www.pipelinepub.com
Volume 18, Issue 7

The Future of Network Slicing

By: Mark Mortensen

3GPP standards provide specifications for communications service providers (CSPs) to "slice" their shared network to support near-guaranteed different QoS characteristics for various service types and service instances. To create and control these networks, vendors have created slicing managers that are managing the full lifecycle of slices for trials with commercial services soon to start in earnest. Still in question is whether these slicing managers represent a new category of OSSs or are just for the trials.

Slicing is just starting

Network slicing is in its technical adolescence but its commercial infancy. Although it was specified in 2017 by 3GPP standards, there are only a few examples deployed in networks today. Most CSPs are experimenting with slicing in POCs, working with leading-edge vendors. These vendors have added slicing to their domain control and cross-domain orchestration systems and to selected network functions. Commercial deployment is expected to start in earnest in 2022 to 2023. Major commercial growth is expected in 2024, focusing on several use cases that are now becoming clear, especially special events venues, private 5G networks, and MVNO slices. Slicing of the network into one or two dozen slices is reasonably easily achievable and will dominate in the next two years. Scaling up to hundreds or thousands (or even tens of thousands) of slices will require massive automation to handle the operational complexity throughout the service lifecycle.

Network slicing requires an overall lifecycle, end-to-end network view that includes design, provisioning, inventory, and assurance functions, all working together (the P-A-I-D functions). Most CSPs desire that these be provided as a "decomposable suite" that provides all the functionality but partitioned into microservices-based components that can be selectively replaced in the architecture by other vendors' components (usually ones that are already in place).

Network slicing is being implemented in a way that requires no changes to the network equipment (whether virtual or physical) already in place. Instead, a combination of specialized network surveillance of sliced resources, slicing markers in an inventory system, and slice provisioning templates are implemented in external software systems—integrated into, or integrated with, domain control systems and cross-domain orchestration systems. All the vendors are implementing a combination of standard interfaces as defined by the GSMA 3GPP, TM forum, IETF, and ETSI standards bodies.

Many network function vendors are ensuring that their elements and their domain control systems can support network slicing and integrate with cross-domain orchestration systems. Vendors such as Infinera, among others, are following this path. Some NF vendors, such as Huawei, are providing a full solution, ensuring that their NFs are fully supported with external software for the DC and CDO functions that supports both their own NFs as well as those from other vendors requested by their CSP customers. Ericsson has this same basic strategy, with more of a professional services-oriented approach than others.

Some NF and independent software vendors (ISVs) are offering comprehensive multivendor slicing management systems, basically CDOs enhanced with added slicing functionality, with a goal of excellent multivendor support and usually incorporating the DCs of other NF vendors (where available) into their architecture. Leaders in this area are Amdocs, Ciena Blue Planet, NEC/Netcracker, Sedona, and Oracle.

Other ISVs are building more targeted software components, usually based on open source, for DC and CDO functionality. ISV service assurance and inventory specialists are working with other vendors to integrate their products into overall slicing management architectures. TEOCO is a leader here.

Other specialized players are providing specialized slicing NFs and DCOs that provide enhanced slicing capabilities in the NFs themselves. Kaloom and Pluribus are standouts in this area.

Network slicing: the state of the art

Nearly all network slicing vendor offerings today involve people engineering the network slice to specified quality of service (QoS) values, then monitoring it via service assurance OSSs to ensure that the QoS values are met. If they are not, then manual methods are employed (usually via engineering changes).

Most network slicing schemes in effect today are static and based on implementing specialized service surveillance. Two years ago, most solutions I had seen were in their infancy and varied quite a bit in their approach. Few were "close to the network" solutions and depended upon external OSS (provisioning, inventory, planning and design, and service assurance system) and even BSS functions. They mostly were what I called *painting the slices onto the network*. In this technique, slice creation means explicitly engineering a set of (or portion of a set of) network resources for the desired characteristics and then monitoring them according to their specific KPIs. Although major failures were to be mitigated against by backup plans, more subtle SLA violations for the slices were handled via manual re-engineering, with a long lifecycle. There was little dynamic behavior exhibited in most cases.

In 2021 and 2022, vendors made great strides in their slicing solutions as slice design and provisioning was added, with substantial automation in cases where the underlying domain controllers were in place. Assurance functions became more sophisticated, with vendors adding latency and other QoS functions to their slicing solutions. Inventory systems, completely refactored or written anew by non-traditional inventory players, became real-time repositories of the state of the network, feeding both the provisioning and assurance part of slicing operations. To feed the provisioning process, slice design functions were written, usually using predefined parameterized slice templates. All of this was done using cloud-native software architecture and delivered on multi-cloud infrastructure with CI/CD processes.

There is little in the way, however, of deployed commercial slicing offers, although the industry has seen many POCs. ACG Research anticipates that the commercial market for vendors' slicing offerings, and CSPs' offerings of sliced services, will begin in earnest in late 2022 and accelerate in 2023 with CSP implementations of a small number of dozens of slices implemented on any one network. These numbers will grow as the operations are improved, scaled up, and automated.

In the future, we also expect that additional functionality will be implemented in the network elements to provide a finer degree of control for network slicing. We also expect that the current GMSA list of standard slicing types will be greatly enhanced.

Network slicing: techniques

It has always been possible to ensure a customer's QoS guarantees are not violated by dedicating equipment to that customer, or selected services for that customer. But the idea of slicing is to have as much of a common network as possible with capacity within the network resources dedicated to certain slices, with known and configurable QoS parameters for the slices, sub-slices, and services. There are two somewhat different techniques for doing this as shown in Figure 1 on the next page. The first is Isolation, how the resources are set aside in the network. The second is Adaptability, how dynamic is the slice (as in how much of the slice operations are automated, allowing them to be quickly changed to meet changing customer requirements or adapted to meet changing network conditions).

Isolation of slice resources

How are the resources set aside in the shared network? Just done statically by the engineering department painting the slices onto the network, with nothing special being done other than monitoring more stringent QoS parameters? Or more sophisticated techniques in the network

Design, engineer and monitor slices

The basic way of doing slicing is via the way that network resources have always been dedicated to individual customers or services. In this basic methodology, the CSP "paints the network" with the slice identifiers in the external OSS systems. This requires no new functions in the equipment, nor any specialized functions in the OSS, domain control, or cross-domain orchestration systems. The resources are flagged in the inventory system as belonging to a particular slice while the service assurance system monitors the QoS parameters, perhaps with more stringent requirements than normal. If SLA violations occur, then the engineering department is notified, and remedial action is planned and put in effect. This can, however, take days or weeks.

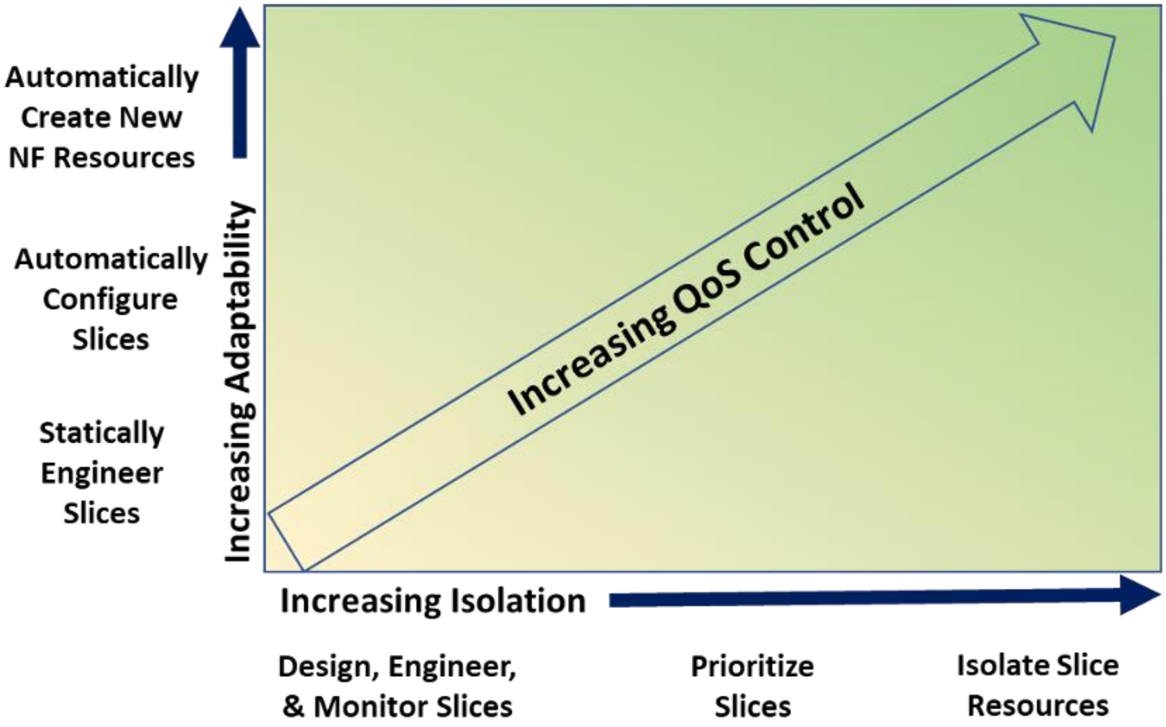
Automatically Create New NF Resources
Automatically Configure Slices
Statically Engineer Slices
Increasing Adaptability
Increasing QoS Control
Increasing Isolation
Design, Engineer, & Monitor Slices
Prioritize Slices
Isolate Slice Resources

Figure 1: Slicing techniques

Prioritize slices

With shared equipment that has priority queueing, admission control, or other features that allow paths to be given priority above other services or slices, slices are given the appropriate QoS parameters. This can mean a low priority for mMTC slices.

Isolate slice resources

Soft isolation of the NF resources for a slice, such as is done in data center networking by companies like Kaloom and Pluribus, can provide additional security and fine-grained control of the resources.

Adaptability of slices

The second technique to gain control of the QoS parameters is to make the slices more adaptable to changing network conditions or customer requirements. This requires capabilities both in the slicing control software as well as the network elements themselves.

Statically engineer slices

By setting aside resources and setting service priorities in the network during the provisioning process, slices can be created. Then, if conditions change, they can be reprovisioned using the standard manual processes.

Automatically configure slices

Providing a way of configuring slices both initially and throughout their lifecycle gives greater control over the QoS, allowing the systems to anticipate and mitigate against expected QoS violations by reconfiguring the parameters of the slice (which resources are used, routes taken, and so on) or even reconfiguring the other slices sharing the same resources.

Automatically create new network function resources

If the network resources that the slice uses have been virtualized, then more NF capacity can be automatically created, and the slice configuration changed to take advantage of the new capacity.

Network slicing solutions basic architecture

The basic architecture of end-to-end slicing managers is shown in Figure 2 on next page. They include the full set of P-A-I-D functions (provisioning, assurance, inventory, and design), interacting with upstream systems that include the additional business functions, and, on the south side, withdomain controllers for each of the domains (or sub-domains). The domain controllers need to have network sub-slice management functions (NSSMF) implemented for configuring the equipment (whether physical or virtual) for slicing, if such functionality is available.

End-to-end service and network slicing orchestration

The job of orchestrating end-to-end slices and verifying that those slices are being provided with the agreed-to QoS parameters requires a full lifecycle management.

Design

The slices are designed to adhere to the standardized offerings and QoS parameters. They are usually templated, with some parameterization available. These templates describe the equipment, the interconnection of the equipment, and the parameters that must be implemented in the equipment.

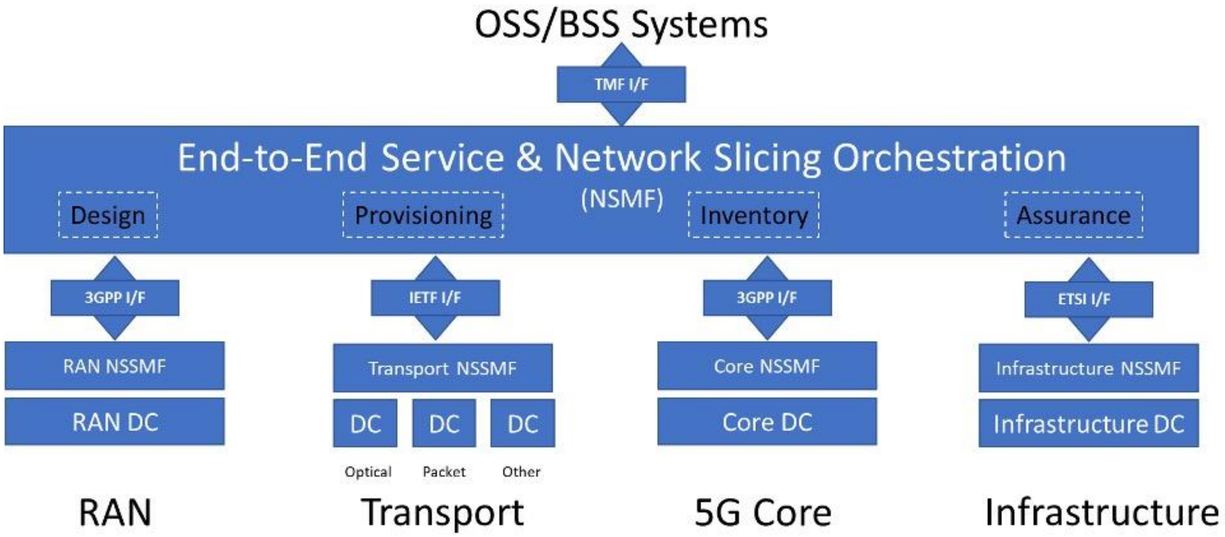

Figure 2: Basic architecture of end-to-end slicing managers

Inventory

The available inventory of equipment needs to be up to date to the minute. It then needs to be allocated to the slice.

Provisioning

The slicing manager then creates a set of processes for implementing the slice in each domain, updating the inventory, and then adding the slicing information to the service assurance system. Increasingly, as domain controllers become more intelligent, the provisioning interface to the domain controller is expected to become more intent-based, with some of the design delegated to the domain controller. In the interim, much of the design needs to be done in the slicing manager.

Assurance

After provisioning and acceptance, the slice must be monitored for proper QoS. This is done by the assurance module. As slicing managers become more intelligent, semi-automated and then automated actions will be taken if the QoS guarantees are not being met to bring these back into conformance.

Vendor solutions – end-to-end network slicing orchestrators

Many vendors have their slicing stories to tell. The major vendors and their full slicing lifecycle management offerings include Amdocs 5G Slicing Manager, Ciena Blue Planet Automation Software, Cisco NSO and Crosswork Automation, Ericsson Dynamic Network Slice Selection, Huawei iMaster, Netcracker Digital OSS, Nokia Network Service Platform, and Oracle Service & Network Orchestration & Open Source MANO.

Conclusion

The future of integrated network slicing is bright, with new innovative services on the horizon. To design, implement, and manage these slices will require the application of tightly integrated, highly automated OSS functions. Will these functions be provided by an assemblage of existing OSSs, updated and more tightly integrated to provide the full lifecycle management? Or will the new generation of fully integrated all-in-one OSS slice managers gain market traction? I will be monitoring and reporting on the market evolution.