



OPERATIONALISING OPEN RAN: DESIGN AND DELIVER

Webinar: Questions and answers

Operationalising Open RAN: Design and deliver

*This document outlines the questions and answers received from the STL Partners webinar, **Operationalising Open RAN: Design and Deliver**, which was hosted on Wednesday 15th June 2022.*

In this document, we seek to address the questions raised in the webinar that we were unable to address in the time available.

***You can also watch the recording of the session, and also access the slides, using the link [here](#).** We have included the following timestamps for the webinar recording:*

- **02:21** for the introduction to our presenters and panellists
- **03:22** for the presentation of the research findings
- **23:57** for the live panel session with:
 - **Timo Jokiaho**, Chief Technologist, Telecom, Media & Entertainment, Red Hat
 - **Sandeep Sharma**, VP and Global Head of portfolio 5G/RAN/ORAN, Tech Mahindra
 - **Dalia Adib**, Director, Consulting, STL Partners

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

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

For more details on the research findings, you can learn more about the [Telco Cloud Tracker](#) by clicking [here](#).

Webinar questions and answers

About STL Partners' Telco Cloud tracker

1. In your telco tracker, what is a "deployment"? For example, you show Verizon with 3. What do deployments comprise?

STL Partners: To be counted in the tracker a deployment has to be a commercial deployment in the live network; we do not include trials. Putting it as plainly as possible, it's:

- live, fully commercial deployments (so not even 'commercial trials', where the solution is powering services that customers are actually paying for, but only on a trial basis).
- a deployment is all the bits that we know about, e.g. the VNF or CNFs involved (e.g. a mobile packet core) and / or some of the other components of the stack, such as the hardware, virtualisation platform, orchestration, etc.
- if the same solution is being deployed across multiple natcos as part of a single project, we don't count it as multiple deployments, just as a single one.
- however, if the same or similar solutions are commissioned by an operator at different times and implemented as separate projects (e.g. different Vodafone natcos rolling out similar NSA 5G cores from the same one, two or three vendors at different times), we'd count those as multiple deployments.

2. Slide 9: Who are the 3 deployments in Middle East ?

STL Partners: The three Open RAN deployments in the Middle East in our tracker are all Open RAN deployments. Two are by Etisalat in the UAE and Afghanistan respectively with a third deployment by Zain in Kuwait.

3. When you mention deployments, perhaps I missed it, do you refer to live deployments carrying traffic?

STL Partners: Refer to the answer to question 1.

4. Do you have a vRAN and/or ORAN vendor breakdown and number of deployments?

STL Partners: Our [Telco cloud tracker](#) does have granularity on the vendors behind these deployments, with detail down to the sub-components of all deployments where information is available. The individual sub-components tracked (along with the vendors providing these) include Baseband units, Central units (CU), Distributed units (DU) and RAN intelligent Controllers (RIC). This level of detail of course also extends beyond Open RAN to the full range of Telco Cloud deployments tracked by our service.

If you are interested in finding out more about getting access to the full tracker and the accompanying analysis, feel free to get in touch with [Dalia Adib](#).

Drivers / opportunities

1. What do you see as the primary drivers for Operators to move to Open RAN.. we talk about innovation capability, removing vendor lock in, cost advantages, etc. Are these all equally distributed or do we see a primary driver?

Timo (Red Hat): I think the primary driver still is the fundamental 3GPP functional split specification. After all, it is a required spec and 3GPP is a standardisation authority in the mobile network. Fairly shortly, all RAN vendors will just need to follow the 3GPP disaggregated, functional split model.

As to why 3GPP have gone down this route, all the benefits mentioned follow from the many deployment models of RAN with this 3 entity split: you can have everything on the cell site or you can centralise 1 or 2 of the entities. It also allows very efficient small cells to be built and deployed in buildings, for example. This variety of deployments can respond to different environments and connectivity characteristics and so is very valuable to operators.

Sandeep (Tech Mahindra): There are three related drivers: the first is supply chain disruption. Operators want more diversity in terms of different ecosystem players to select the best in breed solution that best fits their whole network requirements.

The second, to summarise Timo's point, is the flexibility of deployment options, which allows you to achieve better efficiencies from the assets you deploy into the network.

Thirdly, the newer edge environment that Open RAN unlocks brings intelligence to the network out at the edge. You can do analytics at the edge sites, with many use cases that don't require you to load the whole system back to the core. You can drive intelligence that makes use of data that is already there at the edge of the network. With 5G coming in, this type of intelligence is more and more relevant as we have so many bands coming in and new users with their own expectations from the access network – this will rely on an open and virtualised RAN.

2. For greenfield networks Open RAN could be a good choice, but what about brownfield networks?

STL Partners: The earliest Open RAN deployments were certainly driven by pure Open RAN players (DISH, Rakuten etc) and a significant part of even the most recent deployments are greenfield deployments from more established operators like Verizon, Vodafone and MTN. However, we are already seeing the economics and operational benefits of Open RAN for brownfield deployments starting to be proven and expect the majority of operators to follow once these become established.

Last year, STL Partners **projected that brownfield operators** will switch their buying pattern from legacy RAN to mostly Open RAN by 2030, following four possible pathways:

- Committed Open RAN operators refreshing more restricted initial vRAN deployments to Open-RAN after 5G SA rollout

- Those late with 5G adopting Open-RAN from the outset of their 5G deployments in 2025-27
- Operators that have already made significant early progress with 5G launch and deployed first wave of 5G, working closely with existing vendor(s) to evolve network : Slow to adopt Open RAN from 2026 onwards
- Those with mature 4G/5G and fixed broadbands whose main interest is in Open RAN supporting advanced industrial services and private networks.

3. Can you speak to whether new silicon (ARM/RISC V versus x86) will improve the economics or acceleration of OpenRAN

Sandeep (Tech Mahindra): For any workload that has to be loaded onto cloud infrastructure, there are a certain set of specifications that need to be in place. RAN, being a more critical component that is closer to end consumers, there are particularly stringent requirements that need to be in place. For example, with the amount of data going from RUs to DUs, there needs to be a lot of compute power, and that processing needs to be done at a faster pace. Bigger compute, faster processing and accurate timing and synchronisation needs to be in place on the underlying infrastructure over and above what might be typically available from cloud infrastructure.

Acceleration is also an important aspect, but will not necessarily be required in every scenario – usually only when multiple higher bandwidth applications are processed simultaneously in the RAN. There are some chip vendors leading in this space and a number of new ones coming in .. depending on the type of acceleration you are looking at (in-line or look-aside), smartNICs will play their important role in the ecosystem. A few other key considerations – real-time kernel or having acceleration cards need to be considered when deciding on the underlying infrastructure to run Open RAN workloads.

Timo (Red Hat): All these acceleration and timing or synchronisation considerations apply to distributed units (DUs) specifically. The cloud platform plays a crucial role in providing these functionalities

Technical challenges

1. You mentioned integration complexity. Is there a view on whether the integration costs + product costs are less than the equivalent in traditional RAN ? i.e. how much of a cost is the integration & testing costs?

Sandeep (Tech Mahindra): I think rather than cost we should look at TCO (Total Cost of Ownership) over the lifecycle of the network. The innovations that are coming in will impact the operational efficiencies of a network over a certain period of time, and viewed over time

TCO improvements are certainly proven in Open RAN. The cost optimisation in terms of individual units will come later when we have larger deployments and a more mature stack and ecosystem in place.

Timo (Red Hat): I believe that while in the short term, certain TCO benefits will not be there, it is slightly less relevant. Service providers used to prioritise TCO over everything else, but that is not the case anymore. It used to be the most important thing, and it is of course still important – it will always be important – but it is not the top priority anymore.

Over time, what will help on the TCO side of things is that there will be a set of components that have been interoperability tested and guaranteed to work before deployment, so that when you do the deployment with these specific set of components it won't be that complex anymore and so therefore won't increase the initial integration cost. The set of components that have been interoperability tested with each other in advance will continue to increase until this TCO question is settled.

2. SMO consist of OSS or OSS consist of SMO Part ?

STL Partners: Service Management and Orchestration (SMO) is an automation framework defined by the O-RAN alliance that looks to improve network performance, enhance customer experience and minimize RAN operational costs. This level of optimisation of the RAN can only work at scale using automation to reduce complexity, and so components like AI-enabled RAN Intelligent Controllers are new and crucial to making SMO work.

Functionally, it is defined as a component of OSS and is defined to cover existing functions within OSS, such as service orchestration, inventory/topology and policy control. However, the potential of SMO goes beyond just these functions, with the aim being to automate this RAN optimisation at scale so will touch BSS and other functions. In this way, many see SMO as an evolution of the Self-Organising Network (SON).

3. What about orchestration and ORAN? Do you see initially manual integration or SMO like orchestration usage from Day 1?

Timo (Red Hat): This is likely to be one of the greatest benefits of Open RAN. When you use a similar cloud platform to host DU/CU workloads as those used on the core network (for example Red Hat's containerised OpenShift, which is used for both the core and the RAN), it makes end-to-end automation and management much easier than it has previously been. The industry is obviously going towards automatic provisioning and installation and zero-touch provisioning of the RAN so this is a driver many are waiting on and it contributes very nicely to the final TCO numbers.

If we talk about Service Management and Orchestration (SMO), it will be available eventually, but it's going to be a phased approach - SMO might not be there in the earliest deployments but will follow as O-RAN interfaces become fully defined.

Sandeep (Tech Mahindra): SMO is definitely an important aspect of Open RAN, and in some form or another operators are already using an SMO framework today. Zero touch provisioning is a key part of that and it is the containerisation of RAN workloads that is driving towards this.

4. What about the issue of large scale telemetry required to move data from CUs and DUs for assurance? Is the challenge going to make assurance a second step for ORAN deployment?

STL Partners This is certainly a challenge and also one of the key issues in working towards the RAN Intelligent Controller (RIC) as part of Open RAN. This is the part of Open RAN that will draw out the benefits of automated orchestration as mentioned in the previous question and is part of a secondary phase for Open RAN deployment. The RIC differentiates between near-real-time and non-real-time optimisations, with separate telemetry data streams between DU and RIC, and between CU and RIC, dependent on this.

Short-term, before the RIC standards are finalised and tested, deployments will require a work-around in particular the real-time requirements brought about by the RIC as the CU, DU and RIC may be physically further away from each other. In a traditional deployment, this is less problematic because of the physical proximity: data just has to travel from the top of the tower to the BBU at the bottom. But similar low latencies are intended to be assured in open RAN by distributing the DU close to the RU - either at the foot of each tower or in far-edge sites serving only a handful of cell sites. Ultimately, the goal is to be able to centralise the DU function a lot more.

5. Is it possible to manage the transmission part of the network with the management of fibre or microwave radios?

STL Partners: In reference to backhaul, yes, this should work – just as fibre and microwave links are used in traditional RAN. If it means fronthaul (DU to RU) and midhaul (DU to CU), then yes as well. One of the purposes of open RAN is to enable fronthaul and midhaul to be delivered via Ethernet or IP links (over fibre or microwave), as opposed to via optical wavelengths over dedicated fibre from the BU to the RU as in a traditional RAN.

6. I've reviewed a few Open RAN SI RFPs and I've found them to be weak especially in security, troubleshooting tool requirements and documentation. After years of getting products with these essentially built in, will telcos be surprised and disappointed when their integration efforts result in "immature" and incomplete "solutions"?

Sandeep (Tech Mahindra): In terms of security, Open RAN networks are more observable and have in-built flexibility to take care of the different security mechanisms that are required. From the standards perspective, the O-RAN alliance formed a security-focussed group two years ago and so are actively addressing this.

Open RAN is really an amalgamation of the IT world and the telco world, so we can use the best security practices already available in the network and use them to deploy on IT infrastructure with its own set of best practices. Compared to the traditional way of deploying, it's a way of managing the security that can be describes as "single-handed". In this situation,

operators have a lot more control over what type of security mechanisms they want to run on top of the cloud infrastructure and on top of the end-to-end infrastructure, getting the best out of it. The standards side of it will get there eventually, but already today operators are managing security flexibly in this cloud-native virtualised architecture.

Timo (Red Hat): As Sandeep has said, the O-RAN alliance is doing a lot of work on security: securing the communication between the entities etc. As a platform vendor, we start from basics on the security by securing the supply chain of open-source components used to build the platforms. We use SELinux as a base of the whole thing and harden the platform to be as secure as possible and guarantee that base.

The other view we have as Red Hat is that the security concerns are likely easier to manage in this disaggregated model because the software is divided into different entities – it's not just one monolithic piece of software doing everything. If there is a security breach in any one of these units, it's easier to fix because you only need to fix it in one place. The reality of it is that security breaches are easier to fix because you don't need to touch the whole thing, you just touch where the breach happens.

Sandeep (Tech Mahindra): This highlights the point that there are still many security loopholes even in the traditional way of deploying networks. This is a responsibility we have as an ecosystem together, regardless of it being Open or traditional RAN, working to make these deployments more secure and agile.

7. What technical difficulties have you seen operators currently facing during implementation and how do they cope with that?

STL Partners: Many of the benefits of an Open RAN implementation are indeed sources of challenges in the initial deployment. For example, having an open multi-vendor solution with disaggregation has the benefit of easy replacement of any one part of the solution, faster roll out and a dynamic innovative ecosystem with access to the latest features. However, ensuring that applications are able to take full advantage of new features and capabilities in the hardware is a key challenge and is dependent on backwards compatibility. Operators, vendors, SIs and platform providers today are already working collaboratively on the verification of every individual hardware and software component, 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. The creation of an increasing number of these blueprints is likely to make the ecosystem work.

To sustain these benefits, lines of communication need to be opened up across the ecosystem to maintain seamless integration as different components evolve in this dynamic environment. Here, we have already seen evidence of Open RAN implementations 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.

Other challenges

8. What about lack of mature standards?

Sandeep (Tech Mahindra): Every technology cycle has a standards evolution phase. This is true of 4G/5G as it is for O-RAN. In O-RAN, standards started by identifying certain gaps in the previous approaches to work on and then build from there, so there is clearly a process of maturing of standards happening in the O-RAN alliance. Today's standards are mature enough to deploy a commercial network, it's just that, day by day, there are more and more expectations on the O-RAN alliance, but it is evolving. This iterative process ensures that the needs that are actually there in the market do progressively get addressed.

Timo (Red Hat): I agree with Sandeep but just want to give an example of the maturity level of O-RAN standards. What I consider one of the most important open interfaces is already there: the open fronthaul (interface between the Radio Unit and the Distributed Unit). That is already a solid standard. It is referenceable and people can claim compliance to open fronthaul. There are several companies that have already done that and proved that it works.

When you have a solid open fronthaul implementation, it already enables multivendor RAN because you can pick and choose different Radio Units from different vendors and DU/CU pairs from other vendors. Other interfaces are still evolving but will build on this key standard.

9. Is there any xHaul open standard for the OpenRAN?

STL Partners: Yes, there are standards for fronthaul and midhaul. For example, the open fronthaul standard developed by the TIP OpenRAN group mentioned by Timo. O-RAN alliance also has fronthaul standards and open transport.

10. The success of open RAN is based on economy of scale. Today the cost of an Open RAN site depends mainly on the cost of the O-RRUs. Current providers of O-RRUs cannot compete in price or performance with traditional vendors. For prices to fall, demand needs to increase and that is something that is not happening.... What do you think about this?

Timo (Red Hat): I agree that economies of scale are most important with Radio Units, as these are the highest volume unit in the whole system. Therefore, the open fronthaul standards are key as they are only now enabling this ecosystem of open radio units to be built.

Sandeep (Tech Mahindra): Demand is also related to supply and the more vendors we have joining this movement, the more you will see the benefits of these economies of scale. Today, we are seeing many new vendors coming into this Radio space, increasing diversity in the supply chain that promotes both economies of scale and innovation.

11. What are the top skills we should look for when hiring a team that can deliver integrated Open RAN projects?

STL Partners: You certainly need the existing RAN engineering skillsets. But you also need a raft of cloud and software expertise, as open RAN is essentially a cloud-native network function. And you need the two sides of the equation to work harmoniously together in a single team. So strong leadership from the top and, ideally, some heavyweight experience with open RAN and / or delivering mission-critical VNF / CNF deployments in the recent past.

PARTNERS



Research



Consulting



Events