TEAS Working Group
Internet-Draft
Intended status: Informational
Expires: December 26, 2018

D. King (Ed.)
Old Dog Consulting
Y. Lee (Ed.)
Huawei

June 26, 2018

Applicability of Abstraction and Control of Traffic Engineered Networks (ACTN) to Network Slicing
draft-king-teas-applicability-actn-slicing-03

Abstract

Network abstraction is a technique that can be applied to a network domain to select network resources by policy to obtain a view of potential connectivity

Network slicing is an approach to network operations that builds on the concept of network abstraction to provide programmability, flexibility, and modularity. It may use techniques such as Software Defined Networking (SDN) and Network Function Virtualization (NFV) to create multiple logical (virtual) networks, each tailored for a set of services that are sharing the same set of requirements, on top of a common network.

These logical networks are referred to as transport network slices. A transport network slice does not necessarily represent dedicated resources in the network, but does constitute a commitment by the network provider to provide a specific level of service.

The Abstraction and Control of Traffic Engineered Networks (ACTN) defines an SDN-based architecture that relies on the concepts of network and service abstraction to detach network and service control from the underlying data plane.

This document outlines the applicability of ACTN to transport network slicing in an IETF technology network. It also identifies the features of network slicing not currently within the scope of ACTN, and indicates where ACTN might be extended.

Copyright Notice

Table of Contents

1. Introduction

The principles of network resource separation are not new. For years, separated overlay and logical (virtual) networking have existed, allowing multiple connectivity services to be deployed over a single physical network comprised of single or multiple layers. However, several key differences exist that differentiate overlay and virtual networking from network slicing.

A transport network slice construct provides an end-to-end logical network, often with compute functions and utilising shared underlying (physical or virtual) network resources. This logical network is separated from other, often concurrent, logical networks each with independent control and management, and each of which can be created or modified on demand.

At one end of the spectrum, a virtual private wire or a virtual private network (VPN) may be used to build a network slice. In these cases, the network slices do not require the service provider to isolate network resources for the provision of the service - the service is "virtual".

At the other end of the spectrum there may be a detailed description of a complex service that will meet the needs of a set of applications with connectivity and service function requirements that may include compute resource, storage capability, and access to content. Such a service may be requested dynamically (that is, instantiated when an application needs it, and released when the application no longer needs it), and modified as the needs of the application change.

Each example represents a self-contained network that must be flexible enough to simultaneously accommodate diverse business-driven use cases from multiple players on a common network infrastructure.

This document outlines the application of the ACTN architecture [actn-framework] and enabling technologies to provide transport network slicing in an IETF technology network. It describes how the ACTN functional components can be used to support model-driven partitioning of variable-sized bandwidth to facilitate network sharing and virtualization. Furthermore, the use of model-based interfaces to dynamically request the instantiation of virtual networks could be extended to encompass requesting and instantiation of specific service functions (which may be both physical and/or virtual), and to partition network resources such as compute resource, storage capability, and access to content.

In an IETF context, there are works in progress that have some bearing with network slicing such as Enhanced VPN (VPN+) and DetNet. Both works are an independent work in their own scope while

This document highlights how the ACTN approach might be extended to address these other requirements of network slicing where TE is required.

1.1. Terminology

Resource: Any features that can be delivered, including connectivity, compute, storage, and content delivery.

Service Functions (SFs): Components that provide specific function within a network. SFs are often combined in a specific sequence, service function chain, to deliver services.

Infrastructure Resources: The hardware and necessary software for hosting and connecting SFs. These resources may include computing hardware, storage capacity, network resources (e.g. links and switching/routing devices enabling network connectivity), and physical assets for radio access.

Service Provider: A server network or collection of server networks.

Consumer: Any application, client network, or customer of a network provider.

Service Level Agreement (SLA): An agreement between a consumer and network provider that describes the quality with which features and functions are to be delivered. It may include measures of bandwidth, latency, and jitter; the types of service (such as the network service functions or billing) to be executed; the location, nature, and quantities of services (such as the amount and location of compute resources and the accelerators require).

Network Slice: An agreement between a consumer and a service provider to deliver network resources according to a specific service level agreement. A slice could span multiple technology (e.g., radio, transport and cloud) and administrative domains.

IETF Technology: A TE network slice or transport network slice.

2. Requirements for Network Slicing

The concept of network slicing is considered a key capability for future networks and, to serve customers with a wide variety of different service needs, in term of latency, reliability, capacity, and service function specific capabilities.

This section outlines the key capabilities required, and further discussed in [ngmn-network-slicing], [network-slice-5g], [3gpp.28.801] and [onf-tr526], to realise network slicing in an IETF technology network.

2.1. Resource Slicing

For network slicing, it is important to consider both infrastructure resources and servic functions. This allows a flexible approach to deliver a range of services both by partitioning (slicing) the available network resources to present them for use by a consumer, but also by providing instances of SFs at the right locations and in the correct chaining logic, with access to the necessary hardware, including specific compute and storage resources.

Mapping of resources to slices may 1-to-1, or resources may be shared among multiple slices.

2.2. Network and Function Virtualization

Virtualization is the abstraction of resources where the abstraction is made available for use by an operations entity, for example, by the Network Management Station (NMS) of a consumer network. The resources to be virtualized can be physical or already virtualized, supporting a recursive pattern with different abstraction layers. Therefore, Virtualization is critical for network slicing as it enables effective resource sharing between network slices.

Just as server virtualization makes virtual machines (VMs) independent of the underlying physical hardware, network Virtualization enables the creation of multiple isolated virtual networks that are completely decoupled from the underlying physical network, and can safely run on top of it.

2.3. Resource Isolation

Isolation of data and traffic is a major requirement that must be satisfied for certain applications to operate in concurrent network slices on a common shared underlying infrastructure. Therefore, isolation must be understood in terms of:

- Performance: Each slice is defined to meet specific service requirements, usually expressed in the form of Key Performance Indicators (KPIs). Performance isolation requires that service delivery on one network slice is not adversely impacted by congestion and performance levels of other slices;

- Security: Attacks or faults occurring in one slice must not have an impact on other slices, or customer flows are not only isolated on network edge, but multiple customer traffic is not mixed across the core of the network. Moreover, each slice must have independent security functions that prevent unauthorised entities to have read or write access to slice-specific configuration, management, accounting information, and able to record any of these attempts, whether authorised or not;

- Management: Each slice must be independently viewed, utilised and managed as a separate network.

2.4. Control and Orchestration

Orchestration is the overriding control method for network slicing. We may define orchestration as combining and coordinating multiple control methods to provide an operational mechanism that can deliver services and control underlying resources. In a network slicing environment, an orchestrator is needed to coordinate disparate processes and resources for creating, managing, and deploying the end-to-end service. Two scenarios are outlined below where orchestration would be required:

1. Multi-domain Orchestration: Managing connectivity setup of the transport service, across multiple administrative domains;

2. End-to-end Orchestration: Combining resources for an "end-to-end service (e.g., transport connectivity with firewalling and guaranteed bandwidth and minimum delay for premium radio users (spanning multiple domains).

In addition, 3GPP has also developed Release 14 "Study on management and orchestration of network slicing for next generation network" [3gpp.28.801], which defines an information model where the network slice as well as physical and virtualized network functions belong to the network operator domain, while the virtualized resources belong to another domain operated by a Virtualization infrastructure service provider.

3. Abstraction and Control of Traffic Engineered (TE) Networks (ACTN)

The framework for ACTN [actn-framework] includes a reference architecture that has been adapted for Figure 1 in this document, it describes the functional entities and methods for the coordination of resources across multiple domains, to provide end-to-end services, components include:

- Customer Network Controller (CNC);

- Multi-domain Service Coordinator (MDSC);

- Provisioning Network Controller (PNC).

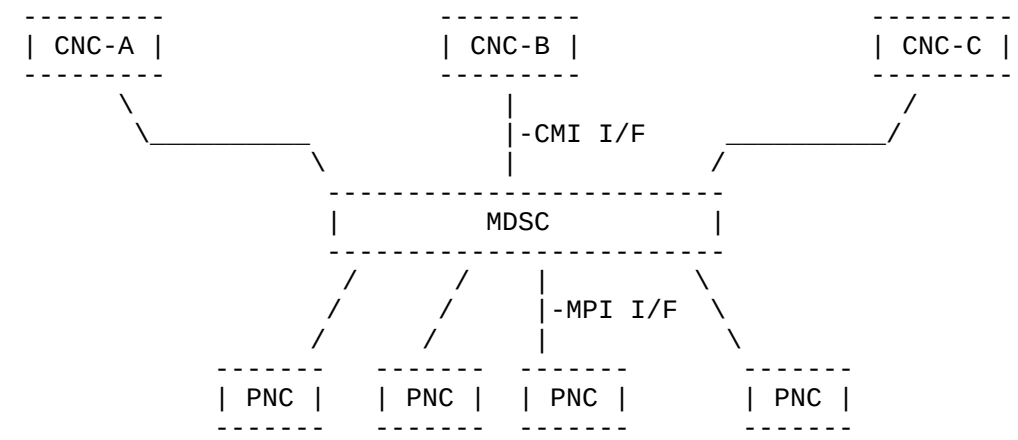

CMI - (CNC-MDSC Interface)
MPI - (MDSC-PNC Interface)

Figure 1: ACTN Hierarchy

ACTN facilitates end-to-end connections and provides them to the user. The ACTN framework highlights how:

- Abstraction of the underlying network resources are provided to higher-layer applications and customers;
- Virtualization of underlying resources, whose selection criterion is the allocation of those resources for the customer, application, or service;
- Creation of a virtualized environment allowing operators to view and control multi-domain networks as a single virtualized network;
- The presentation to customers of networks as a virtual network via open and programmable interfaces.

The ACTN managed infrastructure are traffic engineered network resources, which may include:

- Statistical packet bandwidth;
- Physical forwarding plane sources, such as: wavelengths and time slots;
- Forwarding and cross connect capabilities.

The ACTN type of network virtualization provides customers and applications (tenants) to utilise and independently control

allocated virtual network resources as if resources as if they were physically their own resource. The ACTN network is "sliced", with tenants being given a different partial and abstracted topology view of the physical underlying network. The capabilities that ACTN provides to enable slicing are outlined in Section 2 (Requirements for Network Slicing).

3.1. ACTN Virtual Network as a "Network Slice"

To support multiple clients each with its own view of and control of the server network, a network operator needs to partition (or "slice") the network resources. The resulting slices can be assigned to each client for guaranteed usage which is a step further than shared use of common network resources. See [actn-vn] for detailed ACTN VN and VNS.

An ACTN Virtual Network (VN) is a client view that may be considered a "network slice" of the ACTN managed infrastructure, and is presented by the ACTN provider as a set of abstracted resources.

Depending on the agreement between client and provider various VN operations and VN views are possible.

- Network Slice Creation: A VN could be pre-configured and created via static or dynamic request and negotiation between customer and provider. It must meet the specified SLA attributes which satisfy the customer’s objectives.

- Network Slice Operations: The network slice may be further modified and deleted based on customer request to request changes in the network resources reserved for the customer, and used to construct the network slice. The customer can further act upon the network slice to manage traffic flow across the network slice.

- Network Slice View: The VN topology from a customer point of view. These may be a variety of tunnels, or an entire VN topology. Such connections may comprise of customer end points, access links, intra domain paths and inter-domain links.

Primitives (capabilities and messages) have been provided to support the different ACTN network control functions that will enable network slicing. These include: topology request/query, VN service request, path computation and connection control, VN service policy negotiation, enforcement, routing options. [actn-info]

3.2. Examples of ACTN Delivering Types of Network Slices

In examples below the ACTN framework is used to provide

control, management and orchestration for the network slice life-cycle, the connectivity . These dynamic and highly flexible, end-to-end and dedicated network slices utilising common physical infrastructure, and according to vertical-specific requirements.

The rest of this section provides three examples of using ACTN to achieve different scenarios of ACTN for network slicing. All three scenarios can be scaled up in capacity or be subject to topology changes as well as changes from customer requirements perspective.

3.2.1. ACTN Used for Virtual Private Line Model

ACTN Provides virtual connections between multiple customer locations, requested via Virtual Private Line (VPL) requester (CNC-A). Benefits of this model include:

- Automated: the service set-up and operation is network provider managed;
- Virtual: the private line is seamlessly extended from customers Site A (vCE1 to vCE2) and Site B (vCE2 to vCE3) across the ACTN-managed WAN to Site C;
- Agile: on-demand where the customer needs connectivity and fully adjustable bandwidth.

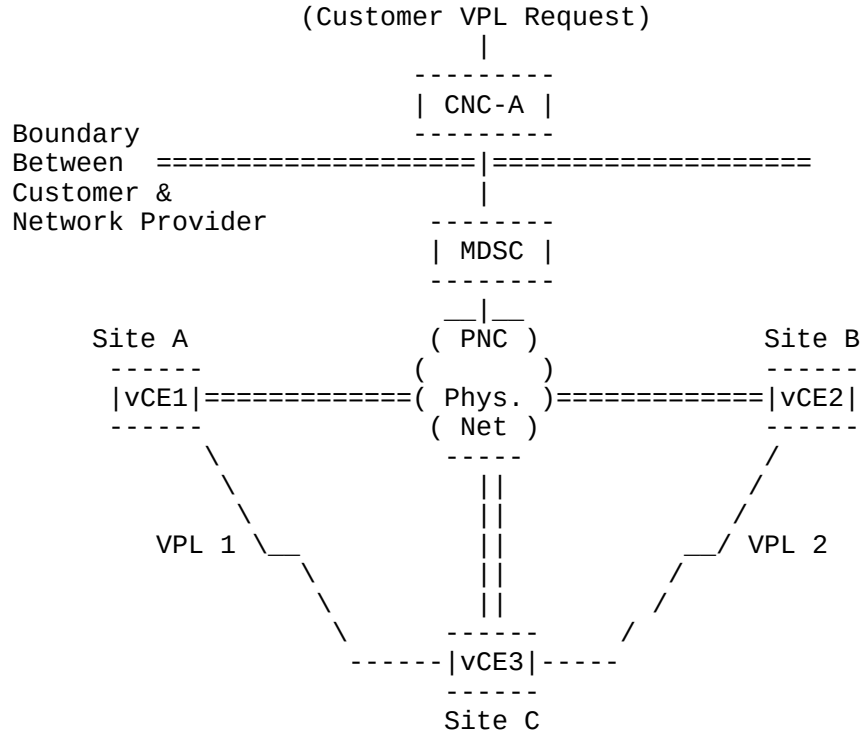

Figure 2: Virtual Private Line Model

3.2.2. ACTN Used for VPN Delivery Model

ACTN Provides VPN connections between multiple sites, requested via a VPN requestor (CNC-A), which is managed by the customer themselves. The CNC will then interact with the network providers MDSC. Benefits of this model include:

- Provides edge-to-edge VPN multi-access connection;
- Mostly network provider managed, with some flexibility delegated to the customer managed CNC.

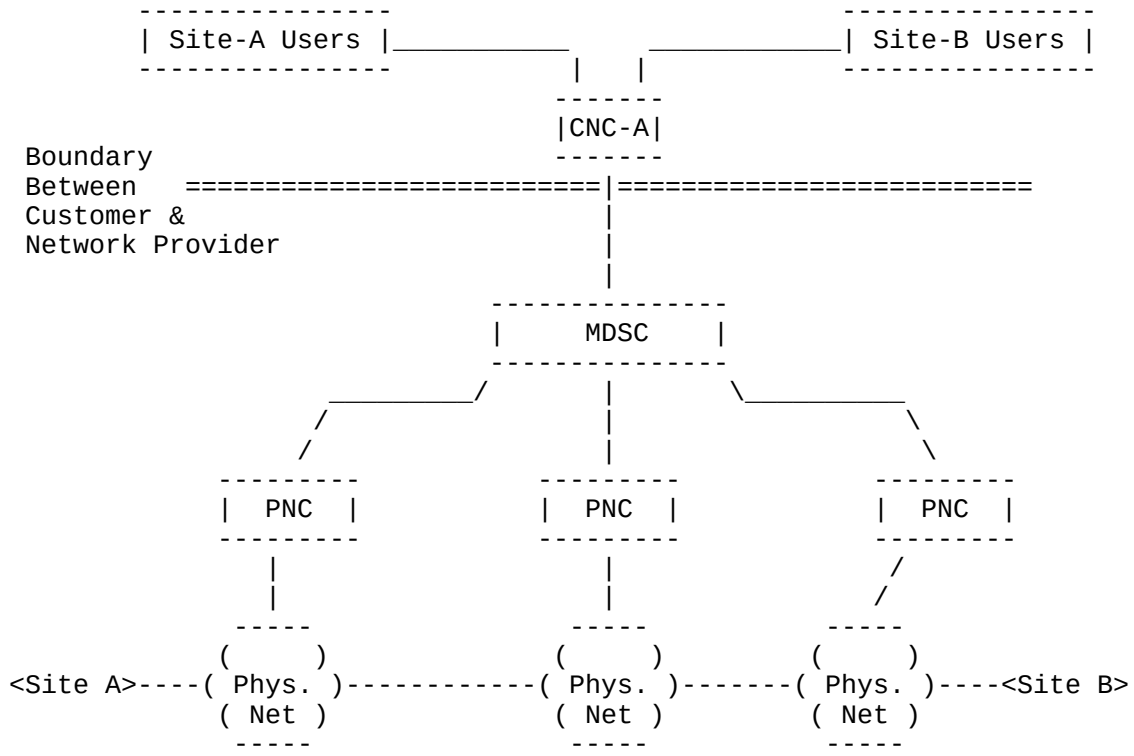

Figure 3: VPN Model

3.2.3. ACTN Used to Deliver a Virtual Customer Network

In this example ACTN provides a virtual network resource to the

customer. This resource is customer managed. Empowering the tenant to control allocated slice (recursively). Benefits of this model include:

- The MDSC provides the topology as part of the customer view so that the customer can control their network slice to fit their needs;

- Resource isolation, each customer network slice is fixed and will not be affected by changes to other customer network slices;

- Applications can interact with their assigned network slice directly, the customer may implement their own network control method and traffic prioritization, manage their own addressing scheme, and further slice their assigned network resource;

- The network slice may also include specific capability nodes, delivered as Physical Network Functions (PNFs) or Virtual Network Functions (VNFs).

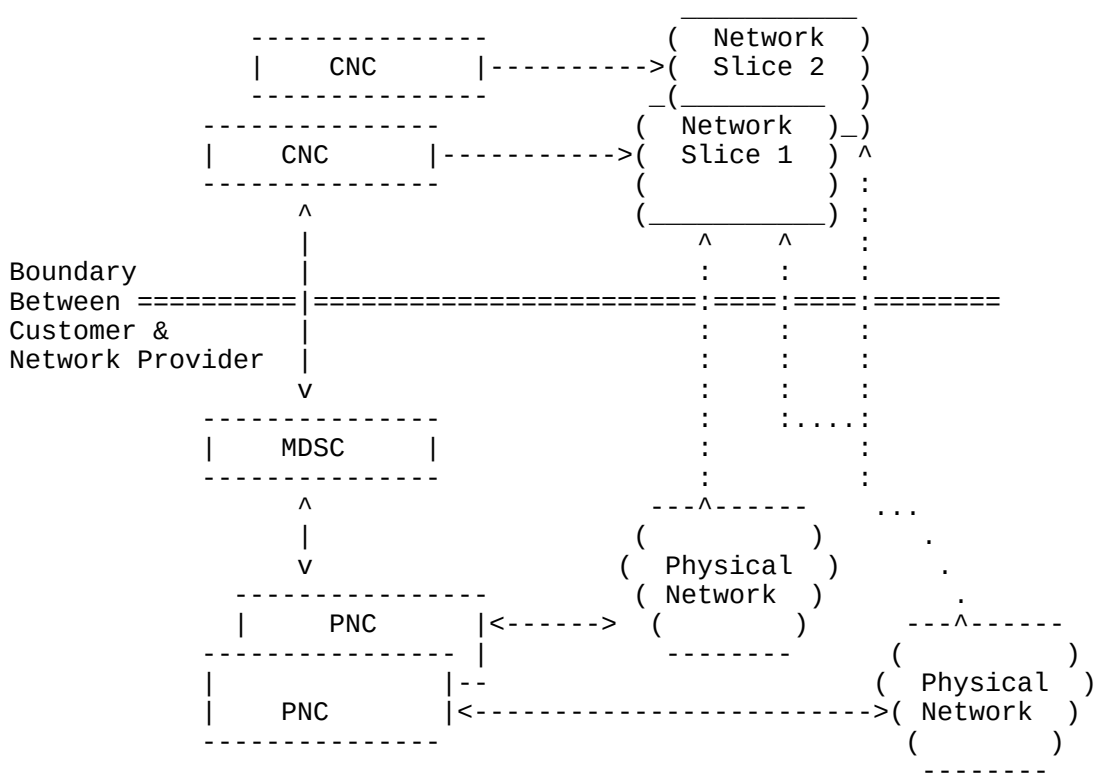

Figure 4: Network Slicing

3.3. Network Slice Service Mapping from TE to ACTN VN Models

The role of TE-service mapping model [te-service-mapping] is to create a binding relationship across a Layer-3 Service Model [l3sm], Layer-2 Service Model and TE Tunnel model, via a generic ACTN Virtual Network (VN) model [actn-vn].

The ACTN VN YANG model is a generic virtual network service model that allows customers (internal or external) to create a VN that meets the customer's service objective with various constraints.

The TE-service mapping model is needed to bind L3VPN specific service model with TE-specific parameters. This binding will facilitate a seamless service operation with underlay-TE network visibility. The TE-service model developed in this document can also be extended to support other services including L2SM, and L1CSM network service models.

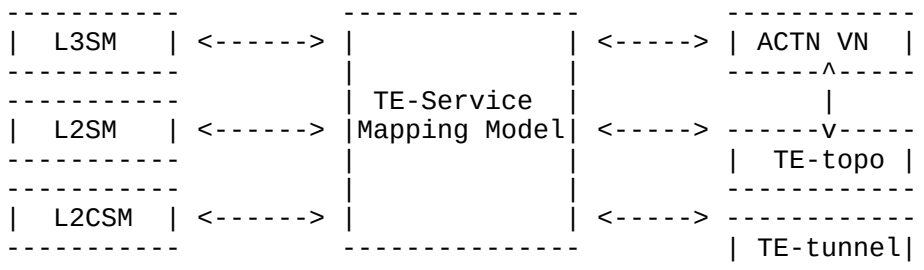

Figure 5: TE-Service Mapping ([te-service-mapping])

Editors note - We plan to provide a list of models available and their relationships/dependencies. We will also provide a vertical hierarchy of how these models may be used between functional components in ACTN.

3.4. ACTN VN KPI telemetry Models

The role of ACTN VN KPI telemetry model [actn-pm-telemetry] is to provide YANG models so that customer can define key performance monitoring data relevant for its VN/network slicing via the YANG subscription model.

Key characteristics of [actn-pm-telemetry] include:

- an ability to provide scalable VN-level telemetry aggregation based on customer-subscription model for key performance parameters defined by the customer;

- an ability to facilitate proactive re-optimization and reconfiguration of VNs/Netork Slices based on network autonomic traffic engineering scaling configuration mechanism.

5. IANA Considerations

This document makes no requests for action by IANA.

6. Security Considerations

Network slicing involves the control of network resources in order to meet the service requirements of consumers. In some deployment models, the consumer is able to directly request modification in the behaviour of resources owned and operated by a service provider. Such changes could significantly affect the service provider's ability to provide services to other consumers. Furthermore, the resources allocated for or consumed by a consumer will normally be billable by the service provider.

Therefore, it is crucial that the mechanisms used in any network slicing system allow for authentication of requests, security of those requests, and tracking of resource allocations.

It should also be noted that while the partitioning or slicing of resources is virtual, the consumers expect and require that there is no risk of leakage of data from one slice to another, no transfer of knowledge of the structure or even existence of other slices, and that changes to one slice (under the control of one consumer) should not have detrimental effects on the operation of other slices (whether under control of different or the same consumers) beyond the limits allowed within the SLA. Thus, slices are assumed to be private and to provide the appearance of genuine physical connectivity.

ACTN operates using the [netconf] or [restconf] protocols and assumes the security characteristics of those protocols. Deployment models for ACTN should fully explore the authentication and other security aspects before networks start to carry live traffic.

7. Acknowledgements

Thanks to Qin Wu, Andy Jones, Ramon Casellas, and Gert Grammel for their insight and useful discussions about network slicing.

8. References

8.1. Normative References

8.2. Informative References

[ngmn-network-slicing]
NGMN, "Description of Network Slicing Concept", 1 2016, <https://www.ngmn.org/uploads/media/160113_Network_Slicing_v1_0.pdf>.

[3gpp.28.801]
3GPP, "Study on management and orchestration of network slicing for next generation network", 3GPP TR 28.801 0.4.0, 1 2017, <http://www.3gpp.org/ftp/Specs/html-info/28801.htm>.

[network-slice-5g]
"Network Slicing for 5G with SDN/NFV: Concepts, Architectures and Challenges", Jose Ordonez-Lucena, Pablo Ameigeiras, Diego Lopez, Juan J. Ramos-Munoz, Javier Lorca, Jesus Folgueira, IEEE Communications Magazine 55, March 2017

[onf-tr526]
ONF TR-526, "Applying SDN Architecture to 5G Slicing", April 2016.

[actn-framework]
Ceccarelli, D. and Y. Lee, "Framework for Abstraction and Control of Traffic Engineered Networks", draft-ietf-teas-actn-framework, work in progress, February 2017.

[te-service-mapping]
Y. Lee, D. Dhody, and D. Ceccarelli, "Traffic Engineering and Service Mapping Yang Model", draft-lee-teas-te-service-mapping-yang, work in progress.

[actn-vn] Y. Lee (Editor), "A Yang Data Model for ACTN VN Operation", draft-lee-teas-actn-vn-yang, work in progress.

[actn-info] Y. Lee, S. Belotti (Editors), "Information Model for Abstraction and Control of TE Networks (ACTN)", draft-ietf-teas-actn-info-model, work in progress.

[actn-pm-elemetry] Y. Lee, et al, "YANG models for ACTN TE Performance Monitoring Telemetry and Network Autonomics", draft-lee- teas-actn-pm-telemetry-autonomics, work in progress.

[l3sm] Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", RFC 8049, February 2017

[netconf] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241.

[restconf] A. Bierman, M. Bjorklund, and K. Watsen, "RESTCONF Protocol", draft-ietf-netconf-restconf, work in progress.

[sf-topology] I. Bryskin, et al, "Use Cases for SF Aware Topology Models", draft-ietf-teas-use-cases-sf-aware-topo-model, work in progress.

[vpn+] S. Bryant and J. Dong, "Enhanced Virtual Private Networks (VPN+)", draft-bryant-rtgwg-enhanced-vpn, work in progress.

9. Contributors

The following people contributed text to this document.

Adrian Farrel
Email: afarrel@juniper.net

Mohamed Boucadair
Email: mohamed.boucadair@orange.com

Sergio Belotti
Email: sergio.belotti@nokia.com

Daniele Ceccarelli
Email: daniele.ceccarelli@ericsson.com

Haomian Zheng
Email: zhenghaomian@huawei.com

Authors' Addresses

Daniel King
Email: daniel@olddog.co.uk

Young Lee
Email: leeyoung@huawei.com