



In partnership with:



# WHY IS NOW THE TIME TO RETHINK EDGE ORCHESTRATION?

Webinar: Q&A

FEBRUARY 2022

# Why is now the time to rethink edge orchestration?

This document outlines the questions and answers received from the STL Partners and Nearby Computing webinar, “**Why is now the time to rethink edge orchestration?**”, which was hosted on Wednesday 2<sup>nd</sup> February 2022.

**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 both prior to and during 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 panelists, please contact:

- Dalia Adib, Director, Consulting (STL Partners) – [dalia.adib@stlpartners.com](mailto:dalia.adib@stlpartners.com)
- Ani Keshishyan, Consultant (STL Partners) – [ani.keshishyan@stlpartners.com](mailto:ani.keshishyan@stlpartners.com)
- Reah Jamnadass, Senior Consultant (STL Partners) – [reah.jamnadass@stlpartners.com](mailto:reah.jamnadass@stlpartners.com)
- David Carrera, CTO & Co-Founder (Nearby Computing) – [dcarrera@nearbycomputing.com](mailto:dcarrera@nearbycomputing.com)
- Sankar Venkatraman, Senior VP (Detasad) – [sankar.venkatraman@detasad.com](mailto:sankar.venkatraman@detasad.com)

We have categorised questions based on topic area. Please use the links below to be directed to each section:

- [Edge orchestration and the telco opportunity](#)
- [Orchestration technology](#)
- [Partners and competitors](#)
- [NearbyOne](#)
- [Detasad’s experience](#)

## Edge orchestration and the telco opportunity

1. "Edge" has many different definitions. Where does your definition end?

*STL Partners: At a high level, edge computing brings the capabilities of the cloud closer to the source of data or end user/device using physical compute infrastructure. This infrastructure is more distributed than the central cloud and can sit anywhere between a device and the cloud/internet. Edge cloud is the virtualised infrastructure and business models that sit on top of that compute.*

*However, "edge" is a broad term, and it's important to define which edge (e.g. location, workload) is being discussed in different scenarios. For example, within the broader "edge computing", there is telco edge computing – distributed compute that is owned and operated by telco. That infrastructure can be used to run consumer/enterprise applications and/or network functions.*

*You can read more about how STL Partners defined edge computing [here](#).*

2. Do you see the edge orchestration solution needing to provide access to end-users to manage their own private workloads, in addition to services managed by the CSP?

*STL Partners: One of the core differentiators between an edge orchestration platform and a traditional network orchestration platform is the multi-tenancy aspect, i.e. allowing customers to manage and control certain parts of the solution, not only telco's internal teams. This could include private network workloads, however the level of control the customers want and need depends on their preferences.*

3. Will telcos be on the receiving or issuing end of edge orchestration services?

*STL Partners: This will depend on the business model that the operator is pursuing to enable the use case requiring edge orchestration. If the operator is only providing its network as a service, it is less likely to be on the issuing side of the edge orchestration order. Alternatively, if the telco takes a more front-seat role and manages the end-to-end delivery of the solution, it is more likely to act as an integrator of the different vendors/solutions at play.*

4. From your view which is more prevalent - cloud native (Kubernetes) or older, traditional workloads at the edge?

*STL Partners: Based on the majority of vendors and customers we are speaking to, cloud-native is the default method, because most deployments of edge are greenfield and need containers to ensure workloads are easier to move and orchestrate. There are exceptions, for example, you could say that most edge deployments in telecoms are actually still virtual machine based (i.e. VNFs rather than CNFs).*

## Orchestration technology

5. Hyperscalers don't currently buy into the ETSI model and appear to suggest workarounds to achieve the same results by using GPS and network probing, what are your thoughts on that?

*STL Partners: I think most vendors and telecoms operators appreciate what ETSI is trying to achieve to improve standardisation across different edges, however the challenge with standards bodies is that they can sometimes run slower than the market. Some are seeing that vendors will continue to innovate independent of ETSI, but will converge with the standards over time, assuming they are ready and still relevant.*

## Partners and competitors

6. The discussion is normally focused on telcos but this can lead to walled gardens that don't help to compete with hyperscalers. Do other players e.g. IXPs (working as orchestrators) help to develop an open environment with many players involved?

*STL Partners: They indeed help, however at this stage in the market it is unsurprising that innovation is happening in islands. IXPs (internet exchange points) will help to solve challenges around making it easier to connect to an edge that can serve multiple telecoms networks. Other initiatives, such as federated edge, also help to ensure interoperability between telco edges.*

7. How do platforms like the GSMA Operator Platform align in the edge orchestrator's roadmap?

*STL Partners: The GSMA Operator Platform is not a platform per se (at least that is our understanding), but a forum to enable the ideal operator platform, i.e. allow telcos to move from network operators to platform providers but avoid challenges related to being fragmented as an industry and less global compared to competitors such as hyperscalers. In fact, the GSMA initiative involves input from multiple platform partners to ensure the APIs, standards, etc. that the forum comes up with is relevant and usable by key parties.*

8. How should telcos partner with orchestration and integration software vendors to deliver edge orchestration for customers?

*STL Partners: Telcos will need to establish which business models they want to pursue, and from there figure out who they need to partner with to provide that service. Different business models will have different orchestration requirements e.g. end-to-end solutions will require everything from the deployment of network functions to full stack SLA monitoring – telcos need to understand these requirements and the role of orchestration in order to identify and evaluate potential partners.*

*The orchestration vendor ecosystem can be confusing, but by identifying which capabilities are needed, it will be easier to have productive conversations with potential partners.*

## Nearby Computing

9. We see lots of specialised stacks for edge infrastructure, how likely are they to be swallowed by hyperscalers' edge cloud infrastructure?

*Nearby Computing: Hyperscalers have something that plays against them: they are not using open technologies to manage infrastructure and they are too closely coupled with their internal stacks that can't open any of their technologies. That results in vendor lock-in. What Hyperscalers do well is integrating others' infrastructure and expose it through their APIs, as well as to integrate software stacks from third parties. When it comes to edge computing they're still trying to understand how to cover that "last mile" in terms of infrastructure, as their usual approach is to own all pieces end to end. So they may well stop at the telco edge cloud level, and leave for others the far-edge and on-prem space from an infra point of view.*

10. Please comment on the aspect of harmonisation of the northbound interface of edge orchestrators thereby enabling orchestrator of orchestrators.

*Nearby Computing: ETSI MEC proposes Mm1 interface as the connection between the E2E orchestrator and the MEAO. But the standard falls short in several aspects related to dynamic management of edge sites (i.e. does not cover infra at all, does not consider service placement policies, does not consider edge services that are not really cellular...). From our experience, NB integration with OSS/E2E orchestration systems is done by pre-designing edge services (infra/app/xNF) in the edge platform and then registering artifacts into the OSS/BSS systems that reference those pre-designed solutions. The interaction, therefore, is rather simple – asking to the Edge DO to deploy the edge service. All policies, infra profiles, etc. are associated to that pre-designed solution.*

11. Is [Nearby One] entirely homemade? If not, which bricks do you use?

*Nearby Computing: It is powered by several open technologies. The packaging is home-made, but it is powered by several open projects: Kubernetes (excluding Kubelet and all the pod-related resources, and API server – so it looks and feels like Kubernetes at the engine level, but allows users to manage multi-cloud platforms, including vSphere and other VM platforms for instance); Helm library to manage releases of internal components (what we call Nearby Blocks, that contain multi-cloud deployment descriptors, management policies, infra descriptors...); GraphQL for the API-driven operation of the platform; and Prometheus is extensively operated as well to aggregate services and infrastructure telemetry.*

12. How does Nearby orchestration manage workloads on hyperscalers infrastructure? Does it substitute to their orchestration? How does it evolve with edge computing's evolution?

*Nearby Computing: Nearby One does not replace any resource management technology (i.e. Kubernetes, vSphere, cloud XaaS). What it delivers is over-the-top management of all these pieces (including infrastructure provisioning), as well as mechanisms to express dependencies and*

*propagate configurations across them. For example, describing a service chain that includes: one edge deployment on top of a K8s cluster hosting a video analytics application and one IoT component; one data center component that provides a windows VM on vSphere to control the video analytics areas of interest; one cloud service on AWS Greengrass or Azure IoT hub to process IoT data; and the necessary configurations to propagate across these 3 components the necessary information so that they can find each other as well as to make automatic placement decisions on each tier of the service based on infra SLIs.*

13. What is the role of AI/DL/ML in [Nearby One's] decision process on edge orchestration?

*Nearby Computing: At this point policies are mainly deterministic based on different inputs coming from service components and the infrastructure. They can be driven by external elements that provide influencing on the decisions. At the innovation level, we're working on several projects to use ML to enhance the policies, but most of the requirements we get today from customers require explainable and predictable decisions.*

14. What triggers the edge orchestrator? Is it a device on ultra-far edge or is it a network element from the telco network?

*Nearby Computing: The edge orchestration engine is usually registered in the OSS/E2E orchestrator of the telco. NB integration with OSS/E2E orchestration systems is done by pre-designing edge services (infra/app/xNF) in the edge platform and then registering artifacts into the OSS/BSS systems that reference those pre-designed solutions. The interaction, therefore, is rather simple – asking the Edge DO to deploy the edge service. All policies, infra profiles, etc. are associated to that pre-designed solution.*

15. Do your edge orchestration services run over the public internet, and does this assume pre-arranged business terms for resources available in each of the data centres, near edge and far edge?

*Nearby Computing: The product runs either on-prem or in the cloud. Just requires a K8s service to host it. Management can be done over the top (internet) or using customers' transport network (or a mix of them). It assumes that any managed cloud hosted services that incur in costs, are pre-arranged in terms of credentials and billing details and information. The product is a platform to operate the edge tiers, but it is the customer that operates it according to its needs. In most cases, actual operation is conducted by SIs.*

16. Will the entity using the edge orchestration services be funnelled back analytics data to be able to see the performance of their end-to-end service? Will these be made available to the entity through a virtual NOC?

*Nearby Computing: We use Prometheus (multiple instances in fact) to aggregate infrastructure and service telemetry. This information is consumed by the orchestration engine to enforce management policies, as well as it can be integrated into the NOC of customers' organizations. We support webhooks that can be integrated with NOC services, as well as we provide Alerta as a front end for that information if needed.*

17. In some of the use cases, a private network could also be a solution. How does the Edge Orchestration approach/architecture differ in that case?

*Nearby Computing: Edge orchestration works for public and private networks. In fact, we see more MPN deployments than edge deployments for commercial/public networks. Operators are moving into the MPN space to simplify the management of networks at the edge, and also many SI/SP around the world start operating them because many countries are allocating spectrum for industrial/utilities/public admin use.*

18. How do you see edge orchestration working in harmony with a CSD/DSP's E2E Service Orchestration layer? And what standards and use cases will dictate such an interworking?

*Nearby Computing: ETSI MEC proposes Mm1 interface as the connection between the E2E orchestrator and the MEAO. But the standard falls short in several aspects related to dynamic management of edge sites (i.e. does not cover infra at all, does not consider service placement policies, does not consider edge services that are not really cellular...). From our experience, NB integration with OSS/E2E orchestration systems is done by pre-designing edge services (infra/app/xNF) in the edge platform and then registering artifacts into the OSS/BSS systems that reference those pre-designed solutions. The interaction, therefore, is rather simple – asking to the Edge DO to deploy the edge service. All policies, infra profiles, etc. are associated to that pre-designed solution.*

## End customers/ Detasad

19. Will KSA be investing in Space Technology?

*Detasad: Not in the near future.*

20. Do you see value in simpler all-in-one software applications that link cloud/edge events?

*Detasad: Absolutely. This will be the trend of the future. For eg: any retail store would like to have an encompassing solution including IoT platform integrated with a Software platform (tailored for Retail) for processing of an event. This would reduce the processing time taken for an event and can be very specific to an industry.*

21. Are you seeing scenarios where the end user / end customer is ready to fund edge modernisation needs which benefit them directly?

*Detasad: In my view Edge Modernization should be offered as a service by a CSP/TSP to a customer. The initial cost of setting up an edge is very high and would not benefit a customer with respect to the RoI. Unless the customer is a very large enterprise who can amortize the cost across several business units.*

# PARTNERS



Research



Consulting



Events