

In partnership with:

VOLT
ACTIVE DATA

 PARTNERS

5G ORCHESTRATION: A CRITICAL ENABLER OF NEW REVENUE STREAMS

Webinar: Questions and Answers

5G Orchestration: A Critical Enabler of New Revenue Streams

This document outlines the questions and answers received from the STL Partners and Volt Active Data webinar, **5G as a growth engine: how do telco IT systems need to evolve?** which was hosted on Tuesday 21st June 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 in the webinar as well as questions that we were unable to address in the time available.

The presentation is based on insights and findings from a research programme that included: a survey with 146 respondents from telcos, vendors, and application providers across Europe, Asia Pacific, North America and Latin America. **We have published an updated version of the report on a CIO's Guide to 5G in June 2022, with an extra chapter on orchestration.**

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:

- Andrew Keene, Senior Director of Product Management, VoltDB, akeene@voltdb.com
- Dalia Adib, Director and Edge Practice Lead, STL Partners, dalia.adib@stlpartners.com
- Ani Keshishyan, Consultant, STL Partners, ani.keshishyan@stlpartners.com
- Patrick Montague Jones, Senior Consultant, STL Partners, patrick.montague-jones@stlpartners.com

When fixed network assets are already deployed in stadiums, what is the business case for 5G & slice orchestration for this particular use case?

Andrew Keene, Volt Active Data: Operators will be fighting wireless & Wi-Fi 6 in the enterprise space. My experience of stadiums is that there is variable Wi-Fi signal - wireless signal that is shared with so many people is unreliable. 5G, particularly with dedicated slices, does offer a way to guarantee premium experience for paying users that cannot be guaranteed on current Wi-Fi systems.

Dalia Adib, STL Partners: Investments into fixed network assets will continue to serve the purpose of running underlying IT and operations around the stadium. The use cases that were discussed in the presentation are mobile driven and so need wireless connectivity. Stadium and events teams are continually seeking to improve customer experience, and while Wi-Fi remains a viable option, most hotspots have certain login requirements. This can act as a barrier to fan accessibility and engagement. Furthermore, it can be said that it is easier to deploy use cases like AR headsets onto a public 5G network than a Wi-Fi network. As an application developer, having to integrate their solution into the enterprise Wi-Fi system is a more bureaucratic process. This is down to various factors, such as privacy concerns. 5G, in comparison, has a lower barrier to entry. Depending on the use case and objective, the application will lend itself to either fixed wireless access or 5G.

What do you think operators can learn from hyperscalers as they move into 5G?

Ani Keshishyan, STL Partners: I think operators could learn from the skills, mindsets and practices that hyperscalers possess. For example, hyperscaler's 'software-centric' thinking is proving beneficial as a foundation for their drive towards automated networks. Hyperscalers work in agile, self-organising and cross-functional teams for innovation projects, whereas operators are less flexible in how they organise themselves. These are just a few examples of what operators could borrow from the hyperscalers.

Why is orchestration for 5G that much more complex than what has come before with 4G?

Dalia Adib, STL Partners: I would echo Andrew's presentation that talked about services becoming more dynamic. I think the transition towards a cloud-native, software-based architecture allows telcos to create these new services. In turn, this will create more complexity. For example, dynamic network slicing requires orchestration of different parts of the network in real-time and that puts demand on orchestration systems that were not present in previous generations. It is very different delivering a linear connectivity service compared with delivering a service to an application developer or enterprise that must meet certain, and potentially fluid, service requirements. Operators will have to ensure that all network resources are able to deliver the requested service against key parameters such as latency, bandwidth or other QoS measures. The other aspect is edge computing which, although is a separate technology, enables 5G and other applications in the network. A lot of telecoms operators see the two as complimentary technologies and as telcos seek to take edge services to market. This shift will bring its own edge-specific orchestration needs. These include monitoring the underlying compute resources and moving workloads to edge nodes closer to the end user - Andrew talked about this in his presentation. The focus for orchestration today is network

and service orchestration, and while these will remain important, 5G & its services will open up new opportunities which will require new types of orchestration.

Andrew Keene, Volt Active Data: The networks are getting more complicated with cloud-native monoliths being broken up into many micro-services. This adds the challenge of scaling and adding to current capacity. There are also many more functions appearing on the 5G core, dealing with actions like mobile edge service discovery. The complexity of 5G networks is greater but there are also more parts that require orchestration, beyond connectivity. 4G networks, even with NFV, are far less dynamic so 5G will also be more complex to orchestrate.

How do you view 5G & edge acting as complimentary technologies in the next few years?

Andrew Keene, Volt Active Data: I think workloads at the edge will have to provide significant additional value that cannot be provided from centralised cloud locations. The more limited compute resource at the edge will force developers to consider how to build their application to function within edge environments. 5G and the dynamic nature of its network are far more geared towards moving workloads around and optimisation at the edge. The end user cares about low-latency, response time etc. and so 5G and edge is a far more compelling offer in reaching these advanced capabilities.

Ani Keshishyan, STL Partners: Here at STL, we've been trying to understand how much demand there is for the joint private 5G and edge proposition. It seems that video analytics has gained a lot of traction in the industry as the killer use case to benefit from this technology predominantly in contexts where real-time decisioning is required. It is also interesting to note that advanced video analytics will be used for a number of applications, such as worker safety or industry 4.0 use cases. These include monitoring production lines, waste management and worker safety which all benefit from this unique marriage between 5G and edge.

Do you believe that there is a need for edge-specific orchestration capabilities to be able to succeed at the edge?

Dalia Adib, STL Partners: We have started to touch on this. Some of the capabilities are similar to running an application in the cloud and on any compute premises because you still need to monitor the application throughout its lifecycle and ensure it is supported with the resources it needs. There are some differences with edge as the environment is a fragmented, disparate group of edge nodes rather than a highly centralised environment like the cloud. The way you ensure resources are allocated correctly and enable lifecycle management is different. The other difference is the need to intertwine network and edge orchestration to ensure application demands are met. Orchestration that ensures the network is available and moving traffic in the right direction so that it can be processed at the edge must be paired with resource allocation and other edge node orchestration. Telecoms operators are in a good position to deliver these capabilities. At the edge, there is a need to move applications around to follow the user. In cloud gaming, for example, the workload may need to follow the user if they move locations to guarantee low-latency. This opens up new edge-specific orchestration requirements that are not currently required for cloud applications.

How specifically will real-time orchestration drive better customer experience or internal efficiencies for operators?

Andrew Keene, Volt Active Data: In some examples, decisions in the network will need to be faster to maintain certain user experience. Edge services will have limited compute capacity so it is important that if workloads are moved to new edge locations, there will be enough compute capacity allocated. Some of the use cases, like Industrial IoT, will require quicker response times and the flexibility to spin up more capacity so that the factory floor is not interrupted. The level of real-time orchestration is going to vary between use cases, but for some it will be particularly critical.

Dalia Adib, STL Partners: Real-time orchestration ensures the application is reliable and certain mission-critical applications like Industrial IoT will need these capabilities. If a manufacturer starts to disaggregate the production plant so that the equipment is running the software at the edge and that piece of equipment is mission critical, the edge must be available and reliable. There may be some need for real-time orchestration in gaming when the workload needs to follow the end-user. However, latency in this example may even be tolerable to a certain extent. There is also need for real-time orchestration in AR/VR use cases that are rendered at the edge. These use cases require sub 50ms latency to enable the best customer experience.

Patrick Montague-Jones, STL Partners: From our research, real-time orchestration is absolutely vital when combined with other network functions. For example, if upgraded policy control works in tandem with orchestration, when a fault occurs in the network, policy rules can dictate that a new service or replacement resources can be spun up in the place of a faulty resource. This can ensure that services remain consistent and end-users are not even aware of a fault in the network.

Do you think that operators that have already invested significantly in NFV would have to start again with fully cloud-native functions to achieve a high level of dynamic services?

Andrew Keene, Volt Active Data: Ultimately, there is no doubt that containerised functions are more dynamic than VNFs. When you consider a virtual machine and how long it can take to start and stop, a policy or charging function could take 10/20 minutes. If you are looking for real-time orchestration then you would want to start replacing VNFs with cloud-native functions but you don't need to start there. It is entirely possible to connect your orchestration platform up to VNF managers and spin up new instances of VNFs.

Dalia Adib, STL Partners: I suppose many operators are implementing that cloud-native functionality as they implement the 5G core, so it's a more greenfield environment. There is also a timing element. We saw in the poll that the consensus was that it will be 2-3 years before we see commercially viable edge services. For telecoms operators, these private networks deployments are a bit of a testbed for these types of services. There have been many trials implementing dynamic slicing for private 5G networks in stadiums to prove out the model before implementing this in public 5G networks. The timing may work out in line with where the investments are being made.

What are the considerations for enterprises thinking about either using public infrastructure or setting up their own private infrastructure?

Andrew Keene, Volt Active Data: It depends on who the end-user is. The stadium example is a B2B2C use case. It makes far more sense to target the public network rather than onboarding them onto a private network. It would be challenging to roam Vodafone users, for example, onto a private business network like a football club. If it's a manufacturing campus that doesn't have access to good public infrastructure, private networks may make the most sense.

Dalia Adib, STL Partners: Another element is availability of the public network. Some areas of demand are moving towards private networks because they are trying to plug in a gap of coverage. In this case, slices would be redundant. There are different drivers like data privacy, security and technology/solution maturity that will determine which infrastructure enterprises will opt for.

How are operators approaching orchestration systems transformation?

Ani Keshishyan, STL Partners: In general, our survey showed that the main motivation behind enhancing their orchestration capabilities is to improve overall business operations and agility, rather than to reduce costs. Depending on the operator and its ambitions, the telco might choose to build its own orchestration capabilities or rely on partner offerings. Larger players are more likely to have the resources to adopt a DIY approach while smaller players are more likely to partner with third party technology vendors.

How will 5G orchestration systems differ for public and private deployments? What timeline would you put on enterprises adopting private 5G networks?

Patrick Montague-Jones, STL Partners: Operators will look to build their networks and supporting networks in as scalable fashion as possible. This means when providing a private network for an enterprise, they will look to build this in as standard a way as possible, sharing architecture blueprints across customers (where possible) to reduce time to plan and serve. The underlying components will broadly be similar, but the demands on the systems will differ by enterprise to achieve the desired outcomes. Widespread adoption of private 5G networks will take time, and require demonstration that 5G can provide material improvement in serving today's use cases over existing fixed and wireless (i.e. Wi-Fi and 4G) networks. Moreover, commercial viability of private 5G deployments for enterprises will be proved out in ability to serve new use cases (e.g. automated guided vehicles, advanced predictive maintenance) that will mature over the next couple of years.

With respect to E2E vertical solutions, what is the share of value/revenue for the connectivity piece, according to your research?

Patrick Montague-Jones, STL Partners: Connectivity has long been the lifeblood of operator business, forming the foundation of all services and generating a large portion of their revenue. Connectivity will continue to be a gateway to providing higher value services but its ability to generate revenue on a standalone basis has diminished. Enterprises and consumers alike expect high-speed connectivity as a baseline requirement, and most of the value is generated from the additional services enabled

by that very same connectivity. Our research suggests connectivity will account for 1-10% of end-to-end vertical solution revenue, dependent on the industry and use case(s) being addressed.

How can operators ensure that network QoS is continuously delivered, network capacity is not reached and service denial is avoided?

Andrew Keene, Volt Active Data: 5G standalone provides better QoS management than 4G, in that any given user of a given service can have the requested quality of service set by the Policy Control Function (PCF), similar to PCRF in 4G. In addition, 5G core functions, such as the Session Management Function (SMF) and User Plane Function (UPF), monitor and report actual experienced QoS (or QoE) and can trigger proactive resolution of issues (e.g. orchestration triggers additional more network capacity). Moreover, use of dedicated network slices allows the CSP to better define the service-specific QoS demands of that slice. For example, the stadium service likely needs bandwidth over ultra-low latency. A user in the stadium won't likely notice 50-100ms delay in a response, but they will notice congested slow bandwidth. Many 5G network function vendors report statistics via Kafka streams, which can be ingested and acted on in real time. Others don't and rely on network probes to monitor whatever quality of experience metrics a given service demands. Either way, the orchestration platform needs to act on these network insights to ensure sufficient capacity is provisioned to meet the service-specific QoS expectations.

How will the OSS/BSS play into the 5G models for public networks?

Andrew Keene, Volt Active Data: OSS/BSS still plays a vital role in 5G public networks with some big architectural changes from 4G. Some functions closer to the 5G core (such as charging) will be more impacted than others (such as CRM). The well-known point to point diameter interfaces that interface with 4G network functions will all become http based, with a Service Based Architecture. Functions, such as charging, are expected to be cloud native and therefore broken into independently scalable micro services etc. Furthermore, it could be argued that OSS/BSS will be more critical in 5G. Many new CSP revenue streams will go beyond connectivity and be driven by wider, value-added services. These services, like those mentioned in my presentation, will need OSS/BSS to work with orchestration and the 5G core functions to define, provision and manage these services in a flexible way. Going forward, OSS and BSS will be increasingly intertwined and cloud-native. This move will bring together IT and OT teams and provide opportunities for products and the required underlying network functionality to be developed in a more aligned manner.

For the common data store with a distributed platform, do you see something in addition to UDM/UDR? Or instead of? Can you expand on this a little?

Andrew Keene, Volt Active Data: While UDRs hold the user data, the network operational data can be married with UDR from individual network functions to create contexts that allow the network to perform more efficiently. Thus, the UDR does not contain all stateful information a 5G orchestration platform needs to make decisions for service orchestration. An orchestration platform needs to be able to act on data held in many parts of the 5G network, should be capable of ingesting real time streams as well as caching other real-time decision-critical data.

We have published an updated version copy of the CIO's Guide to 5G in June 2022, with an additional chapter on orchestration. All registrants to the webinar, will receive a copy of this report for free.

Get in touch with our panellists to learn more:

Andrew Keene, VoltDB, akeene@votldb.com

Dalia Adib, STL Partners, dalia.adib@stlpartners.com

Ani Keshishyan, STL Partners, ani.keshishyan@stlpartners.com

Patrick Montague-Jones, STL Partners, patrick.montague-jones@stlpartners.com

PARTNERS



Research



Consulting



Events