Executive Briefing Service

In co-operation with

JUNIPER NETWORKS

STL PARTNERS

Executive Briefing

# 5G NETWORK SLICING: HOW TO SECURE THE OPPORTUNITY

Network slicing is an opportunity for telcos to disrupt private networking, but it is new and unknown to enterprises. A 'hand holding approach', internal alignment and simplified pilots will be key enablers in its adoption.

Yesmean Luk, Consultant | Matt Pooley, Senior Consultant | Phil Laidler, Partner
yesmean.luk@stlpartners.com | November 2019

# Preface

The document has been prepared by independent research firm STL Partners, and is commissioned by Juniper Networks. It is based on STL Partners' continuous research programme into the future telecoms operator and how to get there.

Mentions or allusions to companies or products in this document are intended as illustrations of market evolution and are not intended as endorsements or product/service recommendations.

# Executive summary

## Slicing and 5G: not quite real… yet

Network slicing has been touted as an integral part of the 5G opportunity. One highly anticipated aspect of 5G is that it will be built on virtualised infrastructure. Network functions will run as software in datacentres, rather than on dedicated appliances as in the past. This will mean that operators can deploy and make changes to functions with far greater flexibility than ever before. It also offers the promise of enabling multiple logical end-to-end networks - each intended to meet specific needs – to be "spun-up", operated and retired as required, over the same shared hardware. This capability is the essence of **network slicing**. In particular, it has been suggested as a mechanism for telcos to provide a new breed of managed networking services to enterprise customers.

The technology of network slicing is not quite a reality today. It is being tested in labs, though its application depends on a full 5G network core which doesn't yet exist commercially anywhere. But assuming the technology can be made to work, is there actually a market for it?

## Do enterprises really want to buy a network slice?

We found that as slicing is still very new as a concept it is largely unheard of by many enterprises. Indeed, many telcos still aren't entirely sure how to think about slicing and how it will work in practice. Clearly, these gaps need to be addressed.

Beyond that, we found that if telcos plan to leverage network slicing as an enterprise offering, they must first address **concerns that enterprises express on network slicing when it's explained to them**, particularly around **security**. For telcos, much of the value of network slicing comes from the potential cost savings and efficiency of hosting services on common public infrastructure, but enterprises often voice concerns about privacy and security on public networks. **We believe that much of this concern is unfounded: many are unaware that the whole point of slicing is to tailor public networks to bespoke enterprise needs – whether that is much higher security standards, or greater reliability through isolation.**

## Building sound foundations

To move slicing forward, the telecoms industry should agree on standards or at least guidelines that will ensure the technology works and is resilient and secure. These include:

- defining "isolation" – this has been widely discussed as crucial for ensuring the performances of individual slices is not affected by intentional or unintentional events and actions, but there isn't technical consensus on how to achieve it
- agreement on the required level of baseline security, for example around authentication protocols and encryption algorithms
- clarity on the roles of telcos and vendors in orchestration and management of slices.

Defining these standards and how they will be met is important to building enterprise trust because slices will only be as secure and reliable as the lowest common denominator.

## Preparing for slicing

Telcos who are considering slicing as an enabler for enterprise networking services should:

1. **Make network slicing part of the private networking services menu.**

    - In order to directly address some of the concerns that enterprises have around security and reliability, network slicing should be positioned as a delivery mechanism for private networking services. This creates the expectations of using dedicated resources. This should not only be about how network slicing is positioned externally to customers but should also shape how telcos organise themselves internally.

    - Figure out how network slicing is internally defined and what the proposition looks like before going to enterprise customers. Rather than selling a network slice, sell a private networking solution enabled by network slicing that addresses enterprises' business needs and problems. There will be a period of transition, but this will help drive adoption because what enterprises really want is an affordable private network, not what technology is used to deliver it.

2. **Be open with enterprise customers and address concerns head-on.**

    - Acknowledge the potential risks of network slicing and be prepared to communicate openly and show that these risks have been addressed and are being proactively managed. This will help to build trust and confidence with customers.

    - Work closely with enterprise customers to understand their needs. Telcos need to adopt a 'hand holding' co-creation approach with customers to help them better understand network slicing and work with them to provide the level of comfort they seek. This requires a strong understanding of enterprise customers' business problems and requirements. This will require different skillset and types of conversations with customers.

3. **Don't wait for maturity to start testing.**

    - The transition to fully-fledged network slicing doesn't have to wait for when 5G is ready. Pilot simplified implementations of network slicing with existing technology for internal use ahead of commercial roll-out. Deploying slicing for internal use will build more experience and provide an opportunity to learn and iron out any problems or details.

    - With these pilots, telcos can work with existing friendly enterprise customers and co-create tailored solutions and these can be upgraded to fully-fledged network slicing when 5G networks are ready.

Although network slicing is still in development, because it takes a while for new ideas to take hold in the market, telcos should start the groundwork now**.** Because of long investment cycles and steep

learning curves for enterprises, telcos should **start working more closely with enterprise customers immediately** to not only understand their needs, but to build trust and confidence for when fully-fledged network slicing technology is ready.

# Table of Contents

# Table of Figures

# Introduction

## Network slicing is central to unlocking the 5G opportunity

There has understandably been a lot of talk and hype about 5G in the telecoms industry. It promises to bring greater speeds, lower latency, greater capacity, ultra-reliability, greater flexibility in the network operations and more. It also pledges to support high device densities and to enable new services, new business and operational models as well as new vertical opportunities.

Given that the rollout of 5G networks is expected to involve a significant investment of hundreds of billions of dollars, there is a need to look at how it might address new business opportunities that previous generations of cellular networks could not. Many, including us[1], have argued that the consumer business case for 5G is limited, and that the enterprise segment is likely to represent the greater opportunity.

One highly anticipated aspect of 5G is that it will be built on virtualised infrastructure. Network functions will run as software in datacentres, rather than on dedicated appliances as in the past. This will mean that operators can deploy and make changes to functions with far greater flexibility than ever before. It also offers the promise of enabling multiple logical end-to-end networks - each intended to meet specific needs – to be "spun-up", operated and retired as required, over the same shared hardware. Traditionally, achieving such a multi-service outcome would have required building dedicated stand-alone networks, which was rarely a viable proposition. This capability is the essence of **network slicing**.

This report will explore the concept of network slicing and what it means for enterprise customers. It will have a particular focus on one aspect of network slicing through the enterprise perspective, that being security. The first section will cover how we define network slicing whilst the second will dive into what the enterprise security-related concerns are. We will then assess the implications of these concerns in the third section, before identifying ways that telcos can address these concerns in order to accelerate the adoption of network slicing.

Our findings in this report are informed by a wider STL Partners research programme that STL Partners has conducted with telcos and enterprises across several verticals, including transport, defence, utilities, logistics and smart cities.

[1] STL Partners report: 5G: 'Just another G' – yet a catalyst of change

# Dynamic, virtualised, end-to-end networks on shared resource

In previous publications[2], we have defined network slicing as "a mechanism to create and dynamically manage logical functionally-discrete end-to-end networks over common physical infrastructure".

Until now, public mobile networks have been built on the principle that all connected devices communicate using the same physical infrastructure (access network, core functions, etc.) as everything else. The operator rolls out a "one-size-fits-all" network which is designed to cover a small set of well-defined use cases (primarily voice calls, SMS and consumer mobile broadband) effectively, but not much else.

This focussed approach has allowed operators to reach users across geographies quickly, while remaining cost-effective. However, it caters poorly for use-cases beyond the design specs. In an ideal world, a network that is designed to support consumer mobile broadband would look very different to a network designed to support an agricultural IoT solution, for example. In the mobile world, however, both would run over the same infrastructure – or the customer would need to deploy a private network to support niche use-cases.

**Figure 1: Limitations of one-size-fits-all networks**


Source: STL Partners

With 5G, this will change. 5G networks differ from previous generations in being built from the ground-up on *virtualised* infrastructure. Network functions will run as software in datacentres, rather than on dedicated physical appliances as in the past. This has many advantages – not least that the operator will be able to spin-up network functions at will without the need to buy expensive bits of dedicated equipment each time. It is this functionality that allows operators to "create and manage logical

[2] STL Partners Report – Network Slicing: The greatest thing since sliced bread?

functionally-discrete end-to-end networks", or in short, "network slices". Figure 2 illustrates how this might look:

**Figure 2: Diagram of network slicing**

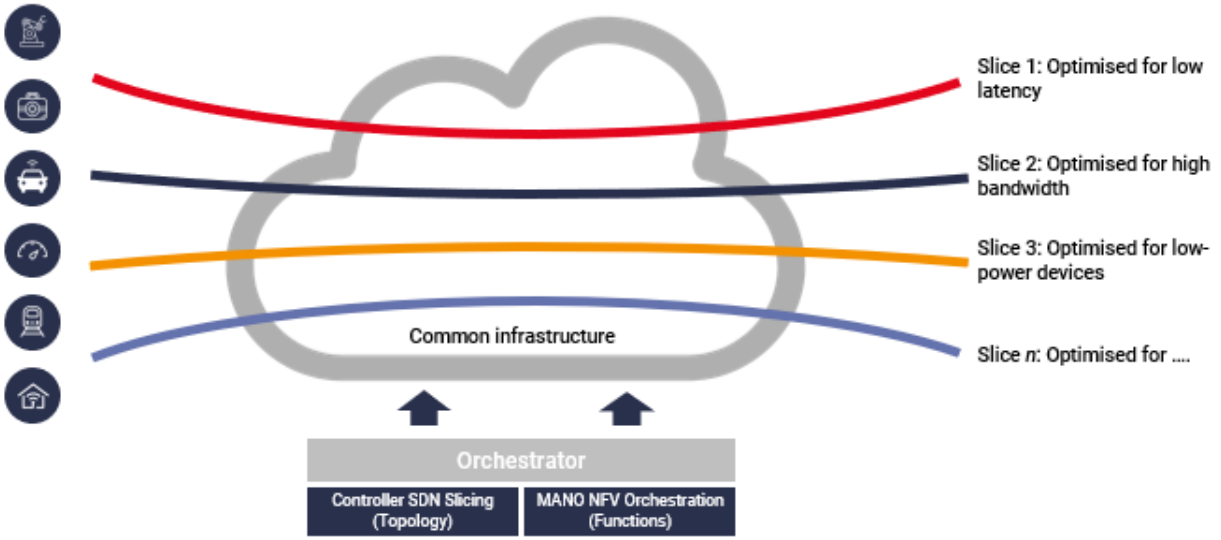


Source: STL Partners

Complete end-to-end network virtualisation (as envisaged with 5G) is necessary for dynamic network slicing. However, this is only one of the many building blocks required. We anticipate the fully-fledged automated network slicing to be deployed 2-3 years after we have significant rollout and coverage of standalone 5G networks: for most countries, not before 2025.

Each end-to-end network slice has the functionality of a complete network, including specific network layer capabilities, operational parameters and network characteristics. Network slices also each have their own set of network functions and required compute, storage and networking resources to meet certain requirements. Once deployed and activated, it is known as a "network slice instance". Each slice has at least one instance, which defines the behaviour of the slice.

Network slices can each have "different requirements on functionality (e.g. priority, charging, policy control, security, and mobility), differences in performance requirements (e.g., latency, mobility, availability, reliability and data rates), or they can serve only specific users".[3] This means that each network slice can be designed to serve either a particular use case, purpose or even an individual customer (or "tenant"). Alternatively, each network slice can support services that have multiple tenants.

[3] 3GPP TS 22.261 version 15.5.0 Release 15

# Slicing might come about in different ways

The 5G standards has been developed with three fundamental use cases in mind: enhanced mobile broadband (eMBB), ultra-reliable low latency communications (URLLC) and mass machine-type communications (mMTC).

It is quite possible that network operators could cater for these by slicing their networks into three parts, each with its own set of network functions.

**Figure 3: Scenario 1 – Many customers on 3 generic slices**

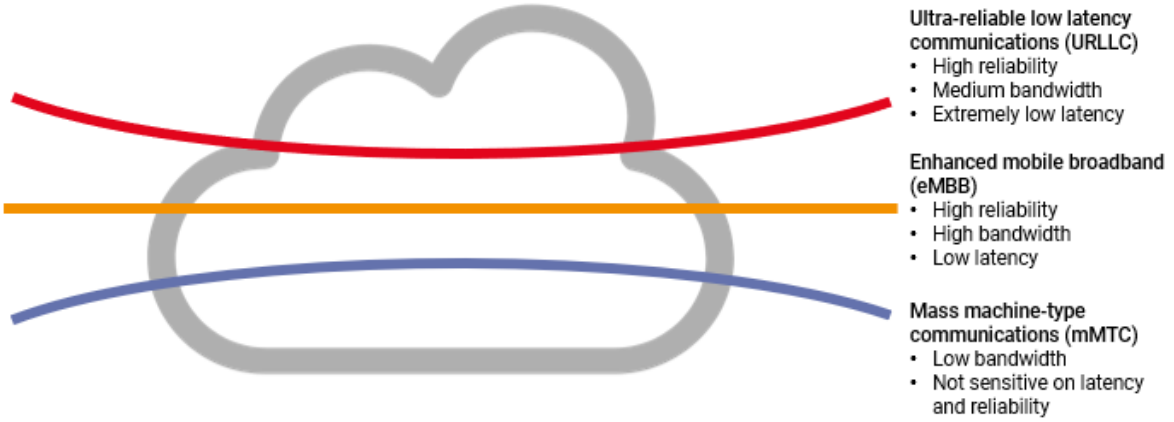


Source: STL Partners

Under this paradigm, a telco would have many customers on each network slice. Customer connections would be assigned to a given slice depending on latency, bandwidth and reliability requirements. For example, customers connecting IoT devices on a mass scale, whether it is a utilities company monitoring its smart meters or a logistics company tracking assets in a warehouse, would be using the mMTC network slice, customers with demanding latency requirements would use the URLLC slice etc.

Having only three generic services serving vastly different requirements and use cases does not quite provide the customisation or "private-feel" that network slicing promises, so instead of defining network slicing within the above scenario, this report focuses on a more forward-looking definition of network slicing that looks at a network slice as a service instance for a **single** customer (or closed user group). Below, we distinguish two different scenarios of how we see customer-specific network slicing working.

The second scenario we see is somewhat of an extension from the first scenario illustrated above. The three fundamental use cases (eMBB, URLLC, mMTC) are seen as generic network slice types, distinguished by the different service requirements. All network slice instances are based on the same network slice templates according to its use case (i.e. eMBB, URLLC, mMTC), which defines the structure, configuration and workflows for the network slice instance throughout its lifecycle[4]. Figure

[4] IETF – Network Slicing Architecture

4: Scenario 2 – Many customer instances of 3 generic slice types illustrates this second scenario below:

**Figure 4: Scenario 2 – Many customer instances of 3 generic slice types**

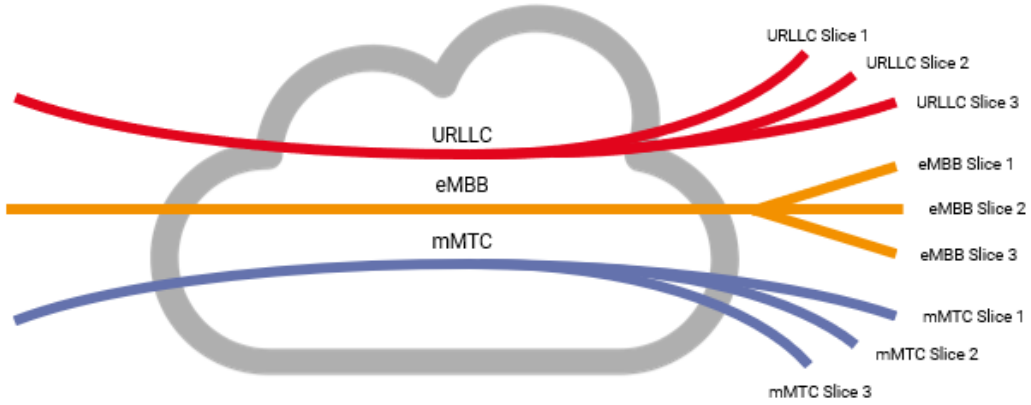


Source: STL Partners

The last scenario is based on complete customisation, where telcos would spin up individual tailored slices per customer or per customer use case (see Figure 5: Scenario 3 – Made-to-order network slicing). This would be much more work-intensive for telcos, in terms of the testing and development efforts required prior to service launch of each slice. It would also require more ongoing maintenance and potentially, more sophisticated automation than the other scenarios.

**Figure 5: Scenario 3 – Made-to-order network slicing**


Source: STL Partners

It may help to think about the three network slicing scenarios using an ice cream-based analogy:

- **Scenario 1:** the three fundamental services are essentially three distinct flavours of ice cream (chocolate, strawberry and vanilla); all customers are served scoops from the same tubs.

- **Scenario 2:** telcos mix and prepare ice-cream flavours for customers on-demand, from a limited menu. These can be standard flavours or variations of each flavour for each customer. For instance, one customer may request chocolate. Although this has been prepared from scratch for this customer, it is essentially the same chocolate recipe enjoyed by everyone else.

- **Scenario 3:** some enterprise customers may require a particular flavour ice cream made to a specific recipe. In extreme cases, the customer may bring its own ingredients, utensils or recipes, along with instructions for the telco. If this flavour has never been mixed before, the telco is going to take longer to learn and get the recipe right and will want to charge extra for this.

Our research has focussed on potential adoption of scenarios 2 and 3.

Scenario 2 is already happening in practice. For example, Deutsche Telekom has deployed a simplified implementation of a network slice for a leading lightbulb manufacturer, Osram, using public and private LTE on one common infrastructure with the aim of moving to fully-fledged 5G network slicing in the future[5].

Many operators, particularly in Europe and North America, expressed concerns regarding Scenario 3. They are, understandably, hesitant to commit to providing fully-customised slices per customer and/or per customer use case, thus potentially resulting in thousands of uniquely different slices, due to the inherent complexity in the management and orchestration at that scale.

The closest example we have seen thus far was from an MVNO that provides connectivity services specifically to the rail and transport industry. It sees itself deploying approximately 8 different network slices, some on multiple CSPs, each designed and customised to meet the specific requirements of its respective identified rail industry use case. Some examples include mission-critical voice for train conductors, remote autonomous train monitoring and assistance and digital services for passenger announcement systems or entertainment. For areas where running a private network is extremely costly and unjustifiable, for example for less busy and more rural train lines, it sees network slicing as potentially offering an improvement on public networks from LTE or other public access technologies.

## Slicing should bring great benefits...

By dramatically reducing barriers to service innovation[6], network slicing has the potential to generate revenue from new services that will help justify the significant network investments that go into 5G.

First, network slicing offers the promise that it will allow operators to deploy new services quickly and cheaply with minimal disruption to existing services, to meet market and customer demands and address new opportunities. It should also enable operators to better serve a wider range of customers

[5] Article: Telekom and OSRAM launch campus network
[6] STL Partners report: Network slicing: the greatest thing since sliced bread

and customer needs. For example, instead of having a private network, the customer could use one or more network slices that are better suited to specific connectivity requirements.

Second, by moving away from the current, static, "one size fits all", in-all-locations paradigm, service providers can cater to diverse needs in different geographies, something that was not previously available with past generations of cellular networks.

Moreover, network slicing should also both enable and require changes in operators' business and operating models in order to support and facilitate the dynamic aspect of slicing: that network slices can be "spun-up", scaled-up, scaled-down and retired when and where needed.

## … but adoption will not be without challenges

With multiple 'networks' being created over common infrastructure, each serving different customers, applications and devices that all need to be managed; the complexity inherent in network slicing will result in various concerns over security. As with any new technology, telcos need to consider the enterprise perspective and what concerns they may have, in order to be able to anticipate and address them appropriately.

This report seeks to answer the question of whether, and to what extent, security-related concerns could be a barrier to enterprise adoption of services that could be enabled by network slicing technology. Later, we outline some strategic recommendations for telcos as to what can be done to achieve and accelerate the adoption of network slicing.

# Enterprise security concerns with network slicing are rooted in the fear of the new and unknown

Network slicing is inherently complex. Multiple networks being created over common infrastructure, each serving different customers, use cases and devices means that management and orchestration of network slices is something that telcos are still grappling with. It not only represents a change in technology but also a shift in the way that the network lifecycle is managed, which is new and unfamiliar to telcos and their enterprise customers. Current security protocols will not necessarily be equipped to cover many of the new dimensions that network slicing brings. This new shift in the way things work will result in various enterprise security concerns. Changes in the network architecture with slicing, with multiple logical networks each having their own resources and sharing others, also poses questions of how the security architecture needs to evolve in order to address new risks.

Enterprise customers define security as not only about preventing services being compromised by intentional malicious attacks, but also about preventing service degradation or disruption due to unintentional operational or technical failures and/or negligence, unplanned breakdowns etc. Due to the interdependence of slices, even if a fault occurrence happens, it could consume resources in one slice, just like an attack would, which would affect the reliability or lifecycle of other network slices that share the same resources. Regardless of how the performance of a slice gets affected, whether it is by a malicious attack, a natural disaster, a bug or unintentional negligence, the consequences are ultimately the same. These are all, in some way, related to security. Therefore, when considering security, we need to think beyond potential intentional malicious attack but also unintentional negligence and unplanned events.

We outline three key questions that enterprises have around their security concerns, as potential tenants of network slices, below.

## What if my network slice gets compromised?

As tenants of a network slice, there are a number of security-related risks for enterprise customers that have been raised in our research, ranging from unauthorised access to the network slice management capabilities to varying levels of security by network slice type. Essentially these risks centre around two sub-questions: "How could my slice potentially be compromised?" and more importantly "What happens if my network slice gets compromised?". By "compromised", we mean 'the failure of a service to operate as intended due to intentional or unintentional disruption on its operation or its resources'.

One potential concern relates to the enterprise's slicing management tools and interfaces. These represent both a potential attack surface and an opportunity for unintended runaway configuration actions. Or straight incompetence. Network slicing provides enterprise customers the choice of being able to manage their own instances of a network slice, a function known as the network slice management (NSM). A network slice manager will have the ability and responsibility to create, manage and terminate network slice instances, either manually or through integration with other systems. There is the risk of impersonation attacks. In the scenario where attackers impersonate the NSM or where other unauthorised users are able to gain access to the APIs for NSM, whether intentionally or unintentionally, they could potentially compromise a critical network function, cause interruptions or terminate a slice and thus affect services running on the slice.

Another security-related concern is around different levels of security controls between slices. Consider the three fundamental scenarios identified for 5G network slicing (eMBB, URLLC, mMTC), these network slice types have different characteristics in terms of availability, latency requirements, throughput type etc. in order to serve different use cases (see Figure 6: Different security levels for each network slice type below).

**Figure 6: Different security levels for each network slice type**

| | Key priority for slice type | Implications on security |
|---|---|---|
| eMBB | Enhanced connectivity with higher capacity | Enhanced security mechanisms currently used for LTE |
| URLLC | Extremely high reliability and low latency | Robust security mechanisms to provide reliability, stringent authentication protocols, cryptographic algorithms and credential management |
| mMTC | Prolonged device battery life | Constrained security protocols – needs quicker and less frequent authentication protocols, lightweight cryptographic algorithms |

Source: STL Partners

As such, they will also have different security requirements and mechanisms, depending on the type of services each slice supports. For example, while eMBB services are likely to adopt enhanced security mechanisms currently in place for LTE, URLLC services require much more robust security mechanisms, such as strong authentication protocol, encryption and credential management, due to the potential mission criticality of some of these use cases.

In contrast, the priority for many mMTC-type services in terms of the overall requirements is to ensure extremely long device battery life. As such, that has made specific choices on the security protocol in network slices for such services, for example the slice may have less frequent/computationally intensive authentication protocols than eMBB or URLLC-type services. With different slices having to make potential trade-offs in order to meet the intended functionality and priority requirements, not just on security, "lower security" network slices can be easier entry points of attack, to compromise resources or to gain insight into infrastructure vulnerabilities, in order to attack "higher security" slices or the network as a whole. Enterprise customers using mMTC-type slices may fear that their slices would be specifically targeted over others, even if they aren't the intended attack target.

The likely differences in security levels between slice types was apparent in our research. For example, we spoke to an MVNO focussed on the rail and transport industries. They expressed the need for managed end-to-end security, covering not only on the network level but also on the application level, and the need for stringent authentication protocols and encryption. The key reason being that safety first is a key motto, therefore security is an extremely high priority. In comparison, a water provider we also spoke to generally sees additional security measures as an overhead (in performance, battery life and cost). As such, it would look for security at the communications layer instead of having to add extra security measures on the devices. However, in general, it expressed no particular security-related concerns in using public networks.

The last point we'd like to make is somewhat removed from the enterprise perspective but is still a genuine concern that could result in one's network slice being compromised. As previously mentioned, network slicing represents a shift away from the current network lifecycle management model and one of the main benefits of network slicing is the ability to create or adapt slices quickly and bring them to production in a fraction of the time than before. However, this benefit also poses a potential procedural risk whereby in the interests of time, enterprises (or telcos) may create slices without bearing the burden of the traditional telco service lifecycle management (e.g. extensive interoperability and standards compliance). In particular, there may be the risk that shortcuts are taken in the testing or configuration phases prior to activating the network slice instance.

This may not be the case for some enterprises, for example in the rail industry, where there are certain requirements depending on the safety integrity level (SIL). Remote autonomous train monitoring and assistance, and the use of GSM-R (Global System for Mobile Communications-Railway) or the replacement Future Rail Mobile Communications System (FRMCS) all need to achieve SIL 4 level, the highest security level, which unsurprisingly requires services to be tested and proven before they are put into service. However, not all industries will have the same measures. Indeed, this ability to adopt different standards and approaches for different slices is one of the key promises of slicing, so long as "different" is still appropriate and sufficient.

## What if another network slice is compromised?

Network slices, by design, share common network infrastructure, underlying hardware and other resources. A Tier 1 European operator have found a general perception across enterprise customers that the more common infrastructure there is, the less secure the service. This concern revolves around the fact that the services enabled by the network slices are all vulnerable to any attack on the network domain or the underlying network infrastructure. By targeting the common network infrastructure, one breach could compromise multiple logical networks. In many ways, this concern echoes many of the concerns initially (and still) raised by enterprises about public cloud services: these have largely diminished in the minds of enterprises.

Building on a point made earlier, if attackers could potentially target one slice in order to compromise another, then that presents more potential vulnerabilities. First, other network slices' management interfaces could also be an attack/incompetence surface. Second, future network slices will potentially support millions of connected devices, which could be a concern for the security of network slices particularly if there are different levels of device authentication between them. As previously mentioned, network slicing is designed so that devices can consume services from multiple slices for different applications. The problem can stem from where many IoT devices are unmanaged, unmonitored and have weak (if any) security protection on them. Each of these devices are then potential entry points for attack and thus dramatically increase the threat surface and therefore drive a need for security mechanisms to address these potential weaknesses. Without sufficient authentication and authorisation controls in place, there are potential risks related to unauthorised access through devices to other network slices and/or sensitive data, breaching both data confidentiality and integrity.

If device connectivity is not properly or securely managed, attackers or other unauthorised users have opportunities to gain access to the slice and to sensitive data through compromised devices. Besides the risk of unauthorised users consuming resources from another network slice, there is the fear that unauthorised users will eavesdrop, steal or even tamper with or publish sensitive data or private information belonging to other users via another slice. The potential threats that could follow with attackers holding such information are all severe concerns. Data confidentiality breaches or theft of customer and enterprise data could have immense consequences on the enterprise's reputation. Many of the enterprises are subject to a lot of public scrutiny and/or hold private and sensitive information about their customers, such as the smart cities bodies, water provider and logistics companies we interviewed, therefore any security breaches would have severe reputational consequences.

## What if another network slice is eating up resources?

There is a risk that comes with potential interdependence of network slices through resource-sharing. Network slices share common finite resources, including hardware-level resources such as memory and compute, as well as network functions. Dynamic slicing brings the risk that one slice could consume resources required by others. With a potentially large number of network slices being created for various use cases supporting different applications, each slice is interdependent with all others that share the same resources, and resources could be eaten up as a result of various things.

One could be due to a denial of service attack (DoS or DDoS), whereby having more network slices expands the attack surface, offering multiple points of entry via any network slice. Moreover, by sharing such resources, each network slice also becomes a potential attack vector for other slices.

Another way that this could happen is if resources in one slice or resources allocated to certain common virtual network functions (VNFs) get exhausted due to any non-malicious event that overloads the capacity of another's network slice. Examples include a firmware upgrade or a mass reboot. This will result in service degradation or failure across all other slices that share those resources. When a capacity of a network slice is overloaded for any reason, that can also inhibit a slice's ability to run their normal security protocol, leaving it vulnerable. When resources in one slice are affected by another slice, there is a concern that enterprises have around how networks will prioritise data in that scenario. This is particularly a concern for enterprises who run more mission-critical services, where public safety is at risk.

Despite the specific concerns we elaborated upon above, the main takeaway from the research was that security is on the agenda when it comes to network slicing, but mainly because slicing is new and still largely unfamiliar to enterprise customers. A global logistics company expressed this exact sentiment in saying that their concerns aren't necessarily about shared hosting, use of public infrastructure or wireless specifically, but it is around the combination of everything and the fact that network slicing technology is still very new. Many enterprises don't yet fully grasp the concept of network slicing and the lack of knowledge leaves them understandably wary.

# Security concerns will slow adoption if not addressed early and transparently

Although there are clear enterprise opportunities for network slicing, as confirmed in our research, there is a general lack of understanding about network slicing and therefore enterprises are understandably wary. Telcos shared that across their customer base, regardless of industry, security is always front of mind. The threats or risks to their business may vary between customer privacy and company data, regulatory compliance, company reputation and employee safety, but all enterprise customers will look to ensure that the connectivity services that they are using meet the security requirements they have.

How enterprises prioritise security largely depends on the nature of their business and industry. For some industries, such as manufacturing, transport and defence, security is a top priority given the potential consequences of a breach or failure in terms of financial cost or impact on employee safety. Other enterprises cite security as being important but further down the pecking order in terms of priorities.

For example, one third-party smart logistics solution provider told us that connectivity is only 3% of its current spend and is just a means to an end – from this company's perspective "it either works or it doesn't". Public networks as-is work well enough in accommodating its requirements so enhanced security at the network level is seen as "nice to have" but given that connectivity is a small part of what third-party logistics companies need, additional security at the network level is not as important.

Given that network slicing is not a commonly understood concept yet, security-related concerns can slow down the speed of network slicing adoption. These concerns are important issues that telcos should not underestimate, but ultimately, they are not a barrier to adoption and can be overcome if telcos can be more open and address them early and head-on.

## Concerns and misconceptions can be addressed through better awareness and understanding

An overwhelming theme of our research is that network slicing is not well understood. Although network slicing has been widely discussed in the telecoms industry, the concept of network slicing is still relatively new to enterprises. Enterprises have either never heard of it or have vaguely heard about it within the context of 5G. A global logistics company stated that the security concerns they had weren't specifically about shared hosting, use of public infrastructure or wireless but more about the combination of these and the fact that network slicing technology is still very new.

Many enterprises who know very little about it are also unsure why network slicing matters when they already have more immediate issues they need to deal with. Others see it as new exciting technology that they could adopt to position themselves as leading innovators in their respective industries, but also don't quite understand what the case for adopting network slicing is.

This is not surprising given that telcos themselves are still at a relatively early stage in defining, let alone implementing network slicing. One particular North American telco said that for them, although network slicing is still beyond their existing horizon (i.e. their current immediate focus), the second horizon where network slicing sits is fast-approaching.

Telcos are also still grappling with a number of aspects of slicing, given there are very few standards defined, particularly around the management and orchestration of network slices. Many are still figuring out various elements such as best practices and commercial guidelines for their customers. With time and more clarity around how telcos will deploy network slicing technology, telcos can better educate enterprise customers about the "how" as well as the "what" of network slicing.

## As a result, enterprises project concerns about public networks' limitations onto slicing

Due to this lack of understanding, enterprises project many of the concerns they have about public networks around security, reliability and service availability onto network slicing. Some enterprises may also equate the cloud characteristics of slicing with public cloud services and hence also extend, to slicing, the same reservations they have about public cloud services.

Although the perception of public networks is dependent upon the enterprise and their current network connectivity requirements, public networks in general aren't seen to be good enough to meet specific requirements that enterprises have. Enterprises that currently use public networks have stated that these are not necessarily meeting their needs, but they work for now and there is sometimes a lack of standard commercially available alternatives. One regional water company in Europe said that public networks will 'suffice for now' but is keen for operators to roll out other alternative networks with lower power requirements for devices. Other enterprises with more demanding requirements or mission critical use cases feel that public networks are "best-efforts" rather than "industry-grade".

Enterprises have also raised concerns about the reliability of public networks, and that public networks or services can fail for any reason. Enterprises also argue that they don't have confidence in the service availability of public networks. Private networks (i.e. dedicated networks such as private LTE, GSM-R or TETRA) are seen to be inherently more reliable and resilient and therefore will always be the first choice if cost is not an issue.

Below are some of the concerns that enterprises have raised about network slicing. For each, we offer to explain how they actually relate to the limitations of public networks and aren't specific to slicing:

**Figure 7: Limitations of public networks vs. network slicing**

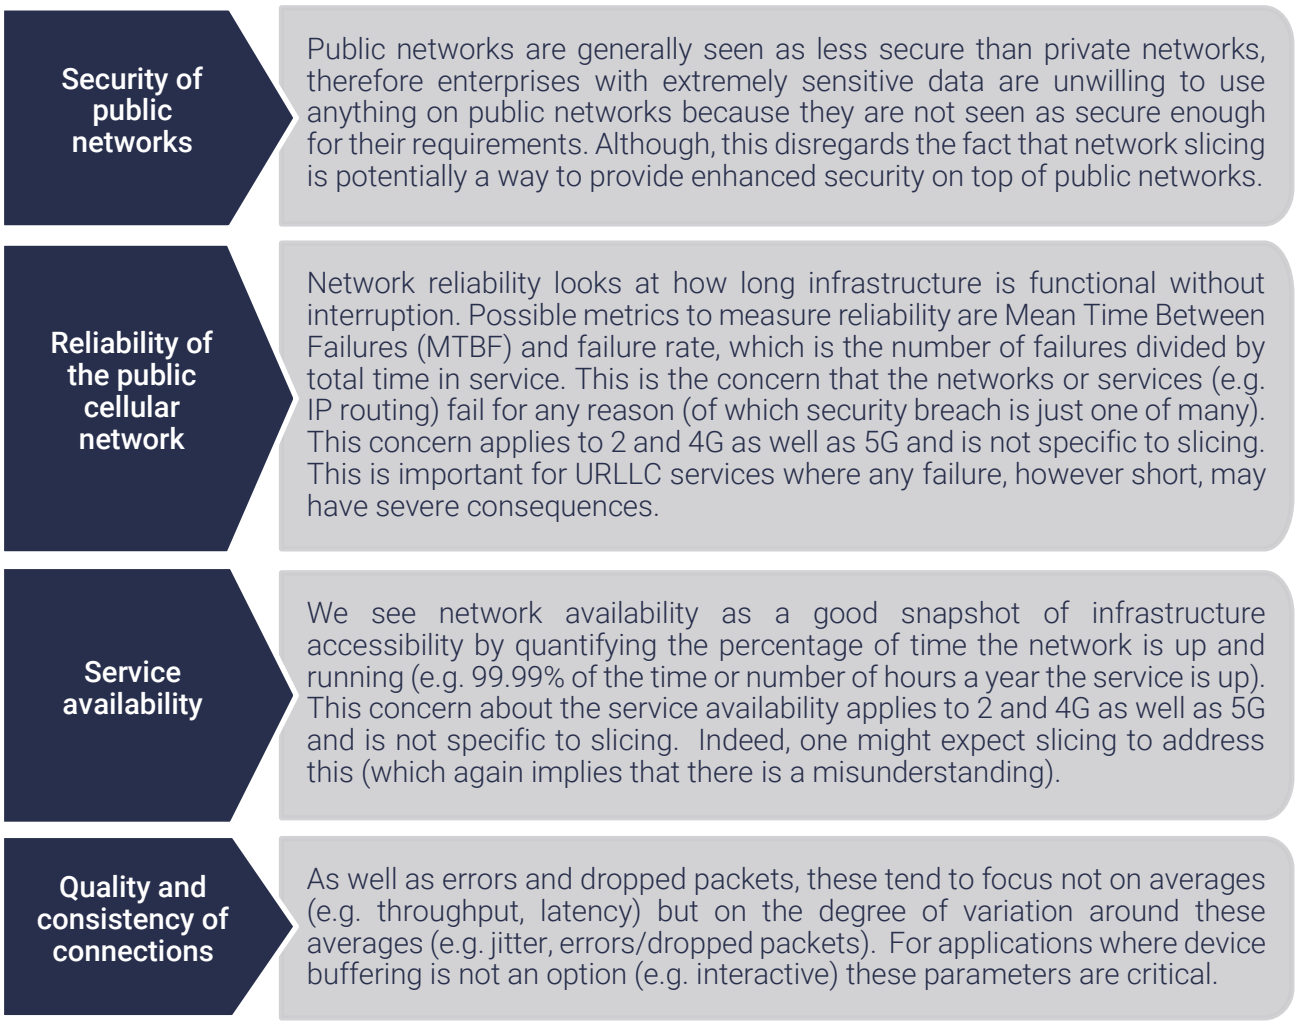


Source: STL Partners

It is clear to us that there is work to be done to get to where we need to be for network slicing to work in terms of understanding and awareness. Many enterprises miss the fact that network slicing inherently seeks to resolve a number of drawbacks of public networks. It enhances operators' ability to meet specific performance characteristics, meet SLAs and provide not only greater stability and reliability in the public networks, but also enhanced security through isolation, which we will explore later in this report.

Enterprises need to be better informed about what network slicing is and by being better informed, enterprises will increasingly understand that these concerns can be mitigated through network slicing. The key proposition of network slicing, that isn't widely understood by enterprises, is that it inherently seeks to resolve a number of drawbacks of public networks. It enhances operators' ability to meet SLAs and provide not only greater stability, predictability and reliability in the public networks, but also enhanced security through isolation, segmentation and other controls (which we will address in the following section). It is also key to note that network slicing (on public network infrastructure) can only do so much to overcome coverage limitations of public networks. This is also a major factor that

enterprises are concerned about. This needs to be addressed by extending slices with dedicated infrastructure e.g. small cells.

Telcos, on the other hand, need to better understand and fully appreciate the requirements of enterprise customers (and the potential impacts of network issues on customers' operational processes) so that they are able to address these concerns appropriately and sufficiently.

## The way that network slicing is designed actually enhances security, and there are additional measures available on top.

Network slicing is actually a way for operators to provide enhanced security and reliability to what are essentially private or quasi-private networks on the same underlying infrastructure that supports public networks. One smart city initiative in Europe sees network slicing as a way to ensure that a reliable and secure portion of the network is always running, particularly for emergency/mission critical services. What is not widely understood is that network slicing is designed to be more secure than public networks, through network segmentation and isolation of network slices, and tighter security protocols based on customer- or use case-specific requirements (authentication, authorisation, encryption etc.)

The first two are of particular interest, but we explore all of them in more detail below:

### Network segmentation

Network slicing allows operators to segment the network infrastructure into different areas based on specific purposes, which acts as a mechanism to restrict access to specific slices to certain people and devices. By creating many networks over one network infrastructure, operators can apply specific security protocols and controls based on the purpose.

Let's take a factory for example. Factories are likely to have multiple dedicated private networks, each using assorted industrial network protocols and serving specific tasks such as real-time control of machinery and equipment. The networks on which each application run on tend to be separate, which means that if one of the networks were to be breached, the impact of the attack would be restricted to the application(s) running on that specific network. However, not all organisations and industries are willing or able to fork out the investment and ongoing costs required to have their own private network for each use case or each of their facilities. This heterogenous private networking fabric also raises major issues around long-term support for "dead-end" technologies. With network slicing, these organisations could have separate slices for different use cases, which differ in terms of mission criticality and connectivity (and security) requirements.

By providing network slices for each use case, network slicing provides the ability to restrict access to each slice only to authorised individuals and authorised devices, thus already making it more secure. If that private network gets attacked or if the network goes down via unintentional failures, then all the services that run on that one network are potentially impacted. By providing each use case with separate slices, that not only provides more security for the services but also greater reliability. If a security event happens on a slice, and if the slice is sufficiently isolated (see below point), only the

services running on that respective slice are affected. There should be no effect on services running on other slices.

## Isolation of network slices

Many security concerns raised relate to the potential impact of one slice on another slice. These concerns ultimately boil down to the level of isolation between network slices.

Network slice isolation refers to the level of sharing of allocated resources between slices. Total isolation essentially means that resources allocated to a slice are entirely separate to the resources allocated to other slices. Network slice isolation can be designed-in and can ensure that the function of one slice does not affect the function of another slice. Proper isolation between network slices ensures the following:

- **The consumption of resources in one network slice instance does not impact the consumption of another network slice.** The 3GPP requirement relating to this specifically states: "Traffic and services in one network slice shall have no impact on traffic and services in other network slices in the same network".[7]

- **Potential network attacks are confined within the affected network slice/network slice instance, therefore preventing spread and denial of service to other slices.** The 3GPP set out the following requirement: "The 3GPP System shall have the capability to provide a level of isolation between network slices which confines a potential cyber-attack to a single network slice".[8][9]

Isolation is a key aspect that has been continuously specified as a requirement, but so far there have been no specific standards set on the level of isolation. It is a potential means of protecting sensitive data and containing attacks or network slice failures. A Tier 1 European telco touted it as the best way to secure a network slice. The level of isolation is dependent upon the baseline security level and also upon customer needs and their willingness to pay – the higher the level of isolation, the more secure and reliable the slice. However, the higher the isolation, the more resource-heavy and costly for the customer due to having more dedicated resources allocated to that slice regardless of workloads.

There are a few dimensions within which operators can isolate slices (see Figure 8: Dimensions of network slice isolation):

[7] 3GPP TS 22.261 version 15.5.0 Release 15
[8] 3GPP TR 22.864,
[9] Wireless World Research Forum white paper: End to End Network Slicing

**Figure 8: Dimensions of network slice isolation**

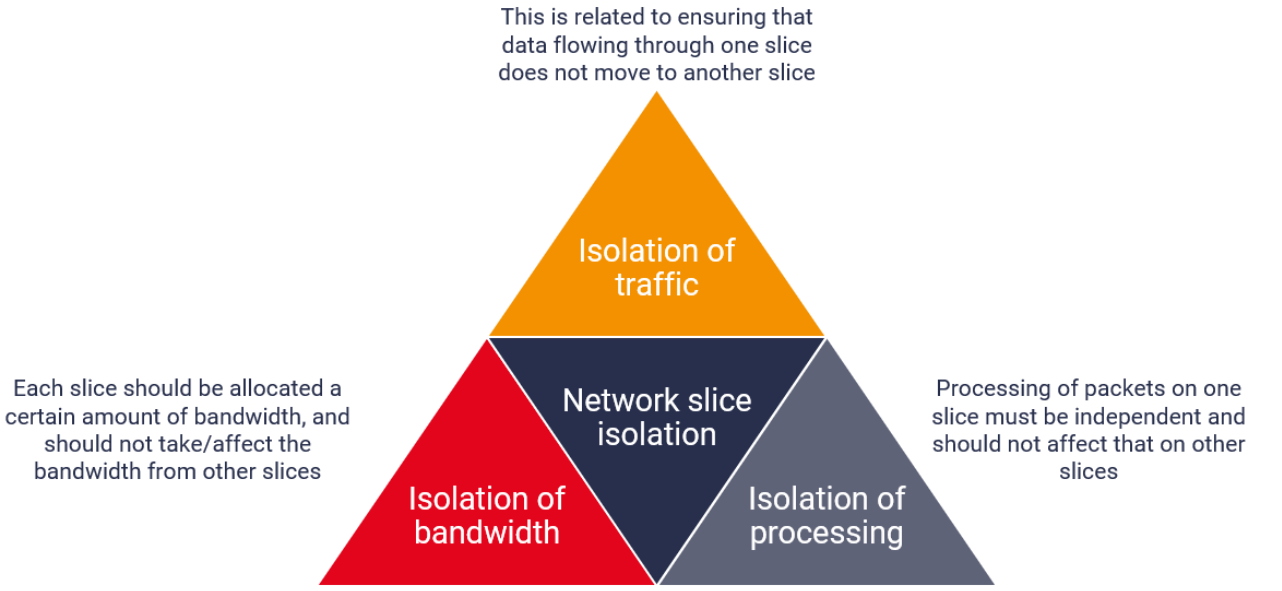


Source: 5G Americas[10], STL Partners

## Network slice authentication and authorisation protocols

Authorisation and authentication are two of the few mechanisms that the 3GPP have outlined[11] as part of the security requirements for network slicing. This relates to authentication and access authorisation of devices, users, virtual network functions within the network slices etc.

Authentication should involve protocols set around separate authentication of devices accessing multiple slices and the frequency of re-authentication.

Having strong authorisation and authentication protocols also helps to control the management of slices by slice tenants, whereby such protocols can prevent impersonation attacks against other network slice managers, or against other network slices themselves. Preventing the former would mitigate any corruption, removal, disclosure and interruption threats.[12]

## Encryption

Encryption is another way to provide customised enhanced security on network slices. Providing cryptographic protection (through applying different algorithms for example) is a way of protecting the privacy of network slice tenant data, whether it is user data or other types of data. It also prevents attackers from eavesdropping or tampering the data from other slices. Any cryptographic algorithms should be configurable based on the needs and requirements of the use case.

[10] 5G Americas – Network Slicing for 5G Networks and Services
[11] 3GPP TS 22.261 version 15.5.0 Release 15 – Sections 8.3 and 8.4
[12] NGMN 5G Security – Network Slicing Version 1.0

# Telcos must act early and work more closely with customers to drive slicing adoption

To overcome customer scepticism, telcos must drive awareness and understanding of network slicing. Telcos shouldn't wait until 5G is rolled out to start the transition towards network slicing. There are a few key things that telcos can do whilst the technology evolves into the fully-fledged network slicing that the industry envisions – we have identified a few of them below and outlined what can be done:

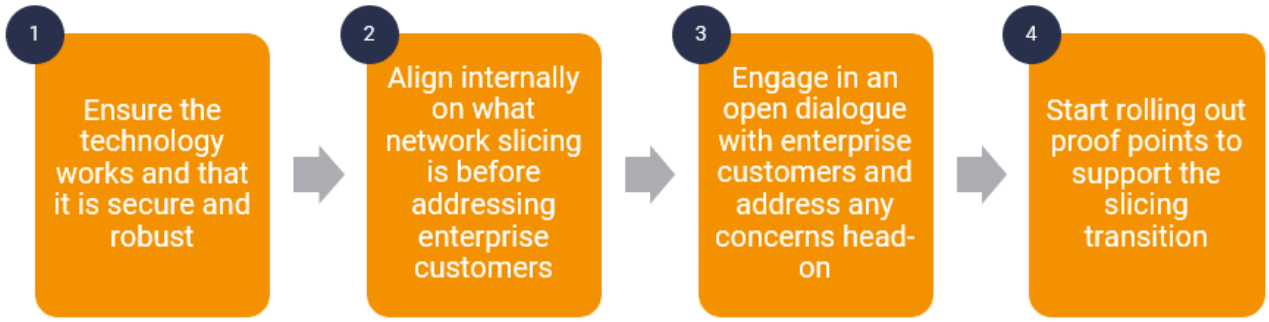


## Ensure that the technology works and that it is secure and robust

This needs to be addressed much more widely at an industry level, and not necessarily specific only to telcos. Telcos have expressed an ongoing need for more standards and guidance from industry bodies and/or technology vendors in defining what roles look like and how network slicing will work in practice. In terms of what telcos need to address in particularly, they must look to implement and test security measures (such as isolation) that provide a sufficient baseline security level to ensure that security parameters in the slices with fewer security requirements do not present risks to other network slices. Network slicing is as secure as the lowest common denominator, therefore the definition for the baseline security level is crucial. The baseline security level across all network slices can be supplemented with further measures depending on customer requirements. Ensuring a robust baseline security level is crucial because although many enterprises (particularly larger ones in certain industries such as in manufacturing, transport and defence) will look to add their own security measures/solutions on top of the security provided at the network level, others see additional security as a resource-hungry overhead and instead are keen to use security at the network level.

Network slicing is still far from being reality and standards have yet to be clearly defined particularly around the "how". Another aspect that many telcos are grappling with is the management and orchestration of network slicing to enable full lifecycle management of network slices and associated network resources. How many slices is too many? Network slicing looks to support numerous use cases, each with their own requirements, so orchestration will become increasingly complex, posing challenges with operationalisation and scale. Telcos need to figure out how to manage the complexity in order for it to not become unmanageable, and still look to their technology partners for support.

Significant testing will also be required to ensure that the security controls in place are robust enough against attacks and breaches. One industry-specific MVNO we spoke to stated that "safety first" is the motto, particularly in the defence and transport industries. In the network slicing context, that essentially means that the network needs to not only be secure, but also resilient. Testing therefore cannot only include preventing security breaches but also ensuring that reliability and performance are not affected by intentional or unintentional actions/events. Although it is not possible to test for every potential eventuality, a robust evolving testing strategy will also help learning and wider operationalisation of network slicing.

## Organise and align internally on what network slicing is and where it fits internally before addressing enterprise customers

Although the concept of network slicing has been around for some time, the industry in general is still in the early phases of network slicing development - there is still large variation as to where different telcos are within that. Some leading telcos are testing simplified implementations for live deployments, whilst others are still figuring out what network slicing could mean for them.

Even within telcos who participated in this research, some are still in the relatively early conceptual phase in thinking about what they might use network slicing for and how they might use it. In contrast, one particular MVNO is already at a much more advanced stage, having clearly defined specific use cases with associated functional requirements in the rail and transport industry, with mission critical voice communication for rail conductors as the first use case that will be deployed. One other Tier 1 telco already has simplified implementations of network slicing up and running ("proto-slices" with dedicated frequencies, some shared infrastructure such as antennas/base stations and some dedicated infrastructure such as core network functions).

The differences in the way that network slicing is defined has implications on how it works and what the technology offers. The whole organisation must have the same definition of what network slicing is in order to clearly communicate this forward to enterprise customers and develop a well-defined business case for it. This definition should also shape the way that telcos organise themselves internally.

There must also be internal alignment on what the network slicing proposition is. Once that is achieved, telcos can address customer problems and identify opportunities where network slicing technology is a feasible solution. Simplicity has been cited as key in building the proposition. Enterprises don't care about the nitty gritty ins and outs of how network slicing technology works, but they do want to know what this means for them. In many ways, a network slice could simply mean that enterprise customers experience better services at lower prices, but either way the proposition needs to be easy to understand from both the internal and customers' point of view. Telcos don't sell a network slice, they sell a solution to enterprises' business needs and problems.

We see three broad internal narratives that operators can align around on network slicing:

1. Slicing as an **evolution of the existing network** and the services that run over it. This narrative emphasises the increasing flexibility of a single overall network. For example, the introduction of NB-IoT on LTE networks as a first step to slicing. The approach emphasises efficiency, scale economies and centralised control/orchestration. Security is reinforced through the operator's rigorous policies and processes. However, it then becomes harder to position slicing as fundamentally different from public network services.

2. Slicing as **private networking**. Under this narrative, the core proposition to enterprises is private networking, with slicing considered a delivery mechanism. Potentially, the same private network could run on dedicated infrastructure (e.g. for on-site coverage) and run on shared infrastructure (e.g. for wide area coverage). There are many potential "hybrid" combinations including only using shared infrastructure (essentially a network slice). The emphasis is on multiple networks rather than the one network. Reliability and security are assured through the isolation implied by private networking. It is easier to position slicing as fundamentally different from public network services, but this approach also creates potential customer expectations of using dedicated resources.

3. Emphasise **needs, performance and services**, not technology. Under this narrative, the focus is on meeting customer needs and performance requirements. The underlying technology is for the operator to determine, not the customer. Slicing is part of the toolkit that operators use to meet customer needs, not part of the proposition. Although this approach might work for consumers and some enterprises, our research suggests that it will not work for most enterprises as it does not address many of their potential concerns.

Based on our research, we propose that network slicing should become a part of the private networking services menu, as part of the second internal narrative described above, in order to directly address some of the concerns that enterprises have around security and reliability. The other benefit of this is to address the clear preference for using private networks over public. The proposition to enterprise customers should be a private networking solution enabled by network slicing that helps to address their key business needs and problems. There will be a period of transition, but this will help in the long-term.

## Engage in an open dialogue with enterprise customers and directly address any concerns via a 'hand holding' approach

Telcos should be upfront with their enterprise customers about the perceived potential risks of network slicing, to demonstrate that they've thought this through, and have evidence to show that these risks have been addressed and are being proactively managed. Telcos have made comparisons between the transition towards network slicing and the transition from dedicated infrastructure to cloud infrastructure, in terms of the initial security concerns and economics. Therefore, if telcos

communicate openly with enterprise customers to address concerns early on, that should ease the transition towards network slicing.

Open communication and co-creation approaches will also allow telcos to gain a better understanding of the customers' needs. Telcos should work more closely with enterprise customers and adopt a 'hand holding' co-creation approach to help customers better understand network slicing and work with them to provide the level of comfort they seek. Two telcos that we spoke to admitted that telcos need to do better in having different conversations with their enterprise customers in order to gain a better understanding. By learning about the vertical-, customer- and use case-specific needs, telcos can more easily promote and apply network slicing when the technology is ready to customers they're already working with.

However, one North American telco cautioned that this new way of engaging with enterprise customers is a bit of a departure from how things are currently done (see Figure 9: Telcos need to change the way they sell below). Telco sales teams may be more accustomed to selling off-the-shelf connectivity solutions. In contrast, network slicing requires telcos and their sales teams to have a strong understanding of what their enterprise customers' business requirements and needs are, in order to be able to identify where slicing can fit in. This requires retraining to build a different skillset to allow telcos to have these conversations with customers, because selling network slices successfully requires a lot of technical expertise and understanding of the customers' commercial business. Telcos should not sell a slice, but instead sell a solution that solves the enterprises' problems through using a network slice.

**Figure 9: Telcos need to change the way they sell**


Source: STL Partners

# Don't wait for maturity to start testing and rolling out pilots to support the transition and learning process

The transition towards network slicing doesn't have to start when 5G is ready; there are ways that telcos can start on that path now, with existing technology. We see the roadmap of network slicing being a broad 3 step process (see Figure 10: Key steps for telcos to drive slicing adoption):

**Figure 10: Key steps for telcos to drive slicing adoption**

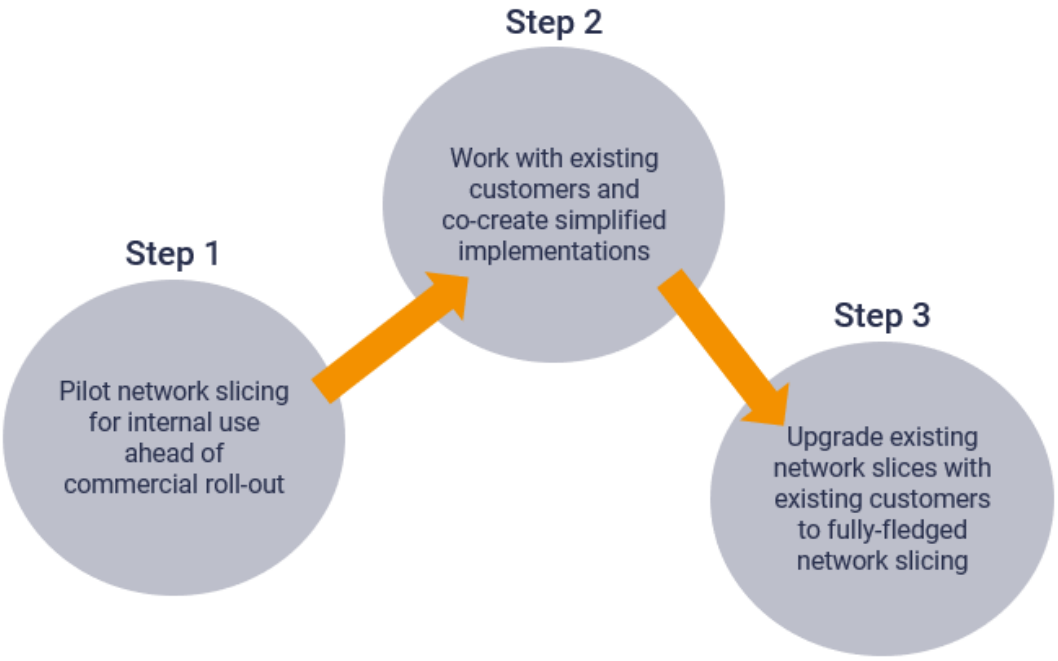


Source: STL Partners

## Network slicing piloting for internal use

Telcos can look to address our first recommendation (i.e. ensuring that the technology works and is secure) through using network slicing for internal use. This could involve deploying pilots for internal "customers" or sister companies to experiment, validate and refine the proposition ahead of commercial roll-out.

Deploying network slicing for internal use is a much lower-risk scenario for telcos to gain a bit more experience, particularly around the management and orchestration elements, and to iron out any minor problems or details in the processes and technicalities.

## Co-creation with existing customers

Before waiting for the vision of dynamic network slicing under 5G to be ready, telcos should be looking to start with simplified implementations, with existing technology on their 4G networks. Telcos need to play a more active role in trying to solve the problems that enterprise customers have today with the current technology available (e.g. NB-IoT or CAT-M1 or regular LTE). Some have admitted that this is something that is not done enough. Although it's a long way from the ultimate vision of network

slicing, the concept of creating customised services and delivering them with specific requirements to a specific customer group or use case is still the same.

As previously mentioned, telcos should take a guided "holding hands" co-creation approach to initial implementations of network slicing with existing customers, where there is a clear interest and immediate fit for network slicing, in order to build confidence and trust not only in slicing itself but also in the telco. Enterprises involved in our research have stated that they would look for investment from a telco partner to work closely together in a partnership both sides are invested in, to get the level of comfort needed around the security aspects, in order to build confidence in the solution.

In our research programme alone, we came across many scenarios with strong opportunities for network slicing. For the rail industry, a leading company in the transport, defence and aerospace industries stated that hybrid slicing can address cases where private networks are too expensive and unjustifiable, particularly for the less busy rural train lines. Another enterprise, a regional water provider, expressed keen interest in using a network slice if it's more customised to its needs than existing 4G networks but would want to have a better idea of how network slices are managed and how prioritisation of different services will work. Many enterprises also see a clear economic benefit for using network slices if it allows them to still enjoy the "private feel" of network slices and more customisation without having expensive private networks. Multiple individuals we spoke to in smart city initiatives in smaller cities across Europe stressed the importance of cost. As much as security is important, at the end of the day smart city initiatives need to be able to justify the price against the opportunity cost of providing other key public services that directly contribute to the wellbeing of the community, particularly in smaller cities. Other smart city initiatives in large cities and some leading enterprises in other industries saw network slicing as a way to position themselves as leaders and innovators in their respective industries.

It is clear that there is a case for network slicing – it can offer an improvement on public from LTE or other public access technologies at a more economic cost, but telcos need to be able to identify where the opportunities for network slicing are and key factors that appeal to each customer in order to offer an appealing alternative to current connectivity solutions.

Fully-fledged network slicing technology isn't quite here yet but there are currently trials by operators – one Tier 1 European operator has live deployments with "prototype slices" with two large enterprise customers whereby the network slices have their own dedicated infrastructure and spectrum but share the antennas and base stations. These simplified implementations should evolve to something much closer to more dynamic genuine network slicing.

While telcos work on refining the technology for fully-fledged network slicing, they can also start engaging with their customers (part of our third recommendation) to ensure that when the network slicing technology is ready, telcos can "upgrade" existing services with current customers. These current customers will then act as proof points for new customers.

## Upgrade existing implementations with existing customers when ready

Once network slicing technology is ready, telcos can upgrade existing simplified implementations of network slicing with current customers to the fully-fledged network slicing, which will be more dynamic. These customers will then act as proof points for new customers or other existing more "conservative" customers. Ultimately, most enterprises that we have spoken to do not have any ambitions of becoming a connectivity provider; they have their own core businesses and continue to look to telcos to provide the connectivity element. Those that have their own private networks often do so because they feel that it is the only way to meet their demanding requirements on reliability, latency and security.

# Conclusion

Although network slicing promises many benefits for both telcos and enterprises, the industry is still a long way out from making network slicing become reality. The network slicing vision relies on end-to-end virtualised networks, which as we know we are still far from, in order for it to be as dynamic, flexible and agile as promised. We anticipate the fully-fledged automated network slicing to be deployed 2-3 years after we have significant rollout and coverage of standalone 5G networks: for most countries, not before 2025.

However, because of long investment cycles and steep learning curves for enterprises, telcos need to act now instead of waiting until the technology is ready for fully-fledged 5G network slicing. There are many takeaways from this report but if there are three things to keep in mind going forward, we would emphasize the importance of the following:

1. **Focus on network slicing as part of private networking services** – Our research has shown that many enterprises have concerns about public networking and therefore, implicitly, network slicing. However, we see network slicing essentially as a way to deliver private networking services on shared infrastructure. For that reason, telcos must focus their efforts on positioning network slicing as part of a menu of private networking services. This refers to internal organisation, proposition development, go-to-market strategy etc.

2. **Be open with your customers** – Network slicing is a term that has been thrown around a lot recently. Based on our research, enterprises do not have sufficient understanding about network slicing and therefore have outstanding questions and concerns about security and other key requirements. Telcos need to engage with enterprise customers early via a 'hand holding' approach to help build their understanding and directly address any concerns through open communication. Pre-empt any questions that may come up and prepare appropriate answers to them – this will help convince enterprise customers that any potential risks are addressed and proactively managed.

3. **Don't wait for technology maturity** – Telcos can start testing and learning from the process now. Telcos can either work with existing friendly enterprise customers or deploy network slicing for internal use to build more experience ahead of commercial roll-out. Some leading telcos are already piloting simplified implementations with existing available technology and will use these implementations as a stepping stone for when fully-fledged network slicing is ready.

STL PARTNERS

Research

Consulting


Events

3200079-001-EN Dec 2019