

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 10, 2020

E. Gray, Editor
Ericsson
J. Drake, Editor
Juniper Networks
R. Rokui
Nokia
D. Dhody
March 09, 2020

Framework for Transport Network Slices
draft-nsdt-teas-ns-framework-00

Abstract

This memo discusses setting up special-purpose transport connections using existing IETF technologies. These connections are called transport slices for the purposes of this memo. The memo discusses the general framework for this setup, the necessary system components and interfaces, and how abstract requests can be mapped to more specific technologies. The memo also discusses related considerations with monitoring and security.

This memo is intended for discussing interfaces and technologies. It is not intended to be a new set of concrete interfaces or technologies. Rather, it should be seen as an explanation of how some existing, concrete IETF VPN and traffic-engineering technologies can be used to create transport slices. Note that there are a rather large of these technologies, and new technologies or capabilities keep being added. This memo is also not intended presume any particular technology choice.

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 September 10, 2020.

Copyright Notice

Copyright (c) 2020 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	2
2.	Requirements	4
3.	Framework	4
3.1.	Management systems or other applications	5
3.2.	Expressing connectivity intents	5
3.3.	Transport Slice Controller (TSC)	7
3.3.1.	Northbound Interface (NBI)	8
3.4.	Mapping	9
3.5.	Underlying technology	9
4.	Applicability of ACTN to Transport Slices	10
5.	Considerations	12
5.1.	Monitoring	12
5.2.	Security Considerations	12
5.3.	Privacy Considerations	12
6.	Acknowledgments	13
7.	References	13
7.1.	Normative References	13
7.2.	Informative References	13
	Authors' Addresses	16

[1.](#) Introduction

This draft provides a framework for discussing transport slices, as defined in [[I-D.nsdt-teas-transport-slice-definition](#)] (to be replaced by [draft-rokui-teas-transport-slice-definition](#)). It is the intention in this document to use terminology consistent with this and other definitions provided in that draft.

In particular, this document uses the following terminology defined in the definitions document: - Transport Slice - Transport Slice

Controller (TSC) - Transport Network Controller (TNC) - Northbound Interface (NBI) - Southbound Interface (SBI)

This framework is intended as a structure for discussing interfaces and technologies. It is not intended to be a new set of concrete interfaces or technologies. Rather, the idea is that existing or under-development IETF technologies (plural) can be used to realize the ideas expressed here.

For example, virtual private networks (VPNs) have served the industry well as a means of providing different groups of users with logically isolated access to a common network. The common or base network that is used to provide the VPNs is often referred to as an underlay network, and the VPN is often called an overlay network. As an example technology, a VPN may in turn serve as an underlay network for transport slices.

Note: It is conceivable that extensions to these IETF technologies are needed in order to fully support all the ideas that can be implemented with slices, but at least in the beginning there is no plan for the creation of new protocols or interfaces.

Driven largely by needs surfacing from 5G, the concept of network slicing has gained traction ([[NGMN-NS-Concept](#)], [[TS23501](#)], [[TS28530](#)], and [[BBF-SD406](#)]). In [[TS23501](#)], Network Slice is defined as "a logical network that provides specific network capabilities and network characteristics", and a Network Slice Instance is defined as "A set of Network Function instances and the required resources (e.g. compute, storage and networking resources) which form a deployed Network Slice". According to [[TS28530](#)], an end-to-end network slice consists of three major types of network segments: Radio Access Network (RAN), Transport Network (TN) and Core Network (CN). Transport network provides the required connectivity between different entities in RAN and CN segments of an end-to-end network slice, with a specific performance commitment. For each end-to-end network slice, the topology and performance requirement on transport network can be very different, which requires the transport network to have the capability of supporting multiple different transport slices.

While network slices are commonly discussed in the context of 5G, it is important to note that transport slices are a narrower concept, and focus primarily on particular network connectivity aspects. Other systems, including 5G deployments, may use transport slices as a component to create entire systems and concatenated constructs that match their needs, including end-to-end connectivity.

A Transport Slice could span multiple technologies and multiple administrative domains. Depending on the consumer's requirements, a transport slice could be isolated from other, often concurrent transport slices in terms of data, control and management planes.

The consumer expresses requirements for a particular transport slice by specifying what is required rather than how the requirement is to be fulfilled. That is, the Transport Slice consumer's view of a transport slice is an abstract one.

Thus, there is a need to create logical network structures with required characteristics. The consumer of such a logical network can require a degree of isolation and performance that previously might not have been satisfied by traditional overlay VPNs. Additionally, the transport slice consumer might ask for some level of control of their virtual networks, e.g., to customize the service paths in a network slice.

This document specifies a framework for the use of existing technologies as components to provide a transport slice service, and might also discuss (or reference) modified and potential new technologies, as they develop (such as candidate technologies described in section 5 of [[I-D.ietf-teas-enhanced-vpn](#)]).

2. Requirements

It is intended that transport slices can be created to meet specific requirements, typically expressed as bandwidth, latency, latency variation, and other desired or required characteristics. Creation is initiated by a management system or other application used to specify network-related conditions for particular traffic flows.

And it is intended that, once created, these slices can be monitored, modified, deleted, and otherwise managed.

It is also intended that applications and components will be able to use these transport slices to move packets between the specified end-points in accordance with specified characteristics.

As an example of additional requirements that might apply to transport slices, see [[I-D.ietf-teas-enhanced-vpn](#)] (in particular, [section 3](#)).

3. Framework

A number of transport slice services will typically be provided over a shared underlying network infrastructure. Each transport slice consists of both the overlay connectivity and a specific set of

dedicated network resources and/or functions allocated in a shared underlay network to satisfy the needs of the transport slice consumer. In at least some examples of underlying network technologies, the integration between the overlay and various underlay resources is needed to ensure the guaranteed performance requested for different transport slices.

**** Align with definitions draft ****

The users of transport slices are:

- o The management system or other application that creates and manages them.
- o The applications and components wishing to use these slices.

Transport users are served by the system that can build transport slices, as follows:

- o A controller takes requests from a management system or other application, which are then communicated via an abstracted northbound interface. This interface transmits data objects that describe the needed transport slices in terms of topology and service level objectives (SLO).
- o These requests are assumed to be translated by one or more underlying systems, which establish specific transport slice instances, using specific implementation techniques (such as IP and MPLS) implemented on top of the underlying network infrastructure.

Section 3 of [[I-D.ietf-teas-enhanced-vpn](#)] provides an example architecture that might apply in using the technology described in that document.

3.1. Management systems or other applications

The transport slice system is used by a management system or other application. These systems and applications may also be a part of a higher level function in the system, e.g., putting together network functions, access equipment, application specific components, as well as the transport slices.

3.2. Expressing connectivity intents

The Transport Slice Controller (TSC) northbound interface (NBI) can be used to communicate between transport slice users (or consumers) and the TSC.

A transport slice user may be a network operator who, in turn, provides the transport slice to another transport slice user or consumer.

Using the NBI, a consumer expresses requirements for a particular slice by specifying what is required rather than how that is to be achieved. That is, the consumer's view of a slice is an abstract one. Consumers normally have limited (or no) visibility into the provider network's actual topology and resource availability information.

This should be true even if both the consumer and provider are associated with a single administrative domain, in order to reduce the potential for adverse interactions between transport slice consumers and other uses/users of the transport network infrastructure.

The benefits of this model can include:

- o Security: because the transport network (or network operator) does not need to expose network details (topology, capacity, etc.) to transport slice consumers the transport network components are less exposed to attack;
- o Layered Implementation: the transport network comprises network elements that belong to a different layer network than consumer applications, and network information (advertisements, protocols, etc.) that a consumer cannot interpret or respond to (note - a consumer should not use network information not exposed via the TSC NBI, even if that information is available);
- o Scalability: consumers do not need to know any information beyond that which is exposed via the NBI.

The general issues of abstraction in a TE network is described more fully in [\[RFC7926\]](#).

This framework document does not assume any particular layer at which transport slices operate as a number of layers (including virtual L2, Ethernet or IP connectivity) could be employed.

Data models and interfaces are of course needed to set up transport slices, and specific interfaces may have capabilities that allow creation of specific layers.

Layered virtual connections are comprehensively discussed in IETF documents and are widely supported. See, for instance, GMPLS-based networks ([\[RFC5212\]](#) and [\[RFC4397\]](#)), or ACTN ([\[RFC8353\]](#) and

[RFC8353]). The principles and mechanisms associated with layered networking are applicable to transport slices.

There are several IETF-defined mechanisms for expressing the need for a desired logical network. The NBI carries data either in a protocol-defined format, or in a formalism associated with a modeling language.

For instance:

- o Path Computation Element (PCE) Communication Protocol (PCEP) [RFC5440] and GMPLS User-Network Interface (UNI) using RSVP-TE [RFC4208] use a TLV-based binary encoding to transmit data.
- o Network Configuration Protocol (NETCONF) [RFC6241] and RESTCONF Protocol [RFC8040] use XML and JSON encoding.
- o gRPC/GNMI [I-D.openconfig-rtgwg-gnmi-spec] uses a binary encoded programmable interface;
- o SNMP ([RFC3417], [RFC3412] and [RFC3414] uses binary encoding (ASN.1).
- o For data modeling, YANG ([RFC6020] and [RFC7950]) may be used to model configuration and other data for NETCONF, RESTCONF, and GNMI - among others; ProtoBufs can be used to model gRPC and GNMI data; Structure of Management Information (SMI) [RFC2578] may be used to define Management Information Base (MIB) modules for SNMP, using an adapted subset of OSI's Abstract Syntax Notation One (ASN.1, 1988).

While several generic formats and data models for specific purposes exist, it is expected that transport slice management may require enhancement or augmentation of existing data models.

3.3. Transport Slice Controller (TSC)

The transport slice controller takes abstract requests for transport slices and implements them using a suitable underlying technology. A transport slice controller is the key building block for control and management of the transport slice. It provides the creation/modification/deletion, monitoring and optimization of transport Slices in a multi-domain, a multi-technology and multi-vendor environment.

A TSC northbound interface (NBI) is needed for communicating details of a transport slice, as well as providing information to a slice

requester/consumer about transport slice status and performance. The details for this NBI are not in scope for this document.

The controller provides the following functions:

- o Provides a technology-agnostic NBI for creation/modification/deletion of the transport slices. The API exposed by this NBI communicates the endpoints of the transport slice, transport slice SLO parameters (and possibly monitoring thresholds), applicable input selection (filtering) and various policies, and provides a way to monitor the slice.
- o Determines an abstract topology connecting the endpoints of the transport slice that meets criteria specified via the NBI. The TSC also retains information about the mapping of this abstract topology to underlying components of the transport slice as necessary to monitor transport slice status and performance.
- o Provides "Mapping Functions" for the realization of transport slices. In other words, it will use the mapping functions that:
 - map technology-agnostic NBI request to technology-specific SBIs.
 - map filtering/selection information as necessary to entities in the underlay network.
- o Via an SBI, the controller collects Telemetry data (e.g. OAM results, Statistics, States etc.) for all elements in the abstract topology used to realize the transport slice.
- o Using the Telemetry data from the underlying realization of a transport slice (i.e. services/paths/tunnels), evaluates the current performance against transport slice SLO parameters and exposes them to the transport slice consumer via the NBI. The TSC NBI may also include a capability to provide notification in case the transport slice performance reaches threshold values defined by the transport slice consumer.

3.3.1. Northbound Interface (NBI)

The Transport Slice Controller provides a Northbound Interface (NBI) that allows consumers of network slices to request and monitor transport slices. Consumers operate on abstract transport slices, with details related to their realization hidden.

The NBI complements various IETF services, tunnels, path models by providing an abstract layer on top of these models.

The NBI is independent of type of network functions or services that need to be connected, i.e. it is independent of any specific storage, software, protocol, or platform used to realize physical or virtual network connectivity or functions in support of transport slices.

The NBI uses protocol mechanisms and information passed over those mechanisms to convey desired attributes for transport slices and their status. The information is expected to be represented as a well-defined data model, and should include at least endpoint and connectivity information, SLO specification, and status information.

To accomplish this, the NBI needs to convey information needed to support communication across the NBI, in terms of identifying the transport slices, as well providing the above model information.

3.4. Mapping

The main task of the transport slice controller is to map abstract transport slice requirements to concrete technologies and establish the required connectivity, and ensuring that required resources are allocated to the transport slice.

3.5. Underlying technology

There are a number of different technologies that can be used, including physical connections, MPLS, TSN, Flex-E, etc.

See [[I-D.ietf-teas-enhanced-vpn](#)] - [section 5](#) - for instance, for example underlying technologies.

Also, as outlined in "applicability of ACTN to Transport Slices" below, ACTN ([[RFC8453](#)]) offers a framework that is used elsewhere in IETF specifications to create virtual network (VN) services similar to Transport Slices.

A transport slice can be realized in a network, using specific underlying technology or technologies. The creation of a new transport slice will be initiated with following three steps:

- o Step 1: A higher level system requests connections with specific characteristics via NBI.
- o Step 2: This request will be processed by a Transport Slice Controller which specifies a mapping between northbound request to any IETF Services, Tunnels, and paths models.

- o Step 3: A series of requests for creation of services, tunnels and paths will be sent to the network to realize the transport slice.

It is very clear that regardless of how transport slice is realized in the network (i.e. using tunnels of type RSVP or SR), the definition of transport slice does not change at all but rather its realization.

4. Applicability of ACTN to Transport Slices

[RFC8453] defined three controllers as per the framework for Abstraction and Control of TE Networks (ACTN) to support virtual network (VN) services - Customer Network Controller (CNC), Multi-Domain Service Coordinator (MDSC) and Provisioning Network Controller (PNC). A CNC is responsible for communicating a customer's virtual network requirements, a MDSC is responsible for multi-domain coordination, virtualization/abstraction, customer mapping/translation and virtual service coordination to realize the virtual network requirement. Its key role is to detach the network/service requirements from the underlying technology. A PNC oversees the configuration, monitoring and collection of the network topology. The PNC is a underlay technology specific controller.

While the ACTN framework is a generic VN framework that is used for various VN service beyond the transport slice, it is still a suitable based to understand how the various controllers interact to realize the Transport slice.

A mapping between the Transport Slice definitions and ACTN definitions is as shown in Figure 1 below.

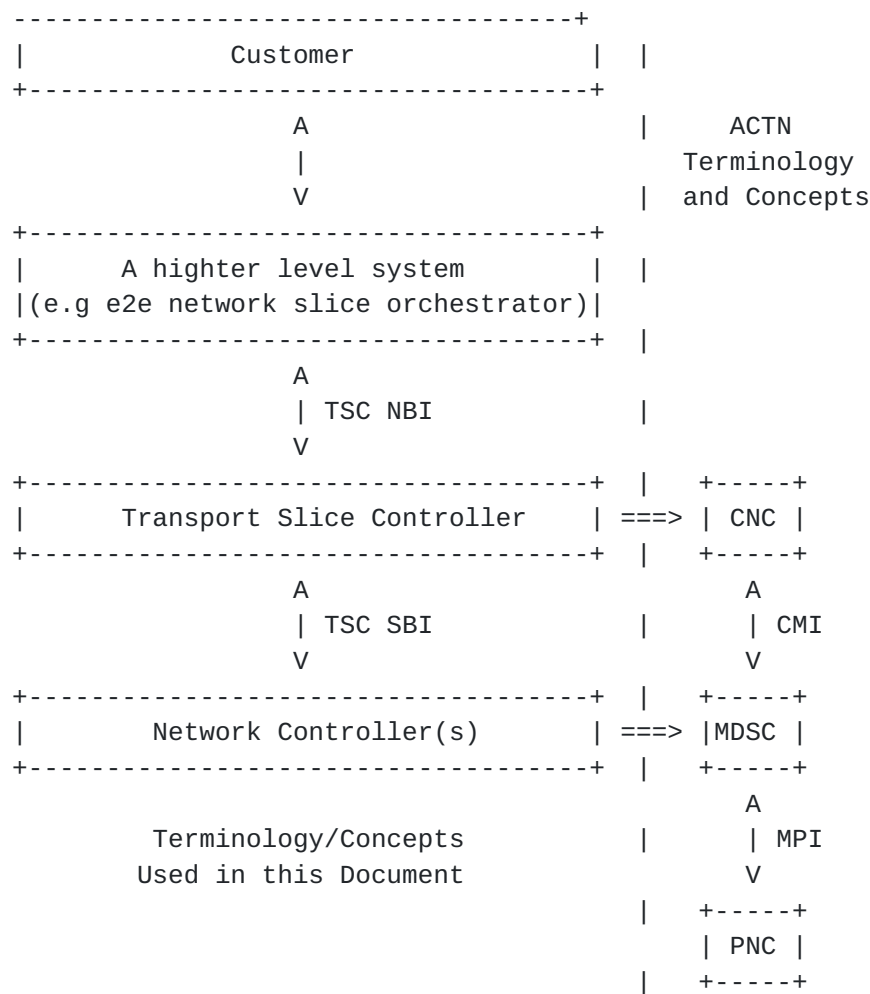


Figure 1

The TSC NBI conveys the generic transport slice requirements. These may then be realized using an SBI.

As per [RFC8453] and [I-D.ietf-teas-actn-yang], the CNC-MDSC Interface (CMI) is used to convey the virtual network service requirements along with the service models and the MDSC-PNC Interface (MPI) is used to realize the service along network configuration models. [I-D.ietf-teas-te-service-mapping-yang] further describe how the VPN services can be mapped to the underlying TE resources.

The Transport Network Controller is depicted as a single block, where as in the ACTN framework this has been decomposed into MDSC and PNC to handle multiple domains and various underlay technologies.

[RFC8453] also describes TE Network Slicing in the context of ACTN as a collection of resources that is used to establish a logically dedicated virtual network over one or more TE networks. In case of

TE enabled underlying network, ACTN VN can be used as a base to realize the transport network slicing by coordination among multiple peer domains as well as underlay technology domains.

5. Considerations

5.1. Monitoring

Transport slice realization needs to be instrumented in order to track how it is working, and it might be necessary to modify the transport slice as requirements change. Dynamic reconfiguration might be needed.

5.2. Security Considerations

Transport slices might use underlying virtualized networking. All types of virtual networking require special consideration to be given to the separation of traffic between distinct virtual networks, as well as some degree of protection from effects of traffic use of underlying network (and other) resources from other virtual networks sharing those resources.

For example, if a service requires a specific upper bound of latency, then that service can be degraded by added delay in transmission of service packets through the activities of another service or application using the same resources.

Similarly, in a network with virtual functions, noticeably impeding access to a function used by another transport slice (for instance, compute resources) can be just as service degrading as delaying physical transmission of associated packet in the network.

While a transport slice might include encryption and other security features as part of the service, consumers might be well advised to take responsibility for their own security needs, possibly by encrypting traffic before hand-off to a service provider.

5.3. Privacy Considerations

Privacy of transport network slice service consumers must be preserved. It should not be possible for one transport slice consumer to discover the presence of other consumers, nor should sites that are members of one transport slice be visible outside the context of that transport slice.

In this sense, it is of paramount importance that the system use the privacy protection mechanism defined for the specific underlying technologies used, including in particular those mechanisms designed

to preclude acquiring identifying information associated with any transport slice consumer.

6. Acknowledgments

The entire TEAS NS design team and everyone participating in related discussions has contributed to this draft. Some text fragments in the draft have been copied from the [[I-D.ietf-teas-enhanced-vpn](#)], for which we are grateful.

Significant contributions to this document were gratefully received from Dhruv Dhody, Xufeng Liu, Jie Dong, and Reza Rokui. Useful comments have also been received from Sergio Belotti and Kiran Makhijani.

7. References

7.1. Normative References

[I-D.nsdtd-teas-transport-slice-definition]
Rokui, R., Homma, S., and K. Makhijani, "IETF Definition of Transport Slice", [draft-nsdt-teas-transport-slice-definition-00](#) (work in progress), November 2019.

7.2. Informative References

[BBF-SD406]
Broadband Forum, ., "End-to-end network slicing", BBF SD-406 , n.d..

[I-D.ietf-teas-actn-yang]
Lee, Y., Zheng, H., Ceccarelli, D., Yoon, B., Dios, O., Shin, J., and S. Belotti, "Applicability of YANG models for Abstraction and Control of Traffic Engineered Networks", [draft-ietf-teas-actn-yang-05](#) (work in progress), February 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+) Services", [draft-ietf-teas-enhanced-vpn-05](#) (work in progress), February 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-03](#) (work in progress), March 2020.

- [I-D.openconfig-rtgwg-gnmi-spec]
Shakir, R., Shaikh, A., Borman, P., Hines, M., Lebsack, C., and C. Morrow, "gRPC Network Management Interface (gNMI)", [draft-openconfig-rtgwg-gnmi-spec-01](#) (work in progress), March 2018.
- [NGMN-NS-Concept]
NGMN Alliance, ., "Description of Network Slicing Concept", https://www.ngmn.org/uploads/media/161010_NGMN_Network_Slicing_framework_v1.0.8.pdf , 2016.
- [RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, [RFC 2578](#), DOI 10.17487/RFC2578, April 1999, <<https://www.rfc-editor.org/info/rfc2578>>.
- [RFC3412] Case, J., Harrington, D., Presuhn, R., and B. Wijnen, "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)", STD 62, [RFC 3412](#), DOI 10.17487/RFC3412, December 2002, <<https://www.rfc-editor.org/info/rfc3412>>.
- [RFC3414] Blumenthal, U. and B. Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", STD 62, [RFC 3414](#), DOI 10.17487/RFC3414, December 2002, <<https://www.rfc-editor.org/info/rfc3414>>.
- [RFC3417] Presuhn, R., Ed., "Transport Mappings for the Simple Network Management Protocol (SNMP)", STD 62, [RFC 3417](#), DOI 10.17487/RFC3417, December 2002, <<https://www.rfc-editor.org/info/rfc3417>>.
- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", [RFC 4208](#), DOI 10.17487/RFC4208, October 2005, <<https://www.rfc-editor.org/info/rfc4208>>.
- [RFC4397] Bryskin, I. and A. Farrel, "A Lexicography for the Interpretation of Generalized Multiprotocol Label Switching (GMPLS) Terminology within the Context of the ITU-T's Automatically Switched Optical Network (ASON) Architecture", [RFC 4397](#), DOI 10.17487/RFC4397, February 2006, <<https://www.rfc-editor.org/info/rfc4397>>.

- [RFC5212] Shiimoto, K., Papadimitriou, D., Le Roux, J.L., Vigoureux, M., and D. Brungard, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", [RFC 5212](#), DOI 10.17487/RFC5212, July 2008, <<https://www.rfc-editor.org/info/rfc5212>>.
- [RFC5278] Livingood, J. and D. Troshynski, "IANA Registration of Enumservices for Voice and Video Messaging", [RFC 5278](#), DOI 10.17487/RFC5278, July 2008, <<https://www.rfc-editor.org/info/rfc5278>>.
- [RFC5440] Vasseur, J.P., Ed. and J.L. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", [RFC 6020](#), DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.
- [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>>.
- [RFC7149] Boucadair, M. and C. Jacquenet, "Software-Defined Networking: A Perspective from within a Service Provider Environment", [RFC 7149](#), DOI 10.17487/RFC7149, March 2014, <<https://www.rfc-editor.org/info/rfc7149>>.
- [RFC7926] Farrel, A., Ed., Drake, J., Bitar, N., Swallow, G., Ceccarelli, D., and X. Zhang, "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", [BCP 206](#), [RFC 7926](#), DOI 10.17487/RFC7926, July 2016, <<https://www.rfc-editor.org/info/rfc7926>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", [RFC 7950](#), DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [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>>.

- [RFC8172] Morton, A., "Considerations for Benchmarking Virtual Network Functions and Their Infrastructure", [RFC 8172](#), DOI 10.17487/RFC8172, July 2017, <<https://www.rfc-editor.org/info/rfc8172>>.
- [RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", [RFC 8345](#), DOI 10.17487/RFC8345, March 2018, <<https://www.rfc-editor.org/info/rfc8345>>.
- [RFC8353] Upadhyay, M., Malkani, S., and W. Wang, "Generic Security Service API Version 2: Java Bindings Update", [RFC 8353](#), DOI 10.17487/RFC8353, May 2018, <<https://www.rfc-editor.org/info/rfc8353>>.
- [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>>.
- [RFC8568] Bernardos, CJ., Rahman, A., Zuniga, JC., Contreras, LM., Aranda, P., and P. Lynch, "Network Virtualization Research Challenges", [RFC 8568](#), DOI 10.17487/RFC8568, April 2019, <<https://www.rfc-editor.org/info/rfc8568>>.
- [TS23501] 3GPP, ., "System architecture for the 5G System (5GS)", 3GPP TS 23.501 , 2019.
- [TS28530] 3GPP, ., "Management and orchestration; Concepts, use cases and requirements", 3GPP TS 28.530 , 2019.

Authors' Addresses

Eric Gray
Ericsson

Email: eric.gray@ericsson.com

John Drake
Juniper Networks

Email: jdrake@juniper.net

Reza Rokui
Nokia

Email: reza.rokui@nokia.com

Dhruv Dhody

Email: dhruv.ietf@gmail.com