

Internet Engineering Task Force (IETF)
Request for Comments: 8453
Category: Informational
ISSN: 2070-1721

D. Ceccarelli, Ed.
Ericsson
Y. Lee, Ed.
Huawei
August 2018

Framework for Abstraction and Control of TE Networks (ACTN)

Abstract

Traffic Engineered (TE) networks have a variety of mechanisms to facilitate the separation of the data plane and control plane. They also have a range of management and provisioning protocols to configure and activate network resources. These mechanisms represent key technologies for enabling flexible and dynamic networking. The term "Traffic Engineered network" refers to a network that uses any connection-oriented technology under the control of a distributed or centralized control plane to support dynamic provisioning of end-to-end connectivity.

Abstraction of network resources is a technique that can be applied to a single network domain or across multiple domains to create a single virtualized network that is under the control of a network operator or the customer of the operator that actually owns the network resources.

This document provides a framework for Abstraction and Control of TE Networks (ACTN) to support virtual network services and connectivity services.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8453>.

Copyright Notice

Copyright (c) 2018 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
2. Overview	4
2.1. Terminology	5
2.2. VNS Model of ACTN	7
2.2.1. Customers	9
2.2.2. Service Providers	9
2.2.3. Network Operators	10
3. ACTN Base Architecture	10
3.1. Customer Network Controller	12
3.2. Multi-Domain Service Coordinator	13
3.3. Provisioning Network Controller	13
3.4. ACTN Interfaces	14
4. Advanced ACTN Architectures	15
4.1. MDSC Hierarchy	15
4.2. Functional Split of MDSC Functions in Orchestrators	16
5. Topology Abstraction Methods	18
5.1. Abstraction Factors	18
5.2. Abstraction Types	19
5.2.1. Native/White Topology	19
5.2.2. Black Topology	19
5.2.3. Grey Topology	20
5.3. Methods of Building Grey Topologies	21
5.3.1. Automatic Generation of Abstract Topology by Configuration	22
5.3.2. On-Demand Generation of Supplementary Topology via Path Compute Request/Reply	22
5.4. Hierarchical Topology Abstraction Example	23
5.5. VN Recursion with Network Layers	25
6. Access Points and Virtual Network Access Points	28
6.1. Dual-Homing Scenario	30

7.	Advanced ACTN Application: Multi-Destination Service	31
7.1.	Preplanned Endpoint Migration	32
7.2.	On-the-Fly Endpoint Migration	33
8.	Manageability Considerations	33
8.1.	Policy	34
8.2.	Policy Applied to the Customer Network Controller	34
8.3.	Policy Applied to the Multi-Domain Service Coordinator	35
8.4.	Policy Applied to the Provisioning Network Controller	35
9.	Security Considerations	36
9.1.	CNC-MDSC Interface (CMI)	37
9.2.	MDSC-PNC Interface (MPI)	37
10.	IANA Considerations	37
11.	Informative References	38
	Appendix A. Example of MDSC and PNC Functions Integrated in a Service/Network Orchestrator	40
	Contributors	41
	Authors' Addresses	42

1. Introduction

The term "Traffic Engineered network" refers to a network that uses any connection-oriented technology under the control of a distributed or centralized control plane to support dynamic provisioning of end-to-end connectivity. TE networks have a variety of mechanisms to facilitate the separation of data planes and control planes including distributed signaling for path setup and protection, centralized path computation for planning and traffic engineering, and a range of management and provisioning protocols to configure and activate network resources. These mechanisms represent key technologies for enabling flexible and dynamic networking. Some examples of networks that are in scope of this definition are optical, MPLS Transport Profile (MPLS-TP) [RFC5654], and MPLS-TE networks [RFC2702].

One of the main drivers for Software-Defined Networking (SDN) [RFC7149] is a decoupling of the network control plane from the data plane. This separation has been achieved for TE networks with the development of MPLS/GMPLS [RFC3945] and the Path Computation Element (PCE) [RFC4655]. One of the advantages of SDN is its logically centralized control regime that allows a global view of the underlying networks. Centralized control in SDN helps improve network resource utilization compared with distributed network control. For TE-based networks, a PCE may serve as a logically centralized path computation function.

This document describes a set of management and control functions used to operate one or more TE networks to construct virtual networks that can be presented to customers and that are built from abstractions of the underlying TE networks. For example, a link in

the customer's network is constructed from a path or collection of paths in the underlying networks. We call this set of functions "Abstraction and Control of TE Networks" or "ACTN".

2. Overview

Three key aspects that need to be solved by SDN are:

- o Separation of service requests from service delivery so that the configuration and operation of a network is transparent from the point of view of the customer but it remains responsive to the customer's services and business needs.
- o Network abstraction: As described in [RFC7926], abstraction is the process of applying policy to a set of information about a TE network to produce selective information that represents the potential ability to connect across the network. The process of abstraction presents the connectivity graph in a way that is independent of the underlying network technologies, capabilities, and topology so that the graph can be used to plan and deliver network services in a uniform way
- o Coordination of resources across multiple independent networks and multiple technology layers to provide end-to-end services regardless of whether or not the networks use SDN.

As networks evolve, the need to provide support for distinct services, separated service orchestration, and resource abstraction have emerged as key requirements for operators. In order to support multiple customers each with its own view of and control of the server network, a network operator needs to partition (or "slice") or manage sharing of the network resources. Network slices can be assigned to each customer for guaranteed usage, which is a step further than shared use of common network resources.

Furthermore, each network represented to a customer can be built from virtualization of the underlying networks so that, for example, a link in the customer's network is constructed from a path or collection of paths in the underlying network.

ACTN can facilitate virtual network operation via the creation of a single virtualized network or a seamless service. This supports operators in viewing and controlling different domains (at any dimension: applied technology, administrative zones, or vendor-specific technology islands) and presenting virtualized networks to their customers.

The ACTN framework described in this document facilitates:

- o Abstraction of the underlying network resources to higher-layer applications and customers [RFC7926].
- o Virtualization of particular underlying resources, whose selection criterion is the allocation of those resources to a particular customer, application, or service [ONF-ARCH].
- o TE Network slicing of infrastructure to meet specific customers' service requirements.
- o Creation of an abstract environment allowing operators to view and control multi-domain networks as a single abstract network.
- o The presentation to customers of networks as a virtual network via open and programmable interfaces.

2.1. Terminology

The following terms are used in this document. Some of them are newly defined, some others reference existing definitions:

Domain: A domain as defined by [RFC4655] is "any collection of network elements within a common sphere of address management or path computation responsibility". Specifically, within this document we mean a part of an operator's network that is under common management (i.e., under shared operational management using the same instances of a tool and the same policies). Network elements will often be grouped into domains based on technology types, vendor profiles, and geographic proximity.

Abstraction: This process is defined in [RFC7926].

TE Network Slicing: In the context of ACTN, a TE network slice is a collection of resources that is used to establish a logically dedicated virtual network over one or more TE networks. TE network slicing allows a network operator to provide dedicated virtual networks for applications/customers over a common network infrastructure. The logically dedicated resources are a part of the larger common network infrastructures that are shared among various TE network slice instances, which are the end-to-end realization of TE network slicing, consisting of the combination of physically or logically dedicated resources.

Node: A node is a vertex on the graph representation of a TE topology. In a physical network topology, a node corresponds to a physical network element (NE) such as a router. In an abstract network topology, a node (sometimes called an "abstract node") is a representation as a single vertex of one or more physical NEs and their connecting physical connections. The concept of a node represents the ability to connect from any access to the node (a link end) to any other access to that node, although "limited cross-connect capabilities" may also be defined to restrict this functionality. Network abstraction may be applied recursively, so a node in one topology may be created by applying abstraction to the nodes in the underlying topology.

Link: A link is an edge on the graph representation of a TE topology. Two nodes connected by a link are said to be "adjacent" in the TE topology. In a physical network topology, a link corresponds to a physical connection. In an abstract network topology, a link (sometimes called an "abstract link") is a representation of the potential to connect a pair of points with certain TE parameters (see [RFC7926] for details). Network abstraction may be applied recursively, so a link in one topology may be created by applying abstraction to the links in the underlying topology.

Abstract Topology: The topology of abstract nodes and abstract links presented through the process of abstraction by a lower-layer network for use by a higher-layer network.

Virtual Network (VN): A VN is a network provided by a service provider to a customer for the customer to use in any way it wants as though it was a physical network. There are two views of a VN as follows:

- o The VN can be abstracted as a set of edge-to-edge links (a Type 1 VN). Each link is referred as a "VN member" and is formed as an end-to-end tunnel across the underlying networks. Such tunnels may be constructed by recursive slicing or abstraction of paths in the underlying networks and can encompass edge points of the customer's network, access links, intra-domain paths, and inter-domain links.
- o The VN can also be abstracted as a topology of virtual nodes and virtual links (a Type 2 VN). The operator needs to map the VN to actual resource assignment, which is known as "virtual network embedding". The nodes in this case include physical endpoints, border nodes, and internal nodes as well as

abstracted nodes. Similarly, the links include physical access links, inter-domain links, and intra-domain links as well as abstract links.

Clearly, a Type 1 VN is a special case of a Type 2 VN.

Access link: A link between a customer node and an operator node.

Inter-domain link: A link between domains under distinct management administration.

Access Point (AP): An AP is a logical identifier shared between the customer and the operator used to identify an access link. The AP is used by the customer when requesting a Virtual Network Service (VNS). Note that the term "TE Link Termination Point" defined in [TE-TOPO] describes the endpoints of links, while an AP is a common identifier for the link itself.

VN Access Point (VNAP): A VNAP is the binding between an AP and a given VN.

Server Network: As defined in [RFC7926], a server network is a network that provides connectivity for another network (the Client Network) in a client-server relationship.

2.2. VNS Model of ACTN

A Virtual Network Service (VNS) is the service agreement between a customer and operator to provide a VN. When a VN is a simple connectivity between two points, the difference between VNS and connectivity service becomes blurred. There are three types of VNSs defined in this document.

- o Type 1 VNS refers to a VNS in which the customer is allowed to create and operate a Type 1 VN.
- o Type 2a and 2b VNS refer to VNSs in which the customer is allowed to create and operates a Type 2 VN. With a Type 2a VNS, the VN is statically created at service configuration time, and the customer is not allowed to change the topology (e.g., by adding or deleting abstract nodes and links). A Type 2b VNS is the same as a Type 2a VNS except that the customer is allowed to make dynamic changes to the initial topology created at service configuration time.

VN Operations are functions that a customer can exercise on a VN depending on the agreement between the customer and the operator.

- o VN Creation allows a customer to request the instantiation of a VN. This could be through offline preconfiguration or through dynamic requests specifying attributes to a Service Level Agreement (SLA) to satisfy the customer's objectives.
- o Dynamic Operations allow a customer to modify or delete the VN. The customer can further act upon the virtual network to create/modify/delete virtual links and nodes. These changes will result in subsequent tunnel management in the operator's networks.

There are three key entities in the ACTN VNS model:

- o Customers
- o Service Providers
- o Network Operators

These entities are related in a three tier model as shown in Figure 1.

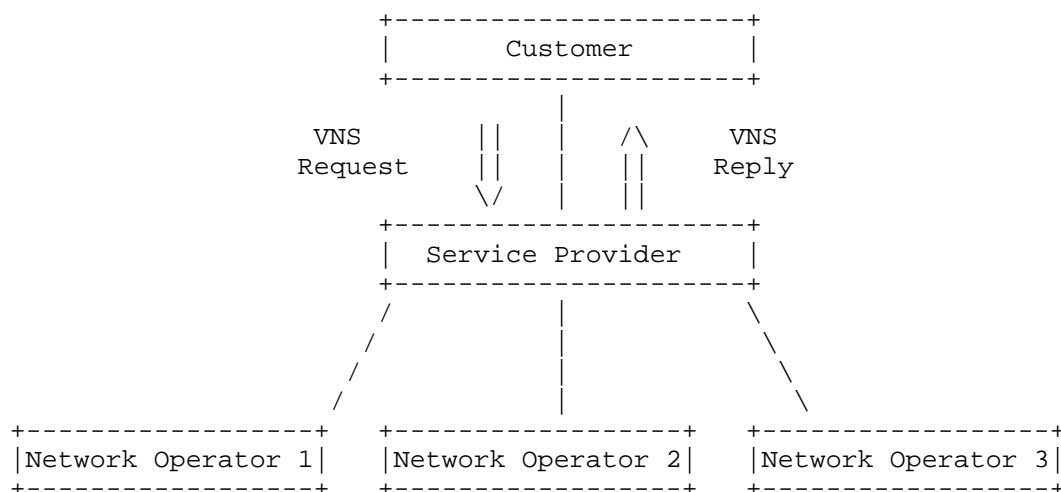


Figure 1: The Three-Tier Model

The commercial roles of these entities are described in the following sections.

2.2.1. Customers

Basic customers include fixed residential users, mobile users, and small enterprises. Each requires a small amount of resources and is characterized by steady requests (relatively time invariant). Basic customers do not modify their services themselves: if a service change is needed, it is performed by the provider as a proxy.

Advanced customers include enterprises and governments. Such customers ask for both point-to point and multipoint connectivity with high resource demands varying significantly in time. This is one of the reasons why a bundled service offering is not enough, and it is desirable to provide each advanced customer with a customized VNS. Advanced customers may also have the ability to modify their service parameters within the scope of their virtualized environments. The primary focus of ACTN is Advanced Customers.

As customers are geographically spread over multiple network operator domains, they have to interface to multiple operators and may have to support multiple virtual network services with different underlying objectives set by the network operators. To enable these customers to support flexible and dynamic applications, they need to control their allocated virtual network resources in a dynamic fashion; that means that they need a view of the topology that spans all of the network operators. Customers of a given service provider can, in turn, offer a service to other customers in a recursive way.

2.2.2. Service Providers

In the scope of ACTN, service providers deliver VNSs to their customers. Service providers may or may not own physical network resources (i.e., may or may not be network operators as described in Section 2.2.3). When a service provider is the same as the network operator, the case is similar to existing VPN models applied to a single operator (although it may be hard to use this approach when the customer spans multiple independent network operator domains).

When network operators supply only infrastructure, while distinct service providers interface with the customers, the service providers are themselves customers of the network infrastructure operators. One service provider may need to keep multiple independent network operators because its end users span geographically across multiple network operator domains. In some cases, a service provider is also a network operator when it owns network infrastructure on which service is provided.

2.2.3. Network Operators

Network operators are the infrastructure operators that provision the network resources and provide network resources to their customers. The layered model described in this architecture separates the concerns of network operators and customers, with service providers acting as aggregators of customer requests.

3. ACTN Base Architecture

This section provides a high-level model of ACTN, showing the interfaces and the flow of control between components.

The ACTN architecture is based on a three-tier reference model and allows for hierarchy and recursion. The main functionalities within an ACTN system are:

- o Multi-domain coordination: This function oversees the specific aspects of different domains and builds a single abstracted end-to-end network topology in order to coordinate end-to-end path computation and path/service provisioning. Domain sequence path calculation/determination is also a part of this function.
- o Abstraction: This function provides an abstracted view of the underlying network resources for use by the customer -- a customer may be the client or a higher-level controller entity. This function includes network path computation based on customer-service-connectivity request constraints, path computation based on the global network-wide abstracted topology, and the creation of an abstracted view of network resources allocated to each customer. These operations depend on customer-specific network objective functions and customer traffic profiles.
- o Customer mapping/translation: This function is to map customer requests/commands into network provisioning requests that can be sent from the Multi-Domain Service Coordinator (MDSC) to the Provisioning Network Controller (PNC) according to business policies provisioned statically or dynamically at the Operations Support System (OSS) / Network Management System (NMS). Specifically, it provides mapping and translation of a customer's service request into a set of parameters that are specific to a network type and technology such that network configuration process is made possible.
- o Virtual service coordination: This function translates information that is customer service related into virtual network service operations in order to seamlessly operate virtual networks while meeting a customer's service requirements. In the context of

ACTN, service/virtual service coordination includes a number of service orchestration functions such as multi-destination load-balancing and guarantees of service quality. It also includes notifications for service fault and performance degradation and so forth.

The base ACTN architecture defines three controller types and the corresponding interfaces between these controllers. The following types of controller are shown in Figure 2:

- o CNC - Customer Network Controller
- o MDSC - Multi-Domain Service Coordinator
- o PNC - Provisioning Network Controller

Figure 2 also shows the following interfaces

- o CMI - CNC-MDSC Interface
- o MPI - MDSC-PNC Interface
- o SBI - Southbound Interface

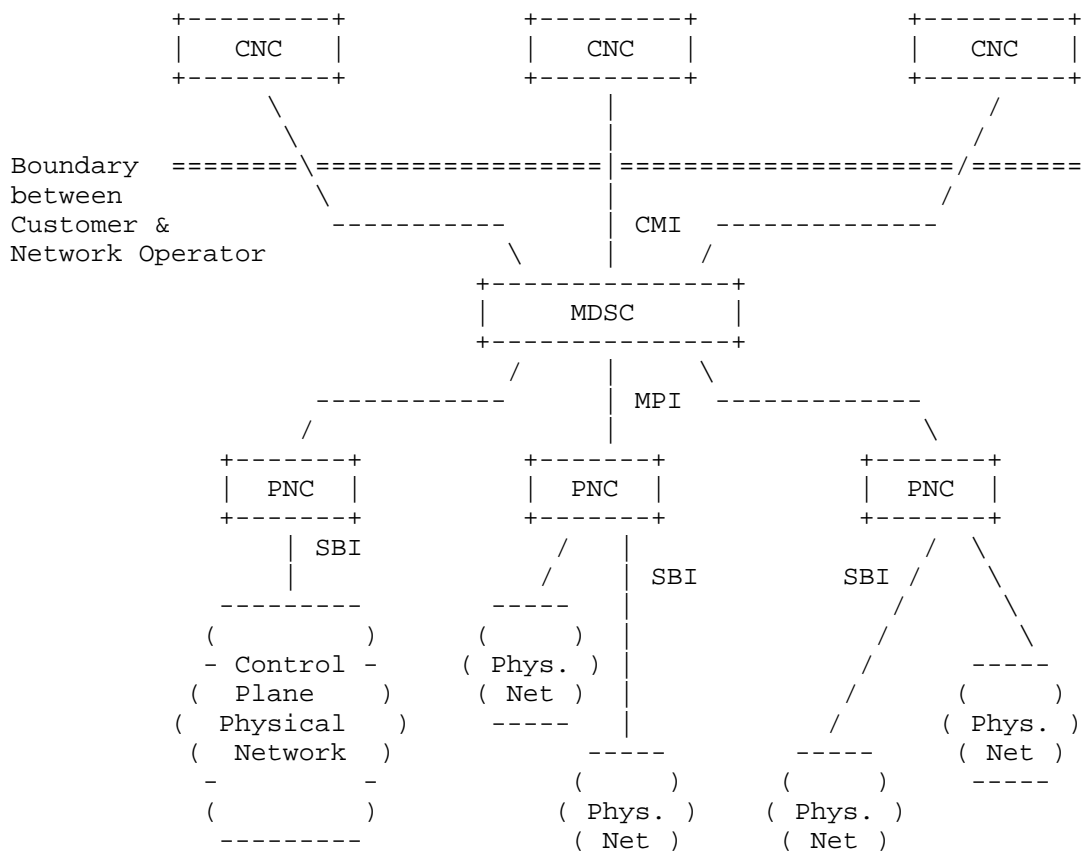


Figure 2: ACTN Base Architecture

Note that this is a functional architecture: an implementation and deployment might collocate one or more of the functional components. Figure 2 shows a case where the service provider is also a network operator.

3.1. Customer Network Controller

A Customer Network Controller (CNC) is responsible for communicating a customer's VNS requirements to the network operator over the CNC-MDSC Interface (CMI). It has knowledge of the endpoints associated with the VNS (expressed as APs), the service policy, and other QoS information related to the service.

As the CNC directly interfaces with the applications, it understands multiple application requirements and their service needs. The capability of a CNC beyond its CMI role is outside the scope of ACTN and may be implemented in different ways. For example, the CNC may, in fact, be a controller or part of a controller in the customer's domain, or the CNC functionality could also be implemented as part of a service provider's portal.

3.2. Multi-Domain Service Coordinator

A Multi-Domain Service Coordinator (MDSC) is a functional block that implements all of the ACTN functions listed in Section 3 and described further in Section 4.2. Two functions of the MDSC, namely, multi-domain coordination and virtualization/abstraction are referred to as network-related functions; whereas the other two functions, namely, customer mapping/translation and virtual service coordination, are referred to as service-related functions. The MDSC sits at the center of the ACTN model between the CNC that issues connectivity requests and the Provisioning Network Controllers (PNCs) that manage the network resources. The key point of the MDSC (and of the whole ACTN framework) is detaching the network and service control from underlying technology to help the customer express the network as desired by business needs. The MDSC envelopes the instantiation of the right technology and network control to meet business criteria. In essence, it controls and manages the primitives to achieve functionalities as desired by the CNC.

In order to allow for multi-domain coordination, a 1:N relationship must be allowed between MDSCs and PNCs.

In addition to that, it could also be possible to have an M:1 relationship between MDSCs and PNCs to allow for network-resource partitioning/sharing among different customers that are not necessarily connected to the same MDSC (e.g., different service providers) but that are all using the resources of a common network infrastructure operator.

3.3. Provisioning Network Controller

The Provisioning Network Controller (PNC) oversees configuring the network elements, monitoring the topology (physical or virtual) of the network, and collecting information about the topology (either raw or abstracted).

The PNC functions can be implemented as part of an SDN domain controller, a Network Management System (NMS), an Element Management System (EMS), an active PCE-based controller [RFC8283], or any other means to dynamically control a set of nodes that implements a

northbound interface from the standpoint of the nodes (which is out of the scope of this document). A PNC domain includes all the resources under the control of a single PNC. It can be composed of different routing domains and administrative domains, and the resources may come from different layers. The interconnection between PNC domains is illustrated in Figure 3.




Figure 3: PNC Domain Borders

3.4. ACTN Interfaces

Direct customer control of transport network elements and virtualized services is not a viable proposition for network operators due to security and policy concerns. Therefore, the network has to provide open, programmable interfaces, through which customer applications can create, replace, and modify virtual network resources and services in an interactive, flexible, and dynamic fashion.

Three interfaces exist in the ACTN architecture as shown in Figure 2.

- o CMI: The CNC-MDSC Interface (CMI) is an interface between a CNC and an MDSC. The CMI is a business boundary between customer and network operator. It is used to request a VNS for an application. All service-related information is conveyed over this interface (such as the VNS type, topology, bandwidth, and service constraints). Most of the information over this interface is agnostic of the technology used by network operators, but there are some cases (e.g., access link configuration) where it is necessary to specify technology-specific details.
- o MPI: The MDSC-PNC Interface (MPI) is an interface between an MDSC and a PNC. It communicates requests for new connectivity or for bandwidth changes in the physical network. In multi-domain environments, the MDSC needs to communicate with multiple PNCs,

each responsible for control of a domain. The MPI presents an abstracted topology to the MDSC hiding technology-specific aspects of the network and hiding topology according to policy.

- o SBI: The Southbound Interface (SBI) is out of scope of ACTN. Many different SBIs have been defined for different environments, technologies, standards organizations, and vendors. It is shown in Figure 3 for reference reason only.

4. Advanced ACTN Architectures

This section describes advanced configurations of the ACTN architecture.

4.1. MDSC Hierarchy

A hierarchy of MDSCs can be foreseen for many reasons, among which are scalability, administrative choices, or putting together different layers and technologies in the network. In the case where there is a hierarchy of MDSCs, we introduce the terms "higher-level MDSC" (MDSC-H) and "lower-level MDSC" (MDSC-L). The interface between them is a recursion of the MPI. An implementation of an MDSC-H makes provisioning requests as normal using the MPI, but an MDSC-L must be able to receive requests as normal at the CMI and also at the MPI. The hierarchy of MDSCs can be seen in Figure 4.

Another implementation choice could foresee the usage of an MDSC-L for all the PNCs related to a given technology (e.g., Internet Protocol (IP) / Multiprotocol Label Switching (MPLS)) and a different MDSC-L for the PNCs related to another technology (e.g., Optical Transport Network (OTN) / Wavelength Division Multiplexing (WDM)) and an MDSC-H to coordinate them.

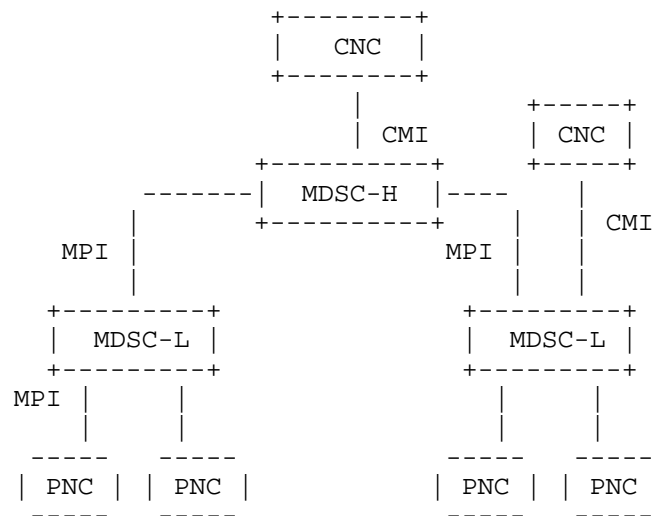


Figure 4: MDSC Hierarchy

The hierarchy of MDSC can be recursive, where an MDSC-H is, in turn, an MDSC-L to a higher-level MDSC-H.

4.2. Functional Split of MDSC Functions in Orchestrators

An implementation choice could separate the MDSC functions into two groups: one group for service-related functions and the other for network-related functions. This enables the implementation of a service orchestrator that provides the service-related functions of the MDSC and a network orchestrator that provides the network-related functions of the MDSC. This split is consistent with the YANG service model architecture described in [RFC8309]. Figure 5 depicts this and shows how the ACTN interfaces may map to YANG data models.

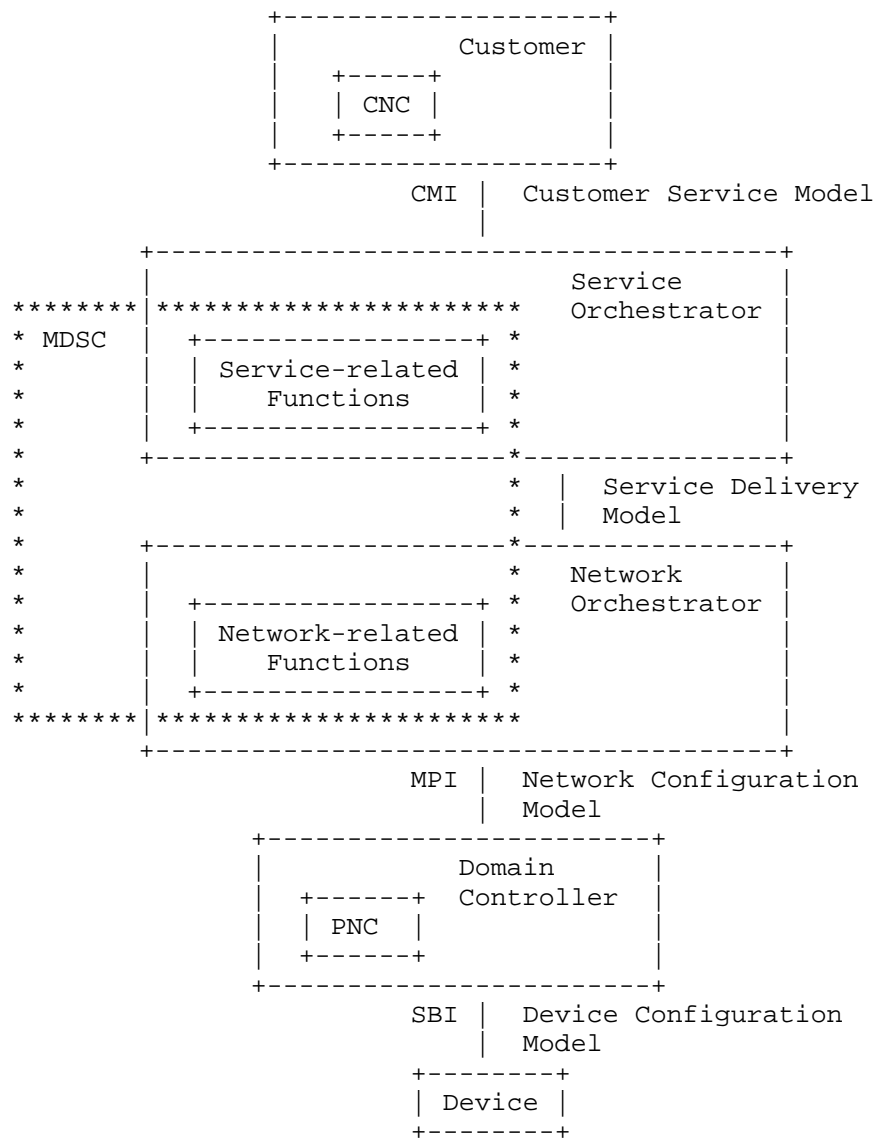


Figure 5: ACTN Architecture in the Context of the YANG Service Models

5. Topology Abstraction Methods

Topology abstraction is described in [RFC7926]. This section discusses topology abstraction factors, types, and their context in the ACTN architecture.

Abstraction in ACTN is performed by the PNC when presenting available topology to the MDSC, or by an MDSC-L when presenting topology to an MDSC-H. This function is different from the creation of a VN (and particularly a Type 2 VN) that is not abstraction but construction of virtual resources.

5.1. Abstraction Factors

As discussed in [RFC7926], abstraction is tied with the policy of the networks. For instance, per an operational policy, the PNC would not provide any technology-specific details (e.g., optical parameters for Wavelength Switched Optical Network (WSO) in the abstract topology it provides to the MDSC. Similarly, the policy of the networks may determine the abstraction type as described in Section 5.2.

There are many factors that may impact the choice of abstraction:

- o Abstraction depends on the nature of the underlying domain networks. For instance, packet networks may be abstracted with fine granularity while abstraction of optical networks depends on the switching units (such as wavelengths) and the end-to-end continuity and cross-connect limitations within the network.
- o Abstraction also depends on the capability of the PNCs. As abstraction requires hiding details of the underlying network resources, the PNC's capability to run algorithms impacts the feasibility of abstraction. Some PNCs may not have the ability to abstract native topology while other PNCs may have the ability to use sophisticated algorithms.
- o Abstraction is a tool that can improve scalability. Where the native network resource information is of a large size, there is a specific scaling benefit to abstraction.
- o The proper abstraction level may depend on the frequency of topology updates and vice versa.
- o The nature of the MDSC's support for technology-specific parameters impacts the degree/level of abstraction. If the MDSC is not capable of handling such parameters, then a higher level of abstraction is needed.

- o In some cases, the PNC is required to hide key internal topological data from the MDSC. Such confidentiality can be achieved through abstraction.

5.2. Abstraction Types

This section defines the following three types of topology abstraction:

- o Native/White Topology (Section 5.2.1)
- o Black Topology (Section 5.2.2)
- o Grey Topology (Section 5.2.3)

5.2.1. Native/White Topology

This is a case where the PNC provides the actual network topology to the MDSC without any hiding or filtering of information, i.e., no abstraction is performed. In this case, the MDSC has the full knowledge of the underlying network topology and can operate on it directly.

5.2.2. Black Topology

A black topology replaces a full network with a minimal representation of the edge-to-edge topology without disclosing any node internal connectivity information. The entire domain network may be abstracted as a single abstract node with the network's access/egress links appearing as the ports to the abstract node and the implication that any port can be "cross-connected" to any other. Figure 6 depicts a native topology with the corresponding black topology with one virtual node and inter-domain links. In this case, the MDSC has to make a provisioning request to the PNCs to establish the port-to-port connection. If there is a large number of interconnected domains, this abstraction method may impose a heavy coordination load at the MDSC level in order to find an optimal end-to-end path since the abstraction hides so much information that it is not possible to determine whether an end-to-end path is feasible without asking each PNC to set up each path fragment. For this reason, the MPI might need to be enhanced to allow the PNCs to be queried for the practicality and characteristics of paths across the abstract node.




Figure 6: Native Topology with Corresponding Black Topology Expressed as an Abstract Node

5.2.3. Grey Topology

A grey topology represents a compromise between black and white topologies from a granularity point of view. In this case, the PNC exposes an abstract topology containing all PNC domain border nodes and an abstraction of the connectivity between those border nodes. This abstraction may contain either physical or abstract nodes/links.

Two types of grey topology are identified:

- o In a type A grey topology, border nodes are connected by a full mesh of TE links (see Figure 7).
- o In a type B grey topology, border nodes are connected over a more-detailed network comprising internal abstract nodes and abstracted links. This mode of abstraction supplies the MDSC with more information about the internals of the PNC domain and allows it to make more informed choices about how to route connectivity over the underlying network.




Figure 7: Native Topology with Corresponding Grey Topology

5.3. Methods of Building Grey Topologies

This section discusses two different methods of building a grey topology:

- o Automatic generation of abstract topology by configuration (Section 5.3.1)
- o On-demand generation of supplementary topology via path computation request/reply (Section 5.3.2)

5.3.1. Automatic Generation of Abstract Topology by Configuration

Automatic generation is based on the abstraction/summarization of the whole domain by the PNC and its advertisement on the MPI. The level of abstraction can be decided based on PNC configuration parameters (e.g., "provide the potential connectivity between any PE and any ASBR in an MPLS-TE network").

Note that the configuration parameters for this abstract topology can include available bandwidth, latency, or any combination of defined parameters. How to generate such information is beyond the scope of this document.

This abstract topology may need to be periodically or incrementally updated when there is a change in the underlying network or the use of the network resources that make connectivity more or less available.

5.3.2. On-Demand Generation of Supplementary Topology via Path Compute Request/Reply

While abstract topology is generated and updated automatically by configuration as explained in Section 5.3.1, additional supplementary topology may be obtained by the MDSC via a path compute request/reply mechanism.

The abstract topology advertisements from PNCs give the MDSC the border node/link information for each domain. Under this scenario, when the MDSC needs to create a new VN, the MDSC can issue path computation requests to PNCs with constraints matching the VN request as described in [ACTN-YANG]. An example is provided in Figure 8, where the MDSC is creating a P2P VN between AP1 and AP2. The MDSC could use two different inter-domain links to get from domain X to domain Y, but in order to choose the best end-to-end path, it needs to know what domain X and Y can offer in terms of connectivity and constraints between the PE nodes and the border nodes.





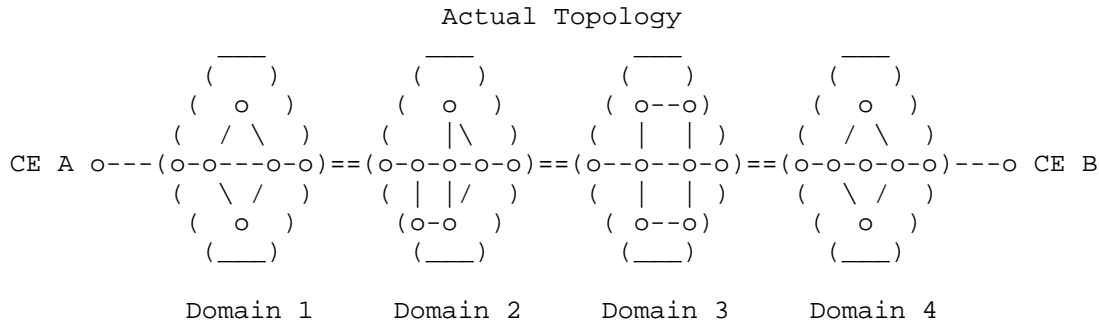
Figure 8: A Multi-Domain Example

The MDSC issues a path computation request to PNC.X asking for potential connectivity between PE1 and border node BrdrX.1 and between PE1 and BrdrX.2 with related objective functions and TE metric constraints. A similar request for connectivity from the border nodes in domain Y to PE2 will be issued to PNC.Y. The MDSC merges the results to compute the optimal end-to-end path including the inter-domain links. The MDSC can use the result of this computation to request the PNCs to provision the underlying networks, and the MDSC can then use the end-to-end path as a virtual link in the VN it delivers to the customer.

5.4. Hierarchical Topology Abstraction Example

This section illustrates how topology abstraction operates in different levels of a hierarchy of MDSCs as shown in Figure 9.





Where

- o is a node
- is a link
- === is a border link

Figure 9: Illustration of Hierarchical Topology Abstraction

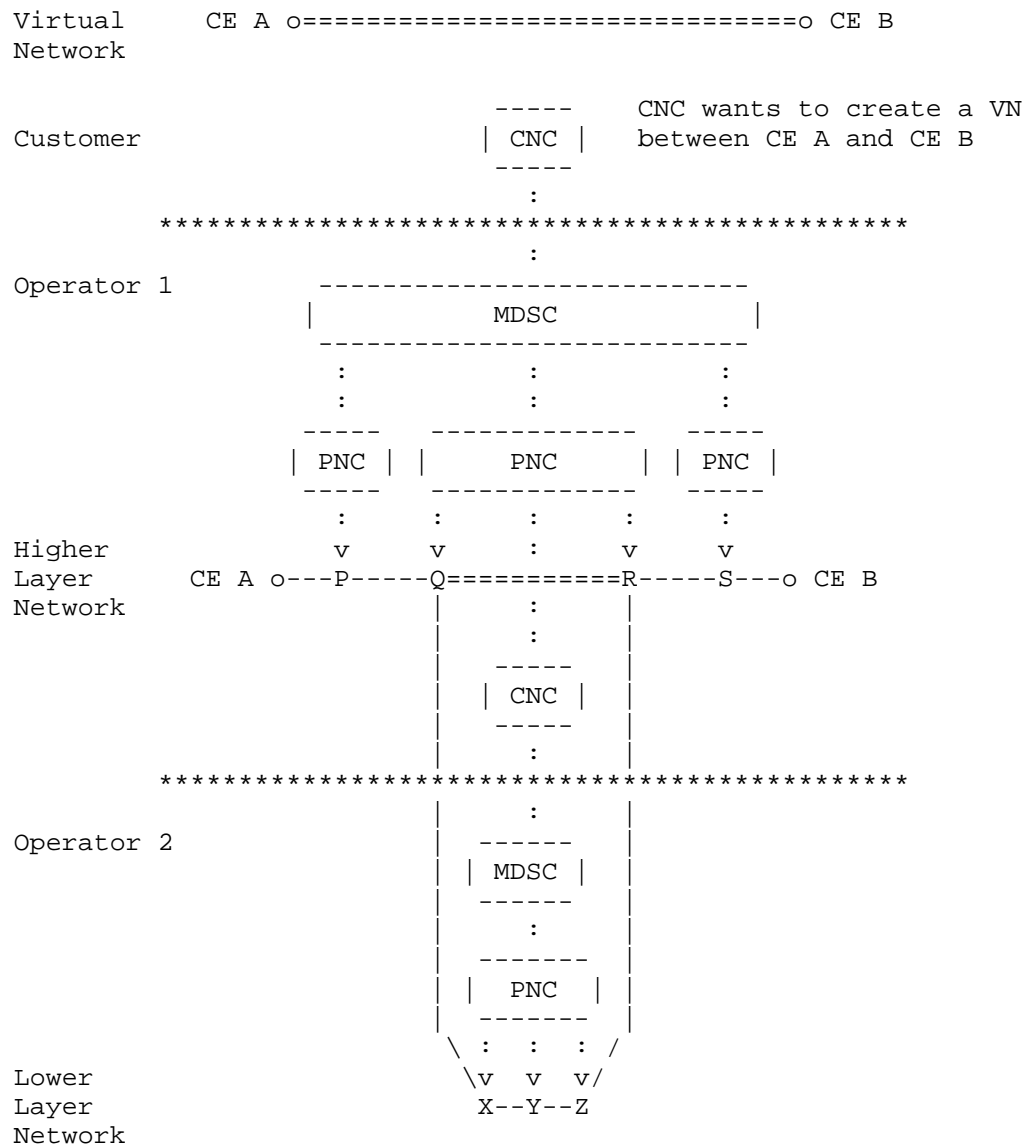
In the example depicted in Figure 9, there are four domains under control of PNCs: PNC1, PNC2, PNC3, and PNC4. MDSC-L1 controls PNC1 and PNC2, while MDSC-L2 controls PNC3 and PNC4. Each of the PNCs provides a grey topology abstraction that presents only border nodes and links across and outside the domain. The abstract topology MDSC-L1 that operates is a combination of the two topologies from PNC1 and PNC2. Likewise, the abstract topology that MDSC-L2 operates is shown in Figure 9. Both MDSC-L1 and MDSC-L2 provide a black topology abstraction to MDSC-H in which each PNC domain is presented as a single virtual node. MDSC-H combines these two topologies to create the abstraction topology on which it operates. MDSC-H sees the whole four domain networks as four virtual nodes connected via virtual links.

5.5. VN Recursion with Network Layers

In some cases, the VN supplied to a customer may be built using resources from different technology layers operated by different operators. For example, one operator may run a packet TE network and use optical connectivity provided by another operator.

As shown in Figure 10, a customer asks for end-to-end connectivity between CE A and CE B, a virtual network. The customer's CNC makes a request to Operator 1's MDSC. The MDSC works out which network resources need to be configured and sends instructions to the appropriate PNCs. However, the link between Q and R is a virtual link supplied by Operator 2: Operator 1 is a customer of Operator 2.

To support this, Operator 1 has a CNC that communicates with Operator 2's MDSC. Note that Operator 1's CNC in Figure 10 is a functional component that does not dictate implementation: it may be embedded in a PNC.



Where

--- is a link

=== is a virtual link

Figure 10: VN Recursion with Network Layers

6. Access Points and Virtual Network Access Points

In order to map identification of connections between the customer's sites and the TE networks and to scope the connectivity requested in the VNS, the CNC and the MDSC refer to the connections using the Access Point (AP) construct as shown in Figure 11.




Figure 11: Customer View of APs

Let's take as an example a scenario shown in Figure 11. CE1 is connected to the network via a 10 Gbps link and CE2 via a 40 Gbps link. Before the creation of any VN between AP1 and AP2, the customer view can be summarized as shown in Figure 12.

+-----+-----+-----+-----+			
	Endpoint	Access Link Bandwidth	
+-----+-----+-----+-----+			
AP id	CE,port	MaxResBw	AvailableBw
+-----+-----+-----+-----+			
AP1	CE1,portX	10 Gbps	10 Gbps
+-----+-----+-----+-----+			
AP2	CE2,portZ	40 Gbps	40 Gbps
+-----+-----+-----+-----+			

Figure 12: AP - Customer View

On the other hand, what the operator sees is shown in Figure 13




Figure 13: Operator View of the AP

which results in a summarization as shown in Figure 14.

+-----+-----+-----+-----+			
	Endpoint	Access Link Bandwidth	
+-----+	+-----+	+-----+	+-----+
AP id	PE,port	MaxResBw	AvailableBw
+-----+	+-----+	+-----+	+-----+
AP1	PE1,portW	10 Gbps	10 Gbps
+-----+	+-----+	+-----+	+-----+
AP2	PE2,portY	40 Gbps	40 Gbps
+-----+	+-----+	+-----+	+-----+

Figure 14: AP - Operator View

A Virtual Network Access Point (VNAP) needs to be defined as binding between an AP and a VN. It is used to allow different VNs to start from the same AP. It also allows for traffic engineering on the access and/or inter-domain links (e.g., keeping track of bandwidth allocation). A different VNAP is created on an AP for each VN.

In this simple scenario, we suppose we want to create two virtual networks: the first with VN identifier 9 between AP1 and AP2 with bandwidth of 1 Gbps and the second with VN identifier 5, again between AP1 and AP2 and with bandwidth 2 Gbps.

The operator view would evolve as shown in Figure 15.

+-----+-----+		+-----+-----+	
Endpoint		Access Link/VNAP Bw	
+-----+-----+			
AP/VNAPid	PE,port	MaxResBw	AvailableBw
+-----+-----+			
AP1	PE1,portW	10 Gbps	7 Gbps
-VNAP1.9		1 Gbps	N.A.
-VNAP1.5		2 Gbps	N.A
+-----+-----+			
AP2	PE2,portY	40Gbps	37 Gbps
-VNAP2.9		1 Gbps	N.A.
-VNAP2.5		2 Gbps	N.A
+-----+-----+			

Figure 15: AP and VNAP - Operator View after VNS Creation

6.1. Dual-Homing Scenario

Often there is a dual-homing relationship between a CE and a pair of PEs. This case needs to be supported by the definition of VN, APs, and VNAPs. Suppose CE1 connected to two different PEs in the operator domain via AP1 and AP2 and that the customer needs 5 Gbps of bandwidth between CE1 and CE2. This is shown in Figure 16.




Figure 16: Dual-Homing Scenario

In this case, the customer will request a VN between AP1, AP2, and AP3 specifying a dual-homing relationship between AP1 and AP2. As a consequence, no traffic will flow between AP1 and AP2. The dual-homing relationship would then be mapped against the VNAPs (since other independent VNs might have AP1 and AP2 as endpoints).

The customer view would be shown in Figure 17.

+-----+-----+-----+-----+-----+				
	Endpoint	Access Link/VNAP Bw		
+-----+-----+-----+-----+-----+				
AP/VNAPid	CE,port	MaxResBw	AvailableBw	Dual Homing
+-----+-----+-----+-----+-----+				
AP1	CE1,portW	10 Gbps	5 Gbps	
-VNAP1.9		5 Gbps	N.A.	VNAP2.9
+-----+-----+-----+-----+-----+				
AP2	CE1,portY	40 Gbps	35 Gbps	
-VNAP2.9		5 Gbps	N.A.	VNAP1.9
+-----+-----+-----+-----+-----+				
AP3	CE2,portX	50 Gbps	45 Gbps	
-VNAP3.9		5 Gbps	N.A.	NONE
+-----+-----+-----+-----+-----+				

Figure 17: Dual-Homing -- Customer View after VN Creation

7. Advanced ACTN Application: Multi-Destination Service

A more-advanced application of ACTN is the case of data center (DC) selection, where the customer requires the DC selection to be based on the network status; this is referred to as "Multi-Destination Service" in [ACTN-REQ]. In terms of ACTN, a CNC could request a VNS between a set of source APs and destination APs and leave it up to the network (MDSC) to decide which source and destination APs to be used to set up the VNS. The candidate list of source and destination APs is decided by a CNC (or an entity outside of ACTN) based on certain factors that are outside the scope of ACTN.

Based on the AP selection as determined and returned by the network (MDSC), the CNC (or an entity outside of ACTN) should further take care of any subsequent actions such as orchestration or service setup requirements. These further actions are outside the scope of ACTN.

Consider a case as shown in Figure 18, where three DCs are available, but the customer requires the DC selection to be based on the network status and the connectivity service setup between the AP1 (CE1) and one of the destination APs (AP2 (DC-A), AP3 (DC-B), and AP4 (DC-C)). The MDSC (in coordination with PNCs) would select the best destination AP based on the constraints, optimization criteria, policies, etc., and set up the connectivity service (virtual network).

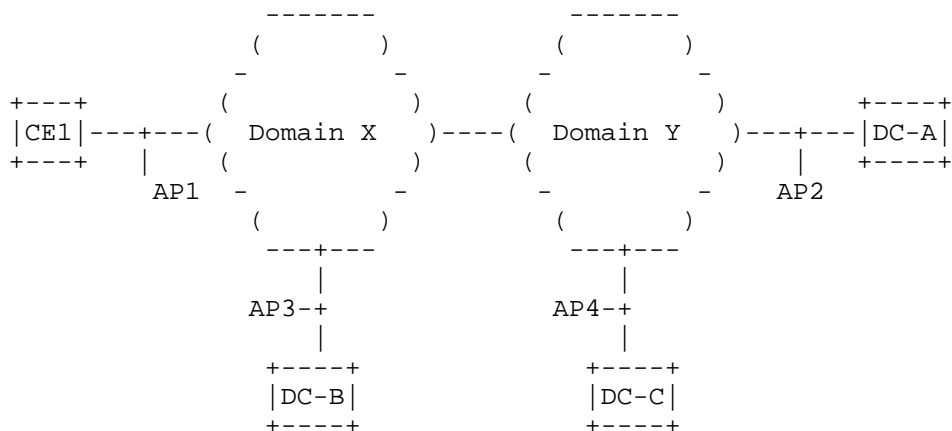


Figure 18: Endpoint Selection Based on Network Status

7.1. Preplanned Endpoint Migration

Furthermore, in the case of DC selection, a customer could request a backup DC to be selected, such that in case of failure, another DC site could provide hot stand-by protection. As shown in Figure 19, DC-C is selected as a backup for DC-A. Thus, the VN should be set up by the MDSC to include primary connectivity between AP1 (CE1) and AP2 (DC-A) as well as protection connectivity between AP1 (CE1) and AP4 (DC-C).

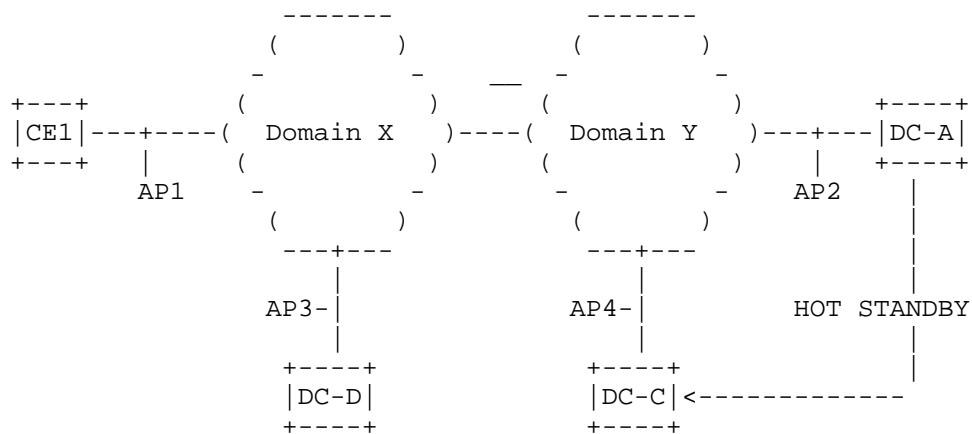


Figure 19: Preplanned Endpoint Migration

7.2. On-the-Fly Endpoint Migration

Compared to preplanned endpoint migration, on-the-fly endpoint selection is dynamic in that the migration is not preplanned but decided based on network condition. Under this scenario, the MDSC would monitor the network (based on the VN SLA) and notify the CNC in the case where some other destination AP would be a better choice based on the network parameters. The CNC should instruct the MDSC when it is suitable to update the VN with the new AP if it is required.

8. Manageability Considerations

The objective of ACTN is to manage traffic engineered resources and provide a set of mechanisms to allow customers to request virtual connectivity across server-network resources. ACTN supports multiple customers, each with its own view of and control of a virtual network built on the server network; the network operator will need to partition (or "slice") their network resources, and manage the resources accordingly.

The ACTN platform will, itself, need to support the request, response, and reservations of client- and network-layer connectivity. It will also need to provide performance monitoring and control of TE resources. The management requirements may be categorized as follows:

- o Management of external ACTN protocols
- o Management of internal ACTN interfaces/protocols
- o Management and monitoring of ACTN components
- o Configuration of policy to be applied across the ACTN system

The ACTN framework and interfaces are defined to enable traffic engineering for virtual network services and connectivity services. Network operators may have other Operations, Administration, and Maintenance (OAM) tasks for service fulfillment, optimization, and assurance beyond traffic engineering. The realization of OAM beyond abstraction and control of TE networks is not discussed in this document.

8.1. Policy

Policy is an important aspect of ACTN control and management. Policies are used via the components and interfaces, during deployment of the service, to ensure that the service is compliant with agreed-upon policy factors and variations (often described in SLAs); these include, but are not limited to connectivity, bandwidth, geographical transit, technology selection, security, resilience, and economic cost.

Depending on the deployment of the ACTN architecture, some policies may have local or global significance. That is, certain policies may be ACTN component specific in scope, while others may have broader scope and interact with multiple ACTN components. Two examples are provided below:

- o A local policy might limit the number, type, size, and scheduling of virtual network services a customer may request via its CNC. This type of policy would be implemented locally on the MDSC.
- o A global policy might constrain certain customer types (or specific customer applications) only to use certain MDSCs and be restricted to physical network types managed by the PNCs. A global policy agent would govern these types of policies.

The objective of this section is to discuss the applicability of ACTN policy: requirements, components, interfaces, and examples. This section provides an analysis and does not mandate a specific method for enforcing policy, or the type of policy agent that would be responsible for propagating policies across the ACTN components. It does highlight examples of how policy may be applied in the context of ACTN, but it is expected further discussion in an applicability or solution-specific document, will be required.

8.2. Policy Applied to the Customer Network Controller

A virtual network service for a customer application will be requested by the CNC. The request will reflect the application requirements and specific service needs, including bandwidth, traffic type and survivability. Furthermore, application access and type of virtual network service requested by the CNC, will be need adhere to specific access control policies.

8.3. Policy Applied to the Multi-Domain Service Coordinator

A key objective of the MDSC is to support the customer's expression of the application connectivity request via its CNC as a set of desired business needs; therefore, policy will play an important role.

Once authorized, the virtual network service will be instantiated via the CNC-MDSC Interface (CMI); it will reflect the customer application and connectivity requirements and specific service-transport needs. The CNC and the MDSC components will have agreed-upon connectivity endpoints; use of these endpoints should be defined as a policy expression when setting up or augmenting virtual network services. Ensuring that permissible endpoints are defined for CNCs and applications will require the MDSC to maintain a registry of permissible connection points for CNCs and application types.

Conflicts may occur when virtual network service optimization criteria are in competition. For example, to meet objectives for service reachability, a request may require an interconnection point between multiple physical networks; however, this might break a confidentiality policy requirement of a specific type of end-to-end service. Thus, an MDSC may have to balance a number of the constraints on a service request and between different requested services. It may also have to balance requested services with operational norms for the underlying physical networks. This balancing may be resolved using configured policy and using hard and soft policy constraints.

8.4. Policy Applied to the Provisioning Network Controller

The PNC is responsible for configuring the network elements, monitoring physical network resources, and exposing connectivity (direct or abstracted) to the MDSC. Therefore, it is expected that policy will dictate what connectivity information will be exchanged on the MPI.

Policy interactions may arise when a PNC determines that it cannot compute a requested path from the MDSC, or notices that (per a locally configured policy) the network is low on resources (for example, the capacity on key links became exhausted). In either case, the PNC will be required to notify the MDSC, which may (again per policy) act to construct a virtual network service across another physical network topology.

Furthermore, additional forms of policy-based resource management will be required to provide VNS performance, security, and resilience guarantees. This will likely be implemented via a local policy agent and additional protocol methods.

9. Security Considerations

The ACTN framework described in this document defines key components and interfaces for managed TE networks. Securing the request and control of resources, confidentiality of the information, and availability of function should all be critical security considerations when deploying and operating ACTN platforms.

Several distributed ACTN functional components are required, and implementations should consider encrypting data that flows between components, especially when they are implemented at remote nodes, regardless of whether these data flows are on external or internal network interfaces.

The ACTN security discussion is further split into two specific categories described in the following subsections:

- o Interface between the Customer Network Controller and Multi-Domain Service Coordinator (MDSC), CNC-MDSC Interface (CMI)
- o Interface between the Multi-Domain Service Coordinator and Provisioning Network Controller (PNC), MDSC-PNC Interface (MPI)

From a security and reliability perspective, ACTN may encounter many risks such as malicious attack and rogue elements attempting to connect to various ACTN components. Furthermore, some ACTN components represent a single point of failure and threat vector and must also manage policy conflicts and eavesdropping of communication between different ACTN components.

The conclusion is that all protocols used to realize the ACTN framework should have rich security features, and customer, application and network data should be stored in encrypted data stores. Additional security risks may still exist. Therefore, discussion and applicability of specific security functions and protocols will be better described in documents that are use case and environment specific.

9.1. CNC-MDSC Interface (CMI)

Data stored by the MDSC will reveal details of the virtual network services and which CNC and customer/application is consuming the resource. Therefore, the data stored must be considered a candidate for encryption.

CNC Access rights to an MDSC must be managed. The MDSC must allocate resources properly, and methods to prevent policy conflicts, resource waste, and denial-of-service attacks on the MDSC by rogue CNCs should also be considered.

The CMI will likely be an external protocol interface. Suitable authentication and authorization of each CNC connecting to the MDSC will be required; especially, as these are likely to be implemented by different organizations and on separate functional nodes. Use of the AAA-based mechanisms would also provide role-based authorization methods so that only authorized CNC's may access the different functions of the MDSC.

9.2. MDSC-PNC Interface (MPI)

Where the MDSC must interact with multiple (distributed) PNCs, a PKI-based mechanism is suggested, such as building a TLS or HTTPS connection between the MDSC and PNCs, to ensure trust between the physical network layer control components and the MDSC. Trust anchors for the PKI can be configured to use a smaller (and potentially non-intersecting) set of trusted Certificate Authorities (CAs) than in the Web PKI.

Which MDSC the PNC exports topology information to, and the level of detail (full or abstracted), should also be authenticated, and specific access restrictions and topology views should be configurable and/or policy based.

10. IANA Considerations

This document has no IANA actions.

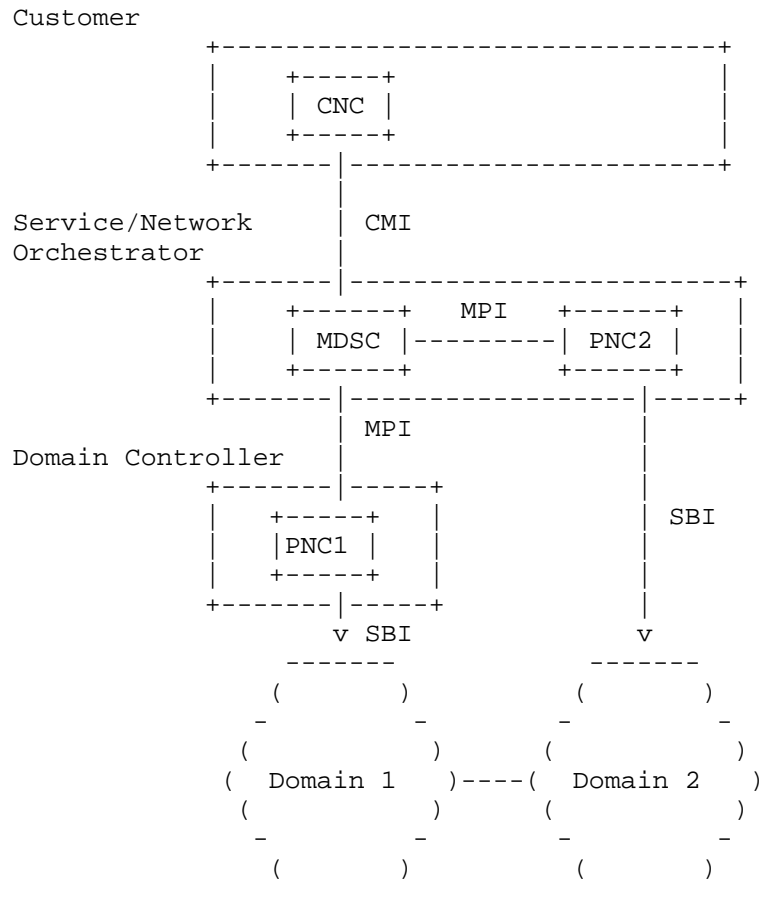
11. Informative References

- [ACTN-REQ] Lee, Y., Ceccarelli, D., Miyasaka, T., Shin, J., and K. Lee, "Requirements for Abstraction and Control of TE Networks", Work in Progress, draft-ietf-teas-actn-requirements-09, March 2018.
- [ACTN-YANG] Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., Yoon, B., Wu, Q., and P. Park, "A Yang Data Model for ACTN VN Operation", Work in Progress, draft-ietf-teas-actn-vn-yang-01, June 2018.
- [ONF-ARCH] Open Networking Foundation, "SDN Architecture", Issue 1.1, ONF TR-521, June 2016.
- [RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, DOI 10.17487/RFC2702, September 1999, <<https://www.rfc-editor.org/info/rfc2702>>.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, DOI 10.17487/RFC3945, October 2004, <<https://www.rfc-editor.org/info/rfc3945>>.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.
- [RFC5654] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, DOI 10.17487/RFC5654, September 2009, <<https://www.rfc-editor.org/info/rfc5654>>.
- [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>>.
- [RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control", RFC 8283, DOI 10.17487/RFC8283, December 2017, <<https://www.rfc-editor.org/info/rfc8283>>.
- [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>>.
- [TE-TOPO] Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and O. Dios, "YANG Data Model for Traffic Engineering (TE) Topologies", Work in Progress, draft-ietf-teas-yang-te-topo-18, June 2018.

Appendix A. Example of MDSC and PNC Functions Integrated in a Service/Network Orchestrator

This section provides an example of a possible deployment scenario, in which Service/Network Orchestrator can include the PNC functionalities for domain 2 and the MDSC functionalities.



Contributors

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

Italo Busi
Huawei
Email: Italo.Busi@huawei.com

Khuzema Pithewan
Peloton Technology
Email: khuzemap@gmail.com

Michael Scharf
Nokia
Email: michael.scharf@nokia.com

Luyuan Fang
eBay
Email: luyuanf@gmail.com

Diego Lopez
Telefonica I+D
Don Ramon de la Cruz, 82
28006 Madrid
Spain
Email: diego@tid.es

Sergio Belotti
Nokia
Via Trento, 30
Vimercate
Italy
Email: sergio.belotti@nokia.com

Daniel King
Lancaster University
Email: d.king@lancaster.ac.uk

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India
Email: dhruv.ietf@gmail.com

Gert Grammel
Juniper Networks
Email: ggrammel@juniper.net

Authors' Addresses

Daniele Ceccarelli (editor)
Ericsson
Torshamnsgatan, 48
Stockholm
Sweden

Email: daniele.ceccarelli@ericsson.com

Young Lee (editor)
Huawei Technologies
5340 Legacy Drive
Plano, TX 75023
United States of America

Email: leeyoung@huawei.com

