing compute & storage technology & building skills still tough | <ul style="list-style-type: none"> - Compute and storage infrastructure costs still high (dedicated infrastructure) | <ul style="list-style-type: none"> - Reliance on hyperscaler business model & pricing - Data sovereignty risk |
| Applicable use cases | <ul style="list-style-type: none"> • Mission critical operations where need to scale is limited and the need for close, real-time control outweighs cost-savings, e.g. in-hospital patient monitoring | <ul style="list-style-type: none"> • Important applications where on-premises investment is not viable but demand could scale and company wants to retain control over compute, e.g. video analytics, disaster recovery | <ul style="list-style-type: none"> • Applications where scalability and cost management is important but responsiveness (inc. latency) & customisation remain important, e.g. cloud gaming, disaster recovery | <ul style="list-style-type: none"> • Business support where the need to reduce cost and manage scalability is compelling and (perceived) risk of data sovereignty is low, e.g. Chatbots, disaster recovery |

General direction of travel for enterprises

As to question of managing these application lifecycles and network integration, there are lessons to be learned from telcos' own internal cloudification. Here, operators, vendors, SIs and hyperscalers today are already working collaboratively on the verification of individual hardware and software components, with exhaustive testing and validation activities to ensure that each component is compatible with others and optimized to run as expected within the overall stack.

To sustain these benefits, lines of communication are being opened up across the ecosystem to maintain seamless integration as different components evolve in this dynamic environment. This increasingly means adopting good practice learnt from the software development world, such as GitOps methodologies and CI/CD pipelines, to ensure automated and programmable configuration tracking and management. We dive into more depth on some of these challenges as specifically encountered in the Private 5G edge in [this report](#).

10. What is Verizon view on the Applications that runs on top of the private networks? What applications are currently in demand from your customers?

Jake Hughes (Verizon): Verizon understands that enterprise customers want to buy/build their own solutions in this emerging marketplace. That being said, Verizon is having success positioning point

solutions that meet specific needs from an operations perspective, in the sports and venue space. There is an increasing demand for solutions that deliver outcomes for COOs and OT Device managers when it comes to in arena logistics and customer experience.

11. **Question for Telstra - what business value would it add to use EDGE as "staging" option yet using cloud for full deployment? Not every application or service is Edge native or Cloud Native - hence the need to refactor the code. Is this additional cost/time truly worth using Edge as a staging point, rather to use Edge for full deployment?**

Deep Mukherjee (Telstra): Not every application is suitable for Edge computing and should be moved to edge. We work with our customers to evaluate the value proposition of edge, whether it is our network edge or a customer premise edge deployment. There are multiple reasons why you want to use edge as a staging option:

- **Reduced latency** especially when the application needs to be distributed in nature. Ex V2X use cases. Initially some of the V2X or Multi-player gaming use cases may be driven from the central cloud, however as the usage grows and you need multiple near real-time interactions in parallel you would want a distributed footprint. Here Staging might help with user acceptance in terms of usability and performance parameters alongside utilizing the cloud.
- **Cost optimization:** Deploying an application entirely on the cloud can sometimes incur higher costs due to data transfer, storage, or computing resources. By using Edge as a staging point, you can offload certain processing tasks or data storage to the Edge, reducing the overall cloud resource consumption and cost. Ex doing pre-processing of data at the edge and removing unwanted data to reduce bandwidth cost on cloud.
- **Reliability:** Edge computing allows you to distribute the workload and processing closer to the edge of the network, reducing the reliance on a centralized cloud infrastructure. By utilizing Edge as a staging point, you can verify the performance and stability of your application in a distributed environment before deploying it to the cloud. This helps identify potential bottlenecks or issues and ensures a more reliable and scalable deployment.
- **There are other benefits like enhanced data privacy and security where you can reduce the exposure to sensitive information to the cloud.** We are working with customers in healthcare, finance or government where data privacy and compliance is critical especially where PII is involved.

However, it's important to consider the trade-offs involved. Refactoring the code to work in a distributed environment and maintaining consistency between the Edge staging and cloud deployment can introduce additional complexity, cost, and time. It's crucial to assess whether the specific benefits gained from using Edge as a staging option outweigh the additional effort and resources required for code refactoring and maintaining dual environments. We at Telstra work with our customers to ensure we suggest the most viable option that best covers the business needs.

PARTNERS



Research



Consulting



Events