

This is a postprint version of the following published document:

Contreras, L.M., López, D.R. (2018). A Network Service Provider Perspective on Network Slicing. *IEEE Softwarization: A Collection of Short Technical Articles*

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

# A Network Service Provider Perspective on Network Slicing

Luis M. Contreras, Diego R. López

*Telefónica Global CTIO Unit*

## Introduction

The feasibility of creating logical, full-functional partitions of network infrastructures (either physical, virtual, or a combination of both), known as slices, will permit network service providers to overcome the great challenge of forthcoming 5G services: how to support and operate different kind of services with very distinct needs onto the same infrastructure. Mixing services like enhanced Mobile Broadband (eMBB), massive Machine-Type Communications (mMTC) and ultra-Reliable and Low Latency Communications (uRLLC), altogether on the same network, makes quite difficult to define a common architecture capable of keeping the requirements of each of them in an ordered and structured form. It is much more convenient to think on segregating them on specialized partitions, designed and optimized for the type of service to be provided.

A second level of segregation, directly related to the business models of network service providers, emerges when considering how (and how many of) those slices are to be provided. In this respect, network slicing can be perceived as the next form of providing wholesale services. That is, it is connected to the allocation of specific resources per vertical customer or tenant. Both levels of segregation go in the direction of allocating physical or virtual resources fitted to the kind of service and appearing to be dedicated per tenant, guaranteeing the isolation among the allocated resources.

Looking at this from a provider mindset, and taking into account the similarities with wholesale product offerings today, different situations can be considered here, leading to slices with different capabilities, specifically in terms of their management and control capabilities, and how much of them the network service provider hands over to the slice tenant:

- Internal slices, understood as the partitions used for internal services of the provider, retaining full control and management of them.
- External slices, being those partitions hosting customer services, appearing to the customer as dedicated networks. In this case, a sub-sequent distinction applies:
  - Provider-managed slices, meaning that the provider keeps the full control and management of the slice. In other words, the customer merely can use the network resources of the provided slice, without any further capability of managing or controlling them.
  - Tenant-managed slices, implying that the customer has full control of the resources and functions allocated. The tenant has access to a (limited) set of operations and/or configuration actions, and the provider just segregates the infrastructure necessary for that purpose. It is worth noting that we use here the term “infrastructure” in the broadest sense, that is,

including all network functions that the customer can deploy and use on the slice.

It is clear that Tenant-managed Slices have to be allocated one per customer, since the customer directly controls it. On the other hand, for Provider-Managed Slices could accommodate different customers sharing the same kind of service in terms of service requirements (bandwidth, latency, number of handled sessions, etc.). Figure 1 graphically presents these different options, illustrating that while the orchestration capability always resides in the provider, the control of the slice or the service on top of such slice could be in the provider or in the tenant depending on the kind of slice as described above.

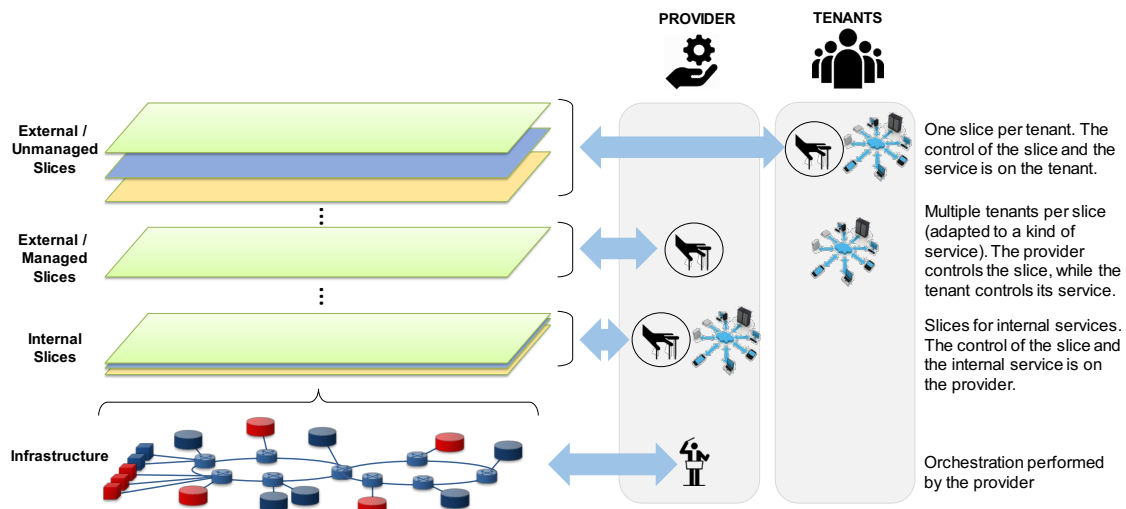


Figure 1. Types of considered slices and control responsibilities

### Environments Facilitating Network Slicing

Network providers are nowadays deploying cloud-like facilities (termed NFV infrastructures: NFVI) that will serve to host virtualized network functions (VNFs). An NFVI allows for deploying VNFs in a dynamic way that can be adapted to the specific needs of each requested service. In the case of Telefónica, those capabilities are based on UNICA [1], the global NFVI being deployed at all its operations. The overarching architecture of UNICA is based on cloud concepts to allow large-scale deployment across multiple sites, but also covers provider needs such as carrier grade performance, scalability, and operational capabilities.

Apart from the elasticity provided by network function virtualization, flexible steering mechanisms are needed to shepherd the traffic flows according to the expected service behavior. In order to address these requirements, network programmability (enabled by Software Defined Networking, SDN), and its integration with the function virtualization described above, are considered as the baseline pieces to approach network slicing [2,3] playing a central technological role in the control, management and operation of future telecom networks.

## Challenges

Both SDN and NFV are progressively being introduced on existing networks, presenting a number of challenges accompanying this technological transition [4]. Network slicing will add an additional degree of complexity. The fundamental aspect of network slicing is that each slice will behave as it was an actual, independent network. This main requirement of *slice isolation* constitutes the essential feature any network service provider has to support in order to offer such services to its customers. The degree of isolation achieved will be critical in determining the ability of a certain provider to address the different classes of slices discussed above, and how they can be requested and used by its customers. Taking this into account, from a provider perspective (extending [5]), a preliminary list of challenges are the following (without considering any particular order):

- Slice operation: Each slice must behave as a dedicated network while sharing underlying resources, physical and virtual. Monitoring and maintenance mechanisms have to be defined in order to show and abstract the proper information for each slice customer. To this respect, external slices have a higher degree of complexity since the information to be exposed, and the access constraints to it, have to be defined or even agreed between the provider and the tenants.
- Scalability: In order to partition network resources in a scalable manner, it is required to clearly define to what extent slice customers can be served or not by a given slice. The application of different SLAs on the offered capabilities of management, control and customization of slices will directly impact the scalability issue. Two dimensions of scalability has to be taken into account. On one hand, there is the dimension related to pure resource allocation and accounting (including those necessary for satisfying protection and availability), directly related with the kind of service requested and negotiated with the tenants. On the other hand, a too much fine-grained offering of slices can provoke an unmanageable number of control and management artifacts to be orchestrated by the provider, making it unpractical. Some kind of aggregation or grouping will be needed for achieving tractability. Scalability can be much impacted by the number of external tenant-managed slices offered.
- Arbitration: In order to resolve conflicts and to ensure negotiated service levels, some arbitration mechanisms are needed that can allow an efficient usage of resources, preventing over-dimensioning. These mechanisms have to be in place not only among the different slices being allocated on the same infrastructure, but within individual slices themselves (like the three main types described in the categorization above), since the relation between customer and slice is not necessarily 1:1 (as in the external provider-managed slices case). Arbitration will have to be applied not only for slice creation or customer activation, but also (and more importantly) when scaling and/or failure events happen, so resources are properly (re-)assigned according to the applicable SLAs.
- Multi-domain: Vertical customers could require no restriction in terms of coverage, service capability, resource constraints, geographical footprint, etc., avoiding any potential limitation of the network provider with which they

maintain a commercial relationship as primary provider. This drives to the necessity of enabling multi-domain slicing, implying functional and commercial interfaces to be normalized for general adoption.

- Slice planning and dimensioning: Over-dimensioning has been the normal way in the past for avoiding any kind of congestion. With slicing the traffic sources and destinations become much less predictable, if at all. This is especially relevant for the case of external tenant-managed slices, where the final decision of where deploying sources and destinations (as well as some intermediate functions that could alter or modify the traffic profile) is on hands of the tenant. Proper planning, dimensioning and enforcement is needed to make sustainable the transition to this new form of service, starting with an appropriate data collection on resource usage, including most significantly virtual ones.
- Orchestration: Slice-based services will not only require the communication between different administrative domains but, what is more challenging, the integration of different technology domains. The complete end-to-end nature of slices involves distinct computing environments and transport technologies, and linking them require from an open and consistent orchestration approach. Furthermore, the scenarios related to external slices will require mediation mechanisms to coordinate requests for the management platforms of the different customers (in the case of tenant-managed slices) or the provider business units (in the case of the provider-managed ones).
- Business layer: The interaction with the vertical customers is key for understanding the needs of the service to be provided. Proper abstractions and templates have to be defined for ensuring the provision of a consistent service portfolio and their integration with the internal network management and orchestration. The support of the different slice classes described above require a consistent set of abstractions at all planes.
- Security: In any shared infrastructure, security is a key element to guarantee proper operation, and especially a fair share of resources to each user. Slice customers must be appropriately authenticated, their rights enforced by authorization mechanisms, and the operations they performed accounted for, so further auditing can be applied in case of any problem is detected. Beyond this, measures have to be in place to proactively detect and address active security attacks, avoiding a security breach in one slice propagates into the infrastructure and/or other slices.
- Slice aging: the same dynamicity and flexibility for the allocation of slices to tenants applies as well to the time-scale on which those slices will be operational. We can expect to have a mix of long- and short-lived slices co-existing on top of the same infrastructure. This fact will impact providers in several ways, from resource planning (slice demand forecast including traffic and resources to be consumed) to security (data preservation per tenant).

These challenges are still a relevant area for research, as explored in projects like 5G-TRANSFORMER [6] and NECOS [7], and additional ones will likely be identified and described in the coming future.

## Conclusions

Network slicing constitutes an extremely promising approach for the provision of network services, able to deal with the demanding 5G requirements, and to support new wholesale offerings. While leveraging recent but well-established technology substrates, such as NFV and SDN, seems a straightforward path to achieve network slicing, there are several important challenges that have to be addressed to make it operationally feasible for network service providers. While a complete list of the open issues related to operational network slicing is still subject for research, the most salient of these challenges have been introduced here, relating them to the fundamental isolation property.

## Acknowledgements

This work has been performed in the framework of the H2020 project 5G-TRANSFORMER (Grant Agreement no. 761536) and EUB project NECOS (Grant Agreement no. 777067), which are partly funded by the European Commission. This information reflects the consortia views, but neither the consortia nor the European Commission are liable for any use that may be done of the information contained therein.

## References

- [1] Analysys Mason, “Telefónica’s UNICA architecture strategy for network virtualization”, white paper, July 2017, available at: <http://www.analysismason.com/telefonica-UNICA-architecture-strategy-for-network-virtualisation-report>
- [2] J. Ordóñez-Lucena, P. Ameigeiras, D. Lopez, J.J. Ramos-Munoz, J. Lorca, J. Folgueira, “Network Slicing for 5G with SDN/NFV: Concepts, Architectures and Challenges”, *IEEE Communications Magazine*, Vol. 55, Issue 5, pp. 80-87, 2017.
- [3] M.-P. Odini, “SDN and NFV Evolution Towards 5G”, *IEEE Softwarization*, September 2017, available at: <https://sdn.ieee.org/newsletter/september-2017/sdn-and-nfv-evolution-towards-5g>
- [4] L.M. Contreras, P. Doolan, H. Lønsethagen, D.R.López, “Operation, organization and business challenges for network providers in the context of SDN and NFV”, *Elsevier Computer Networks*, Vol. 92, pp. 211-217, 2015.
- [5] N. Davies, P. Thompson, “Challenges of Network Slicing”, *IEEE Softwarization*, January 2017, available at: <https://sdn.ieee.org/newsletter/january-2017/challenges-of-network-slicing>
- [6] 5G-TRANSFORMER Project. <http://5g-transformer.eu/>
- [7] NECOS Project: <https://intrig.dca.fee.unicamp.br/2017/09/05/necos-2-year-eu-brazil-collaborative-project-starting-in-nov2017/>