

TEAS Working Group
Internet-Draft
Intended status: Informational
Expires: October 2, 2021

D. King
Old Dog Consulting
J. Drake
Juniper Networks
H. Zheng
Huawei Technologies
A. Farrel
Old Dog Consulting
March 31, 2021

Applicability of Abstraction and Control of Traffic Engineered Networks
(ACTN) to Network Slicing
draft-king-teas-applicability-actn-slicing-10

Abstract

Network abstraction is a technique that can be applied to a network domain. It utilizes a set of policies to select network resources and obtain a view of potential connectivity across the network.

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 or virtual networks, each tailored for a set of services that share the same set of requirements.

Abstraction and Control of Traffic Engineered Networks (ACTN) is described in RFC 8453. It defines an SDN-based architecture that relies on the concept of network and service abstraction to detach network and service control from the underlying data plane.

This document outlines the applicability of ACTN to network slicing in a Traffic Engineering (TE) network that utilizes IETF technology. It also identifies the features of network slicing not currently within the scope of ACTN, and indicates where ACTN might be extended.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 2, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | |
|---|----|
| 1. Introduction | 3 |
| 1.1. Terminology | 4 |
| 2. Requirements for Network Slicing | 5 |
| 2.1. Resource Slicing | 6 |
| 2.2. Network Virtualization | 6 |
| 2.3. Service Isolation | 6 |
| 2.4. Control and Orchestration | 7 |
| 3. Abstraction and Control of Traffic Engineered (TE) Networks (ACTN) | 7 |
| 3.1. ACTN Virtual Network as a Network Slice | 8 |
| 3.2. ACTN Virtual Network for Network Slice Aggregation | 9 |
| 3.3. Management Components for ACTN and Network Slicing | 9 |
| 3.4. Examples of ACTN Delivering Types of Network Slices | 10 |
| 3.4.1. ACTN Used for Virtual Private Line | 10 |
| 3.4.2. ACTN Used for VPN Delivery Model | 12 |
| 3.4.3. ACTN Used to Deliver a Virtual Consumer Network | 13 |
| 4. YANG Models | 15 |
| 4.1. Network Slice Service Mapping from TE to ACTN VN Models | 15 |
| 4.2. Interfaces and Yang Models | 16 |
| 4.3. ACTN VN Telemetry | 17 |
| 5. IANA Considerations | 18 |
| 6. Security Considerations | 18 |
| 7. Acknowledgements | 19 |
| 8. Contributors | 19 |

| | |
|-------------------------------------|----|
| 9. Informative References | 19 |
| Authors' Addresses | 22 |

1. Introduction

The principles of network resource separation are not new. For years, the concept of separated overlay and logical (virtual) networking has existed, allowing multiple 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 network slice is a virtual (that is, logical) network with its own network topology and a set of network resources that are used to provide connectivity that conforms to a specific Service Level Agreement (SLA) or set of Service Level Objectives (SLOs). The network resources used to realize a network slice belong to the network that is sliced. The resources may be assigned and dedicated to an individual slice, or they may be shared with other slices enabling different degrees of service guarantee and providing different levels of isolation between the traffic in each slice.

[I-D.ietf-teas-ietf-network-slice-definition] provides a number of useful definitions for network slicing in the context of IETF network technologies. In particular, that document defines the term "IETF network slice" to be the generic network slice concept applied to a network that uses IETF technologies. An IETF network slice could span multiple technologies (such as IP, MPLS, or optical) and multiple administrative domains. The logical network that is an IETF network slice may be kept separate from other concurrent logical networks each with independent control and management: each can be created or modified on demand. Since this document is focused entirely on IETF technologies, it uses the term "network slice" as a more concise expression. Further discussion on the topic of IETF network slices can be found in [I-D.ietf-teas-ietf-network-slice-framework].

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. This type of service is called an enhanced VPN and is described in more detail in [I-D.ietf-teas-enhanced-vpn]. It is often based on Traffic Engineering (TE) constructs in the underlay network.

Abstraction and Control of TE Networks (ACTN) [RFC8453] is a framework that facilitates the abstraction of underlying network resources to higher-layer applications and that allows network operators to create virtual networks for their customers through the abstraction of the operators' network resources.

As noted in [I-D.ietf-teas-ietf-network-slice-framework], ACTN is a toolset capable of delivering network slice functionality. This document outlines the application of ACTN and associated enabling technologies to provide network slicing in a network that utilizes IETF technologies such as IP, MPLS, or GMPLS. It describes how the ACTN functional components can be used to support model-driven partitioning of resources into variable-sized bandwidth units to facilitate network sharing and virtualization. Furthermore, the use of model-based interfaces to dynamically request the instantiation of virtual networks can be extended to encompass requesting and instantiation of specific service functions (which may be both physical or virtual), and to partition network resources such as compute resource, storage capability, and access to content. Finally, this document highlights how the ACTN approach might be extended to address the requirements of network slicing where the underlying network is TE-capable.

1.1. Terminology

As far as is possible, this document re-uses terminology from [I-D.ietf-teas-ietf-network-slice-definition], [I-D.ietf-teas-enhanced-vpn] and [I-D.ietf-teas-ietf-network-slice-framework]. The terms defined below are give context and meaning for use in this document only and do not force wider applicability. As other work matures, it is hoped that the terminology will converge.

Service Provider: A server network or collection of server networks. The persons or organization responsible for operating such networks.

Consumer: As defined in [I-D.ietf-teas-ietf-network-slice-definition], a consumer is the component or entity that requests and uses a network slice. This may be any application, client network, or customer of a service

provider. In the ACTN framework [RFC8453] the consumer of a network service is termed a 'customer' because it will often be the case that a VPN consumer is a customer of the operator of the core network that delivers the service. In the context of a network slice, the consumer may well be a customer, but might also be a client network of the service provider (which could also be an internal organization of the service provider), or an application that engineers traffic in the network.

Service Functions (SFs): Components that provide specific functions within a network. SFs are often combined in a specific sequence called a service function chain to deliver services [RFC7665].

Resource: Any feature including connectivity, bufferage, compute, storage, and content delivery that forms part of or can be accessed through a network. Resources may be shared between users, applications, and clients, or they may be dedicated for use by a unique consumer.

Infrastructure Resources: The hardware and 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 Level Agreement (SLA): Per [I-D.ietf-teas-ietf-network-slice-definition], an SLA is an explicit or implicit contract between the consumer of a network slice and the provider of the slice. The SLA is expressed in terms of a set of Service Level Objectives (SLOs) and may include commercial terms as well as the consequences of violating the SLOs. The SLA 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 firewalls or billing) to be provided; the location, nature, and quantities of services (such as the amount and location of compute resources and the accelerators required).

Network Slice Service: An agreement between a consumer and a service provider to deliver network resources according to a specific service level agreement.

2. Requirements for Network Slicing

According to [I-D.ietf-teas-ietf-network-slice-framework] the consumer expresses requirements for a particular IETF network slice by specifying what is required rather than how the requirement is to be fulfilled. That is, the IETF network slice consumer's view of a IETF network slice is an abstract one.

The concept of network slicing is a key capability to serve consumers with a wide variety of different service needs expressed as SLOs in term of latency, reliability, capacity, and service function specific capabilities.

This section outlines the key capabilities required to realize network slicing in a TE-enabled IETF technology network.

2.1. Resource Slicing

Network resources need to be allocated and dedicated for use by a specific network slice, or they may be shared among multiple slices. This allows a flexible approach that can deliver a range of services by partitioning (that is, slicing) the available network resources to make them available to meet the consumer's SLA.

2.2. Network Virtualization

Network virtualization enables the creation of multiple virtual networks that are operationally decoupled from the underlying physical network, and are run on top of it. Slicing enables the creation of virtual networks as consumer services.

2.3. Service Isolation

A consumer may request, through their SLA, that changes to the other services delivered by the service provider do not have any negative impact on the delivery of the service. This quality is referred to as "isolation" [I-D.ietf-teas-ietf-network-slice-definition] [I-D.ietf-teas-enhanced-vpn].

Delivery of such service isolation may be achieved in the underlying network by various forms of resource partitioning ranging from dedicated allocation of resources for a specific slice, to sharing or resources with safeguards.

Although multiple network slices may utilize resources from a single underlying network, isolation should be understood in terms of the following three categorisations.

- o Performance isolation requires that service delivery for one network slice does not adversely impact congestion or performance levels of other slices.
- o Security isolation means that attacks or faults occurring in one slice do not impact on other slices. Moreover, the security functions supporting each slice must operate independently so that an attack or misconfiguration of security in one slice will not

prevent proper security function in the other slices. Further, privacy concerns require that traffic from one slice is not delivered to an end point in another slice, and that it should not be possible to determine the nature or characteristics of a slice from any external point.

- o Management isolation means that each slice must be independently viewed, utilized, and managed as a separate network. Furthermore, it should be possible to prevent the operator of one slice from being able to control, view, or detect any aspect of any other network slice.

2.4. Control and Orchestration

Orchestration combines and coordinates multiple control methods to provide a single mechanism to operate one or more networks to deliver services. In a network slicing environment, an orchestrator is needed to coordinate disparate processes and resources for creating, managing, and deploying the network slicing service. Two aspects of orchestration are required:

- o Multi-domain Orchestration: Managing connectivity to set up a network slice across multiple administrative domains.
- o End-to-end Orchestration: Combining resources for an end-to-end service (e.g., underlay connectivity with firewalling, and guaranteed bandwidth with minimum delay).

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

ACTN facilitates end-to-end connectivity and provide virtual connectivity services (such as virtual links and virtual networks) to the user. The ACTN framework [RFC8453] introduces three functional components and two interfaces:

- o Customer Network Controller (CNC)
- o Multi-domain Service Coordinator (MDSC)
- o Provisioning Network Controller (PNC)
- o CNC-MDSC Interface (CMI)
- o MDSC-PNC Interface (MPI)

RFC 8453 also highlights how:

- o Abstraction of the underlying network resources is provided to higher-layer applications and consumers.
- o Virtualization is achieved by selecting resources according to criteria derived from the details and requirements of the consumer, application, or service.
- o Creation of a virtualized environment is performed to allow operators to view and control multi-domain networks as a single virtualized network.
- o A network is presented to a consumer as a single virtual network via open and programmable interfaces.

The ACTN managed infrastructure consists of traffic engineered network resources. The concept of traffic engineering is broad: it describes the planning and operation of networks using a method of reserving and partitioning of network resources in order to facilitate traffic delivery across a network (see [I-D.ietf-teas-rfc3272bis] for more details). In the context of ACTN, traffic engineering network resources may include:

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

The ACTN network is "sliced" with consumers each being given a different partial and abstracted topology view of the physical underlay network.

3.1. ACTN Virtual Network as a Network Slice

To support multiple consumers, each with its own view of and control of a virtual network constructed using a server network, a service provider needs to partition the server network resources to create network slices assigned to each consumer.

An ACTN Virtual Network (VN) is a consumer view of a slice of the ACTN-managed infrastructure. It is a network slice that is presented to the consumer by the ACTN provider as a set of abstracted resources. See [I-D.ietf-teas-actn-vn-yang] for a detailed description of ACTN VNs and an overview of how various different types of YANG model are applicable to the ACTN framework.

Depending on the agreement between consumer and provider, various VN operations are possible:

- o Network Slice Creation: A VN could be pre-configured and created through static configuration or through dynamic request and negotiation between consumer and service provider. The VN must meet the network slice requirements specified in the SLA to satisfy the consumer's objectives.
- o Network Slice Operations: The VN may be modified and deleted based on consumer requests. The consumer can further act upon the VN to manage the consumer's traffic flows across the network slice.
- o Network Slice View: The VN topology is viewed from the consumer's perspective. This may be the entire VN topology or a collection of tunnels that are expressed as consumer end points, access links, intra domain paths and inter-domain links.

[RFC8454] describes a set of functional primitives that support these different ACTN VN operations.

3.2. ACTN Virtual Network for Network Slice Aggregation

Scaling considerations for IETF network slicing are an important consideration. If the service provider must manage and maintain network state for every network slice then this will quickly limit the number of customer services that can be supported.

The importance of network slice aggregation is discussed in [I-D.ietf-teas-enhanced-vpn] and further in [I-D.dong-teas-enhanced-vpn-vtn-scalability]. That work notes the importance of aggregating network slices into groups of similar slices before realizing those aggregates in the network.

The same consideration applies to ACTN VNs. But fortunately, ACTN VNs may be arranged hierarchically by recursing the MDSCs so that one VN is realised over another VN. This allows the VNs presented to the customer to be aggregated before they are instantiated in the physical network.

3.3. Management Components for ACTN and Network Slicing

The ACTN management components (CNC, MDSC, and PNC) and interfaces (CMI and MPI) are introduced in Section 3 and described in detail in [RFC8453]. The management components for network slicing are described in [I-D.ietf-teas-ietf-network-slice-framework] and are known as the consumer orchestration system, the IETF network slice controller (NSC), and the network controller. The network slicing

management components are separated by the network slice controller northbound interface (NSC NBI) and the network slice controller southbound interface (NSC SBI).

[I-D.ietf-teas-ietf-network-slice-framework] describes the mapping between network slicing management components and ACTN management components. This is presented visually in Figure 1 and provides a useful reference for understanding the material in Section 3.4 and Section 4.




Figure 1: Mapping Between IETF Network Slice and ACTN Components

3.4. Examples of ACTN Delivering Types of Network Slices

The examples that follow build on the ACTN framework to provide control, management, and orchestration for the network slice life-cycle. These network slices utilize common physical infrastructure, and meet specific service-level requirements.

Three examples are shown. Each uses ACTN to achieve a different network slicing scenario. All three scenarios can be scaled up in capacity or be subject to topology changes as well as changes of consumer requirements.

3.4.1. ACTN Used for Virtual Private Line

In the example shown in Figure 2, ACTN provides virtual connections between multiple consumer locations (sites accessed through Customer Edge nodes - CEs). The service is requested by the consumer (via

CNC-A) and delivered as a Virtual Private Line (VPL) service. The benefits of this model include:

- o Automated: the service set-up and operation is managed by the network provider.
- o Virtual: the private line connectivity is provided from Site A to Site C (VPL1) and from Site B to Site C (VPL2) across the ACTN-managed physical network.
- o Agile: on-demand adjustments to the connectivity and bandwidth are available according to the consumer's requests.




Figure 2: Virtual Private Line Model

3.4.2. ACTN Used for VPN Delivery Model

In the example shown in Figure 3, ACTN provides VPN connectivity between two sites across three physical networks. The requirements for the VPN are expressed by the users of the two sites who are the consumers. Their requests are directed to the CNC, and the CNC interacts with the network provider's MDSC. The benefits of this model include:

- o Provides edge-to-edge VPN multi-access connectivity.

- o Most of the function is managed by the network provider, with some flexibility delegated to the consumer-managed CNC.




Figure 3: VPN Model


3.4.3. ACTN Used to Deliver a Virtual Consumer Network

In the example shown in Figure 4, ACTN provides a virtual network to the consumer. This virtual network is managed by the consumer. The figure shows two virtual networks (Network Slice 1 and Network Slice 2) each created for a different consumer under the care of a different CNC. There are two physical networks controlled by separate PNCs. Network Slice 2 is built using resources from just

one physical network, while Network Slice 1 is constructed from resources from both physical networks.

The benefits of this model include:

- o The MDSC provides the topology to the consumer so that the consumer can control their network slice to fit their needs.
- o Applications can interact with their assigned network slices directly. The consumer may implement their own network control methods and traffic prioritization, and manage their own addressing schemes.
- o Consumers may further slice their virtual networks so that this becomes a recursive model.
- o Service isolation can be provided through selection of physical networking resources through a combination of efforts of the MSDC and PNC.
- o The network slice may include nodes with specific capabilities. These can be delivered as Physical Network Functions (PNFs) or Virtual Network Functions (VNFs).



Key: --- ACTN control connection
... Virtualization/abstraction through slicing

Figure 4: Network Slicing

4. YANG Models

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

The role of the TE-service mapping model [I-D.ietf-teas-te-service-mapping-yang] is to create a binding relationship across a Layer 3 Service Model (L3SM) [RFC8299], Layer 2 Service Model (L2SM) [RFC8466], and TE Tunnel model [I-D.ietf-teas-yang-te], via the generic ACTN Virtual Network (VN) model [I-D.ietf-teas-actn-vn-yang].

The ACTN VN model is a generic virtual network service model that allows consumers to specify a VN (i.e., network slice) that meets the consumer's service objectives with various constraints on how the service is delivered.

The TE-service mapping model [I-D.ietf-teas-te-service-mapping-yang] is used to bind the L3SM with TE-specific parameters. This binding facilitates seamless service operation and enables visibility of the underlay TE network. The TE-service model developed in that document can also be extended to support other services including L2SM, and the Layer 1 Connectivity Service Model (L1CSM) [I-D.ietf-ccamp-llcsm-yang] L1CSM network service models.

Figure 5 shows the relationship between the models discussed above.




Figure 5: TE-Service Mapping

4.2. Interfaces and Yang Models

Figure 6 shows the three ACTN components and two ACTN interfaces as listed in Section 3. The figure also shows the Southbound Interface (SBI) between the PNC and the devices in the physical network. That interface might be used to install state on every device in the network, or might instruct a "head-end" node if a control plane is used within the physical network. In the context of [RFC8309], the SBI uses one or more device configuration models.

The figure also shows the Network Slice Service Interface. This interface allows a consumer of a service to make requests for delivery of the service, and it facilitates the consumer modifying and monitoring the service. In the context of [RFC8309], this

"northbound interface (NBI)" is a customer service interface and uses a service model.

When an ACTN system is used to manage the delivery of network slices, a network slice resource model is needed. This model will be used for instantiation, operation, and monitoring of network and function resource slices. The YANG model defined in [I-D.wd-teas-transport-slice-yang] provides a suitable basis for requesting, controlling, and deleting, network slices.




Figure 6: The Yang Interfaces in Context

4.3. ACTN VN Telemetry

The ACTN VN KPI telemetry model [I-D.ietf-teas-actn-pm-telemetry-autonomics] provides a way for a consumer to define performance monitoring relevant for its VN/network slice via the NETCONF subscription mechanisms [RFC8639], [RFC8640], or using the equivalent mechanisms in RESTCONF [RFC8641], [RFC8650].

Key characteristics of [I-D.ietf-teas-actn-pm-telemetry-autonomics] include:

- o An ability to provide scalable VN-level telemetry aggregation based on consumer subscription model for key performance parameters defined by the consumer.
- o An ability to facilitate proactive re-optimization and reconfiguration of VNs/network 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, as mentioned in Section 2.3 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.

Some service providers may offer secure network slices as a service. Such services may claim to include edge-to-edge encryption for the consumer's traffic. However, a consumer should take full responsibility for the privacy and integrity of their traffic and should carefully consider using their own edge-to-edge encryption.

ACTN operates using the NETCONF [RFC6241] or RESTCONF [RFC8040] 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, Gert Grammel, and Kiran Makhijani for their insight and useful discussions about network slicing.

8. Contributors

The following people contributed text to this document.

Young Lee
Email: younglee.tx@gmail.com

Mohamed Boucadair
Email: mohamed.boucadair@orange.com

Sergio Belotti
Email: sergio.belotti@nokia.com

Daniele Ceccarelli
Email: daniele.ceccarelli@ericsson.com

9. Informative References

- [I-D.dong-teas-enhanced-vpn-vtn-scalability]
Dong, J., Li, Z., Qin, F., and G. Yang, "Scalability Considerations for Enhanced VPN (VPN+)", draft-dong-teas-enhanced-vpn-vtn-scalability-01 (work in progress), November 2020.
- [I-D.ietf-ccamp-llcsm-yang]
Lee, Y., Lee, K., Zheng, H., Dios, O., and D. Ceccarelli, "A YANG Data Model for L1 Connectivity Service Model (L1CSM)", draft-ietf-ccamp-llcsm-yang-13 (work in progress), November 2020.

- [I-D.ietf-teas-actn-pm-telemetry-autonomics]
Lee, Y., Dhody, D., Karunanithi, S., Vilata, R., King, D.,
and D. Ceccarelli, "YANG models for VN/TE Performance
Monitoring Telemetry and Scaling Intent Autonomics",
draft-ietf-teas-actn-pm-telemetry-autonomics-04 (work in
progress), November 2020.
- [I-D.ietf-teas-actn-vn-yang]
Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., and B.
Yoon, "A YANG Data Model for VN Operation", draft-ietf-
teas-actn-vn-yang-10 (work in progress), November 2020.
- [I-D.ietf-teas-enhanced-vpn]
Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A
Framework for Enhanced Virtual Private Networks (VPN+)
Service", draft-ietf-teas-enhanced-vpn-06 (work in
progress), July 2020.
- [I-D.ietf-teas-ietf-network-slice-definition]
Rokui, R., Homma, S., Makhiyani, K., Contreras, L., and J.
Tantsura, "Definition of IETF Network Slices", draft-ietf-
teas-ietf-network-slice-definition-00 (work in progress),
January 2021.
- [I-D.ietf-teas-ietf-network-slice-framework]
Gray, E. and J. Drake, "Framework for IETF Network
Slices", draft-ietf-teas-ietf-network-slice-framework-00
(work in progress), March 2021.
- [I-D.ietf-teas-rfc3272bis]
Farrel, A., "Overview and Principles of Internet Traffic
Engineering", draft-ietf-teas-rfc3272bis-10 (work in
progress), December 2020.
- [I-D.ietf-teas-te-service-mapping-yang]
Lee, Y., Dhody, D., Fioccola, G., WU, Q., Ceccarelli, D.,
and J. Tantsura, "Traffic Engineering (TE) and Service
Mapping Yang Model", draft-ietf-teas-te-service-mapping-
yang-05 (work in progress), November 2020.
- [I-D.ietf-teas-yang-te]
Saad, T., Gandhi, R., Liu, X., Beeram, V., and I. Bryskin,
"A YANG Data Model for Traffic Engineering Tunnels, Label
Switched Paths and Interfaces", draft-ietf-teas-yang-te-25
(work in progress), July 2020.

- [I-D.wd-teas-transport-slice-yang]
Bo, W., Dhody, D., Han, L., and R. Rokui, "A Yang Data Model for Transport Slice NBI", draft-wd-teas-transport-slice-yang-02 (work in progress), July 2020.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<https://www.rfc-editor.org/info/rfc7665>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8299] Wu, Q., Ed., Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", RFC 8299, DOI 10.17487/RFC8299, January 2018, <<https://www.rfc-editor.org/info/rfc8299>>.
- [RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", RFC 8309, DOI 10.17487/RFC8309, January 2018, <<https://www.rfc-editor.org/info/rfc8309>>.
- [RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <<https://www.rfc-editor.org/info/rfc8453>>.
- [RFC8454] Lee, Y., Belotti, S., Dhody, D., Ceccarelli, D., and B. Yoon, "Information Model for Abstraction and Control of TE Networks (ACTN)", RFC 8454, DOI 10.17487/RFC8454, September 2018, <<https://www.rfc-editor.org/info/rfc8454>>.
- [RFC8466] Wen, B., Fioccola, G., Ed., Xie, C., and L. Jalil, "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery", RFC 8466, DOI 10.17487/RFC8466, October 2018, <<https://www.rfc-editor.org/info/rfc8466>>.
- [RFC8639] Voit, E., Clemm, A., Gonzalez Prieto, A., Nilsen-Nygaard, E., and A. Tripathy, "Subscription to YANG Notifications", RFC 8639, DOI 10.17487/RFC8639, September 2019, <<https://www.rfc-editor.org/info/rfc8639>>.

- [RFC8640] Voit, E., Clemm, A., Gonzalez Prieto, A., Nilsen-Nygaard, E., and A. Tripathy, "Dynamic Subscription to YANG Events and Datastores over NETCONF", RFC 8640, DOI 10.17487/RFC8640, September 2019, <<https://www.rfc-editor.org/info/rfc8640>>.
- [RFC8641] Clemm, A. and E. Voit, "Subscription to YANG Notifications for Datastore Updates", RFC 8641, DOI 10.17487/RFC8641, September 2019, <<https://www.rfc-editor.org/info/rfc8641>>.
- [RFC8650] Voit, E., Rahman, R., Nilsen-Nygaard, E., Clemm, A., and A. Bierman, "Dynamic Subscription to YANG Events and Datastores over RESTCONF", RFC 8650, DOI 10.17487/RFC8650, November 2019, <<https://www.rfc-editor.org/info/rfc8650>>.

Authors' Addresses

Daniel King
Old Dog Consulting

Email: daniel@olddog.co.uk

John Drake
Juniper Networks

Email: jdrake@juniper.net

Haomian Zheng
Huawei Technologies

Email: zhenghaomian@huawei.com

Adrian Farrel
Old Dog Consulting

Email: adrian@olddog.co.uk