



5G TELCO CLOUDS: WHERE WE ARE AND WHERE WE ARE HEADED

Questions and Answers

Telco Cloud: Questions and Answers

This document outlines the questions and answers received from the STL Partners & Juniper webinar, **5G Telco Clouds: Where we are and where we are headed**, which was hosted on Wednesday 8th May 2020. You can watch the recording of the session, and also access the slides, from the site [here](#).

If you have any questions not addressed in the webinar or this Q&A document, or want to hear more about our latest research [report](#) on Telco Cloud or our work in the Telco Cloud space, please contact:

- Matt Pooley, Senior Consultant & Telco Cloud Practice Lead – matt.pooley@stlpartners.com

Telco Cloud is a key practice area for us, across both research and consulting, and we would be keen to work together with those exploring and pursuing the opportunities within the sector.

- Q1) When talking about telco cloud, do we have a chicken and egg problem? (i.e. where the chicken is telco cloud and the egg is innovation)

STL: A simple answer would be yes. But many organisations outside of telecoms have managed to solve this problem. We work extensively on this topic – see our Transformation Stream for some particularly good case studies. We have written about a “simplicity” approach to tackling innovation in short cycles: build something for a specific need, but do a bit of extra legwork to ensure you can expand sideways on the next way round.

Juniper: Telco cloud is a pillar of digital transformation which requires an investment in people, processes and technology together to achieve cloud outcomes like agility, flexibility, scale, lower TCO, short life cycles, and dramatically improved TTM. By taking a platform-first approach to design, develop and deploy your telco cloud, the platform becomes a strategic asset from which to drive innovation.

- Q2) Where does real cloud native (something like what CNCF advocates) fit into the telco cloud? NFV is not and never will be cloud native like the clouds the webscale companies run

STL: The answer to this question depends whether you mean that NFV *can never* be cloud native at hyperscale, or whether it *will never* be cloud native at hyperscale. If the former, I disagree: NFV *could* and arguably *should* be cloud native to deal with the scale of traffic that is being predicted over the next decade.

But it is equally true that the Cloud-Native Function is at risk of becoming the next over-hyped thing in networking, that its importance has been exaggerated, and that telcos will not adopt it as many in the industry would hope.

Juniper: Many operators we work with have a vision of standardizing upon a cloud native cloud computing environment. Juniper is heavily involved with both CNTT and CNCF TUG, acting as both co-lead and contributor to several working groups, to help define the reference architectures, use cases, and APIs to facilitate this evolution. The operative word here is 'evolution' as operators have deployed VM-based, NFV environments today requiring strategies to deliver and support hybrid cloud platforms. Juniper's Contrail Cloud supports both OpenStack and Kubernetes today to support both VM- and container-based workloads.

Q3)

Q4) How were the number of deployments counted in the presentation? (Slide 11)

STL: We have built our own database, known as the **NFV Deployment tracker**, that is part of our Telco Cloud research stream. It tracks commercial NFV and SDN deployments by telcos globally and is updated on a quarterly basis.

Q5) There are a lot of "buzz" clouds out there (edge, centralised, distributed/edge, telco, data clouds to mention a few). What are your views on a single cloud that fits all vs. clouds with different characteristics?

STL: We are pragmatic on this question.

Juniper: To capitalize on their fleet of valuable edge real estate, operators need the agility to develop, deploy, and auto scale any new service or application in a consistent way and launch that service across multiple or even simultaneous locations to deliver a ubiquitous experience to their end users. Developing, distributing, and managing multiple, VNF-specific vertical clouds throughout an operator's edge network is simply untenable. What differentiates telco clouds is the requirement to support the operator's infrastructure VNFs in addition to more typical data centre workloads. As a result, the telco cloud architecture needs a common IP fabric to provide access to a diverse set of custom ASICs, merchant silicon, GPU, and other hardware-assisted forwarding platforms in addition to the traditional x86 servers used for data centre workloads.

Q6) Do you think telco operators can work with hyperscale cloud providers, such as AWS and Azure to provide a hybrid telco cloud to the enterprise customers?

Juniper: Yes. More to the point, we believe that it's essential that Telco Operators facilitate the hybrid cloud to their incumbent base of enterprise customers – it's all about choice. No single hyperscaler has a lock on all the needs of every enterprise, and each is evolving at a pace and in a direction that will drive a necessary & relatively frequent recalibration of usage by each and every enterprise – who better positioned than the Telco Operator to broker that ongoing challenge.

- Q7) Is there a general sense of anticipation for specific 5G attributes such as front-end improvements or network slicing or URLLC, mMTC etc. or is it still pretty generic?

STL: Within the industry, yes. But my sense from conversations with most outside of telecoms is simply that "5G is going to be faster than 4G". Network slicing, in particular, is not a commonly-understood term in other industries (although once explained, the concept is generally received with interest).

- Q8) Has any operator implemented a true CI/CD pipeline that involves vendors dropping their "packages" on a continuous basis (or at least better than quarterly?)

Juniper: Yes. Juniper is the prime integrator and blueprint developer for a major Tier-1 provider who is in trial with a set CI/CD pipelines developed to provide end to end automation. We will be providing more detailed information later this year.

- Q9) What is the role of ONAP in service orchestration?

STL: This is what ONAP exists for, and more than one operator has explicitly told us that ONAP has provided the skeleton required allowing them to get their clouds up and running quickly. Frameworks and initiatives like this will be very important in accelerating adoption of MANO. But we cannot speak as to ONAP's relative strengths/weaknesses (as opposed to, say, OSM).

- Q10) Other than partnering with AWS or Azure, are aggregation platforms emerging for Edge (such as MobileEdgeX and Edge Gravity)? Are these being deployed? Any use case worth mentioning?

Juniper: Juniper has announced a major financial investment in **StackPath**. StackPath designs, builds and operates a distributed edge cloud connected with a global high-speed backbone network delivering managed edge cloud services. StackPath offers many edge cloud services including a global CDN service. Our current COVID19 quarantine has dramatically increased bandwidth demands on our operator's transit networks putting a spotlight on the value and benefit of distributing CDN services to the operator's network edge to free up valuable transit capacity and improve end user experience. Juniper and StackPath

are collaborating to deliver a fully integrated, managed telco cloud pod with StackPath's CDN service as an embedded compute workload. Juniper will be providing more public detail on the partnership and the innovative products and services under development for our service provider, enterprise and cloud customers in the coming months.

Q11) Traditional telco infrastructure was used to 5 9's availability whilst cloud infrastructure seems to have 3 9's availability. What are some of the challenges in designing cloud infrastructure for telco service to meet 5 9's availability?

STL: The perceived need to maintain "carrier-grade" standards is one of the key barriers to adoption of telco cloud. Indeed, licenced operators have regulatory obligations regarding their service that do not apply elsewhere – and they take this very seriously.

Juniper: It's about building a strategy that presumes there will be some amount of failure coupled with a thorough understanding of how the set of virtualized functions operate under load, ensuring that there is an available supply of the correct resources needed in the right locations and being able to respond to changes in the environment, in real-time.

Q12) Can you define failings in automation?

Juniper: Believing that simply automating a manual workflow is the answer or believing that building automation without a strategy towards maintaining it is the right answer

Q13) What are your thoughts on the pending acquisition of Affirmed Networks by Microsoft?

STL: Too early to tell for sure what will come of it – if you look at Microsoft's other recent high-profile acquisitions (GitHub, LinkedIn), it is easy to guess where interest might originate but difficult to know exactly what (if any) integration or consolidation is going on in the background. But it is a sure sign that telecoms is on the radar of the big IT players.

Q14) Would you see a world where the hyperscalers move to the edge of the telco edge?

Juniper: Hyperscalers clearly see both the need and opportunity to penetrate the telco edge and we've seen some service providers partner with public cloud providers to outsource their edge cloud architectures. The relationship with public cloud providers delivers an immediate time to market advantage but it is a risky approach given the operator is opening their valuable edge real estate to a potential competitive threat. As mentioned, Juniper is working with StackPath to offer operators an alternative model to monetize their edge position, under their

own control, while achieving the same TTM advantage and immediate cloud capabilities but without risk.

- Q15) With regards to edge cloud, should we talking about numbers of sites, as though geographies of operators are equivalent, or should we be talking about latency requirements and accept that Luxembourg and Brazil deployments may well be different on-site numbers?

STL: Our modelling shows that numbers and locations of edge sites will vary widely by operator – even within the same geography. Edge cloud strategy is driven by a host of very operator-specific factors including existing infrastructure capabilities, network centralisation/distribution programmes, priority use-cases, and more. So compaying “numbers of sites” may not be a helpful metric. But nor is “latency” – as operators will deploy edge capacity to support very different use-cases, with very different latency sensitivities.

Juniper: It’s a valid point. The discussion about the number of sites is really driven by the application and service needs – low latency, high bandwidth, short transmit periods for IOT devices, etc. Applications like 4K/8K streaming and gaming services, connected vehicles, remote surgery, in-vehicle infotainment, AR/VR and mixed reality can only be delivered from a highly distributed edge. In some cases, the ‘edge’ is extremely distributed into the cabinets and ‘street furniture’ meaning the number of locations could be massive. In the end, it is the service and applications that will drive the physical locations and number of sites required by each operator.

- Q16) Do you think that telcos are worried that if they don’t build their telco clouds in time, the hyperscalers will come and do it instead?

STL: Yes, but not worried enough to move faster! The truth is that there is a general lack of understanding of the hyperscale offer in this space. Meanwhile the hyperscalers are talking more and more about how they can support telco cloud infrastructure – they assume that they will win, and historically that assumption has come true!

Juniper: [Covered by question 14, please refer back for the Juniper response.]

PARTNERS



Research



Consulting



Events