TEAS Working Group
Internet-Draft
Intended status: Informational
Expires: 1 March 2025

D. King
Old Dog Consulting
J. Drake
Independent
H. Zheng
Huawei Technologies
A. Farrel
Old Dog Consulting

28 August 2024

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

Abstract

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

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 Engineered (TE) network that utilizes IETF technologies. 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/.

Copyright Notice

Table of Contents

1. Introduction

The principles of network resource separation are not new. For years, the concepts of separated overlay and logical (virtual) networking have existed, allowing multiple services to be deployed over a single physical network comprised of a single or multiple layers. However, several key aspects 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 specific Service Level Agreements (SLAs) or a 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.

[RFC9543] provides a definition for network slicing in the context of IETF network technologies. In particular, that document defines the term "IETF Network Slice" as 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 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 and details of how an IETF Network Slice service may be requested and realized as an IETF Network Slice can be found in [RFC9543].

Within this document, the terms "network slice", "network slice service", and "network slice controller" refer to network slicing of networks built using IETF technologies as described in [RFC9543].

At one end of the spectrum, a Virtual Private Wire (VPW) 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 network service that will meet the needs of a set of applications with connectivity and service function requirements that may include compute resources, storage capabilities, 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. An example of such a type of service can be provided using an enhanced VPN described 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 and supply virtual networks for their customers through the abstraction of the operators' network resources.

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 TE-based technologies. 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 resources, storage capability, and access to content. In Section 3, the 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

This document re-uses terminology from [RFC8453], [RFC9543] and [I-D.ietf-teas-enhanced-vpn].

Service Provider: See "Provider" in [RFC9543].

Consumer: See [RFC9543].

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, buffers, 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 customer.

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): See [RFC9543].

Service Level Expectation (SLE): See [RFC9543].

Service Level Objective (SLO): See [RFC9543].

IETF Network Slice Service: See [RFC9543].

2. Overview of Key Requirements for Network Slicing

According to Section 6.2 of [RFC9543] "Expressing Connectivity Intents", the customer expresses requirements for a particular network slice by specifying what is required rather than how the requirement is to be fulfilled. That is, the customer’s view of a network slice is an abstract one expressed as a network slice service request.

The concept of network slicing is a key capability to serve a customer with a wide variety of different service needs expressed as SLOs/SLEs in terms of, e.g., 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 Partitioning

Network resources can be allocated and dedicated for use by a specific network slice service, or they may be shared among multiple slice services. 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 customer’s SLA.

2.2. Network Topology Customization and 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 customer services.

2.3. Service Isolation

A customer 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" in (Section 8 of [RFC9543]).

Delivery of 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 of 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 categorizations.

* Performance isolation requires that service delivery for one network slice does not adversely impact congestion, packet drop, or performance levels perceived by the users of other slices.

* 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 endpoint in another slice, and that it should not be possible to determine the nature or characteristics of a slice from any external point.

* 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. Orchestration

An orchestrator is used to coordinate disparate processes and resources for creating, managing, and deploying the network slicing service in a network. The following aspects of orchestration should be considered:

* Multi-domain Orchestration: Managing connectivity to set up a network slice across multiple administrative domains.

* 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): Overview of Key Components

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

* Customer Network Controller (CNC)

* Multi-domain Service Coordinator (MDSC)

* Provisioning Network Controller (PNC)

* CNC-MDSC Interface (CMI)

* MDSC-PNC Interface (MPI)

RFC 8453 also highlights how:

* Abstraction of the underlying network resources is provided to higher-layer applications and customers.

* Virtualization is achieved by selecting resources according to criteria derived from the details and requirements of the customer, application, or service.

* Creation of a virtualized environment is performed to allow operators to view and control multi-domain networks as a single virtualized network.

* A network is presented to a customer as a single virtual network via open and programmable interfaces.

The ACTN infrastructure resources include traffic-engineered network capabilities. 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 [RFC9522] for more details).

In the context of ACTN, traffic engineered infrastructure resources may include Statistical Packet Bandwidth, which refers to using statistical methods instead of assigning fixed bandwidth. This approach allocates bandwidth based on how data is flowing and statistical multiplexing. ACTN traffic engineered network resources also consider the physical parts of the network, such as optical channels and time slots, which facilitates the best use of the network’s resources by matching bandwidth with real-time traffic demands.

Therefore, an ACTN network may be "sliced", with each customer 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 customers, each with its own view and control of a virtual network constructed using an underlay network, a service provider needs to partition the network resources to create network slices assigned to each customer.

An ACTN Virtual Network (VN) is a customer view of a slice of the ACTN infrastructure resources. It is a network slice that is presented to the customer 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 models are applicable to the ACTN framework.

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

* Network Slice Creation: A VN could be pre-configured and created through static configuration or through a dynamic request and negotiation between a customer and service provider. The VN must meet the network slice requirements specified in the SLA to satisfy the customers objectives.

* Network Slice Operations: The VN may be modified or deleted based on direct customer requests. Also, the way that the VN is engineered can be adjusted by the operator to continuously ensure that the delivered service complies with the requested SLA. The customer can further act upon the VN to manage their traffic flows across the network slice.

* Network Slice View: A VN topology is viewed from the customer's perspective. This may be the entire VN topology, or a collection of tunnels that are expressed as customer endpoints, access links, intra-domain paths and inter-domain links.

Section 3, "Virtual Network Primitives", in [RFC8454] describes a set of functional primitives that support these different ACTN VN operations.

3.2. ACTN Virtual Network and Scaling Network Slices

If the service provider must manage and maintain state in the core of the network for every network slice, then this will quickly limit the number of customer services that can be supported.

The importance of scalability for network slices is discussed in [I-D.ietf-teas-enhanced-vpn] and further in [I-D.ietf-teas-nrp-scalability]. That work notes the importance of collecting network slices or their composite connectivity constructs into groups that require similar treatment in the network before realizing those groups 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 realized 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 [RFC9543] and are known as the customer orchestration system, the IETF Network Slice Controller (NSC), and the network controller. The network slicing management components are separated by the Network Slice Service Interface and the Network Configuration Interface, modeling the architecture described in [RFC8309].

The mapping between network slicing management components and ACTN management components is presented visually in Figure 1 and provides a reference for understanding the material in Section 3.4 and Section 4.

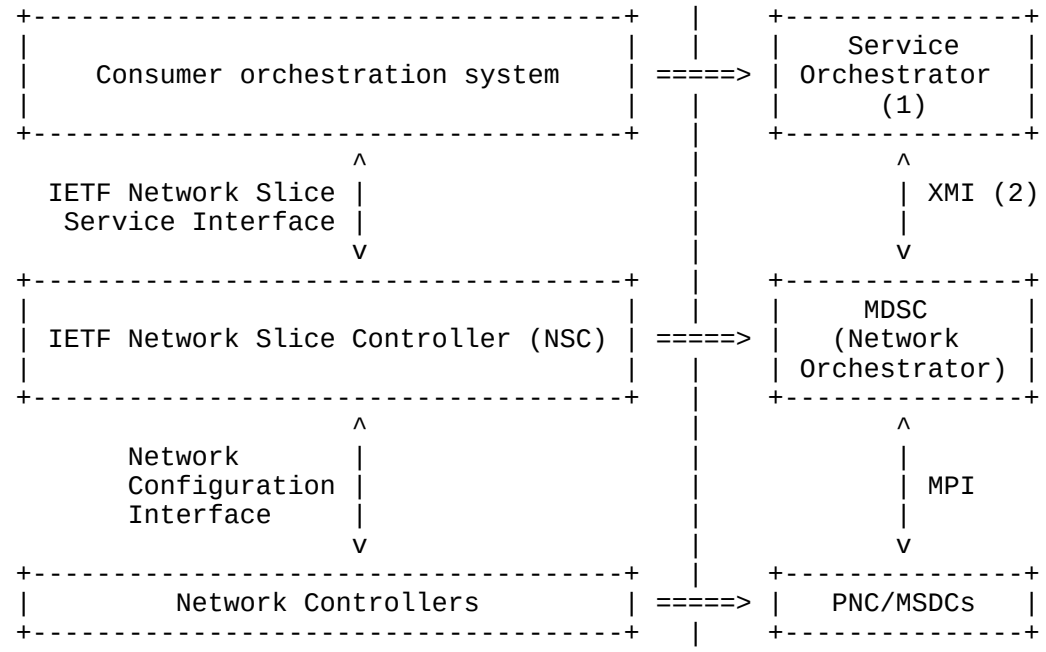

Figure 1: Mapping Between IETF Network Slice and ACTN Management Components

Note 1 - The Service Orchestrator may also contain some MDSC service-related functions, as described in section 4.2 of [RFC8453].

Note 2 - The Service Orchestrator-to-MDSC Interface (XMI) is an interface between two MDSC functional elements encompassing different MDSC service-related functions which is not defined in [RFC8453].

Note 2 - The Service Orchestrator-to-MDSC Interface (XMI) is an interface between two MDSC functional elements encompassing different MDSC service-related functions which is not defined in [RFC8453]. Depending on the function being delivered, the XMI might be realised by the Layer 2 VPN Network Management YANG model [RFC9291] or the Layer 3 VPN Network Management YANG model [RFC9182].

3.4. Examples of ACTN Delivering Network Slice Services

The following examples 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 in customer requirements.

3.4.1. ACTN Used for Virtual Private Line

In the example shown in Figure 2, ACTN provides virtual connections between multiple customer locations (sites accessed through Customer Edge nodes - CEs). The service is requested by the customer (via CNC-A) and delivered as a Virtual Private Line (VPL) service. The characteristics of this model include the following benefits.

* Programmable: The service setup and operation are managed by the network provider via APIs.

* 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 infrastructure resources (physical network).

* Flexible: On-demand adjustments to the connectivity and bandwidth are available according to the customer’s requests, which may be automated.

In terms of the network slicing concept defined in [RFC9543], in this example the customer requests a single network slice with two pairs of point-to-point connectivity constructs between the service demarcation points CE1 and CE3, and CE2 and CE3 with each pair comprising one connectivity construct in each direction.

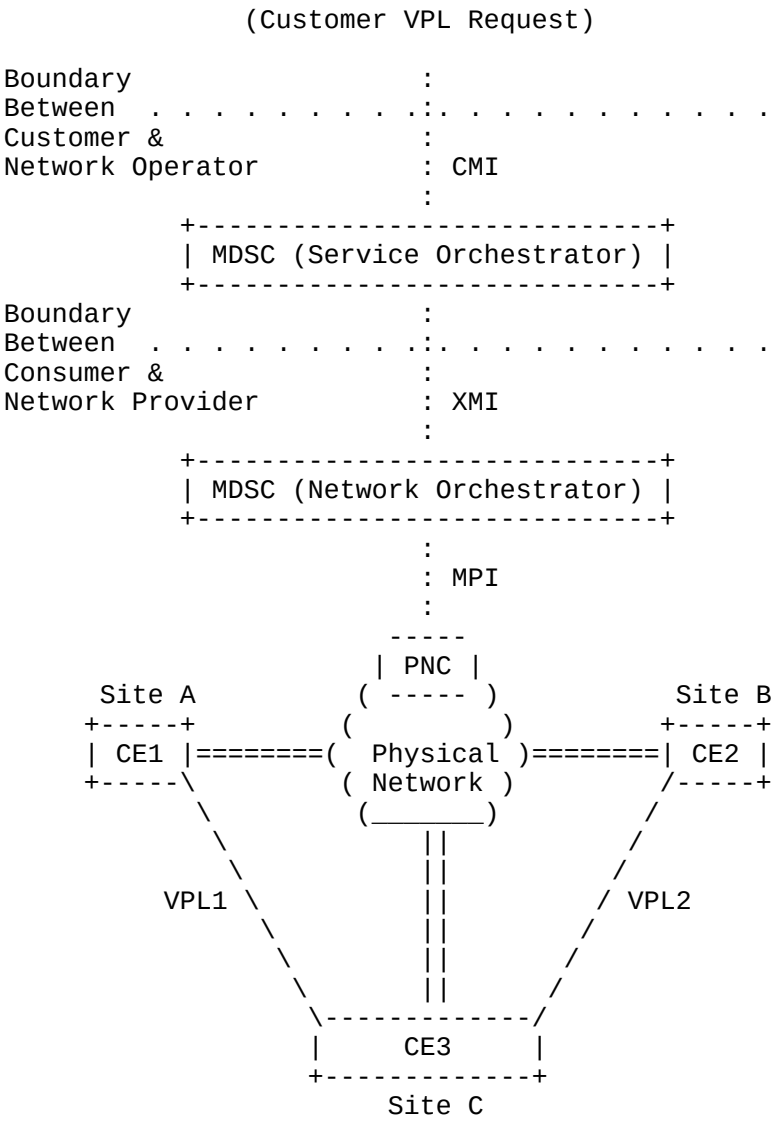

Figure 2: Virtual Private Line Model

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 users of the two sites express the requirements for the VPN. The request is directed to the CNC, and the CNC interacts with the network provider's MDSC. The main characteristics of this model are as follows.

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

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

In terms of the network slicing concept defined in [RFC9543], in this example, the customer requests a single network slice with a pair of point-to-point connectivity constructs (one in each direction) between the service demarcation points at site A and site B. The customer is unaware that the service is delivered over multiple physical networks.

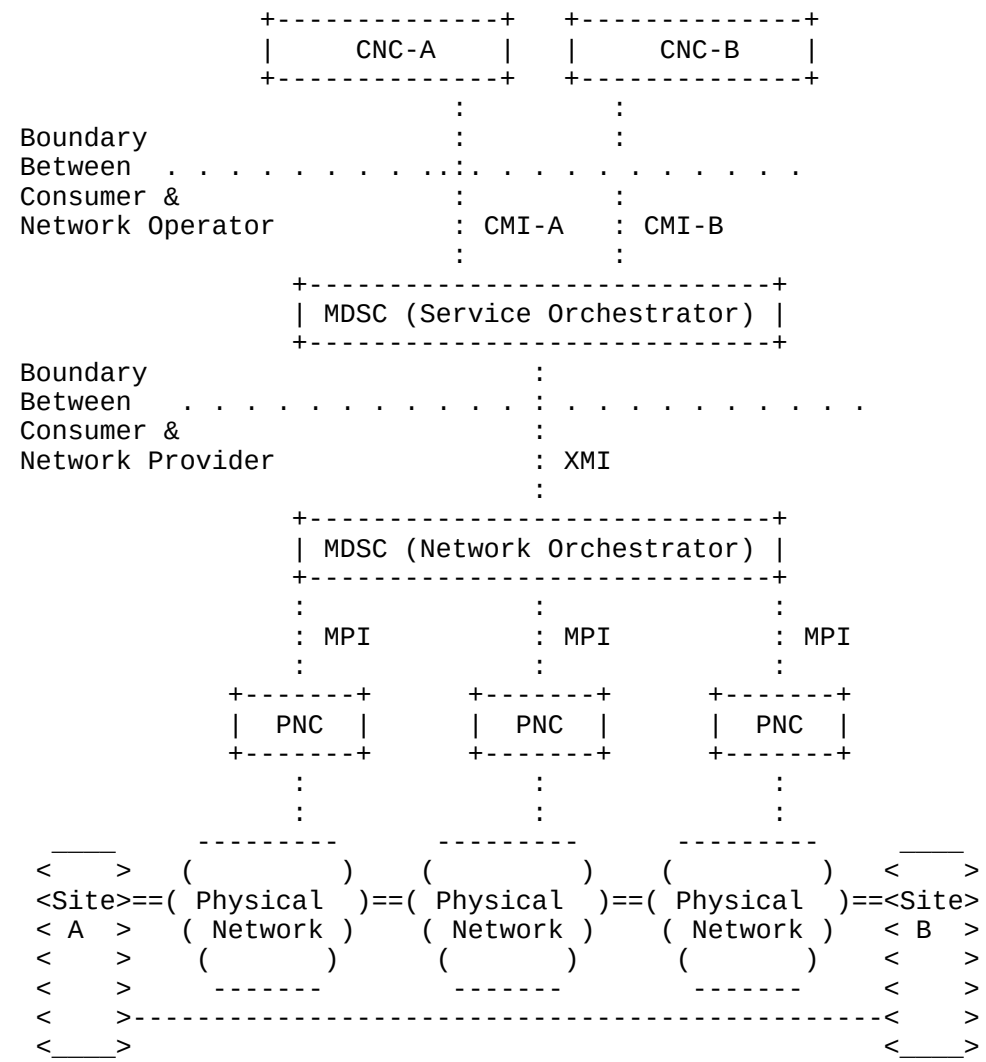

Key: ... ACTN control connectivity
=== Physical connectivity
--- Logical connectivity

Figure 3: VPN Model

3.4.3. ACTN Used to Deliver a Virtual Customer Network

In the example shown in Figure 4, ACTN provides a virtual network to the customer. This virtual network is managed by the customer. The figure shows two virtual networks (Network Slice 1 and Network Slice 2) each created for different customers under the care of different CNCs. 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 using resources from both physical networks.

The characteristics of this model include the following.

* The MDSC provides the topology to the customer so that the customer can control their network slice to fit their needs.

* Customers may interact with their assigned network slices directly. The customer may implement their own network control methods and traffic classification, mapping, prioritization, and manage their own addressing schemes.

* Customers may further slice their virtual networks so that this becomes a recursive model.

* Service isolation can be provided through selection of physical networking resources through a combination of efforts of the MSDC and PNC.

* The network slice may include nodes with specific capabilities. These can be delivered as Physical Network Functions (PNFs) or Virtual Network Functions (VNFs).

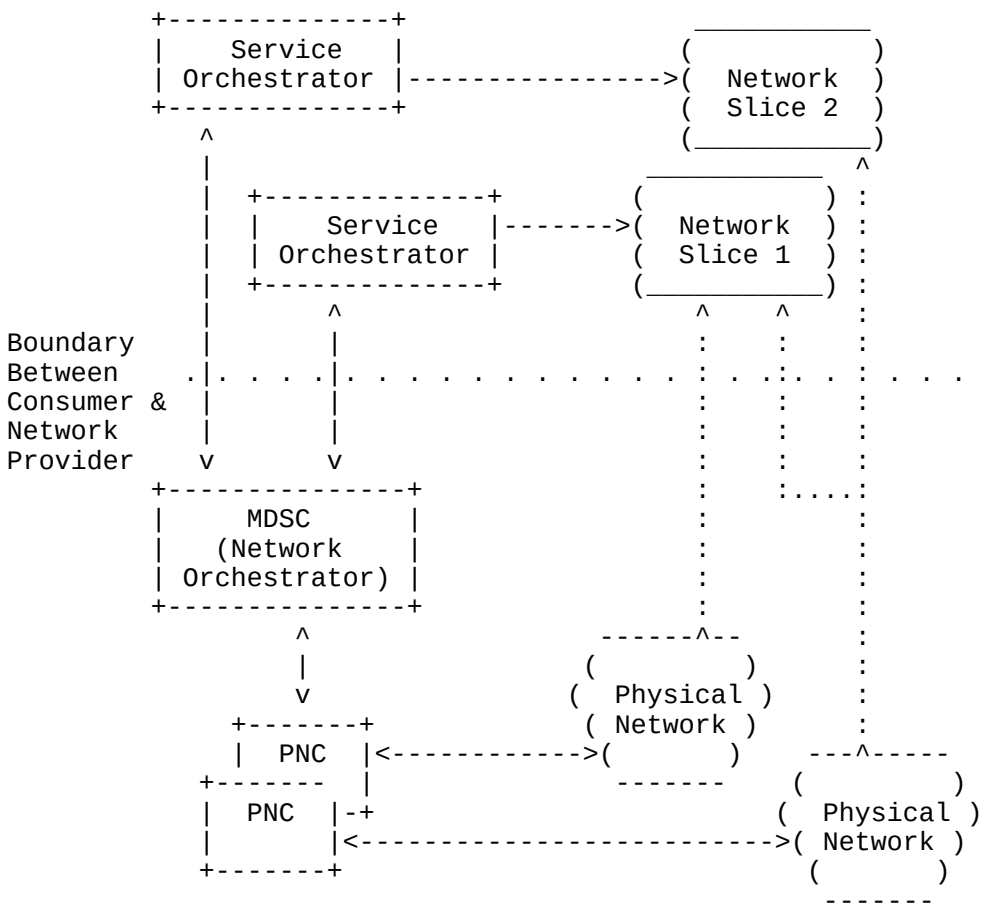

Figure 4: Network Slicing

4. YANG Models

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

The TE-service mapping model [I-D.ietf-teas-te-service-mapping-yang] creates a binding relationship across a L3VPN Service Model (L3SM) [RFC8299], L2VPN 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].

When necessary, it must be possible to map between a slice service request and an ACTN VN model. The ACTN VN model is a generic virtual network service model that allows customers to specify a VN that meets the customer's service objectives with various constraints, which could be included in the initial request, and how the service is delivered. Therefore, a request for a network slice service may be mapped directly to a request for a VN.

The TE-service mapping model [I-D.ietf-teas-te-service-mapping-yang] binds 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-l1csm-yang] L1CSM network service models.

Figure 5 shows the relationship between the YANG models discussed above.

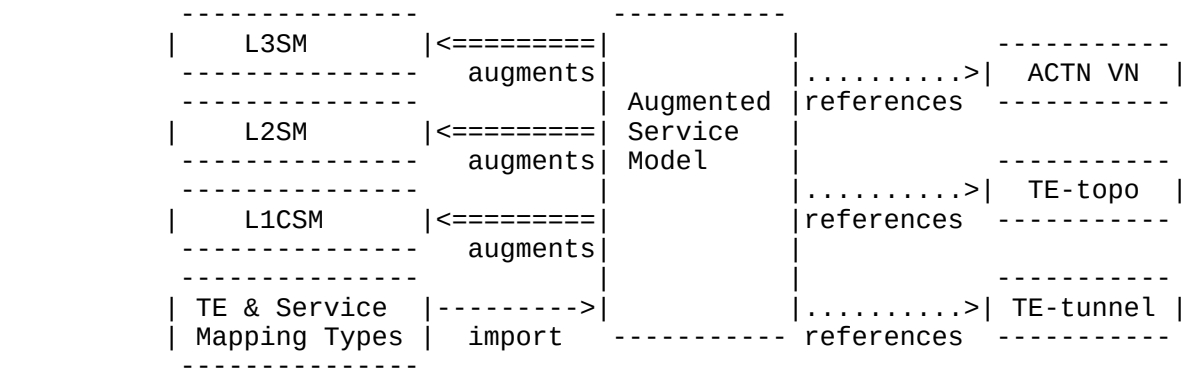

Figure 5: Relationships Between YANG Models

Work is still needed to define YANG models to help map network slice services to Traffic Engineering (TE) models. For example, [I-D.dhody-teas-ietf-network-slice-mapping] shows how the Virtual Network (VN) model and the TE Tunnel model can support network slice services.

4.2. Interfaces and YANG Models

Figure 6 shows the two ACTN components (MDSC and PNC) and one ACTN interface (MPI), as listed in Section 3. The figure also shows the Device Configuration Interface 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 when a control plane is used within the physical network. In the context of [RFC8309], the Device Configuration Interface uses one or more device configuration models.

Figure 6 also shows the Network Slice Service Interface. This interface allows a customer to make requests for delivery of the service, and it facilitates the customer modifying and monitoring the service. In the context of [RFC8309], this 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.ietf-teas-ietf-network-slice-nbi-yang] provides a suitable basis for requesting, controlling, and deletion, of a Network Slice Service.

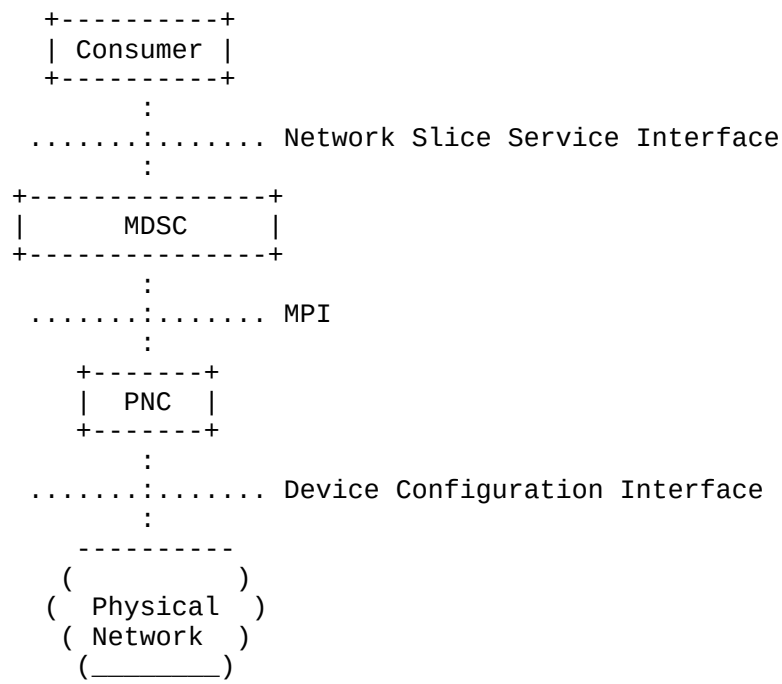

Figure 6: The YANG Interfaces in Context

4.3. ACTN VN Telemetry

The ACTN VN telemetry model [I-D.ietf-teas-actn-pm-telemetry-autonomics] provides a way for a customer to define performance monitoring relevant for the 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 the following:

* An ability to provide scalable VN-level telemetry aggregation based on a customer subscription model for key performance parameters defined by the customer.

* An ability to facilitate proactive re-optimization and reconfiguration of VNs/network slices based on autonomic network traffic engineering scaling configuration mechanisms.

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 customers. In some deployment models using ACTN, the customer may directly request a 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 customers. Furthermore, the resources allocated for or consumed by a customer will typically 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 customers expect and require that there is no risk of data leakage from one slice to another, and no transfer of knowledge of the structure or even existence of other slices. Further, in some service requests, there is an expectation that changes to one slice (under the control of one customer) should not have detrimental effects on the operation of other slices (whether under control of different or the same customers) even within limits allowed within the SLA. Thus, slices are assumed to be private and to provide the appearance of genuine physical connectivity.

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

Further security considerations and recommendations may be found in Section 9 of [RFC8453] and Section 10 of [RFC9543], with the latter document providing additional privacy considerations in Section 11.

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 Italo Busi, Qin Wu, Andy Jones, Ramon Casellas, Gert Grammel, Joe Clarke, Peter Yee, Alvaro Retana, テ詠ic Vyncke, Linda Dunbar and Kiran Makhijani for their reviews, insight, and useful discussions about network slicing.

This work is partially supported by the European Commission under Horizon 2020 grant agreement number 101015857 Secured autonomic traffic management for a Tera of SDN flows (Teraflow).

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: dceccare@cisco.com

9. Normative References

[I-D.ietf-teas-actn-pm-telemetry-autonomics]
Lee, Y., Dhody, D., Vilalta, R., King, D., and D. Ceccarelli, "YANG models for Virtual Network (VN)/TE Performance Monitoring Telemetry and Scaling Intent Autonomics", Work in Progress, Internet-Draft, draft-ietf-teas-actn-pm-telemetry-autonomics-12, 16 March 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-pm-telemetry-autonomics-12>.

[I-D.ietf-teas-actn-vn-yang]
Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., and B. Y. Yoon, "A YANG Data Model for Virtual Network (VN) Operations", Work in Progress, Internet-Draft, draft-ietf-teas-actn-vn-yang-29, 22 June 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-vn-yang-29>.

[I-D.ietf-teas-enhanced-vpn]
Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Network Resource Partition (NRP) based Enhanced Virtual Private Networks", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-20, 14 June 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-enhanced-vpn-20>.

[I-D.ietf-teas-ietf-network-slice-nbi-yang]
Wu, B., Dhody, D., Rokui, R., Saad, T., and J. Mullooly, "A YANG Data Model for the RFC 9543 Network Slice Service", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slice-nbi-yang-15, 27 August 2024, <https://datatracker.ietf.org/api/v1/doc/document/draft-ietf-teas-ietf-network-slice-nbi-yang/>.

[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 Data Model", Work in Progress, Internet-Draft, draft-ietf-teas-te-service-mapping-yang-15, 16 March 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-te-service-mapping-yang-15>.

[I-D.ietf-teas-yang-te]
Saad, T., Gandhi, R., Liu, X., Beeram, V. P., and I. Bryskin, "A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths and Interfaces", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-36, 2 February 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-36>.

[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>.

[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>.

[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>.

[RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <https://www.rfc-editor.org/info/rfc9543>.

10. Informative References

[I-D.dhody-teas-ietf-network-slice-mapping]
Dhody, D. and B. Wu, "IETF Network Slice Service Mapping YANG Model", Work in Progress, Internet-Draft, draft-dhody-teas-ietf-network-slice-mapping-05, 5 July 2024, <https://datatracker.ietf.org/doc/html/draft-dhody-teas-ietf-network-slice-mapping-05>.

[I-D.ietf-ccamp-l1csm-yang]
Lee, Y., Lee, K., Zheng, H., de Dios, O. G., and D. Ceccarelli, "A YANG Data Model for L1 Connectivity Service Model (L1CSM)", Work in Progress, Internet-Draft, draft-ietf-ccamp-l1csm-yang-26, 11 April 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-l1csm-yang-26>.

[I-D.ietf-teas-nrp-scalability]
Dong, J., Li, Z., Gong, L., Yang, G., and G. S. Mishra, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-05, 5 July 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-05>.

[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>.

[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>.

[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>.

[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>.

[RFC9182] Barguil, S., Gonzalez de Dios, O., Ed., Boucadair, M., Ed., Munoz, L., and A. Aguado, "A YANG Network Data Model for Layer 3 VPNs", RFC 9182, DOI 10.17487/RFC9182, February 2022, <https://www.rfc-editor.org/info/rfc9182>.

[RFC9291] Boucadair, M., Ed., Gonzalez de Dios, O., Ed., Barguil, S., and L. Munoz, "A YANG Network Data Model for Layer 2 VPNs", RFC 9291, DOI 10.17487/RFC9291, September 2022, <https://www.rfc-editor.org/info/rfc9291>.

[RFC9522] Farrel, A., Ed., "Overview and Principles of Internet Traffic Engineering", RFC 9522, DOI 10.17487/RFC9522, January 2024, <https://www.rfc-editor.org/info/rfc9522>.

Authors' Addresses

Daniel King
Old Dog Consulting
Email: daniel@olddog.co.uk

John Drake
Independent

Email: je_drake@yahoo.com

Haomian Zheng
Huawei Technologies
Email: zhenghaomian@huawei.com

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk