

OPEN NETWORKING
FOUNDATION

TR-526 Applying SDN Architecture to 5G Slicing

Issue 1
April 2016



ONF Document Type: TR (Technical Recommendation), non-normative,
type 2 ONF Document Name: Applying SDN Architecture to 5G Slicing

Disclaimer

THIS SPECIFICATION IS PROVIDED “AS IS” WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Any marks and brands contained herein are the property of their respective owners.

Open Networking Foundation
2275 E. Bayshore Road, Suite 103, Palo Alto, CA 94303
www.opennetworking.org

©2016 Open Networking Foundation. All rights reserved.

Open Networking Foundation, the ONF symbol, and OpenFlow are registered trademarks of the Open Networking Foundation, in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Table of Contents

1	Abstract	4
2	Scope and Introduction	5
	2.1 SDN Architecture	5
	2.2 5G	6
3	Industry Requirements for 5G Slicing	7
	3.1 Introduction to the term Slicing.....	7
	3.2 Relevance of Slicing for 5G.....	8
	3.3 General Requirements for SDN-based Slicing.....	8
4	SDN Architecture as an Enabler for 5G Slicing	10
	4.1 Relevant SDN Architectural Concepts	10
	4.2 Application to 5G and Slicing	12
5	Conclusions	15
6	References	16
7	Abbreviations and Conventions	17
	LIST OF CONTRIBUTORS	18
	Appendix A: Examples of SDN-based resource views	19

List of Figures

Figure 1: Core concepts of the SDN Architecture (adapted from [2])	10
Figure 2: SDN-based Slice abstraction	13
Figure 3: Examples of simple and more detailed SDN-based resource views for Network Slices	19

1 Abstract

The architecture of software defined networking (SDN) is defined in ONF TR-502 [1] and TR-521 [2]. By providing a complete view of all resources required to serve a business purpose, the SDN architecture supports the key principles of Slicing in a network. As a conceptual framework for a standardized platform supporting Network Slicing, it can serve as one of the technical building blocks to fulfill the business requirements for the fifth generation of mobile technology (5G). In the 5G White Paper [4], NGMN has laid out business demands and introduced the service deployment concept of Network Slicing.

The SDN architecture enables a common infrastructure to efficiently support multiple client network instances (tailored and optimized for services with different, diverse requirements). Its key concepts (resource virtualization, recursion) support the envisioned high degree of flexibility through network slicing. Common abstractions enable the presentation of all relevant resources (networking, processing and storage types) included in a slice to fulfill a business purpose, while open and programmable interfaces allow dynamic control and automation of its creation and operation, i.e. dynamic slicing.

The purpose of this document is to describe how key functional aspects of the SDN architecture apply for the enablement of the business-driven concept of network slicing, one of the key concepts for 5G, as envisioned by NGMN.

Applying the SDN architecture, an SDN controller's Client Context provides the complete abstract set of resources and supporting control logic for constituting a slice, including the complete collection of related client service attributes. As such, a 5G slice is comparable to, if not the same as, an SDN client context, isolated by the controller's virtualization and client policy functions and continuously optimized by the orchestration and global policy functions.

In particular, the SDN architecture supports the ability for network slice blueprints to be predefined in terms of the necessary services and abstract resources, with network slice instances created on demand on dynamically determined endpoints and resources.

2 Scope and Introduction

The purpose of this document is to describe how the SDN architecture can be applied to enable 5G slicing. It aims to show that applying SDN can provide a foundation for a standardized platform supporting slicing, in fulfillment of 5G business requirements.

The document refers to the business perspective for the term and understanding of 5G slicing, with the goal of a synthesis to its technical representation using SDN architectural concepts, attempting to contribute to discussions and work emerging/ongoing in the industry, including other standards development organizations. Recognizing that multiple industry-wide efforts for 5G development are in progress at the time of this writing, this document makes use of publicly available documents, listed in the reference clause.

While more detailed, use-case and application dependent, profiling of slices, as well as considerations for the integration and migration of existing systems, are intentionally out of scope of this document and for further study, we note that, at the time of this writing, relevant work has been initiated in the ONF Mobile Network Working Group and other organizations, that this work intends to facilitate.

This scope of this document is technical, focusing on concepts and components on an architectural level, not at an implementation-related specification level. The remainder of this section provides an introduction, for readers interested in the foundations, to both the SDN architecture and 5G. Section 3 addresses Slicing as a key concept, its relevance and related requirements, approaching it from the business perspective. As the part mainly relevant to technical architects and developers, section 4 identifies core principles of the SDN Architecture and provides guidelines on its application to 5G slicing, concluded by a summary in section 5.

2.1 SDN Architecture

The Architecture for Software-Defined Networking (SDN) is defined in ONF TR-502 [1] and TR-521 [2]. SDN applies the flexibility of software to the entirety of the networking space; including business relationships, geographic span, and everything from service creation and delivery to operation and maintenance.

According to [1] and [2], SDN is based on the following three principles:

- Decoupling of control from traffic forwarding and processing
- Logically centralized control
- Programmability of network services.

The SDN architecture enables a common infrastructure to efficiently support multiple client network instances (tailored to and optimized for services with diverse requirements). It is technology independent and hence able to support wireline, wireless and mobile technologies and related abstractions.

By abstracting physical resources, SDN enables to create “soft partitions”, allowing an easier (re-)grouping and (re-)programming, including the ability to provide service guarantees/SLA while making the best use of shared resources.

Major components of SDN are resources and controllers. Service delivery (send, receive, transmit, transform data) makes use of resources; provisioning, management and control of services and related resources is executed via the controller.

Resources and the views onto them are not limited to just network resources. The concept of resources is unbounded, and explicitly includes storage and processing as well as forwarding domains. Resources are managed-controlled through a controller.

Controllers are the central entities in the SDN architecture. Via so-called Northbound Interfaces (NBI), they offer client entities (applications, users, tenants) the ability to interact with network control and to manage-control (i.e. program) services. The extent of management-control rights and capabilities depends on the role of the respective entity using an NBI.

The term *virtualization* is used to describe the function of a controller to aggregate and abstract the underlying resources it manages-controls. Views onto such virtualized resources, or resource groups dedicated to particular clients, are provided to clients/applications/users via NBI. Virtual resources may be statically configured by administrators or allocated dynamically as needed, according to business agreements and policy.

The term *orchestration* is used to describe the responsibility of the controller to dispatch resources in a way that simultaneously satisfies service demands from all of its clients as cost-effectively as possible. This includes dynamic and highly granular sharing of resources when appropriate, while preserving policy commitments to isolate each client's traffic from that of all others. This is an important touch point of SDN with the needs of 5G slicing. It is important to recognize that orchestration is a continuing real-time process, as network resources, traffic load, and service demands change. Both because of scale and information hiding between trust domains, orchestration is performed at every level of scale and partitioning.

Recursion (hierarchical and federated) allows that a controller may at the same time use resources from a "lower level" controller and provide services to a "higher level" controller, in a hierarchical relationship. Federation can be provided by SDN controllers peering on an equal level.

SDN service agility is enhanced by Network Functions Virtualization (NFV) and its ability to rapidly create, scale or relocate virtual resources. As described by ONF TR-518 [3], to an SDN controller, Virtualized Network Functions (VNF) are resources, i.e. functions in a network graph with known connectivity points and known and controllable transfer function.

The SDN architecture allows for long-term coexistence with existing BSS/OSS/NMS/EMS and signaling functions, with incremental migration as warranted by ongoing business case analysis.

2.2 5G

The Next Generation Mobile Networks (NGMN) Alliance and other worldwide initiatives have started to define business requirements and principles for the so-called fifth generation of mobile networking (5G). The vision described in the NGMN White Paper for 5G [4] serves as a guideline for 5G definition and design. It also shows areas for further exploration by NGMN, industry stakeholders and standardization bodies.

While the NGMN 5G White Paper, being used as a reference in this document, was written with a mobile system perspective, its principles and requirements apply to a broad ecosystem context, including the wireline segment, to fully live up to its 5G industry vision. NGMN envisions an end-to-end system that includes all aspects of a network that is designed with a high level of convergence, leveraging today's and tomorrow's access mechanisms.

NGMN considers 5G to be a highly heterogeneous environment characterized by the existence of multiple types of access technologies, multi-layer networks, multiple types of devices, and multiple types of user interactions. While extending the performance envelope of mobile networks is of key importance, the NGMN 5G White Paper requests to include by design embedded flexibility to optimize the network usage, allowing to accommodate a wide range of use cases, as well as business and partnership models. Fundamental is, hence, a modular architecture and functional design that allows to scale on demand and to accommodate various use cases in an agile and cost efficient manner.

NGMN has developed and analyzed use cases for 5G, aiming to identify representative examples, and grouped them into eight use case families. The use cases are not meant to be exhaustive, but rather as a tool to ensure that the level of flexibility required in 5G is well captured. The use case families described in the NGMN 5G White Paper [4] are: Broadband Access in Dense Areas, Broadband Access Everywhere, Higher User Mobility, Massive Internet of Things (IoT), Extreme Real-Time Communications, Lifeline Communication, Ultra-reliable Communications, and Broadcast-like Services. NGMN categorized them according to the most significant requirement. Use cases assigned to the families may overlap, leading to the introduction of diverse, sometimes orthogonal requirements.

Network slicing is a key element in NGMN's vision for 5G to support these use case families, covering a great diversity of requirements, and to enable the deployment of multiple logical networks as independent business operations on a common physical infrastructure platform. SDN and NFV are outlined as potential technical enablers to address 5G system requirements, particularly with regard to operational sustainability and business agility.

3 Industry Requirements for 5G Slicing

3.1 Introduction to the term Slicing

Network slicing is a key architectural approach in NGMN's vision for 5G, particularly for accommodating new and diversified business demands of the 5G era in a cost-efficient way. Slicing enables the (concurrent) deployment of multiple logical, self-contained networks on a common infrastructure platform. NGMN defines slicing as an end-to-end concept, including core network (with specific functionality or by partitioning (or sharing) it) as well as access network (partitioned). From a business viewpoint, a 5G slice is a combination of all relevant network resources / functions / assets required to fulfill a specific business case, including OSS and BSS processes. The establishment of a 5G slice is supposed to be business-driven and not technology-driven.

NGMN has recently published a document with a refined definition and description of its network slicing concept [5].

From the technical infrastructure viewpoint, slicing requires the partitioning and assignment of a set of resources that can be used in an isolated, disjunctive or shared manner. A set of such dedicated resources can be called a slice instance. Examples of resources to be partitioned or shared, understanding they can be physical or virtual, would be: bandwidth on a network link, forwarding tables in a network element (switch, router), processing capacity of servers, processing capacity of network elements. As it can be assumed that slice instances will often contain a combination/group of the above resources, appropriate resource abstractions as well as the exposure of abstract resources towards clients are needed for the operation of slices.

The SDN architecture does currently not have a specific terminology defined for slicing, but the views that the SDN controller provides via its Northbound Interface fulfill the above descriptions and definitions for slices. Furthermore, SDN and NFV are expected to allow optimizing and customizing related (virtualized) function instances, and functional chains thereof, per use case, with a goal of fulfilling purpose-specific requirements.

3.2 Relevance of Slicing for 5G

Slicing allows to concurrently support a large variety of use cases with diverging requirements on a common infrastructure platform. Given the much greater diversity of the requirements in 5G networks expected by the industry, new and more flexible approaches are needed, to efficiently support business flexibility and future innovations, to accelerate the integration and to manage the complexity of supporting a large set of use cases and business models.

Slicing is regarded as a key concept to enable this transformation, to better leverage underlying network infrastructures, i.e. making it more universal bottom-up as well as more user-customizable and application-aware from top-down. Furthermore, it can enable the owners of resources to assign groups of resources to users and/or applications which can use them without harmful interference between each other.

While current definitions do not detail whether a slice is statically pre-provisioned or dynamically created, enabling the dynamic creation and operation of slices (“dynamic slicing”) including networking, storage and compute resources, with an eventual support for “Slicing-As-A-Service”, would be of great benefit.

3.3 General Requirements for SDN-based Slicing

The 5G vision and the requirements derived from it exceed what network systems are able to provide today. It not only raises considerably increased performance expectations, beyond existing technologies and standards, it provides direction toward which ecosystem players need to develop to in order to gain long term operational sustainability and business agility. The envisioned use cases and use case families demonstrate the need to be able to create and maintain a large variety of relationships between the operator’s network infrastructure and service providers on one side and their customers on the other side. A wide range of business models, including partnership roles, will need to be accommodated.

Improvements are required to deliver a broad range of services concurrently over a common infrastructure. Slicing is considered to be an enabler for this. Each slice is expected to purpose-specifically support: a) connectivity between endpoints; b) resources to process traffic where required, c) processes for network/operations management and business support.

In addition, the services are expected to be deployed and scaled on demand in an agile and cost efficient manner. This covers aspects ranging from devices, (mobile/fixed) infrastructure, network functions, value enabling capabilities and management-control functions. Network operators and service providers operate in a highly heterogeneous environment, characterized by the existence of multiple types of access technologies, multi-layer networks, multiple types of devices, multiple types of user interactions, etc.

Network usage must be flexibly optimized and functional modifications must be feasible and applied within short time frames. Operation should be automated and programmable, as much as possible, reducing the need for manual interference and configuration to an absolute minimum.

Open, well-defined APIs are required for services to become programmable. They expose the resource view and provide network functionality according to service level agreements. It must be possible to customize what is being exposed to applications and clients/customers depending on the respective use case, business model and administrative policies.

Slice instantiation needs to be done in a structured, use case and client specific way. A multitude of different slice instantiations will have to be dynamically accommodated on a common infrastructure, potentially involving more than one operator and/or provider.

Use case, application and deployment dependent resource variants need to be supported in slices, including the coexistence of both standardized generic and user-specific network functionality. Common abstractions and resource representations are needed in order to avoid silos, e.g. avoiding deviations in how physical and virtual resources are managed-controlled (including configuring and programming resources for actual usage).

Slice-supporting manage-control processes need to be able to provision network connectivity (local, wide-area, inter-area) between physical and virtual endpoints as well as the full life cycles of all relevant physical and virtual functional resources, in an end-to-end manner, supporting the operation of X-as-a-Service models.

In order to provide the agility and dynamicity required, systems and APIs for service and resource provisioning need to be designed as modular, re-usable building blocks with model-driven and callable interfaces, versus hard-coded and protocol-based. Existing management systems were largely designed with an understanding that services and resource allocations were largely static and although automated to the extent possible, were never far away from human administrators. Instead of static management adapters, a flexible approach to network services modeling needs to be supported based on model-driven concepts and data modeling languages.

For the creation and operation of dynamic network slices, customer-driven or application-driven variations in service parameters, real-time dynamics and policy driven real-time service variation need to be supported.

While statically configured slices are of course possible, the ability to roll-out new services quickly and to use network resources efficiently are requirements where SDN is expected to help by offering slicing dynamism.

4 SDN Architecture as an Enabler for 5G Slicing

4.1 Relevant SDN Architectural Concepts

The SDN architecture enables a common infrastructure to efficiently support multiple client network instances (tailored and optimized for services with diverse requirements). Common abstractions enable the presentation of all relevant resources to be included in a slice for fulfilling a business purpose. Open and programmable interfaces allow dynamic control and automation of their creation and operation. The SDN architecture is technology independent and hence able to support wireline, wireless and mobile technologies and related abstractions. Through its main functions of virtualizing underlying resources and orchestrating resources to be used by services, SDN enables dynamic resource control.

There are two distinct perspectives on the architecture. The resource-oriented perspective takes the bottom-up view of a provider allocating resources to clients, and focusses on the provider's administrative needs. The service-oriented perspective is the top-down view of a customer who invokes and manages services received from the provider.

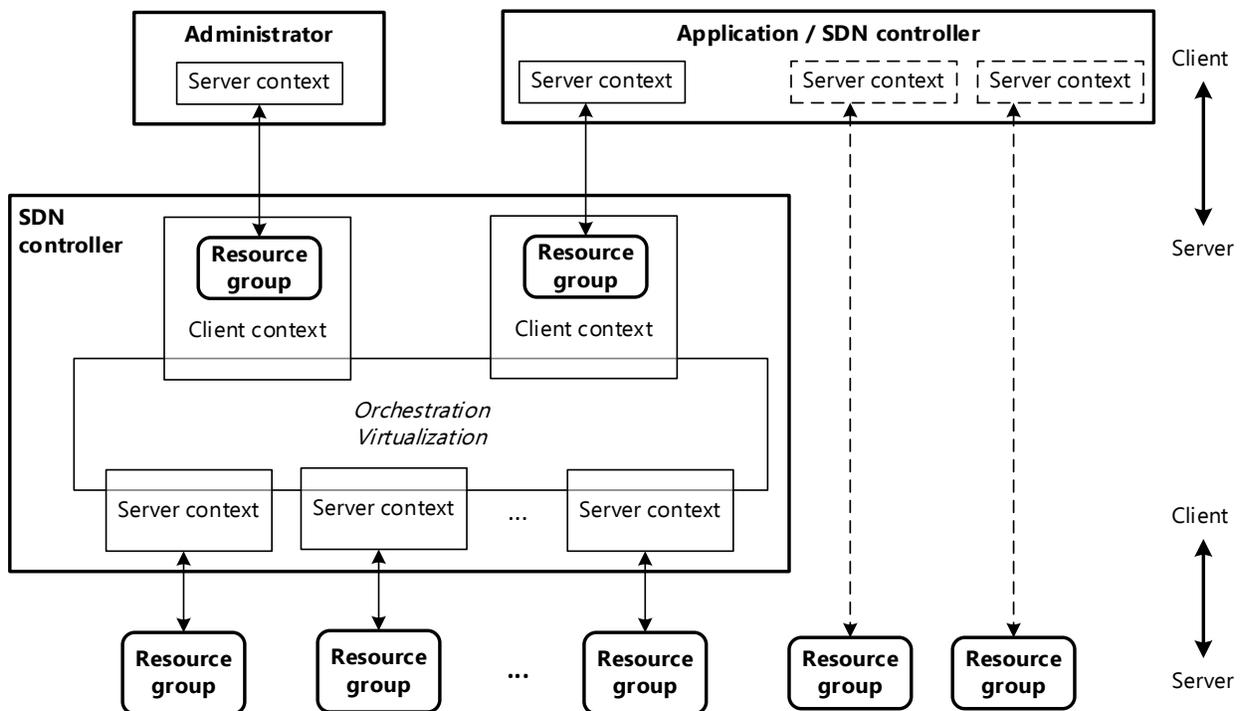


Figure 1: Core concepts of the SDN Architecture (adapted from [2])

The interfaces between the SDN architecture components (resources and controller, and controller and applications) display client-server relationships (Figure 1).

The SDN controller is at the center of a feedback loop: it mediates client requirements with resource availability, supporting policy-driven real-/run-time optimization of changes in network state, service parameters, service and traffic flow.

According to [2] and illustrated in Figure 1, an SDN controller works with two major types of resource views: It interacts with its client via a (client specific) client context, and with its resources via a (server specific) server context.

A *client context* represents the necessary and sufficient material in the SDN controller to support a given client, where a client may be a customer, partner, or even another entity within the same administration that owns the controller. It includes all of the attributes of a service as requested by the client, and may contain service-specific information necessary to map service attributes into the realization of the service.

A *server context* is the symmetric counterpart to a client context. It contains everything necessary and sufficient to interact with a group of underlying resources, which could be, for example, a discrete network element or the virtual resources contracted from a partner domain.

Resources fall into and may be combinations of the categories network, storage and compute, with resource capabilities therefore ranging from simple, e.g. switching between ports, to complex, e.g. firewall with DPI or a video transcoder.

The state of the resources owned by the controller is continually adapted in compliance with policies provided by the administrator. These policies include parameters to satisfy commitments to clients. For all services requested by clients, resources are provided and maintained throughout their lifecycles. Especially where services are delivered over a common infrastructure, the contention of multiple services for resources is dynamically handled by the controller taking into account the aggregated requirements and (administrative) policy constraints.

Some subset of the resources used by SDN may be supplied by NFV (see also [3]). SDN and NFV, complementing each other, enable the required modularity and abstraction of network functions. SDN focuses especially on the use of resources of any and all kinds and locations, referring to resources generically without restriction as to type. NFV focuses especially on the creation and support of resources that can be largely decoupled from hardware realization. The common purpose is to convey information amongst arbitrary entities. Traffic may be modified on its journey, either to support the communications process or in the form of added-value services.

SDN end-to-end global span necessitates hierarchical or federated partitioning, according to business, organization, technology, scale, legacy, or other boundaries, but can also address a basic case of single layer of provider-tenant relationship and administrative domains that likewise do not exceed a depth of one provider-tenant relationship.

The administration of an SDN controller, shown in Figure 1 as administrator, may be performed by an OSS or an OSS application. It is responsible for internal configuration and (administrative) policies and constraints of the controller, including all client and server contexts. The SDN controller will seek to dynamically reconfigure the resources, e.g. to apply varying service chain policies depending on the application and subscriber type to effect the desired policy. It can be expected that this administrator application is in the same trust domain as the controller it manages-controls.

A working SDN controller is created from some substrate. It may be placed in a Virtualized Network Function (VNF). By default an SDN controller is created with an administrator client context that has unrestricted visibility and authority to perform all other operations. The administrator configures the controller with server contexts to access underlying resources, and

updates them from time to time as needed. The underlying resources are themselves configured by their own administrators at their own (underlying) levels. The administrator then creates a client context for each of its clients, which includes allocation of underlying resources to the client in the process called virtualization, as well as supplementary configuration. The administrator configures each client context with policy that defines the actions and bounds permitted to the client. An administrator may modify a client context during its lifetime, and may destroy a client context if the client relationship terminates. [2]

In order to transfer data traffic through the network, both the actual data-plane connectivity and an association between the client organization and the serving organization (owning the respective controller) must exist or be established. The association is part of the initial agreement between the two entities to have a business relationship. Once this relationship exists, the respective client context will be created, which is followed by a management-control session between client and SDN controller to request or modify services. A continuing management-control session is normally not required during delivery of the service.

There are trust and, potentially, ownership borders across interfaces between resource group and server context. These interfaces expose views on resources available in the client context and related services. The views and services visible to each client depend on the association (business agreement) between service user and service provider, and normally include information hiding, namespace and functional mapping and translation, and service level agreements.

4.2 Application to 5G and Slicing

The SDN architecture enables flexible, dynamic and scalable support of a wide range of business needs. All business use cases, each with its own specific set of performance, scalability and availability requirements, can be active concurrently on a common infrastructure, utilizing resources in isolated or shared manner.

NGMN defines slicing from a business-driven (service-related) top-down view [4] [5], similar to the level of abstraction a SDN controller provides to applications through client contexts. The document "Description of Network Slicing Concept" [5] published by NGMN in 2016Q1 details network slicing, intending to serve as reference for further technical design and standardization.

Applying the SDN architecture, the client context provides the complete abstract set of resources (as Resource Group) and supporting control logic that constitutes a slice, including the complete collection of related client service attributes. As such, a 5G slice is comparable to, if not the same as, an SDN client context, isolated by the controller's virtualization and client policy functions and continuously optimized by the orchestration and global policy functions. Going a step further, a 5G controller is an SDN controller, or vice versa. Both 5G and SDN controller manage-control combinations of all relevant network resources / functions / assets required to serve client-specific purposes. The client context also offers to the client functions to manage-control the slice resources, including OAM related-functions, as visible by / available to the client according to administrative policy.

Recursion in the SDN architecture allows for a high-level client context to be virtualized and orchestrated over a number of lower-level resource groups, spanning geographic, business and technology boundaries as needed. Each lower-level resource group may itself be exposed by the

client context of a lower-level SDN controller. The SDN architecture thereby naturally supports a recursive composition of slices.

The SDN controller, at the center of a feedback loop and acting autonomously according to administratively configured policies, enables dynamic allocation, modification and optimization of resource usage. The resources provided to the controller’s clients are virtualized (abstracted views and services) from the underlying resources and presented to slices through resource groups.

Figure 2 shows a mapping of the SDN architecture, terminology and abstractions to 5G slicing.

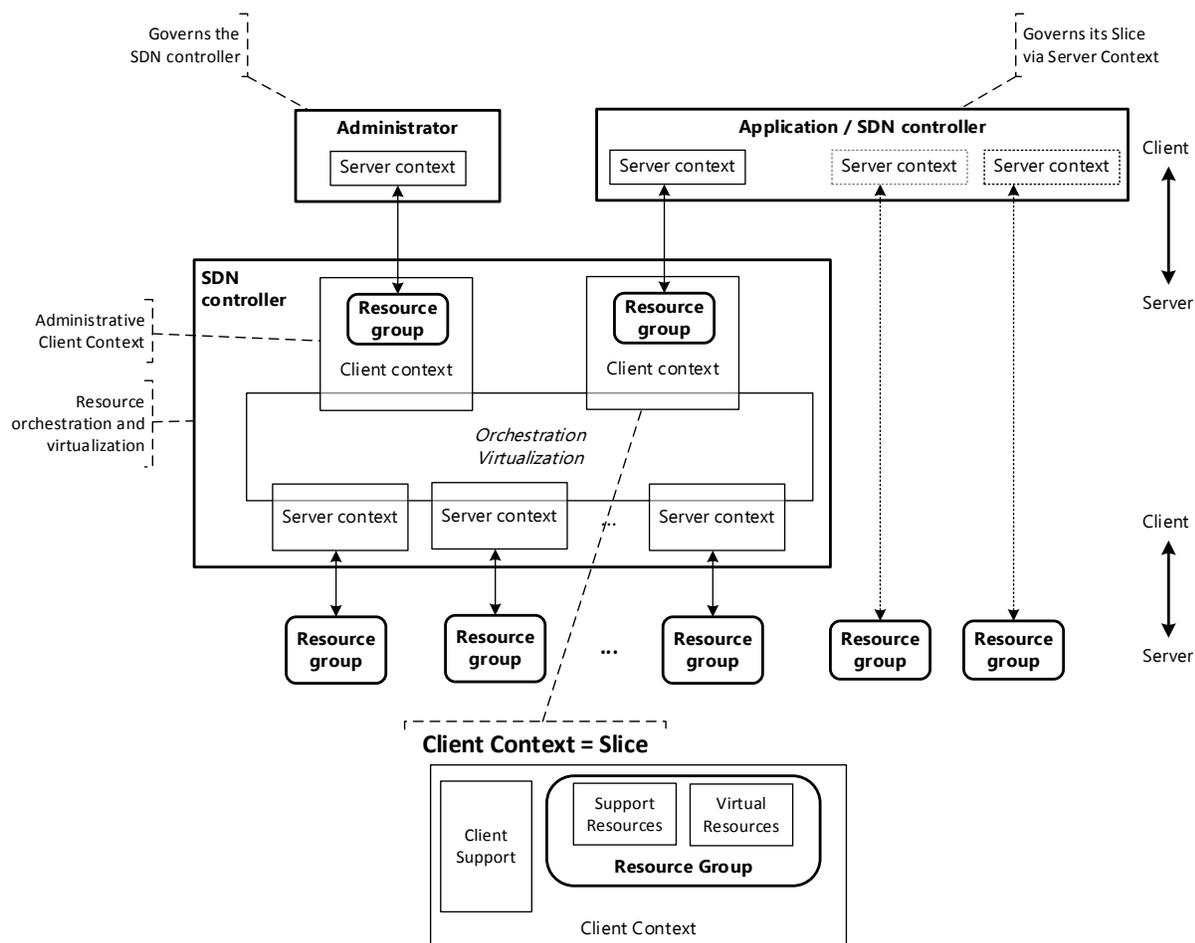


Figure 2: SDN-based Slice abstraction

Each client context represents a set of resources managed-controlled by a controller and is directly applicable to 5G slicing. Resources may be statically predefined by the administrator, or allocated dynamically by the controller’s orchestration function. A (higher-level) controller may slice resources further to support different clients, or combine resource groups to create more complex services.

Through a client context, the SDN controller provides a complete policy-dependent/policy-adapted resource view for a slice instance, as well as related orchestration logic/lifecycle and

business policy. The resource group is a major part of it, representing the resources granted to the slice instance, including network, compute, and storage resources, making use of required resource groups available to the controller at southbound. Virtual resources represent infrastructure resources that are created from the SDN controller's underlying resources through the process of virtualization, and that are exposed to the client by way of a mapping function. Support resources represent functions hosted in the SDN controller itself. Their purpose is to enable or facilitate interaction with the client, including attributes of services as requested by the client, and information necessary to map these service attributes into the realization of the respective service. The client support block may contain additional functions, resources and information needed, but not exposed to the client.

The administration (which may be performed by an OSS or an OSS application) configures the policies to be enforced by the SDN controller. This includes creating, modifying, deleting client contexts for clients (customers) that may interact via the NBI with the controller to manage-control the operation of "their" resources (restricted by the respective views and services available through the NBI). This means, that once the controller is set up and configured by the administration, it functions autonomously. The administration is responsible for the overall framework within which the controller operates. The application/translation of client-requested services and service attributes is the responsibility of the controller. For creating a new slice (i.e. a client context inside a controller), the administration is always involved.

Assuming a typical operator network including multiple business entities, resource owners must establish: what virtual resources are exposed to each client, how they are identified, how they are made available at contractually agreed level, and how they are isolated from other resource users. The SDN controller takes care of isolation in the control and data planes so that each network slice instance can be operated as a distinct and logically separate network.

While more detailed, use-case and application dependent, profiling of slices, as well as considerations for the integration and migration of existing systems, are intentionally out of scope of this document and for further study, we note that, at the time of this writing, relevant work has been initiated in the ONF and other organizations, that this document intends to facilitate.

5 Conclusions

The SDN architecture provides a foundation for a standardized platform supporting slicing, including a unified control plane for an adaptive data plane, as an enabler to support the 5G business requirements. While the SDN architecture and its resource abstractions apply to wireline, wireless and mobile networking contexts, it is applicable for 5G, particularly as the technical concept of SDN enables realizing the business purpose of slicing in a network system.

The vision for 5G and motivation for SDN match. For future business opportunities system flexibility is a key requirement, i.e. to not only support concurrent and/or currently emerging use cases, but even those not yet envisioned. This is fully inline with the original motivation for SDN and for programmable networking.

From an SDN point of view, there is more value in a slice than is captured by the idea of an administratively-created, pre-allocated set of resources. Resources can be used more efficiently if a slice can be instantiated on demand, with per-service instance granularity. At the same time, flexible definition and on-demand construction of slices create space for competitive differentiation in service offerings, and in fine-grained customization of individual service instances.

The SDN architecture enables a common infrastructure to efficiently support multiple slices (tailored and optimized for services with different, diverse requirements). Its key concepts (resource virtualization, continuing orchestrated optimization, recursion) support the 5G envisioned high degree of flexibility through network slicing. An SDN controller provides resource isolation for multiple slices running on a common infrastructure. In summary, the SDN architecture is well aligned with 5G slicing requirements:

- Slices can be defined according to any criteria appropriate for their purpose, and instantiated through the SDN architecture.
- A single end-end slice can be defined and instantiated over the entire world, if desired. Slices can also be partitioned for scale, ownership, technology or other criteria as is convenient.
- NGMN describes a slice instance as independent of other slices. The SDN architecture supports address, namespace, traffic load, and other isolation according to client agreements.
- Slices can be instantiated statically when appropriate. Roaming devices and intermittent service demands may be served more efficiently if slices are instantiated as needed.
- The 5G will necessarily evolve in a brownfield. The SDN architecture provides for gradual migration and long-term coexistence with current management and signaling systems and network devices.

We recognize that, at the time of this writing, 5G development is ongoing with industry-wide efforts. This document, based on the ONF architecture at present state, intends to facilitate this development and related ongoing work, at ONF and other bodies, recognizing that a more detailed understanding may also identify gaps or needed extensions.

6 References

- [1] ONF TR-502 “SDN Architecture 1.0”, June 2014.
(https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR_SDN_ARCH_1.0_06062014.pdf)
- [2] ONF TR-521 “SDN architecture - Issue 1.1”, February 2016.
(https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-521_SDN_Architecture_issue_1.1.pdf)
- [3] ONF TR-518 “Relationship of SDN and NFV”, October 2015.
(https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/onf2015.310_Architectural_comparison.08-2.pdf)
- [4] NGMN Alliance 5G White Paper, March 2015.
(http://www.ngmn.de/fileadmin/ngmn/content/downloads/Technical/2015/NGMN_5G_White_Paper_V1_0.pdf)
- [5] NGMN Alliance document “Description of Network Slicing Concept”, January 2016.
(https://www.ngmn.org/uploads/media/160113_Network_Slicing_v1_0.pdf)

7 Abbreviations and Conventions

This document uses the following abbreviations and acronyms:

5G	Fifth Generation of Mobile Networking
API	Application Programming Interface
BSS	Business Support System
NBI	Northbound Interface
NFV	Network Functions Virtualization
NGMN	Next Generation Mobile Networks Alliance
OSS	Operations Support System
SDN	Software Defined Networking
SLA	Service Level Agreement
VNF	Virtualized Network Function

LIST OF CONTRIBUTORS

Manuel Paul, DT

Sibylle Schaller, NEC

Malcolm Betts, ZTE

Dave Hood, Ericsson

Meral Shirazipour, Ericsson

Diego Lopez, Telefonica

John Kaippallimalil, Huawei

Appendix A: Examples of SDN-based resource views

The SDN architecture allows flexible structuring of domains and administrative hierarchies depending on the preferences of the owner of the infrastructure. It is possible that an infrastructure owner/operator manages-controls all its resources using just a single administrative context, i.e. a single administrator application is used with all controllers in the environment. Of course, where a client application is a controller in its own terms (i.e. crossing trust and administrative boundary between server and client domains), the ownership of these two controllers is different and each is governed within its respective administrative context.

Figure 3 below provides examples of simple and more detailed SDN-based resource views (not showing administrative context):

1. If all business use cases are provided by the same organization, then all controllers shown may be governed by a single administrator.
2. If more organizations are involved to provide a use case, several administrative entities are involved. Controllers belonging to different organizations are governed by the owning administration.

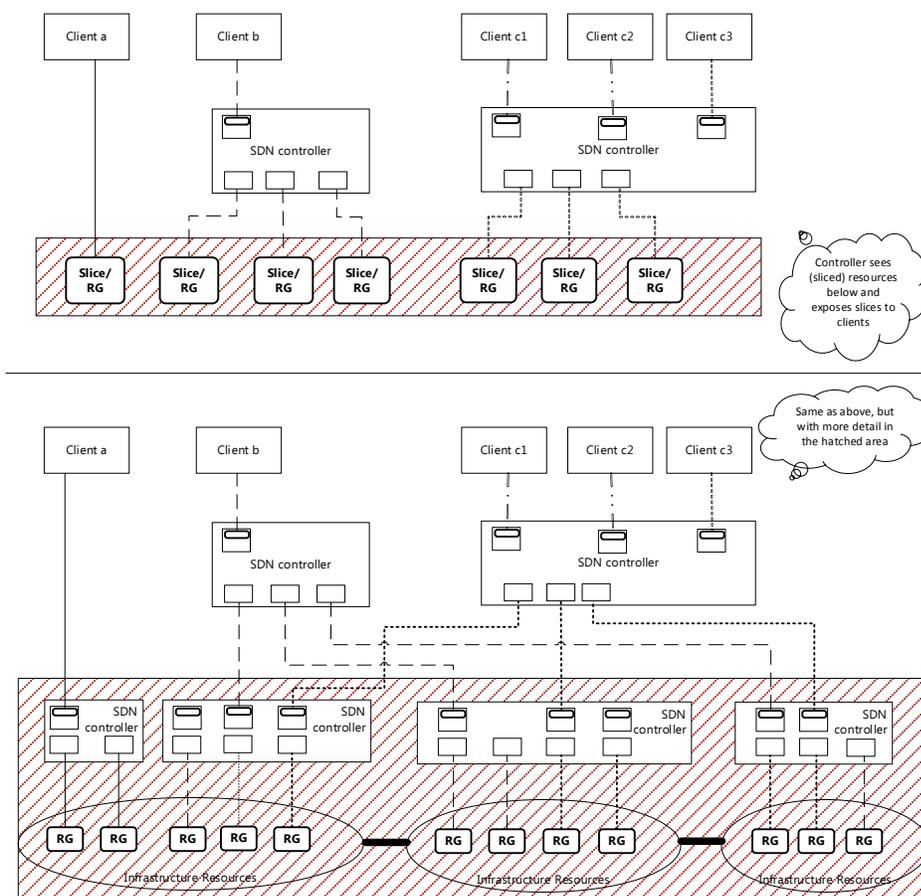


Figure 3: Examples of simple and more detailed SDN-based resource views for Network Slices