Unleashing the economic potential of network slicing

Network slicing has the potential of enabling operators to address new customer segments with highly personalized services. However, the costs incurred and the revenue benefits that could be gained remain largely unquantified.

Bell Labs Consulting performed a detailed analysis of network slicing to understand its business viability for operators, the technology’s practical limitations and how these can be overcome.

NOKIA

Contents

NOKIA

Executive summary

5G network slicing is a new network approach that can provide highly tailored services to specific customer groups and even individual customers. Virtualization and cloud technologies enable the same underlying network infrastructure to be used to create virtual network instances with performance characteristics that fit the needs of new customer groups, something that conventional one-size-fits-all networks cannot achieve.

These network slices can be set up quickly and managed dynamically to deliver services that can command a premium revenue. There is a large potential base of customers within public safety, industrial automation, healthcare, Internet of Things (IoT) and many other areas that would value such services. However, creating a dedicated network for each customer segment is impractical due to capital and operational expenses.

The technology of 5G network slicing has been researched extensively and is being standardized, yet its economic aspects are still unclear. Exactly what new revenue could be generated? What are the consequences of network slicing on network operational and investment costs? How do these costs change as more network slices are added? Is there a sound business model for network slicing and how should an operator target potential customers?

To answer such questions, Bell Labs Consulting has run an exhaustive analysis of network slicing. The results reveal how network slicing can enable operators to offer premium services to address new customer segments and charge premium rates. Key findings include:

- Network slicing boosts operating margins by more than 5 percentage points as the number of network slices is scaled up, while margins from dedicated networks stagnate as the number of networks is increased
- Bell Labs Consulting estimates that 15 percent of subscribers are interested in a premium connectivity service and that every 1 percentage point increase in such customers results in a 3 percentage points increase in revenue
- As the number of premium customers grow, the Total Cost of Ownership (TCO) of sliced networks becomes more and more efficient compared to dedicated networks
- Network management automation is critical to minimize operational costs and enable more network slices to capture more connections and thus customer revenues
- The accurate preparation of network slices and their operation is critical to ensure that premium services are supported by the network. Nokia 5G Services provide the expertise and processes to achieve this goal

5G network slicing is the future of networks

5G networks will be able to support diverse and extreme requirements for latency, throughput, capacity and availability. 5G technology will enable a plethora of services impractical with previous wireless technologies. With network slicing, operators and enterprises will be able to address precisely the specific needs of different customer segments.

For example, they will connect the factory of the future and help create a fully automated and flexible production system. In healthcare, hospitals will be able to arrange remote robotic surgery and on the roads, self-driving cars and smart infrastructure will reduce accidents. City governments could be served for real time traffic management, emergency response and video surveillance with customized quality and security requirements. Even within a city, separate slices may be created for different municipalities and organizations for administrative reasons, even though they may have technically similar requirements.

A good example of how such a deployment could work is being explored in a trial for the Hamburg Port Authority, in which Deutsche Telekom and Nokia are testing key aspects of 5G, including network slicing, in a large 80 km^2 industrial area in the Port of Hamburg. 5G network slicing will be tested with use cases like traffic light management, data processing from mobile sensors and virtual reality.

Meanwhile at Mobile World Congress 2018, Nokia and BMW demonstrated the creation of a slice in a run-time environment for automobiles, while maintaining other slices, such as public safety, running with specific Service Level Agreements (SLAs).

Network slices are network instances for individual customers and are enabled by Software Defined Networking (SDN), Network Function Virtualization (NFV), cloud technologies, automation, end-to-end service provisioning and orchestration. Virtualization enables separation of the software from the hardware to implement many functions on common infrastructure. SDN enables dynamic capacity planning, routing and service chaining, based on real-time needs. Finally, orchestration allows end-to-end slice management during its life-cycle.

Network slices are composed of capabilities from multiple network segments from the access to the core, as well as applications. Only an end-to-end approach ensures that performance requirements of various services for different customers can be met with certainty. Furthermore, slicing allows the same infrastructure to be dynamically shared by different tenants, thus optimizing the Total Cost of Ownership (TCO).

Figure 1. Network slicing partitions common network infrastructure into multiple, logical, end-to-end, virtual network instances to provide customized virtual private services

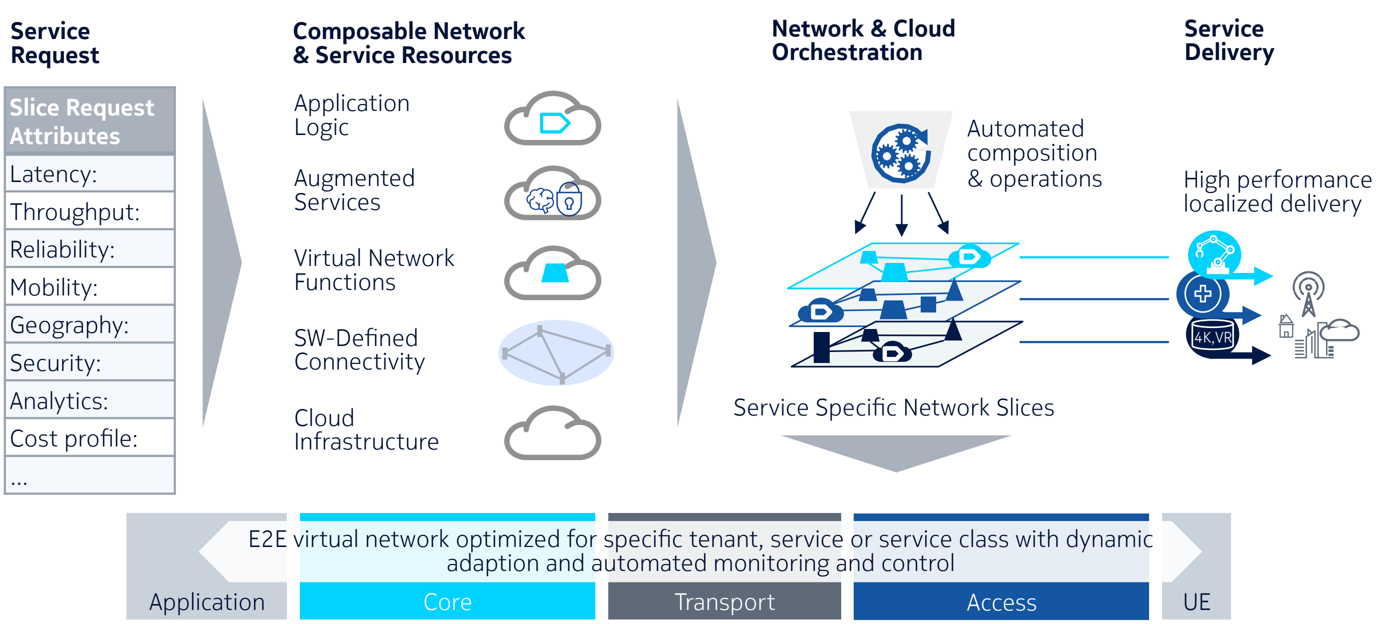

What is the true value of network slicing?

While the nuts and bolts of 5G network slicing, that is the underlying technology and how slices are created, have been discussed for some time, little research has yet been done into its economic value for operators. To address this shortfall, Bell Labs Consulting analyzed the costs of implementing and operating network slicing compared to the potential revenue that could be generated.

Network slicing creates revenue opportunities for operators in two ways:

- Serving existing high-end customers with higher grade service: services tailored to the needs of a current customer or customer segment are more valuable and can win a premium Average Revenue Per User (ARPU)
- Reaching new customers for whom premium service is critical: new customer segments can be addressed, creating completely new, additional revenue streams.

These benefits are clear, but the potential level of new revenue available to an operator is not so obvious.

Equally, while there are cost benefits of network slicing, the added costs for serving more and higher quality customers have not been well understood. The cost impacts are reflected in both Capital Expenses (CAPEX) as well Operational Expenses (OPEX) and each has positive and negative components:

- CAPEX
  - Positive: Network slicing enables custom design, provisioning and maintenance of logical and isolated network instances for each customer. Network automation enables finer tracking and capacity scaling of each network slice. All these aspects, including statistical multiplexing of different services on shared infrastructure, avoid over-provisioning of network elements, thereby saving equipment hardware and installation costs.
  - Negative: On the other hand, software costs could increase because of the need to implement customized service chaining which requires dedicated software and corresponding installation and engineering costs.
- OPEX
  - Positive: Efficiency is achieved through streamlined processes and improved service life-cycle management with automation.
  - Negative: Monitoring and support costs may increase as the number of dedicated slices with different Quality of Service (QoS) requirements increases.

Three network scenarios provide comparative costs

Networks will evolve from a single physical or virtual network to a network with dedicated slices enabled by virtualization and finally to a full-fledged cloud-optimized network featuring dynamic slicing. The analysis models snapshots of the business cases for these stages using the following scenarios:

Single network: A single virtualized network that provides best-effort support for three service categories – extreme mobile broadband (eMBB), Massive Machine-Type Communications (mMTC) and Critical Machine-Type Communications (cMTC). This is the conventional one-size-fits-all network.

Dedicated networks: Customized, dedicated virtualized networks are built for customers willing to pay a premium for customized QoS or dedicated network capacity. In this case, networks are designed and deployed exclusively for respective customers and thus, include dedicated hardware and software systems over and on top of a deployed single network. There is limited automation of network management and limited cross-domain orchestration. Life cycle management of network functions is mostly manual and network functions may be under-used as they are duplicated across required dedicated networks.

End-to-end Network Slicing: Full network slicing comprising a single virtualized network, with slices for customers with specialized technical, performance, regulatory, security, confidentiality, geographic or other requirements. This single virtualized network is based on cloud-native architecture comprising microservices, layered and modular architecture, end-to-end service chaining with SDN programmable transport and a highly automated network management layer. Automation provides capabilities for instantiating the required network functions and services (orchestration), for connecting these functions together (SDN), for customizing and personalizing the use of these functions (service chaining), for automating the configuration of these functions (SON), for optimizing the customer experience when using these functions (CEM) and for providing the cognition necessary for automated decision making (analytics).

Bell Labs Consulting calculated the incremental costs and revenue for each scenario as the number of dedicated networks and virtualized network slices was increased to provide services to different customers.

Comparing the costs and revenue in each case provides an insight into the economic potential of network slicing.

Figure 2. Three network deployment scenarios were considered for the techno-economic modeling

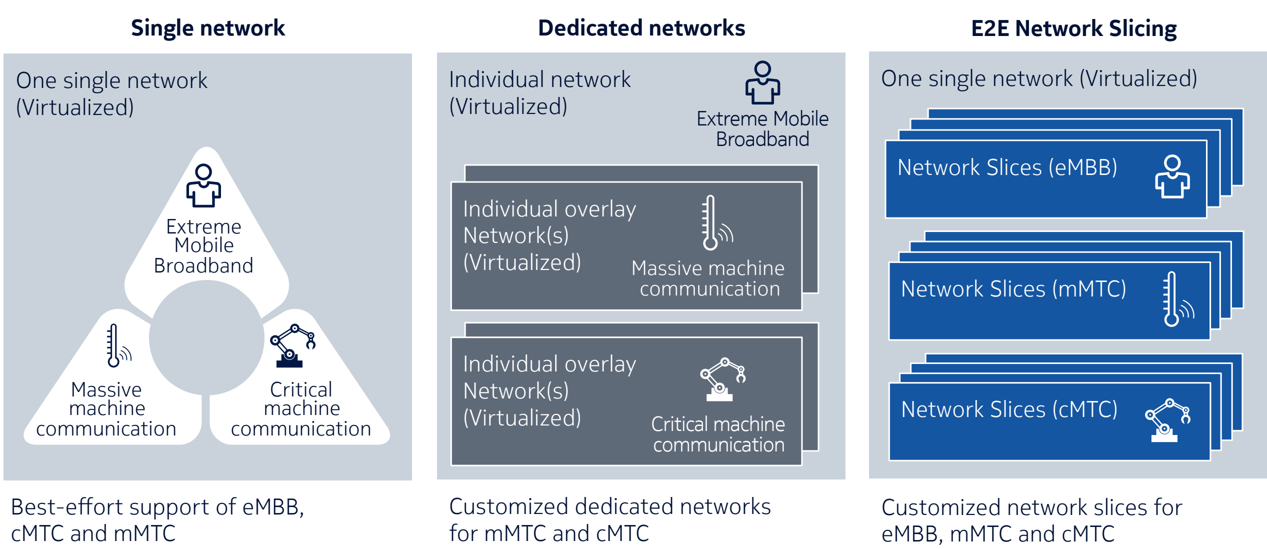

Cost comparisons for the three scenarios

The single network can be regarded as a starting point with the other scenarios adding dedicated networks or slices to address new customer segments. For clarity, the comparisons are shown in two separate charts, each revealing the impact on hardware, software and operational costs as the number of dedicated networks or slices increases. Please note that more connections are enabled by the greater number of dedicated networks or network slices.

The left chart shows how the incremental cost of dedicated networks grows compared to single networks. The right chart shows the same comparison between sliced and single networks.

Figure 3. The TCO for dedicated networks rises substantially as more networks are introduced. Network slicing on the other hand shows less sharply rising TCO

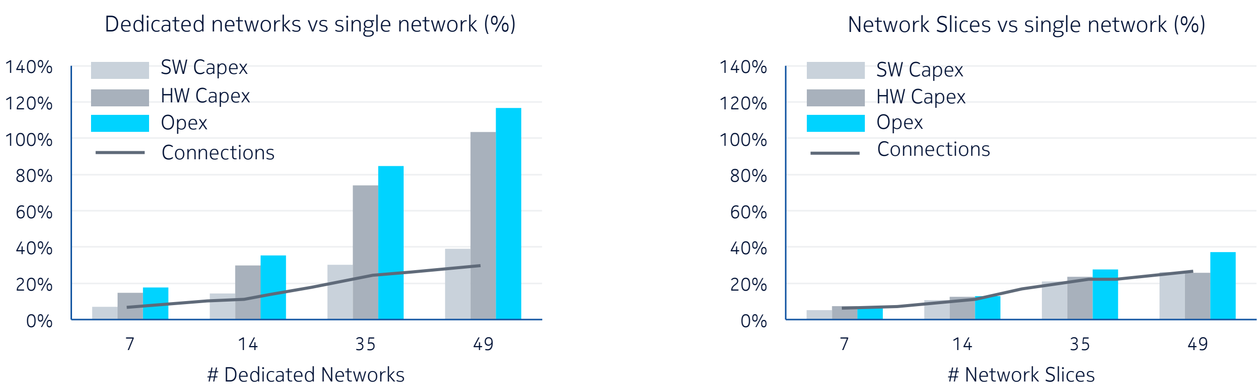

As the number of connections grows, dedicated networks incur substantially higher hardware CAPEX and OPEX because the infrastructure is not shared, while the software CAPEX grows with the number of software licences.

With network slicing, however, both CAPEX and OPEX rise linearly as the number of connections and network slices increases. Bell Labs Consulting analyzed up to 50 slices in this study considering current evolutions of IoT and private networks. This does not mean that we do not see a potential of hundreds or thousands of slices once the technology is fully mastered and operators find new ways to fully monetize the potential of slicing.

Figure 4. The growth of connections and TCO for dedicated and slicing scenarios

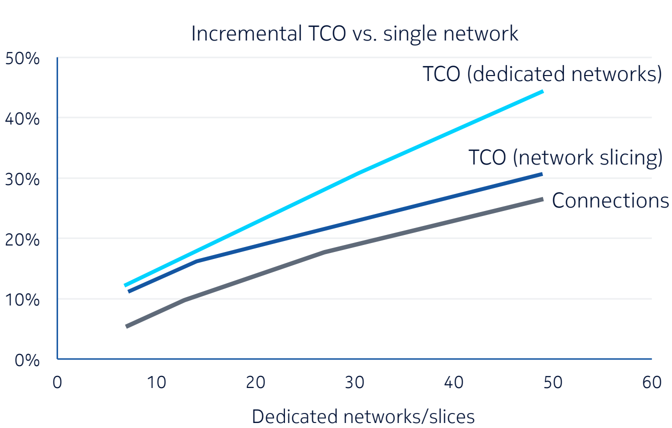

The TCO of dedicated networks increases more rapidly than with network slicing. In practice, dedicated networks will not scale beyond a limited number of networks, such as ten, based on the current model as their costs rise sharply. While the costs for network slicing also rise, the increases are much less severe, enabling them to support more network slices. Also, higher-grade services cannot be enabled by the single network.

Although the cost increases for both scenarios with the number of connections, the cost of slicing follows the connection growth almost linearly while the cost of dedicated networks starts to diverge with the higher number of connections.

For the remainder of the cost discussions, we focus only on network slicing and drop the dedicated networks scenario as Future X networks are expected to support thousands and even millions of customers and customer segments with customized QoS and security requirements, scenarios for which dedicated networks are simply impractical.

The need for automation: managing the cost of increased connections

The goal of slicing is to serve a very high number of customers with customized service requirements which can only be realized by dynamically created slices. These slices will have to be taken down as soon as they are no longer needed to minimize the cost of hardware and software resources. As shown in Figure 3, operational costs are the biggest component of the TCO elements that grow with the number of connections. As connections are the sources of revenue, we want to determine how to control operational costs without jeopardizing connection growth.

Figure 5. OPEX components for comparing the costs of single network vs slicing

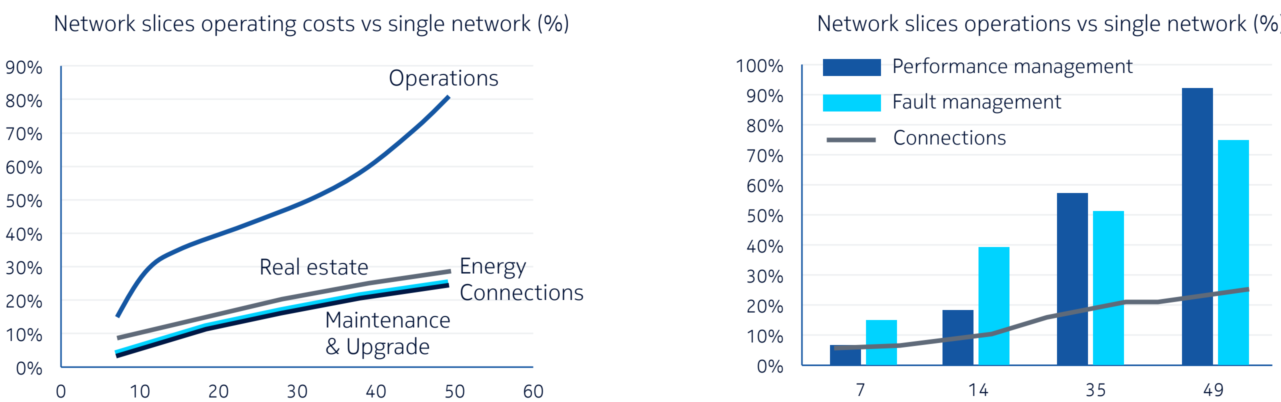

The left-hand side of the picture shows that OPEX components related to real estate, maintenance and energy tend to grow linearly with the number of connections, but there are other OPEX elements that push up the total OPEX at a much higher rate. The right-hand side of the chart makes it clear what those components are – elements that require a lot of human intervention like Performance Management and Fault Management. These components increase not only due to increased connections but with increased network slices, since each network slice customer requires customized and individual support for fault and performance management, security monitoring, service life-cycle management and network optimization.

Operations comprises aspects related to fault management, performance management, service life-cycle management, security monitoring, etc.

Increasing the level of network management automation can substantially reduce the OPEX for each network slice. Bell Labs Consulting assumed 50 percent level of automation as challenging but achievable. By level of automation, we imply amount of human labor that can be reduced through the automation of required tasks in terms of either the number of hours required and/or number of operational staff needed to accomplish certain tasks.

In effect, even high numbers of slices could be managed without needing to increase the number of operational staff. The more automation applied, the lower the operating costs. At 100 percent automation, there is virtually no cost increase with the number of slices. Granted this is a long-term goal and impractical in the short to medium term, yet even 50 percent automation will bring very significant benefits.

Figure 6. Automation is essential to minimize the costs of operations and enable higher numbers of network slices to be deployed to capture more connections and customer revenues

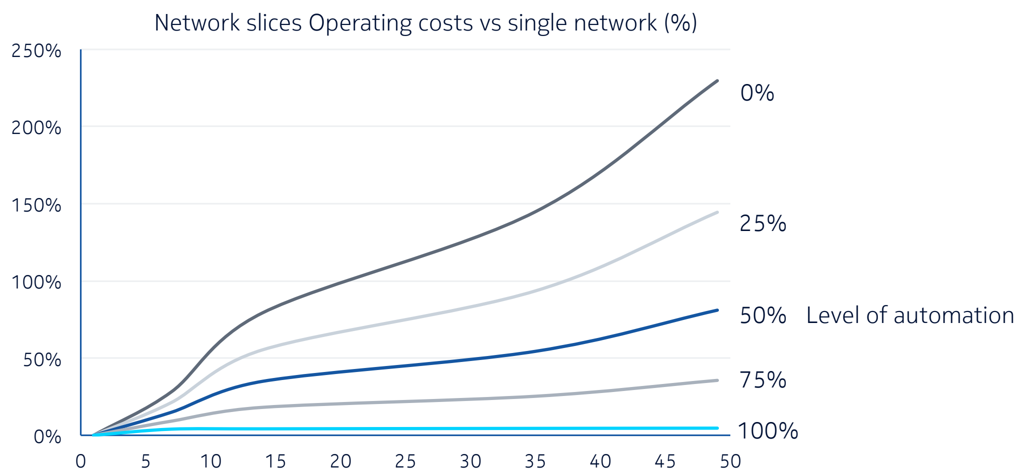

Analyzing the revenue impact

Bell Labs Consulting modeled a large US-based operator to provide a baseline for calculating costs and revenue. The first stage was to identify the operator's current services, revenues and market share in mobile broadband, IoT and enterprises to give the analysis a reference starting point.

The second stage was to identify for all segments the potential growth in market share and revenues when being able to offer a customized connectivity service. High on the list is public safety that often has specific regulatory or security needs that can be competitively addressed by the operator. Also, IoT segments with stringent technical requirements. The key was to estimate potential increase of the operator's market share in these segments. This increase can be enabled by the operator's differentiation in providing a better technical solution and more innovative business models.

In the third stage, estimations are made for increased Average Revenue Per User (ARPU). It is expected that quality conscious customers will be willing to pay a premium for value-added tailored services. The analysis assessed the potential number of organizations within each market segment that would value dedicated network services and be willing to pay for them. This enabled the possible additional ARPU available to the operator to be established.

The impact of faster time to market has not been captured in this model. Adding this component is expected boost the case for slicing even further.

Figure 7. Process of estimating additional revenue enabled by slicing

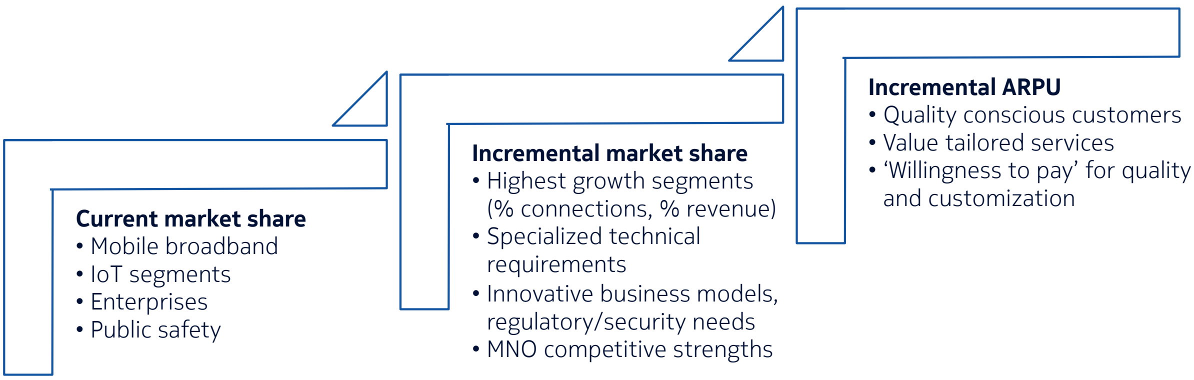

Key finding

An early lesson from the modelling work is how critical it is for the operator's sales teams to understand deeply the value of 5G network slicing. Sales must be dedicated to finding customers with both the specialized requirements and an appreciation of the benefits of a service based on network slicing, indicating they would be willing to pay a premium price.

With regard to the relative revenue growth with slicing or dedicated networks compared to a single virtualized network, Bell Labs Consulting assumes that operators can achieve the same revenue growth in both scenarios, disregarding the fast time-to-market that can be achieved with slicing.

Figure 8. Relative revenue growth with number of connections

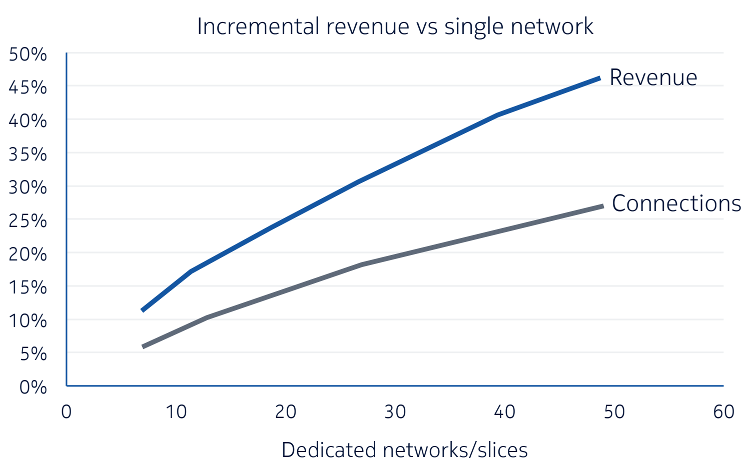

Revenue increases depend on premium customers

Increasing the number of network slices or dedicated networks enables more specialized services to be offered to more niche customers willing to pay premium rates. This additional revenue is not available to a single, one-size-fits-all virtualized network because it cannot provide the dedicated performance parameters.

A sensitivity analysis underlines the importance of targeting the right customers willing to pay for premium services delivered by network slicing. In the sensitivity analysis, three factors – percentage of subscribers requiring tailored services, incremental market take rate and increase in ARPU– are each independently varied +/-20 percent to see their impact on the revenue.

In the base case, with 50 network slices, 15 percent of customers requiring tailored services results in a 46 percent increase in revenue. As the percentage of these subscribers varies between 12 percent and 18 percent (a 6 percent change), the incremental revenue varies between 37 percent and 55 percent (an 18 percent change). Thus, every 1 percentage point increase in customers requiring tailored connections, results in a 3 percentage point increase in revenue. A similar effect is found when varying the market uptake from 40 percent to 60 percent. However, varying the ARPU from 48 percent to 72 percent (a 24 percent change) varies the revenue only from 43 percent to 50 percent (a 7 percent change) thus resulting in a smaller relative variability. This shows that it is more important to target the right customers than it is to increase the ARPU per customer.

Figure 9. A sensitivity analysis shows the importance of targeting the right customers that have the appetite and need for premium services

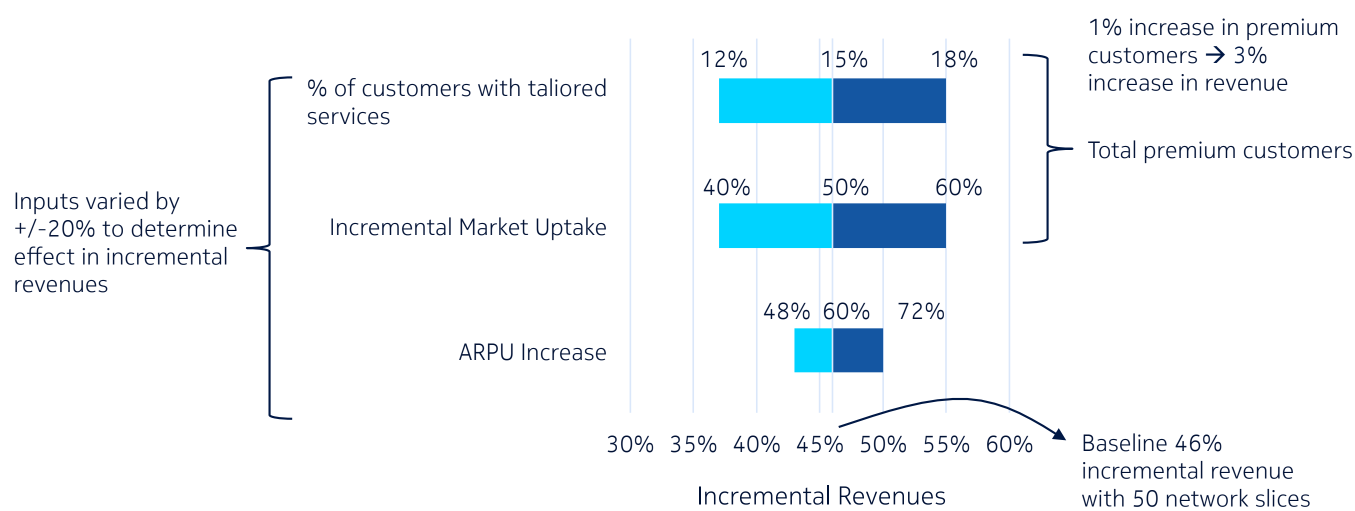

Putting it all together to derive what matters most: profitability

So far, we saw the costs and revenue impacts individually on slicing. Looking at their combined effect on overall profitability, Earnings Before Interest and Taxes (EBIT) for slicing clearly increases faster in comparison to dedicated networks. This is because as revenue increases with the increase in the number of dedicated networks or network slices, the cost for dedicated networks increases substantially while increases in the cost of fully-sliced network is minimal.

With thousands of slices expected in Future X networks, the advantage of slicing is thus established beyond doubt, against both single as well as dedicated networks.

Figure 10. Network slicing delivers greater operating margin than dedicated networks when more than about ten slices are created. The operating margin seems to stay flat or even begins to decrease for dedicated networks whereas it continues to grow with network slicing

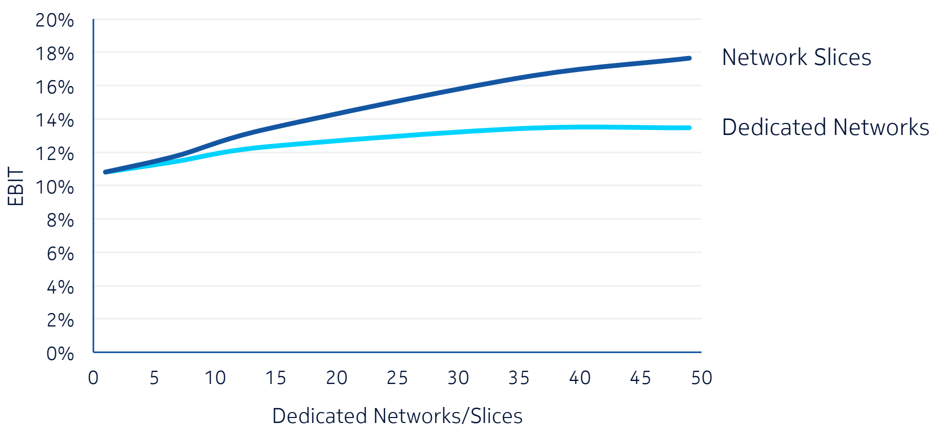

Setting up and optimizing network slices for high performance

Network slicing opens new business opportunities for operators by enabling them to provide specialized services that deliver specific performance parameters. Guaranteeing stringent KPIs enables operators to charge premium rates to customers that value such performance. The flip side is that such agreements will inevitably come with tough contractual obligations and penalties when the agreed KPIs are not met.

This means that the setting up and operation of network slices must ensure they deliver the required performance, requiring expertise in many areas. Nokia 5G Services provides the necessary expertise.

One of the most important steps is the preparation phase in which the network slicing is designed to support end-user services. This requires a detailed understanding of the use cases and the traffic profiles assigned to the users, encompassing many aspects such as the required geographical coverage, throughput performance, latency needs and the number of users.

This establishes a solid basis for the design of the network slice on top of the existing network that will not only deliver the required performance for the slice users, but not adversely affect the existing users of the network. The design will define the slicing methodology, such as QoS based, and assess and plan the network in all the domains (radio, mobile access, enhanced packet core) in terms of capacity required, topology and architecture. This will assess the need for technologies such as Multi-access Edge Computing (MEC) to provide ultra-low latency. The definition of each network slice is described in a template that will be used during the deployment phase.

Following the preparation and design phases, the network slices can be instantiated and activated ensuring the required service quality is monitored and maintained. A key aspect that needs to be managed here is the required elasticity. Cloud resources are finite and actions may be needed at the physical network level and/or the network management level to avoid bottlenecks occurring.

Finally, at the end of the service life, proper termination and deactivation of the network slices must be performed to ensure that decommissioning does not cause any adverse effects on other slices or the wider network in general.

Network slicing provides additional value, but also increases network complexity. For that reason, proven know-how in end-to-end network planning combined with expertise in network management systems and tools are essential to the successful deployment and operations of network slicing.

During the preparation phase, user traffic maps can be created using geolocation techniques and data analytics to correlate network data, crowd-sourcing and social media. Nokia has studied and developed these methodologies to enable a better estimation of user profile segmentation, users' locations in 3D and forecasting of new end-user applications in future use cases, based on machine learning models guaranteeing the right inputs for a correct design.

Approaches based on advanced data analytics and machine learning can be applied in the optimization phase to identify issues in the complex environment of multiple network slices for multiple domains and quickly and accurately propose solutions.

Figure 11. Nokia 5G Services provide a well-defined process for establishing and running network slices that will deliver the required performance for high-value customers without affecting other users of the network

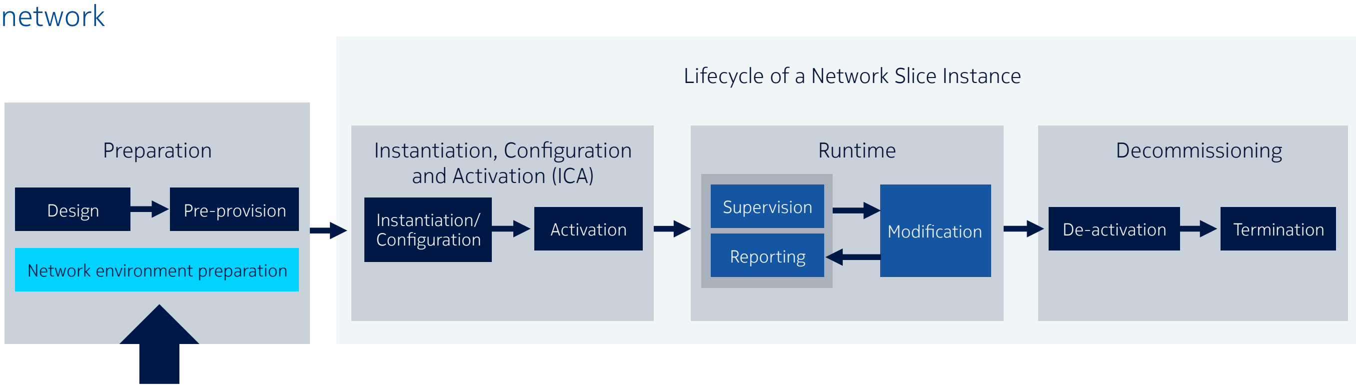

Conclusion: Network slicing delivers higher operating margins

Bell Labs Consulting has run a thorough techno-economic analysis of the revenue and cost impacts of 5G network slicing on an operator's business. Comparing costs of dedicated networks and network slicing against a baseline of a single virtualized network revels some key findings:

- The commercial viability of network slicing depends on the operator being able to identify potential customers that have a need for the capabilities that network slicing offers. Furthermore, the customers must recognize the value of the services delivered and be willing to pay a premium for them
- Network slicing has the potential to reach new customer segments and generate substantial extra revenue for an operator. These customer needs cannot be met with single, one-size-fits-all networks
- Dedicated networks can address the needs of existing and some new customer segments, but they do not scale efficiently. Beyond a limited number of networks based on the current model, costs become unsustainable
- End-to-end network slicing opens further customer segments with costs that scale more or less in line with the additional demand, making them commercially viable
- As the number of slices becomes very large, network slicing experiences sharply higher operational costs because of the complexity of managing the multiple slices and customers
- Implementing automation effectively brings down the operational costs of network slices, enabling an operator to address more customers profitably. With full automation, the TCO of even high numbers of network slices becomes acceptable
- The accurate preparation of network slices and their operation is critical to ensure that premium services are supported by the network. Nokia 5G Services provide the expertise and processes to achieve this goal.

With network slicing, operators have a new means of expanding their business and winning new revenue by delivering premium services focused on the requirements of customer segments, such as high network performance in key areas or to meet specific regulatory or compliance needs.

Even within one segment, different customers will have varying needs. By applying automation to control network slice operating costs, operators have the potential to use network slicing to win business in many new areas.

Nokia Oyj
Karaportti 3
FI-02610 Espoo, Finland
Tel. +358 (0) 10 44 88 000

Document code: SR1801020902EN (April)