7 LESSONS FROM OUR NETWORK SLICING JOURNEY

How to create a disaggregated 5G network to deliver network-as-a-service

Capgemini engineering

CONTENTS

INTRODUCTION

As part of our extensive work supporting the setup and delivery of Network-as-a-Service offers, Capgemini recently worked with a major telco, a cloud company, and multiple network technology and software providers from radio to core, to build a proof of concept (PoC) for a disaggregated 5G network. The goal of the project was to prove that Network-as-a-Service (NaaS) was possible using multiple existing vendors and current standards. The project successfully launched, managed and closed multiple network slices, designed around real customer use cases.

This whitepaper discusses the rationale for moving to NaaS, the trends that make this critical for network operators, the PoC itself, and the lessons we learned along the way.

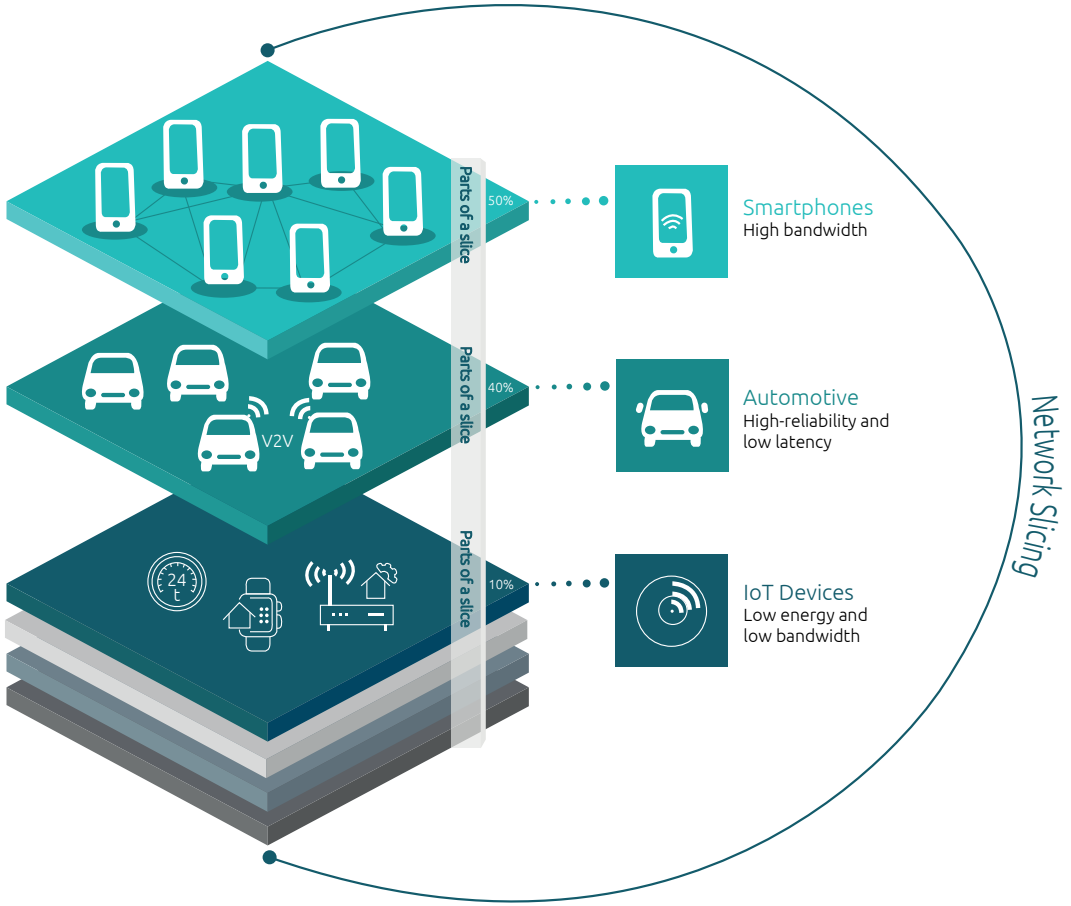

> Capgemini recently worked with a major telco, cloud company, and multiple network technology and software providers from radio to core, to build a proof of concept for a disaggregated 5G network."

WHY NETWORK-AS-A-SERVICE?

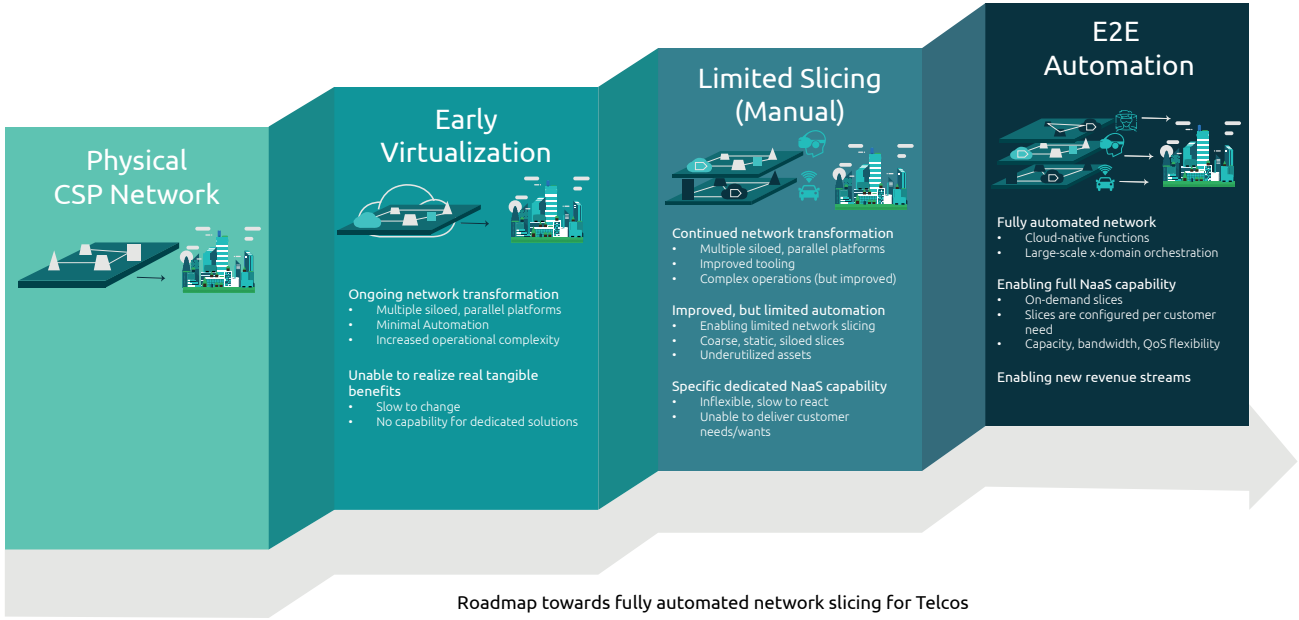

Roadmap towards fully automated network slicing for Telcos

An organisation's traditional network setup involves lots of physical kit - switches, routers, etc. – which must be built and configured. That requires upfront capital to set up, plus ongoing monitoring and upgrades. Such physical networks are hard to evolve, making it difficult to keep up with emerging data-intensive uses – like high-res video streaming, AR/VR, Industrial IoT, and autonomous vehicles – which create more network traffic and place ever more complex demands on the network.

NaaS is a model where the infrastructure and management are delivered by the network operator, as part of an ongoing contract. It is largely software-defined, with most functions running in the cloud, making it easy to scale capacity up and down, and quickly deploy new use cases.

NaaS is particularly interesting, now that 5G networks are here, since these offer high bandwidth that can support new high data use cases. They can also offer network slicing, where a section of the network can be quickly specified and deployed to support new use cases as and when needed without the cost and time of physically deploying a network. This can be done over the top of the existing network, without needing to modify it, and thus insulating the existing network and its customers from disruption.

For example, imagine a company that runs multiple shopping centres wants to add a new video analytics-based security system. The company simply logs into its network provider's self-service portal, defines the parameters it needs for that use case (eg. bandwidth, latency, and jitter) and activates the slice, ready to use across all its venues.

That opens many new use cases that enterprise customers can buy one-by-one as services, or to cover temporary needs such as providing additional data-sharing capabilities during large annual sporting events, without any onsite infrastructure deployment. That, in turn, lets operators monetize their investments in 5G by selling network services in use-case-focused packages, which are more appealing to customers, and easier to get budget sign-off for.

However, setting up these enterprise networks has proven technically challenging. This is, in large part, because many operators are freeing themselves from vendor lock-in by leveraging the disaggregation and virtualization inherent in 5G architecture. That lets them create far more sophisticated and functional networks, and reduce costs. But the price of this freedom is integrating disparate (and not immediately compatible) hardware and software from multiple vendors.

Network Equipment Providers (NEPs) have already demonstrated 5G network PoCs that can deliver NaaS using their own technology. But as yet, there are no successful PoCs that demonstrate how this could be done whilst taking advantage of multi-vendor disaggregated networks which is the likely setup for most future networks. We believe we are the first to do this. In this paper, we will discuss the project, and the lessons it holds for other operators setting up disaggregated 5G private networks with a view to offering network-as-a-service.

BENEFITS TO THE CUSTOMER OF USING NAAS

Before we come on to our PoC, we will quickly address some of the advantages to your customers of NaaS. These are compelling customer benefits, which make NaaS more attractive than traditional network offers, for those that can get the setups right to deliver it well.

Flexibility

NaaS allows network segmentation according to user requirements, such as latency, bandwidth and speed. This rational use of resources enables rapid, seamless resource reallocation, according to the requirements of the enterprise. It also provides flexibility as needs change, by providing an 'elasticity' of service. For example, networks can scale up or down to cope with one-off customer needs (like special events). In the network slice model, new slices can be added or removed as the customer's business evolves, without new infrastructure.

A smooth transition

This flexibility can also smooth the transition to NaaS from a traditional 'owned' enterprise infrastructure. Enterprises can define new business and functionality needs and get them set up quickly without disrupting existing services. In this way, they can gradually move towards NaaS without any major disruption.

Decreased costs and resource requirements

NaaS reduces a range of costs for the customer associated with operating their own network. For example, doing away with the need to purchase, set up and own their network equipment avoids CAPEX. Outsourcing network management and maintenance – which can largely be handled in the cloud by the network provider or its partners - reduces the need for expensive in-house expertise. By moving spending from CAPEX to OPEX at a subscription and/or 'pay as you go' model, NaaS can provide more predictable costs.

Reduced risk

NaaS also reduces risk when introducing and implementing new services. This applies to financial risk – which can be mitigated by knowing exactly how much that new network capability will cost; to service risk, allowing new services to be rolled out independently of existing networks, protecting existing customers from any impact of that service; and to security risk – by assuring the compatibility of new components, thus minimising the likelihood that they will create vulnerabilities within the network.

CHALLENGES OF MAKING NETWORK-AS-A-SERVICE A REALITY

Behind the scenes, setting up the infrastructure to offer NaaS, that is flexible enough to deliver all the above customer benefits, is hard. Three challenges in particular stand out.

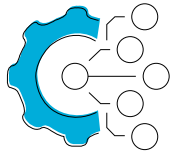

1. **Technology expertise**
   Making the network truly flexible and optimised requires taking advantage of open networks and multi-vendor technologies. Those multi-vendor technologies include physical radio infrastructure, cloud platforms (Azure, AWS, etc), network management systems, software-based services (eg. video conferencing, predictive maintenance, fleet management), and endpoint devices (IoT, robotics, etc.). As networks open and these technologies proliferate, network operators should establish ongoing awareness and partnerships to be able to benefit from the best technologies and create the best as-a-service offers.

2. **Standardisation**
   All these new (and old) technology vendors may adhere to different standards, or interpretations of standards, or no standards at all. And new standards arise and evolve all the time. Network setups, therefore, require standardisation between different technologies to ensure everything properly interoperates. All of this creates a requirement for high-skilled staff who stay current with the rapidly evolving technological and standards landscape.

3. **Alleviating customer concerns around visibility, security, control, and connectivity**
   As enterprises move from physically owned networks to NaaS, there may need to be a mindset shift amongst some customer IT departments around relinquishing control of setups and security, and whether there will be adequate connectivity to deliver services that involve transmitting lots of data back and forth to be run remotely in the cloud. There will be a need to build trust and assurance around these new setups – both through good customer communication and through highly reliable network design.

LESSONS FROM OUR PROOF-OF-CONCEPT

As part of our work supporting companies to deliver NaaS, Capgemini recently delivered an end-to-end PoC of a disaggregated network for a multinational telco.

To date, there have been a number of successful NaaS PoCs. However, all (to our knowledge) have been single-vendor. Since many future NaaS implementations are likely to use a mix of suppliers, we wanted to demonstrate that a disaggregated, multi-vendor network was a viable concept, and learn the lessons of setting one up.

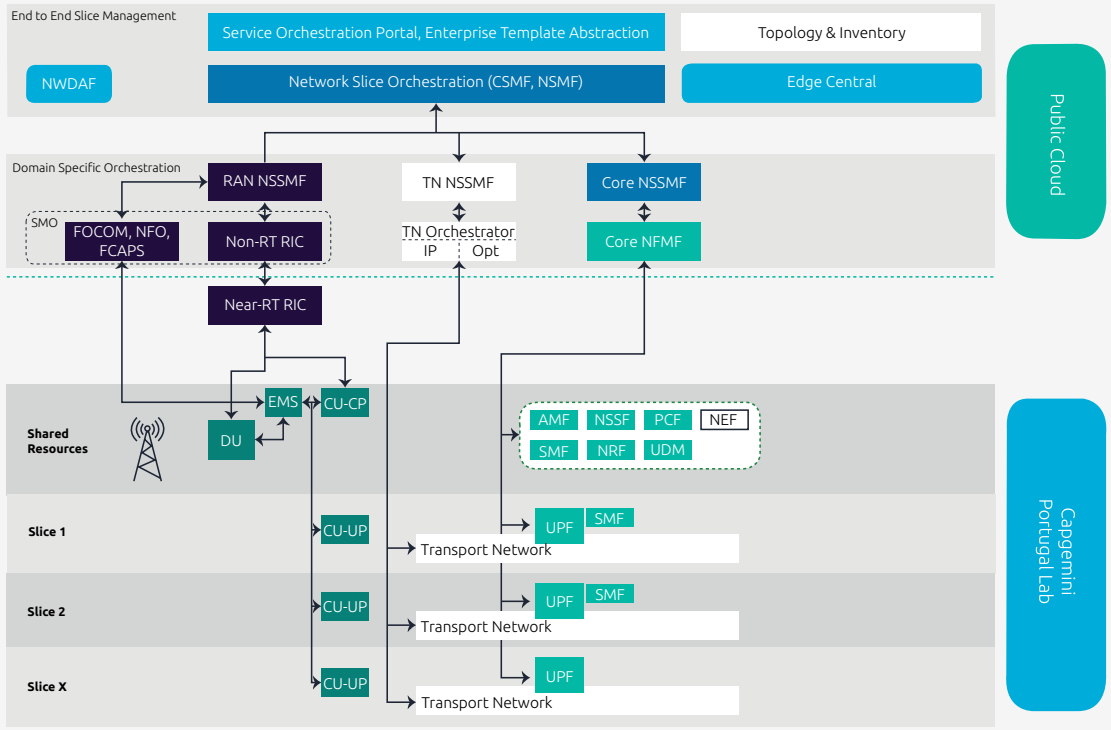

What we did

In this case, the customer wanted to prove the concept of multi-vendor NaaS in order to increase revenue from their investment in 5G. The solution had to be end-to-end (E2E) in order to facilitate this, ie cover an entire fully integrated network across all the technology domains (all the way from the radio that provides the network's connectivity to the core that provides its processing) and prove the maturity of the tech and the standards.

The network PoC was built at the Capgemini lab in Portugal. The majority of the network was run on software and used physical devices only in the radio domain, as would be expected. End-to-end service orchestration, as well as orchestration of each domain, was handled remotely in a public cloud.

The setup used separate vendors for overarching network orchestration and slice management, and multiple vendors within individual domains, including radio domain orchestration, radio network functions (ORAN compliant), core orchestration, and core network functions. We did not use a transport network, since it was located at one site.

This is representative of what a disaggregated network would look like, and what a CSP stitching it together would have to deal with.

Whilst we encountered a number of challenges, these were resolved and the network delivered against its objectives, proving that NaaS across disaggregated networks is possible.

Lessons learned

Through delivering the PoC, we learned a number of valuable lessons which should be taken forward when building disaggregated networks in the real world.

1. Become an expert on standards and integration

There are a variety of standards on how to design technologies to integrate with 5G networks, provided by different bodies, like TMForum, ORAN, GSMA and 3GPP. These specs may work 'on paper', but until an attempt is made to integrate a solution, it's hard to tell if these specifications have been interpreted uniformly by the technology providers. The proprietary integration of non-standard-compliant components, in the development of slicing and slice management, is challenging. Lots of discovery and planning is needed to figure out what is necessary.

Some vendors were at greater levels of maturity and compliance, but even they are still interpreting new and evolving standards. Some use open standards and some closed. Others (some startups for example) have valuable applications but have not yet given any thought to standards and will need a lot of help making their tech compliant.

Whilst standards are still immature, the only way to know if any two applications are compliant is to plug them together in a real use case and see if they can talk to each other. That's when you find out the differences in interpretations of the standards.

Sometimes that creates compatibility problems, which need to be diagnosed and solved. We need to map one vendor's proprietary extensions to the others. We may need to make changes to the technologies' software, or build workarounds (like APIs) that translate one product's inputs and outputs from (and into) the language of the overarching network architecture, or software containers that allow applications to be isolated into discrete units but drop into the cloud environment.

The radio side – the network on the ground that translates the instructions from the cloud – is the hardest bit to get right. The PoC was built to be Open Radio Access Network (ORAN) compliant. Although as yet these standards are immature and there is lots of room for interpretation, we believe this is the future standard with which vendors should be aiming to comply, since it opens new opportunities for CSPs to differentiate. It is important, therefore, to build PoCs to these standards to support vendors and CSPs to move in this direction.

The core – where the software functions are deployed – is less of a risk, as it was built within a cloud environment with mature software capability.

2. Understand new technical challenges that come with disaggregation

Disaggregated 5G private networks present a number of technical challenges that need to be overcome for everything to work together. These will not all be completely new, but they are problems that are more likely to arise in multi-vendor networks. For example, many end-user technologies were not built with 5G networks in mind. A major learning of the POC was that smartphones, and other equipment which use voice, are often not fully compatible with standalone 5G Networks. They are setup to lock onto a network that allows voice, so if the 5G network does not allow it, they will disconnect and lock onto another one, such as a 4G public network. 5G networks should therefore have an IMS core that will allow voice, in addition to the packet core, so voice-reliant devices do not constantly drop off the network.

Another lesson was the need to use the same centralized tool for IP address management across the network, including public cloud and internal network. Otherwise you end up with clashes between the same IP addresses in the public and private cloud which can't talk to each other. The best approach is likely to be leveraging the IP address management tool from the CSP's ecosystem across the whole network.

We also observed multiple CI/CD (continuous integration/continuous delivery) toolsets used by different vendors to achieve automation both at network and software implementation layers. In a disaggregated network, we found a clear benefit from a centralized CI/CD toolset in terms of bringing synergy and control to the network, even when diverse toolsets are used. This provides CSPs with a more robust way of managing the software and network function lifecycle, but needs upfront work to combine and optimize multi-vendor toolsets.

Finally, we were pleased to note the E2E Orchestrator we worked with on the PoC provided flexibility in enabling Capgemini to create adaptors that can convert between APIs and data structures, for both discovery and implementation. This is critical for being able to stitch the individual slice topology to the actual network resources, allowing it to access the communication and data transfer tools it needed to perform. The project nonetheless identified a need for new toolsets to improve this process, which are now being developed.

3. Test the network

Once the network is connected, it requires rigorous testing to check all the different vendor technologies and network layers work together. That means performing a request and tracing it through the different layers of the network architecture, to ensure each part performs as was anticipated by the design. If at any interface, the output doesn't match the receiver, you will either get an error back, or the wrong outcome. That must be followed up and fixed.

Of course, checking all of the above is easy when you built the entire network. But, in a disaggregated network it requires transparency and cooperation between partners, and giving at least one trusted party access to all systems for the purpose of testing and oversight.

Additionally, similar tools should be evaluated, selected and setup to perform slice testing before it is handed over to the customer, to be sure the slice is activated correctly and to prove SLAs and KPIs for the slice are being met at handover and while in service.

A final related point here – for moving beyond the PoC – is that building a disaggregated network should be done whilst minimising changes to existing networks, so as not to affect existing services running over those networks/elements.

4. Look beyond deployment to full lifecycle management and 'clean up'

Vendors need to support the creation of network functions to run the slice in active state and support activation and deactivation of slices onto the network as required. This needs tools in the slice layer to enable the slice to perform its functions as designed, but also in shared network functions – NSSF, AMF, etc - to support slice activation and deactivation, including scheduling of slice creation.

Long term planning is also important. A focus on setting up the network can mean overlooking problems that occur in use. One challenge is that cloud orchestrators sometimes don't properly 'clean up' after closing a slice. This can result in resources being taken away from other parts of the network by a now inactive system. In our PoC, we deployed multiple different slices, which each delivered traffic to support the planned quality of service, to demonstrate full lifecycle management: showing a slice working, measuring its traffic, closing it, and then relaunching it to demonstrate that we had properly 'cleaned up' the network.

5. Price the slice

Once technical possibility is resolved, you will have an idea of what deploying and maintaining a network slice entails. That should allow you to assess what is possible with regard to customer demands, and to start building a business case for what to charge, and what level of SLA you are willing to commit to.

Once agreed, you then need to reserve the infrastructure on the network for the slice and design the slice topology (workload placement to meet the relevant SLAs). Lessons can be learned from doing this manually, which should ultimately be used to automate this process.

6. Build a proof of concept

To be able to do all of the above well, it is advisable to conduct your own PoC before launching a real-world disaggregated network. Disaggregated networks are hard to deliver without the 'been there done that' lessons, and mistakes will inevitably be made the first time. A PoC can provide your team with the experience of dealing with a NaaS in advance of fully deploying one.

The PoC can be used to learn how new technology works, what it is capable of, how it integrates with network architectures and other technologies, and the standards you need to learn – as well as common compatibility problems. That can also help better specify an RFP when provisioning the service. Put another way, the lessons learned from a NaaS PoC are paramount to putting it into production.

A PoC also provides valuable lessons about the business processes and operating models needed to offer NaaS. These will vary between organisations, but will likely include a move from siloed technology towers within each domain, each delivering against, for example a quarterly traffic forecast, to a far more integrated and dynamic approach with cross-domain teams and the ability to rapidly make changes to the network as short-term demands arise.

7. Have a blueprint for your target architecture

A strategic objective for NaaS is the drive to a fully automated digital network which can enable simple interactions resulting in new business outcomes. In order to facilitate this, a target architecture model is required to drive the evolution of the network. This target architecture model should also be utilised to facilitate any PoCs including refactoring the target architecture model leveraging the results of those PoCs.

The target architecture should provide the capabilities and capacity to service all of the target market opportunities, for example providing the models to support the different services ranging from:

- large numbers of subscribers/sessions to small
- high to low bandwidth
- different levels of resilience
- N6 services, Edge/MEC integration etc

The architecture should also support the scheduling of slices and reservation of the required infrastructure capacity (to ensure it is there when needed).

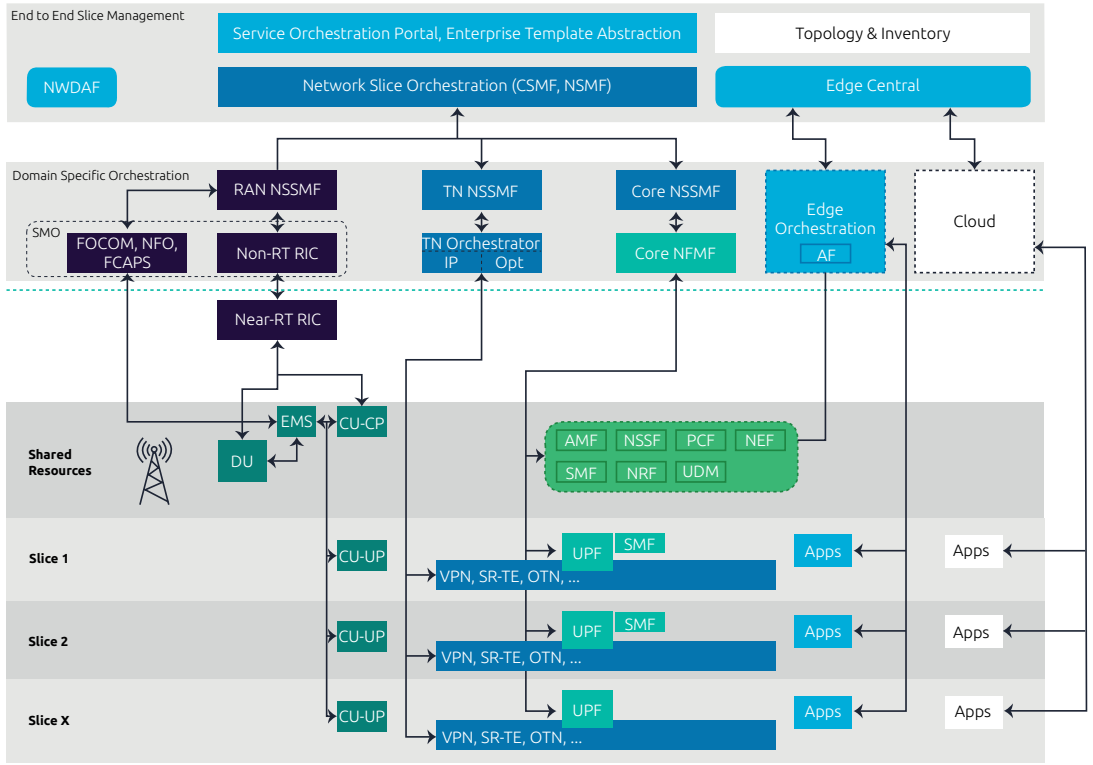

Conclusion

Mastery of network slicing in 5G offers great rewards but is no easy task. The technological challenge is significant - for a start, the technology is still evolving. As such, what is less clear is how to design, deploy and manage such advanced networks within tight customer budgets and timeframes, and to meet a substantial (and growing) number of use cases - all whilst aggregating various disaggregated commercial network systems with varying levels of compatibility.

Failure to implement can result in a range of expensive consequences, including delays and cost overruns. Even for skilled engineers, there is a great deal to learn about the range of available technologies and their use cases, the proliferating and evolving set of standards, and the lessons of stitching them together that are only learned with experience.

Why Capgemini?

Capgemini work with a wide range of partners across the entire Telco industry ecosystem, performing due diligence, system integration, customization and support – this gives us unique and deep insight into the art of the possible.

Technical expertise

We have the unique experience of delivering a disaggregated NaaS PoC to a major telco. Our 20,000 network specialists worldwide have the expertise and experience to ensure the proper security, state, and health of a network and its systems. We have successfully implemented over 1,100 advanced networks (5G and edge projects) since the start of 2020, allowing customers to take advantage of reusable microservices, orchestration, and analytics capabilities.

We offer end-to-end technology consultative services - from network planning and deployment, to managed operation. We possess a deep knowledge of the complex multi-vendor landscape and are able to select (and combine) the right radio, core and transport technologies for the best effect. We provide a world-leading software framework for the custom development of network components, and a library of microservices (AR/VR, image recognition, etc.) that can be deployed quickly as a foundation for new use cases on top of the network, as it evolves.

Business acumen

We can also provide a vision of how the connectivity infrastructure enables long-term returns through use cases. For example, which use cases are available for a specific customer, and how could they benefit from NaaS connectivity? We can also help to develop performance KPIs that help our clients focus on their domain, ie. the use cases and differentiation they can deliver to their customers.

Speed

Our solutions can also be pre-integrated to reduce deployment time. Technology selection and trials are significantly shortened, bringing down the required time to deploy, trial and get data points to decide on network scale. Our network blueprints can cater to all variations of network deployed in enterprise and telco, and can reduce 70% of new network planning, design and test cycle before the first network deployment.

About the authors

Shamik Mishra
Vice President and CTO, Connectivity

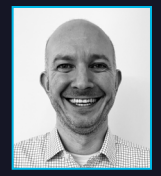

Simon Dumbleton
Solution Director

Chandrashekhar Thakare
Solution Architect

Capgemini engineering

About Capgemini Engineering

World leader in engineering and R&D services, Capgemini Engineering combines its broad industry knowledge and cutting-edge technologies in digital and software to support the convergence of the physical and digital worlds. Coupled with the capabilities of the rest of the Group, it helps clients to accelerate their journey towards Intelligent Industry. Capgemini Engineering has more than 55,000 engineer and scientist team members in over 30 countries across sectors including Aeronautics, Space, Defense, Naval, Automotive, Rail, Infrastructure & Transportation, Energy, Utilities & Chemicals, Life Sciences, Communications, Semiconductor & Electronics, Industrial & Consumer, Software & Internet.

Capgemini Engineering is an integral part of the Capgemini Group, a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided every day by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of over 340,000 team members in more than 50 countries. With its strong 55-year heritage and deep industry expertise, Capgemini is trusted by its clients to address the entire breadth of their business needs, from strategy and design to operations, fueled by the fast evolving and innovative world of cloud, data, AI, connectivity, software, digital engineering and platforms. The Group reported in 2021 global revenues of €18 billion.

For more information please visit:
www.capgemini.com

Contact us at:
engineering@capgemini.com

Team Name_Date_Designer