TEAS Working Group
Internet Draft
Intended status: Informational
Expires: December 2017

Daniele Ceccarelli (Ed)
Ericsson
Young Lee (Ed)
Huawei

June 13, 2017

Framework for Abstraction and Control of Traffic Engineered Networks

draft-ietf-teas-actn-framework-06

Abstract

Traffic Engineered 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.

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 Traffic Engineered Networks (ACTN).

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on December 13, 2017.

Copyright Notice

Copyright (c) 2017 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 (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction....3
1.1. Terminology....5
2. Business Model of ACTN....9
2.1. Customers....9
2.2. Service Providers....10
2.3. Network Providers....12
3. Virtual Network Service....12
4. ACTN Base Architecture....13
4.1. Customer Network Controller....15
4.2. Multi Domain Service Coordinator....16
4.3. Physical Network Controller....17
4.4. ACTN Interfaces....17
5. Advanced ACTN architectures....18
5.1. MDSC Hierarchy for scalability....18
5.2. Functional Split of MDSC Functions in Orchestrators....19
6. Topology Abstraction Method....21
6.1. No abstraction (native/white topology)....22
6.2. One Virtual Node (black topology)....23
6.3. Abstraction of TE tunnels for all pairs of border nodes (grey topology)....24

6.3.1. Grey topology type A: border nodes with a TE links between them in a full mesh fashion.....25
6.3.2. Grey topology Type B.....25
6.4. Topology Abstraction Granularity Level example.....25
7. Access Points and Virtual Network Access Points.....27
7.1. Dual homing scenario.....29
8. Advanced ACTN Application: Multi-Destination Service.....30
8.1. Pre-Planned End Point Migration.....31
8.2. On the Fly End Point Migration.....32
9. Advanced Topic.....32
10. Manageability Considerations.....32
10.1. Policy.....33
10.2. Policy applied to the Customer Network Controller.....34
10.3. Policy applied to the Multi Domain Service Coordinator..34
10.4. Policy applied to the Physical Network Controller.....34
11. Security Considerations.....35
11.1. Interface between the Customer Network Controller and Multi Domain Service Coordinator (MDSC), CNC-MDSC Interface (CMI)...36
11.2. Interface between the Multi Domain Service Coordinator and Physical Network Controller (PNC), MDSC-PNC Interface (MPI)...36
12. References.....37
12.1. Informative References.....37
13. Contributors.....38
Authors' Addresses.....39
APPENDIX A - Example of MDSC and PNC functions integrated in Service/Network Orchestrator.....40
APPENDIX B - Example of IP + Optical network with L3VPN service..40
APPENDIX C - Example of ODU service connectivity.....41

1. Introduction

Traffic Engineered networks have a variety of mechanisms to facilitate separation of data plane and control plane 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.

The term Traffic Engineered network is used in this document to refer 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. Some examples of networks that are in scope of this definition are optical networks, MPLS Transport Profile (MPLS-TP) networks [RFC5654], and MPLS Traffic Engineering (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 of the control plane from the data plane has been already achieved with the development of MPLS/GMPLS [GMPLS] and the Path Computation Element (PCE) [RFC4655] for TE-based networks. 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, PCE is essentially equivalent to a logically centralized path computation function.

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

. Separation of service requests from service delivery so that the orchestration of a network is transparent from the point of view of the customer but remains responsive to the customer's services and business needs.

. 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 domain. The process of abstraction presents the connectivity graph in a way that is independent of the underlying network technologies, capabilities, and topology so that it can be used to plan and deliver network services in a uniform way

. Coordination of resources across multiple domains and multiple layers to provide end-to-end services regardless of whether the domains use SDN or not.

As networks evolve, the need to provide separated service request/orchestration and resource abstraction has emerged as a key requirement for operators. In order to support multiple clients each with its own view of and control of the server network, a network operator needs to partition (or "slice") the network resources. The resulting slices can be assigned to each client for guaranteed usage which is a step further than shared use of common network resources.

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

We call the set of management and control functions used to provide these features Abstraction and Control of Traffic Engineered Networks (ACTN).

Particular attention needs to be paid to the multi-domain case, 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) as a single virtualized network.

The ACTN framework described in this document facilitates:

. Abstraction of the underlying network resources to higher-layer applications and customers [RFC7926].

. Virtualization of particular underlying resources, whose selection criterion is the allocation of those resources to a particular customer, application or service [ONF-ARCH].

. Network slicing of infrastructure to meet specific customers' service requirements.

. Creation of a virtualized environment allowing operators to view and control multi-domain networks as a single virtualized network.

. The presentation to customers of networks as a virtual network via open and programmable interfaces.

1.1. Terminology

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

. Network Slicing: In the context of ACTN, network slicing is a collection of resources that are used to establish logically dedicated virtual networks over TE networks. It allows a network provider to provide dedicated virtual networks for application/customer over a common network infrastructure. The logically dedicated resources are a part of the larger common network infrastructures that are shared among various network slice instances which are the end-to-end realization of 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). 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. Just as network slicing and network abstraction may be applied recursively, so a node in a topology may be created by applying slicing or abstraction on 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 RFC 7926 for details). Network slicing/virtualization and network abstraction may be applied recursively, so a link in a topology may be created by applying slicing and/or abstraction on the links in the underlying topology.

. CNC: A Customer Network Controller is responsible for communicating customer's virtual network service requirements to network provider. It has knowledge of the end-point associated with virtual network service, service policy, and other QoS information related to the service it is responsible for instantiating.

. PNC: A Physical Network Controller is responsible for controlling devices or NEs under its direct control. 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 or any other means to dynamically control a set of nodes and that is implementing an NBI compliant with ACTN specification.

. PNC domain: 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 can be a link or a node.

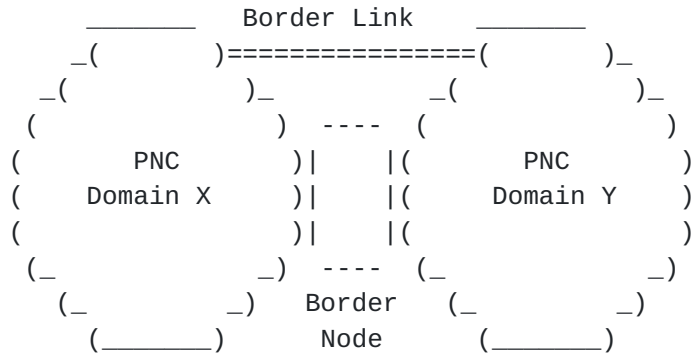

Figure 1: PNC Domain Borders

. MDSC: A multi-domain Service Coordinator is a functional block that implements all four ACTN main functions, i.e., multi domain coordination, virtualization/abstraction, customer mapping/translation, and virtual service coordination. The first two functions of the MDSC, namely, multi domain coordination and virtualization/abstraction are referred to as network-related functions while the last two functions, namely, customer mapping/translation and virtual service coordination are referred to as service-related functions. See details on these functions in Section 4.2. In some implementation, PNC and MDSC functions can be co-located and implemented in the same box.

. A Virtual Network (VN) is a customer view of the TE network. Depending on the agreement between client and provider various VN operations and VN views are possible as follows:

  o VN Creation - VN could be pre-configured and created via offline negotiation between customer and provider. In other cases, the VN could also be created dynamically based on a request from the customer with given SLA attributes which satisfy the customer's objectives.

  o Dynamic Operations - The VN could be further modified or deleted based on a customer request. The customer can further act upon the virtual network resources to perform end-to-end tunnel management (set-up/release/modify).

These changes will result in subsequent LSP management at the operator's level.

o VN Type:

a. The VN can be seen as set of end-to-end tunnels from a customer point of view, where each tunnel is referred as a VN member. Each VN member can then be formed by recursive slicing or abstraction of paths in underlying networks. Such end-to-end tunnels may comprise of customer end points, access links, intra-domain paths, and inter-domain links. In this view, VN is thus a set of VN members (which is referred to as Type 1 VN)

b. The VN can also be seen as a topology comprising of physical, sliced, and abstract nodes and links. This VN is referred to as Type 2 VN. The nodes in this case include physical customer end points, 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. With VN type 2, it is still possible to view VN member-level.

. Virtual Network Service (VNS) is requested by the customer and negotiated with the provider. There are three types of VNS defined in this document. Type 1 VNS refers to VNS in which customer is allowed to create and operate a Type 1 VN. Type 2a and 2b VNS refers to the VNS in which customer is allowed to create and operates a Type 2 VN. With Type 2a VNS, once the VN is statically created at service configuration time, the customer is not allowed to change the topology (i.e., adding or deleting abstract nodes/links). Type 2b VNS is the same as Type 2a VNS except that the customer is allowed to change topology dynamically from the initial topology created at service configuration time. See Section 3 for details.

. Abstraction. This process is defined in [RFC7926].

. Abstract Link: The term "abstract link" is defined in [RFC7926].

. 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.

. Access link: A link between a customer node and a provider node.

. Inter-domain link: A link between domains managed by different PNCs. The MDSC is in charge of managing inter-domain links.

. Access Point (AP): An access point is used to keep confidentiality between the customer and the provider. It is a logical identifier shared between the customer and the provider, used to map the end points of the border node in both the customer and the provider NW. The AP can be used by the customer when requesting VN service to the provider.

. VN Access Point (VNAP): A VNAP is defined as the binding between an AP and a given VN and is used to identify the portion of the access and/or inter-domain link dedicated to a given VN.

2. Business Model of ACTN

The Virtual Private Network (VPN) [RFC4026] and Overlay Network (ON) models [RFC4208] are built on the premise that the network provider provides all virtual private or overlay networks to its customers. These models are simple to operate but have some disadvantages in accommodating the increasing need for flexible and dynamic network virtualization capabilities.

There are three key entities in the ACTN model:

- Customers
- Service Providers
- Network Providers

These are described in the following sections.

2.1. Customers

Within the ACTN framework, different types of customers may be taken into account depending on the type of their resource needs, and on

their number and type of access. For example, it is possible to group them into two main categories:

Basic Customer: Basic customers include fixed residential users, mobile users and small enterprises. Usually, the number of basic customers for a service provider is high: they require small amounts of resources and are characterized by steady requests (relatively time invariant). A typical request for a basic customer is for a bundle of voice services and internet access. Moreover, basic customers do not modify their services themselves: if a service change is needed, it is performed by the provider as a proxy and the services generally have very few dedicated resources (such as for subscriber drop), with everything else shared on the basis of some Service Level Agreement (LSA), which is usually best-efforts.

Advanced Customer: Advanced customers typically include enterprises, governments and utilities. Such customers can ask for both point-to point and multipoint connectivity with high resource demands varying significantly in time and from customer to customer. 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 virtual network service.

Advanced customers may own dedicated virtual resources, or share resources. They 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 provider domains, they have to interface to multiple providers and may have to support multiple virtual network services with different underlying objectives set by the network providers. To enable these customers to support flexible and dynamic applications they need to control their allocated virtual network resources in a dynamic fashion, and that means that they need a view of the topology that spans all of the network providers. Customers of a given service provider can in turn offer a service to other customers in a recursive way.

2.2. Service Providers

Service providers are the providers of virtual network services (see Section 3 for details) to their customers. Service providers may or may not own physical network resources (i.e, may or may not be network providers as described in Section 2.3). When a service provider is the same as the network provider, this is similar to existing VPN models applied to a single provider. This approach

works well when the customer maintains a single interface with a single provider. When customer spans multiple independent network provider domains, then it becomes hard to facilitate the creation of end-to-end virtual network services with this model.

A more interesting case arises when network providers only provide infrastructure, while distinct service providers interface to the customers. In this case, service providers are, themselves customers of the network infrastructure providers. One service provider may need to keep multiple independent network providers as its end-users span geographically across multiple network provider domains.

The ACTN network model is predicated upon this three tier model and is summarized in Figure 2:

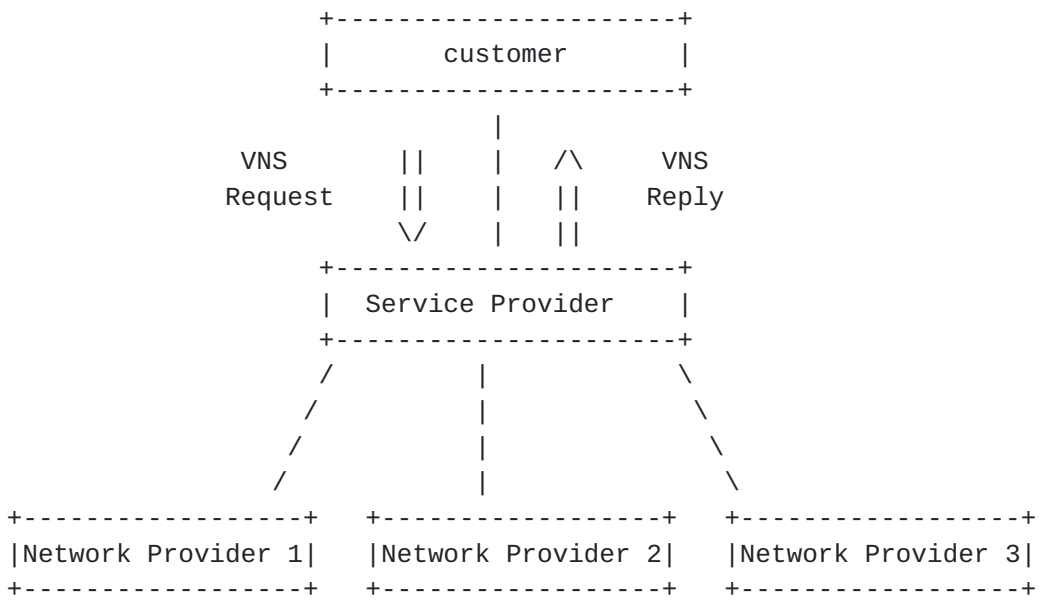

Figure 2: Three tier model.

There can be multiple service providers to which a customer may interface.

There are multiple types of service providers, for example:

- Data Center providers can be viewed as a service provider type as they own and operate data center resources for various WAN customers, and they can lease physical network resources from network providers.
- Internet Service Providers (ISP) are service providers of internet services to their customers while leasing physical network resources from network providers.
- Mobile Virtual Network Operators (MVNO) provide mobile services to their end-users without owning the physical network infrastructure.

2.3. Network Providers

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

3. Virtual Network Service

Virtual Network Service (VNS) is requested by the customer and negotiated with the provider. There are three types of VNS defined in this document.

Type 1 VNS refers to VNS in which customer is allowed to create and operate a Type 1 VN. Type 1 VN is a VN that comprises a set of end-to-end tunnels from a customer point of view, where each tunnel is referred as a VN member. With Type 1 VNS, the network operator does not need to provide additional abstract VN topology associated with the Type 1 VN.

Type 2a VNS refer to VNS in which customer is allowed to create and operates a Type 2 VN, but not allowed to change topology once it is configured at service configuration time. Type 2 VN is an abstract VN topology that may comprise of virtual/abstract nodes and links. The nodes in this case may include physical customer end points, border nodes, and internal nodes as well as abstracted nodes. Similarly, the links may include physical access links, inter-domain links, and intra-domain links as well as abstract links.

Type 2b VNS refers to VNS in which customer is allowed to create and operate a Type 2 VN and the customer is allowed to dynamically change abstract VN topology from the initially configured abstract VN topology at service configuration time.

From an implementation standpoint, Type 2a VNS and Type 2b VNS differentiation might be fulfilled via local policy.

In all types of VNS, customer can specify a set of service related parameters such as connectivity type, VN traffic matrix (e.g., bandwidth, latency, diversity, etc.), VN survivability, VN service policy and other characteristics.

4. 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 aligned with the ONF SDN architecture [ONF-ARCH] and presents a 3-tiers reference model. It allows for hierarchy and recursiveness not only of SDN controllers but also of traditionally controlled domains that use a control plane. It defines three types of controllers depending on the functionalities they implement. The main functionalities that are identified are:

. Multi-domain coordination function: This function oversees the specific aspects of the 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.

. Virtualization/Abstraction function: 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.

. Customer mapping/translation function: This function is to map customer requests/commands into network provisioning requests that can be sent to the Physical Network Controller (PNC) according to business policies provisioned statically or dynamically at the OSS/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.

. Virtual service coordination function: This function translates customer service-related information 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, guarantees of service quality, bandwidth and throughput. It also includes notifications for service fault and performance degradation and so forth.

Figure 3 depicts the base ACTN architecture with three controller types and the corresponding interfaces between these controllers. The types of controller defined in the ACTN architecture are shown in Figure 3 below and are as follows:

. CNC - Customer Network Controller
. MDSC - Multi Domain Service Coordinator
. PNC - Physical Network Controller

Figure 3 also shows the following interfaces:

. CMI - CNC-MDSC Interface
. MPI - MDSC-PNC Interface
. SBI - South Bound Interface

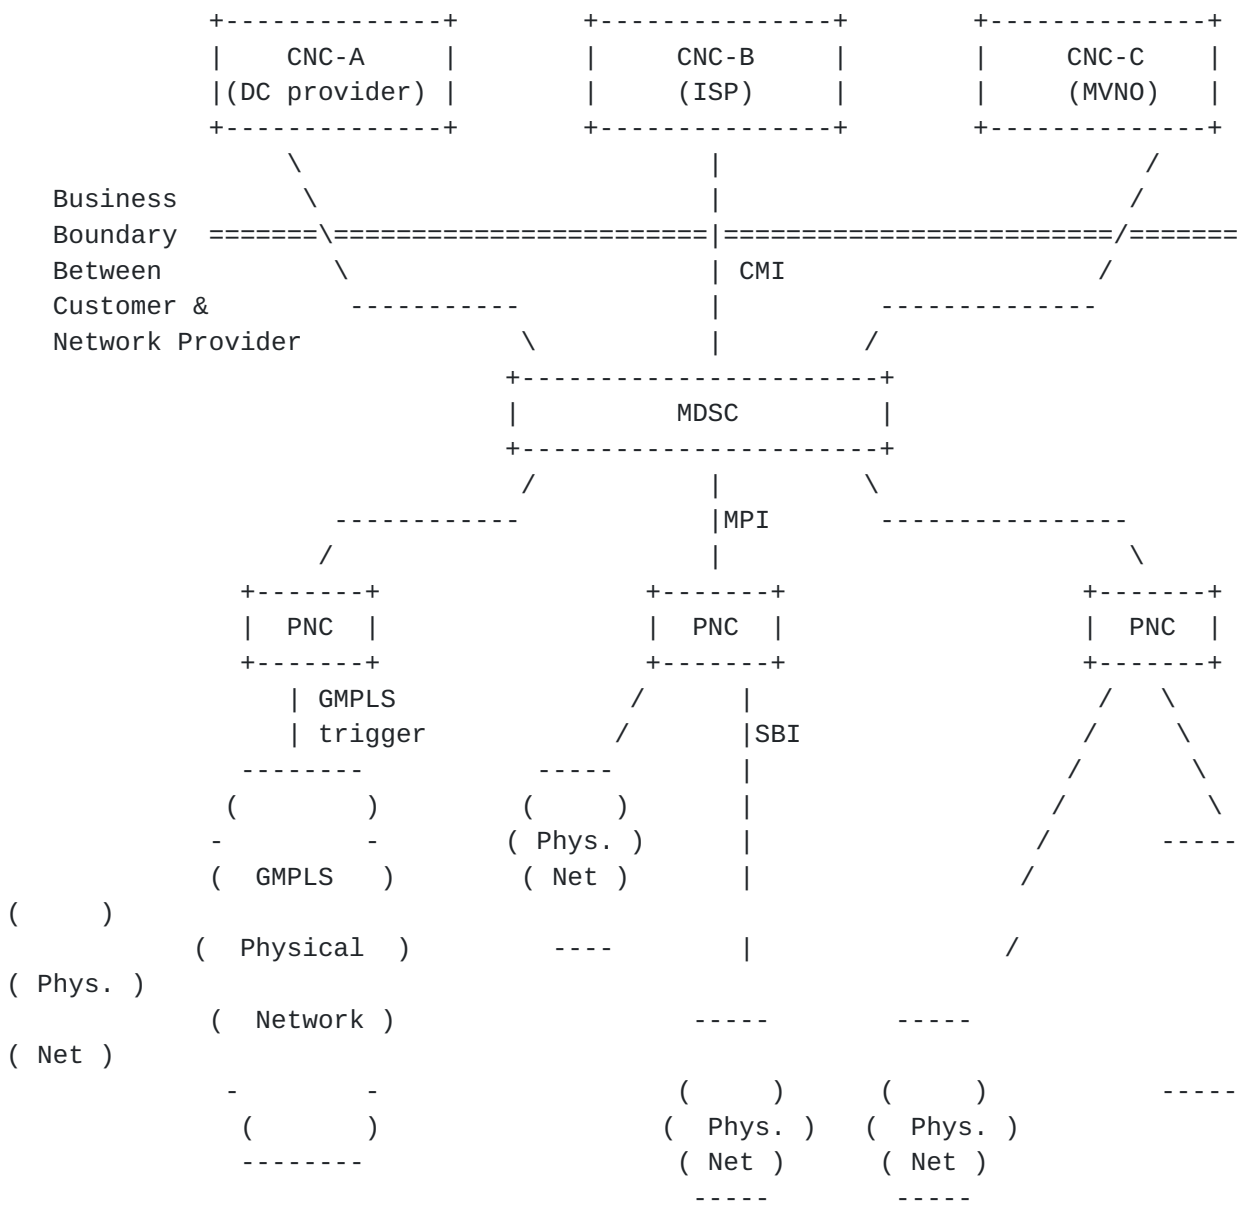

Figure 3: ACTN Base Architecture

4.1. Customer Network Controller

A Virtual Network Service is instantiated by the Customer Network Controller via the CNC-MDSC Interface (CMI). As the Customer Network Controller directly interfaces to the applications, it understands multiple application requirements and their service needs. It is assumed that the Customer Network Controller and the MDSC have a

common knowledge of the end-point interfaces based on their business negotiations prior to service instantiation. End-point interfaces

refer to customer-network physical interfaces that connect customer premise equipment to network provider equipment.

4.2. Multi Domain Service Coordinator

The Multi Domain Service Coordinator (MDSC) sits between the CNC that issues connectivity requests and the Physical Network Controllers (PNCs) that manage the physical network resources. The MDSC can be collocated with the PNC.

The internal system architecture and building blocks of the MDSC are out of the scope of ACTN. Some examples can be found in the Application Based Network Operations (ABNO) architecture [RFC7491] and the ONF SDN architecture [ONF-ARCH].

The MDSC is the only building block of the architecture that can implement all four ACTN main functions, i.e., multi domain coordination, virtualization/abstraction, customer mapping/translation, and virtual service coordination. The first two functions of the MDSC, namely, multi domain coordination and virtualization/abstraction are referred to as network-related functions while the last two functions, namely, customer mapping/translation and virtual service coordination are referred to as service-related functions.
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 between MDSCs and PNCs (i.e. 1 parent MDSC and N child MDSC or 1 MDSC and N PNCs).

In addition to that, it could also be possible to have an M:1 relationship between MDSCs and PNC to allow for network resource partitioning/sharing among different customers not necessarily connected to the same MDSC (e.g., different service providers).

4.3. Physical Network Controller

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

The internal architecture of the PNC, its building blocks, and the way it controls its domain are out of the scope of ACTN. Some examples can be found in the Application Based Network Operations (ABNO) architecture [RFC7491] and the ONF SDN architecture [ONF-ARCH]

The PNC, in addition to being in charge of controlling the physical network, is able to implement two of the four main ACTN main functions: multi domain coordination and virtualization/abstraction function.
Note that from an implementation point of view it is possible to integrate one or more MDSC functions and one or more PNC functions within the same controller.

4.4. ACTN Interfaces

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 while having no impact on other customers. Direct customer control of transport network elements and virtualized services is not perceived as a viable proposition for transport network providers due to security and policy concerns among other reasons. In addition, the network control plane for transport networks has been separated from the data plane and as such it is not viable for the customer to directly interface with transport network elements.

. CMI Interface: The CNC-MDSC Interface (CMI) is an interface between a CNC and an MDSC. As depicted in Figure 3, the CMI is a business boundary between customer and network provider. It is used to request virtual network services required for the applications. Note that all service related information such as specific service properties, including virtual network service type, topology, bandwidth, and constraint information are conveyed over this interface. Most of the information over this interface is technology agnostic; however, there are some cases, e.g., access link configuration, where it should be

possible to explicitly request for a VN to be created at a given layer in the network (e.g. ODU VN or MPLS VN).

. MPI Interface: The MDSC-PNC Interface (MPI) is an interface between an MDSC and a PNC. It communicates the creation requests for new connectivity or for bandwidth changes in the physical network. In multi-domain environments, the MDSC needs to establish multiple MPIs, one for each PNC, as there is one PNC responsible for control of each domain. The MPI could have different degrees of abstraction and present an abstracted topology hiding technology specific aspects of the network or convey technology specific parameters to allow for path computation at the MDSC level. Please refer to CCAMP Transport NBI work for the latter case [Transport NBI].

. SBI Interface: This interface is out of the scope of ACTN. It is shown in Figure 3 for reference reason only.

Please note that for all the three interfaces, when technology specific information needs to be included, this info will be add-ons on top of the general abstract topology. As far as general topology abstraction standpoint, all interfaces are still recursive in nature.

5. Advanced ACTN architectures

This section describes advanced forms of ACTN architectures as possible implementation choices.

5.1. MDSC Hierarchy for scalability

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 higher-level MDSC (MDSC-H) the lower-level MDSC (MDSC-L) and the interface between them is basically of a recursive nature of the MPI. An implementation choice could foresee the usage of an MDSC-L for all the PNCs related to a given network layer or technology (e.g. IP/MPLS) a different MDSC-L for the PNCs related to another layer/technology (e.g. OTN/WDM) and an MDSC-H to coordinate them.

Figure 4 shows this case.

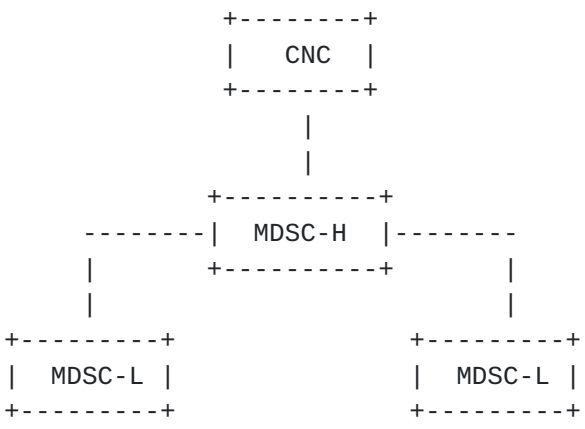

Figure 4: MDSC Hierarchy

Note that both the MDSC-H and the MDSC-L in general cases implement all four functions of the MDSC discussed in Section 3.2.

5.2. Functional Split of MDSC Functions in Orchestrators

Another implementation choice could foresee the separation of MDSC functions into two groups (i.e., one group for service-related functions and another group for network-related functions) which will result in a service orchestrator for providing service-related functions of MDSC and other non-ACTN functions and a network orchestrator for providing network-related functions of MDSC and other non-ACTN functions. Figure 5 shows this case and it also depicts the mapping between ACTN architecture and the YANG service model architecture described in [Service-YANG]. This mapping is helpful for the readers who are not familiar with some TEAS specific terminology used in this document. A number of key ACTN interfaces exist for deployment and operation of ACTN-based networks. These are highlighted in Figure 5 (ACTN Interfaces).

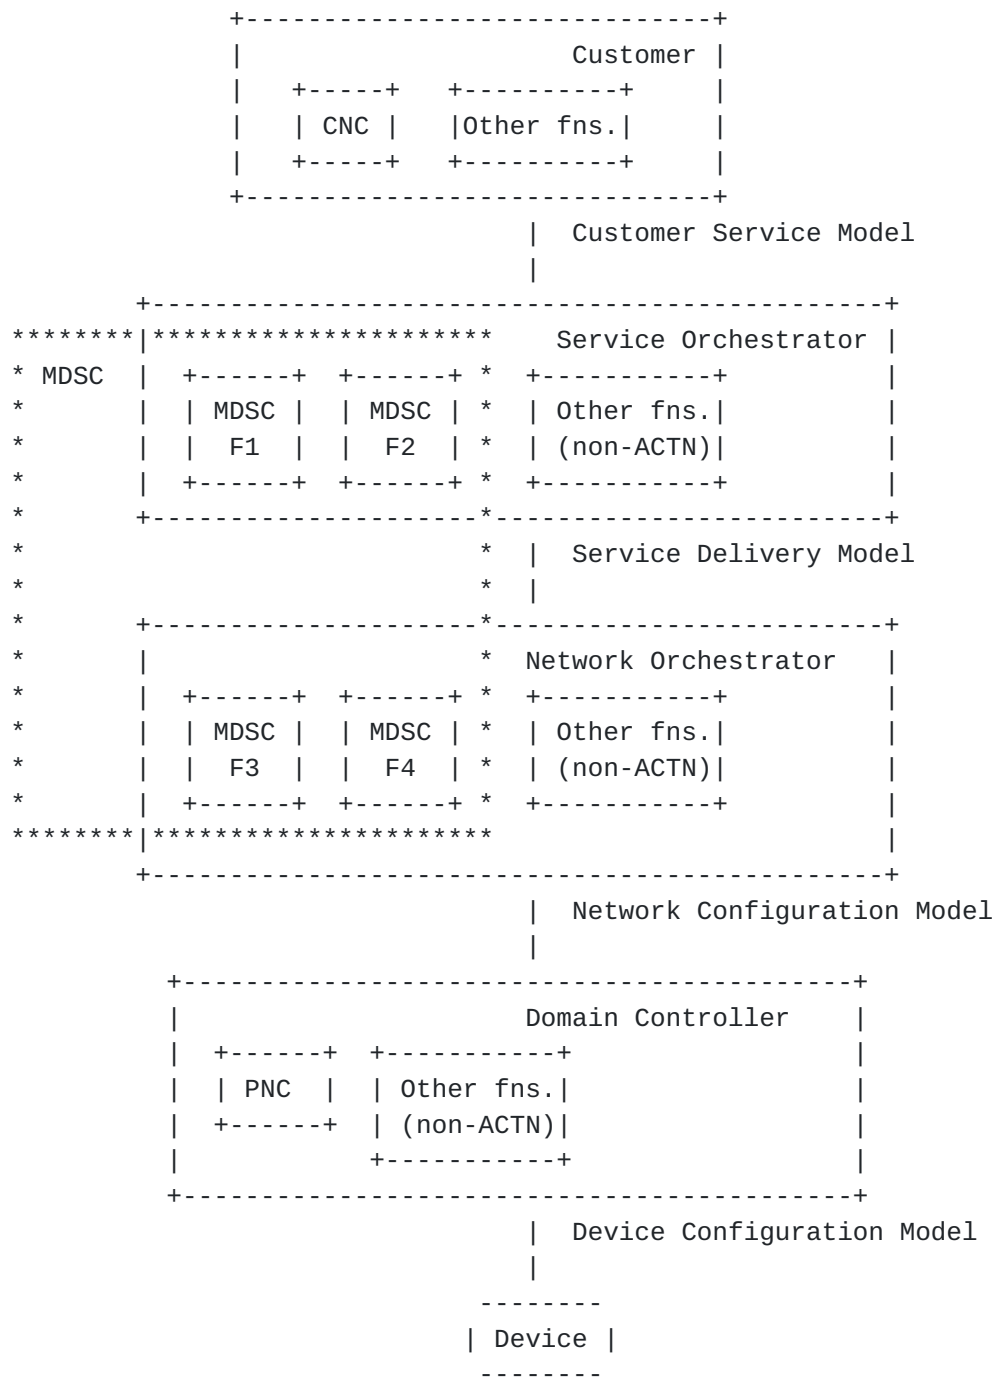

Figure 5: ACTN Architecture in the context of YANG Service Models

In Figure 5, MDSC F1 and F2 correspond to customer mapping/translation, and virtual service coordination, respectively, which are the MDSC service-related functions as defined in Section 4. MDSC F3 and F4 correspond to multi domain coordination, virtualization/abstraction, respectively, which are the MDSC network-related functions as defined in Section 4. In some implementation, MDSC F1 and F2 can be implemented as part of a Service Orchestrator which may support other non-ACTN functions. Likewise, the MDSC F3 and F4 can be implemented as part of a Network Orchestrator which may support other non-ACTN MDSC functions.

Also note that the PNC is not same as domain controller. Domain controller in general has a larger set of functions than that of PNC. The main functions of PNC are explained in Section 3.3. Likewise, Customer has a larger set of functions than that of the CNC.

Customer service model describes a service as offer or delivered to a customer by a network operator as defined in [Service-YANG]. The CMI is a subset of a customer service model to support VNS. This model encompasses other non-TE/non-ACTN models to control non-ACTN services (e.g., L3SM).

Service delivery model is used by a network operator to define and configure how a service is provided by the network as defined in [Service-YANG]. This model is similar to the MPI model as the network-related functions of the MDSC, i.e., F3 and F4, provide an abstract topology view of the E2E network to the service-related functions of the MDSC, i.e., F1 and F2, which translate customer's request at the CMI into the network configuration at the MPI.

Network configuration model is used by a network orchestrator to provide network-level configuration model to a controller as defined in [Service-YANG]. The MPI is a subset of network configuration model to support TE configuration. This model encompasses the MPI model plus other non-TE/non-ACTN models to control non-ACTN functions of the domain controller (e.g., L3VPN).

Device configuration model is used by a controller to configure physical network elements.

6. Topology Abstraction Method

This section discusses topology abstraction types and their context in ACTN architecture. Topology abstraction is useful in ACTN architecture as a way to scale multi-domain network operation. As

the MDSC needs to coordinate with the PNCs in multi-domain networks for determining a feasible domain sequence and provisioning the end-to-end tunnel based on the determined domain sequence, topology abstraction provides an efficient network scaling mechanism. Note that this is the abstraction performed by the PNC to the MDSC or by the MDSC-L to the MDSC-H, and that this is different from the VN Type 2 topology (that is created and negotiated between the CNC and the MDSC as part of the VNS). The purpose of topology abstraction discussed in this section is for an efficient internal network operation based on abstraction principle.

Refer to [ACTN-Abstraction] for detailed discussions on factors that affect topology abstraction, how to build topology abstraction, various considerations and which method to pick based on abstraction types and the nature of underlying transport technology (e.g., MPLS, OTN, WSON, etc.).

6.1. No abstraction (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 as shown in Figure 6. In this case, the MDSC has the full knowledge of the underlying network topology and as such there is no need for the MDSC to send a path computation request to the PNC. The computation burden will fall on the MDSC to find an optimal end-to-end path and optimal per domain paths.

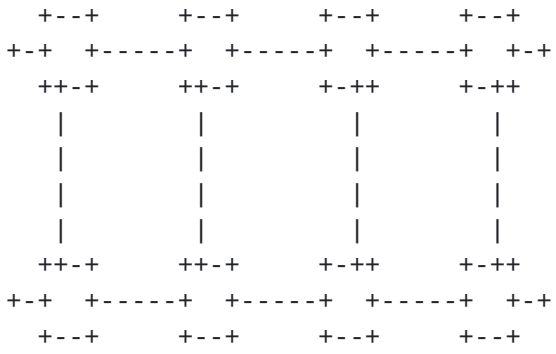

Figure 6a: The native/white topology

6.2. One Virtual Node (black topology)

The entire domain network is abstracted as a single virtual node (see the definition of virtual node in [RFC7926]) with the access/egress links without disclosing any node internal connectivity information.

Figure 6b 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 path computation requests to the PNCs before it can determine an end-to-end path. If there are a large number of inter-connected domains, this abstraction method may impose a heavy coordination load at the MDSC level in order to find an optimal end-to-end path.

The black topology would not give the MDSC any critical network resource information other than the border nodes/links information and as such it is likely to have a need for complementary communications between the MDSC and the PNCs (e.g., Path computation Request/Reply).

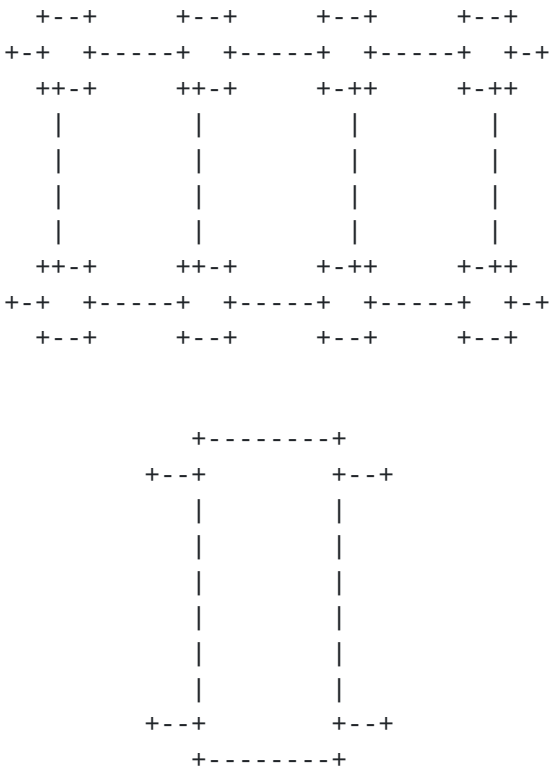

Figure 6b: The native topology and the corresponding black topology with one virtual node and inter-domain links

6.3. Abstraction of TE tunnels for all pairs of border nodes (grey topology)

This abstraction level, referred to a grey topology is between black topology and white topology from a granularity point of view. As shown in Figures 7a and 7b, we may further differentiate from a perspective of how to abstract internal TE resources between the pairs of border nodes:

- Grey topology type A: border nodes with a TE links between them in a full mesh fashion (See Figure 7a)
- Grey topology type B: border nodes with some internal abstracted nodes and abstracted links (See Figure 7b)

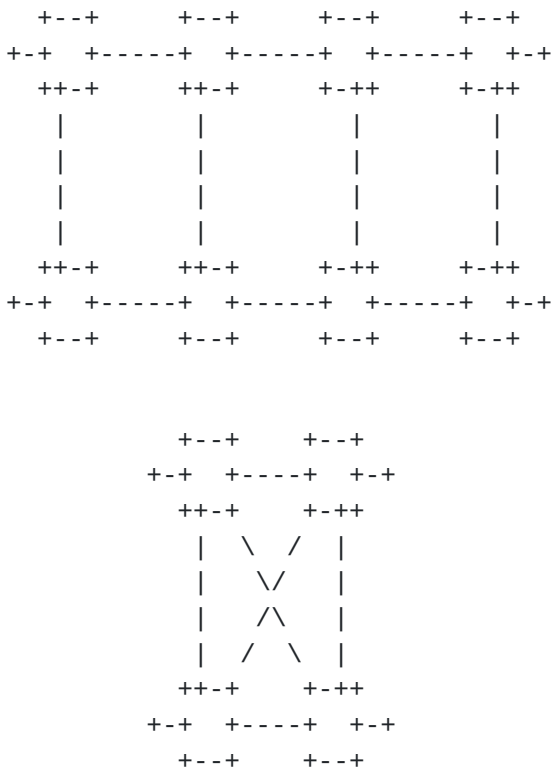

Figure 7a: The native topology and the corresponding grey topology type A with TE links between border nodes

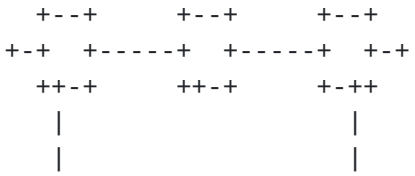

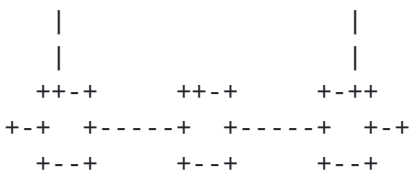

Figure 7b: The grey topology type B with abstract nodes/links between border nodes

6.3.1. Grey topology type A: border nodes with a TE links between them in a full mesh fashion

For each pair of ingress and egress nodes (i.e., border nodes to/from the domain), TE link metric is provided with TE attributes such as max bandwidth available, link delay, etc. This abstraction depends on the underlying TE networks.
Note that this grey topology can also be represented as a single abstract node with the connectivity matrix defined in [TE-Topology], abstracting the internal connectivity information. The only thing might be different is some additional information about the end points of the links of the border nodes (i.e., links outward customer-facing) as they cannot be included in the connectivity matrix's termination points.

6.3.2. Grey topology Type B

The grey abstraction type B would allow the MDSC to have more information about the internals of the domain networks by the PNCs so that the MDSC can flexibly determine optimal paths. The MDSC may configure some of the internal virtual nodes (e.g., cross-connect) to redirect its traffic as it sees changes from the domain networks.

6.4. Topology Abstraction Granularity Level example

This section illustrates how topology abstraction operates in different level of granularity over a hierarchy of MDSCs which is shown in Figure 8 below.

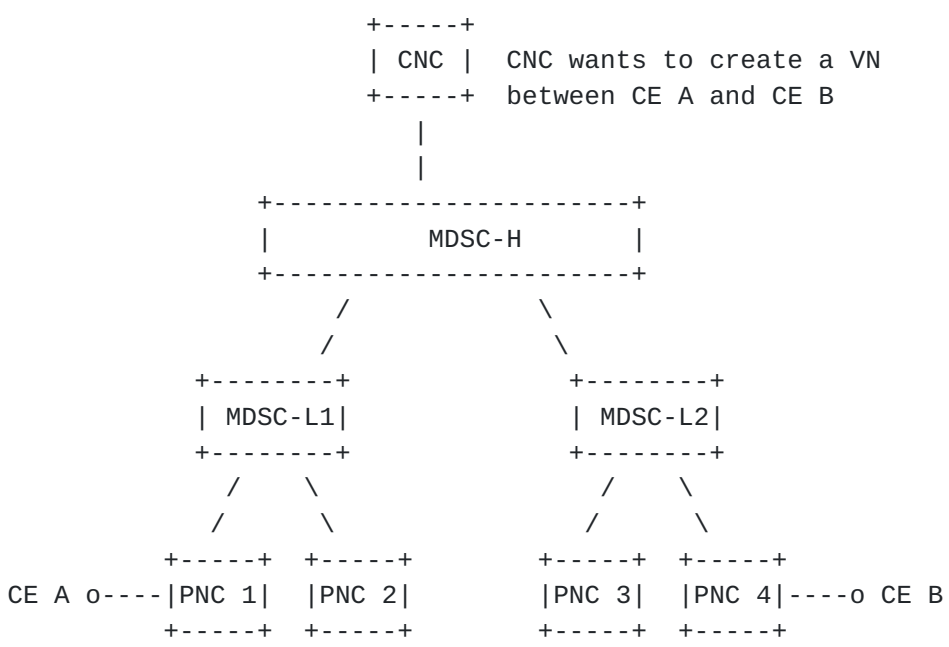

Topology operated by MDSC-H

Topology operated by MDSC-L1

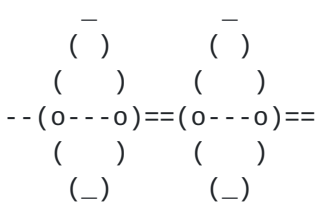

Topology operated by MDSC-L2

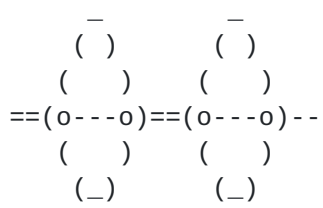

Actual Topology

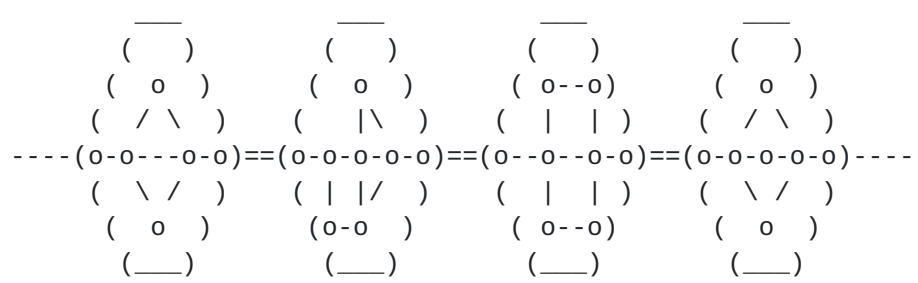

Where o is a node and -- is a link and === a border link

Figure 8: Illustration of topology abstraction granularity levels

In the example depicted in Figure 8, there are four domains under control of the respective PNCs, namely, PNC 1, PNC 2, PNC3 and PNC4. Assume that MDSC L-1 is controlling PNC 1 and PNC 2 while MDSC L-2 is controlling PNC 3 and PNC 4. Let us assume that each of the PNCs provides a grey topology abstraction in which to present only border nodes and links within and outside the domain. The abstract topology MDSC-L1 would operate is basically a combination of the two topologies the PNCs (PNC 1 and PNC 2) provide. Likewise, the abstract topology MDSC-L2 would operate is shown in Figure 8. Both MDSC-L1 and MDSC-L2 provide a black topology abstraction in which each PNC domain is presented as one virtual node to its top level MDSC-H. Then the MDSC-H combines these two topologies updated by MDSC-L1 and MDSC-L2 to create the abstraction topology to which it operates. MDSC-H sees the whole four domain networks as four virtual nodes connected via virtual links. This illustrates the point discussed in Section 5.1: The top level MDSC may operate on a higher level of abstraction (i.e., less granular level) than the lower level MSDCs.

7. Access Points and Virtual Network Access Points

In order not to share unwanted topological information between the customer domain and provider domain, a new entity is defined which is referred to as the Access Point (AP). See the definition of AP in Section 1.1.

A customer node will use APs as the end points for the request of VNS as shown in Figure 9.

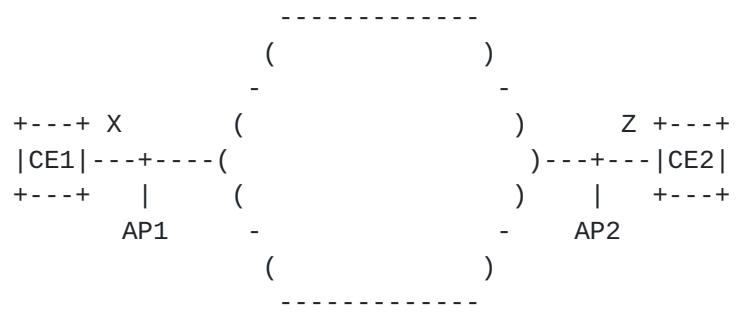

Figure 9: APs definition customer view

Let's take as an example a scenario shown in Figure 7. CE1 is connected to the network via a 10Gb link and CE2 via a 40Gb link. Before the creation of any VN between AP1 and AP2 the customer view can be summarized as shown in Table 1:

	End Point	Access Link Bandwidth	
AP id	CE,port	MaxResBw	AvailableBw
AP1	CE1,portX	10Gb	10Gb
AP2	CE2,portZ	40Gb	40Gb

Table 1: AP - customer view

On the other hand, what the provider sees is shown in Figure 10.

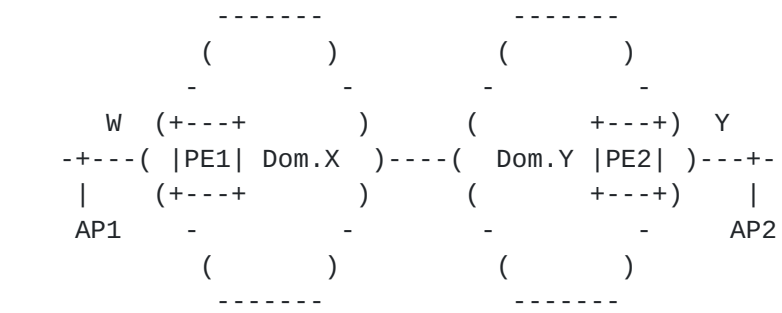

Figure 10: Provider view of the AP

Which results in a summarization as shown in Table 2.

	End Point	Access Link Bandwidth	
AP id	PE,port	MaxResBw	AvailableBw
AP1	PE1,portW	10Gb	10Gb
AP2	PE2,portY	40Gb	40Gb

Table 2: AP - provider view

A Virtual Network Access Point (VNAP) needs to be defined as binding between the AP that is linked to a VN and that is used to allow for 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 the simple scenario depicted above we suppose we want to create two virtual networks. The first with VN identifier 9 between AP1 and AP2 with bandwidth of 1Gbps, while the second with VN id 5, again between AP1 and AP2 and with bandwidth 2Gbps.

The provider view would evolve as shown in Table 3.

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

Table 3: AP and VNAP - provider view after VNS creation

7.1. 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 5Gbps of bandwidth between CE1 and CE2. This is shown in Figure 11.

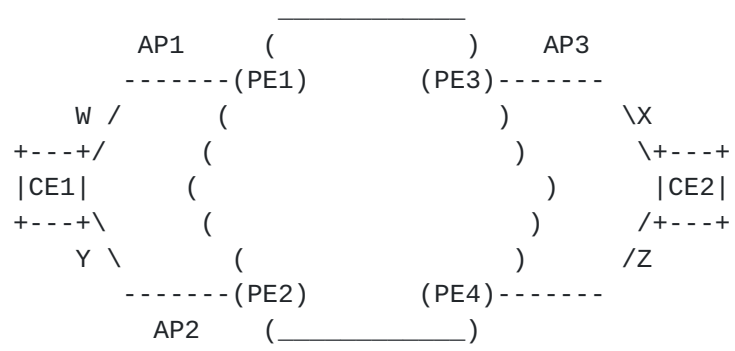

Figure 11: Dual homing scenario

In this case, the customer will request for 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 end points).

The customer view would be shown in Table 4.

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

Table 4: Dual homing - customer view after VN creation

8. Advanced ACTN Application: Multi-Destination Service

A further advanced application of ACTN is in the case of Data Center selection, where the customer requires the Data Center selection to be based on the network status; this is referred to as Multi-Destination in [ACTN-REQ]. In terms of ACTN, a CNC could request a connectivity service (virtual network) between a set of source Aps and destination APs and leave it up to the network (MDSC) to decide which source and destination access points to be used to set up the connectivity service (virtual network). The candidate list of source and destination APs is decided by a CNC (or an entity outside of ACTN) based on certain factors which 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 12, where three data centers are available, but the customer requires the data center 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 setup the connectivity service (virtual network).

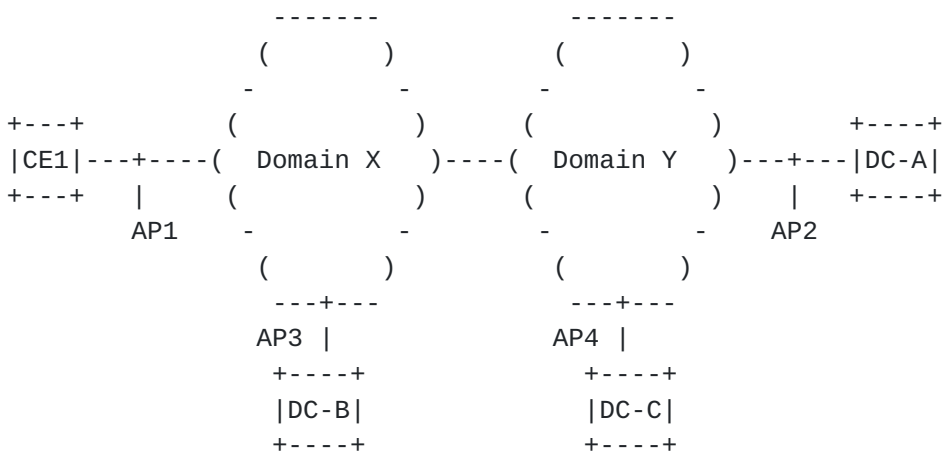

Figure 12: End point selection based on network status

8.1. Pre-Planned End Point Migration

Further in case of Data Center selection, customer could request for 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 13 DC-C is selected as a backup for DC-A. Thus, the VN should be setup 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).

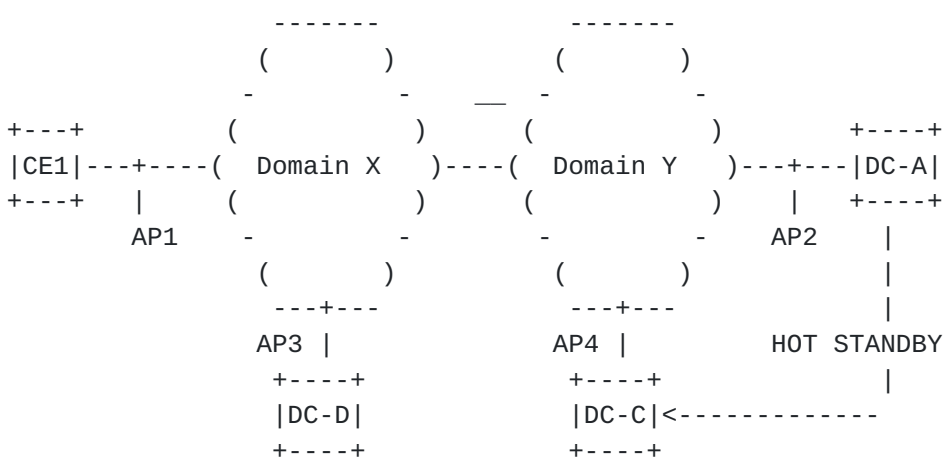

Figure 13: Pre-planned end point migration

8.2. On the Fly End Point Migration

Compared to pre-planned end point migration, on the fly end point selection is dynamic in that the migration is not pre-planned 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 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.

9. Advanced Topic

This section describes how ACTN architecture supports some deployment scenarios. See Appendix A for details on MDSC and PNC functions integrated in Service/Network Orchestrator and Appendix B for IP + Optical with L3VPN service.

10. Manageability Considerations

The objective of ACTN is to manage traffic engineered resources, and provide a set of mechanism to allow clients to request virtual connectivity across server network resources. ACTN will support multiple clients each with its own view of and control of the server network, the network operator will need to partition (or "slice") their network resources, and manage them 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 traffic engineered resources. The management requirements may be categorized as follows:

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

10.1. Policy

It is expected that a policy will be an important aspect of ACTN control and management. Typically, policies are used via the components and interfaces, during deployment of the service, to ensure that the service is compliant with agreed policy factors (often described in Service Level Agreements - SLAs), these include, but are not limited to: connectivity, bandwidth, geographical transit, technology selection, security, resilience, and economic cost.

Depending on the deployment the ACTN deployment 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:

- 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.

- A global policy might constrain certain customer types (or specific customer applications) to only use certain MDSCs, and be restricted to physical network types managed by the PNCs. A global policy agent would govern these types of policies.

This 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.

10.2. Policy applied to the Customer Network Controller

A virtual network service for a customer application will be requested from the CNC. It will reflect the application requirements and specific service policy 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.

10.3. Policy applied to the Multi Domain Service Coordinator

A key objective of the MDSC is to help the customer express the application connectivity request via its CNC as set of desired business needs, therefore policy will play an important role.

Once authorised, 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 connectivity end-points, use of these end-points should be defined as a policy expression when setting up or augmenting virtual network services. Ensuring that permissible end-points are defined for CNCs and applications will require the MDSC to maintain a registry of permissible connection points for CNCs and application types.

It may also be necessary for the MDSC to resolve policy conflicts, or at least flag any issues to administrator of the MDSC itself. Conflicts may occur when virtual network service optimization criterion 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 confidentially policy requirement of specific type of end-to-end service. This type of situation may be resolved using hard and soft policy constraints.

10.4. Policy applied to the Physical Network Controller

The PNC is responsible for configuring the network elements, monitoring physical network resources, and exposing connectivity (direct or abstracted) to the MDSC. It is therefore expected that policy will dictate what connectivity information will be exported between the PNC, via the MDSC-PNC Interface (MPI), and MDSC.

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 become 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 virtual network service performance, security and resilience guarantees. This will likely be implemented via a local policy agent and subsequent protocol methods.

11. Security Considerations

The ACTN framework described in this document defines key components and interfaces for managed traffic engineered networks. Securing the request and control of resources, confidentially 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 as a rule implementations should consider encrypting data that flow between components, especially when they are implemented at remote nodes, regardless if these are external or internal network interfaces.

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

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

- Interface between the Multi Domain Service Coordinator and Physical 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.

11.1. Interface between the Customer Network Controller and Multi Domain Service Coordinator (MDSC), CNC-MDSC Interface (CMI)

The role of the MDSC is to detach the network and service control from underlying technology to help the customer express the network as desired by business needs. It should be noted that data stored by the MDSC will reveal details of the virtual network services, and which CNC and application is consuming the resource. The data stored must therefore be considered as a candidate for encryption.

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

A CNC-MDSC protocol interface will likely be an external protocol interface. Again, 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.

11.2. Interface between the Multi Domain Service Coordinator and Physical Network Controller (PNC), MDSC-PNC Interface (MPI)

The function of the Physical Network Controller (PNC) is to configure network elements, provide performance and monitoring functions of the physical elements, and export physical topology (full, partial, or abstracted) to the MDSC.

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.

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.

12. References

12.1. Informative References

[RFC2702] Awduche, D., et. al., "Requirements for Traffic Engineering Over MPLS", RFC 2702, September 1999.

[RFC4026] L. Andersson, T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[RFC4208] G. Swallow, J. Drake, H.Ishimatsu, Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", IETF RFC 4655, August 2006.

[RFC5654] Niven-Jenkins, B. (Ed.), D. Brungard (Ed.), and M. Betts (Ed.), "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.

[RFC7149] Boucadair, M. and Jacquenet, C., "Software-Defined Networking: A Perspective from within a Service Provider Environment", RFC 7149, March 2014.

[RFC7926] A. Farrel (Ed.), "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", RFC 7926, July 2016.

[GMPLS] Manning, E., et al., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[ONF-ARCH] Open Networking Foundation, "SDN architecture", Issue 1.1, ONF TR-521, June 2016.

[RFC7491] King, D., and Farrel, A., "A PCE-based Architecture for Application-based Network Operations", RFC 7491, March 2015.

[Transport NBI] Busi, I., et al., "Transport North Bound Interface Use Cases", draft-tnbidt-ccamp-transport-nbi-use-cases, work in progress.

[ACTN-Abstraction] Y. Lee, et al., "Abstraction and Control of TE Networks (ACTN) Abstraction Methods", draft-lee-teas-actn-abstraction, work in progress.

13. Contributors

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

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

Khuzema Pithewan
Infinera
Email: kpithewan@infinera.com

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

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, USA
Phone: (469)277-5838
Email: leeyoung@huawei.com

Luyuan Fang
Microsoft
Email: luyuanf@gmail.com

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

Sergio Belotti
Alcatel Lucent
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

APPENDIX A - Example of MDSC and PNC functions integrated in Service/Network Orchestrator

This section provides an example of a possible deployment scenario, in which Service/Network Orchestrator can include a number of functionalities, among which, in the example below, PNC functionalities for domain 2 and MDSC functionalities to coordinate the PNC1 functionalities (hosted in a separate domain controller) and PNC2 functionalities (co-hosted in the network orchestrator).

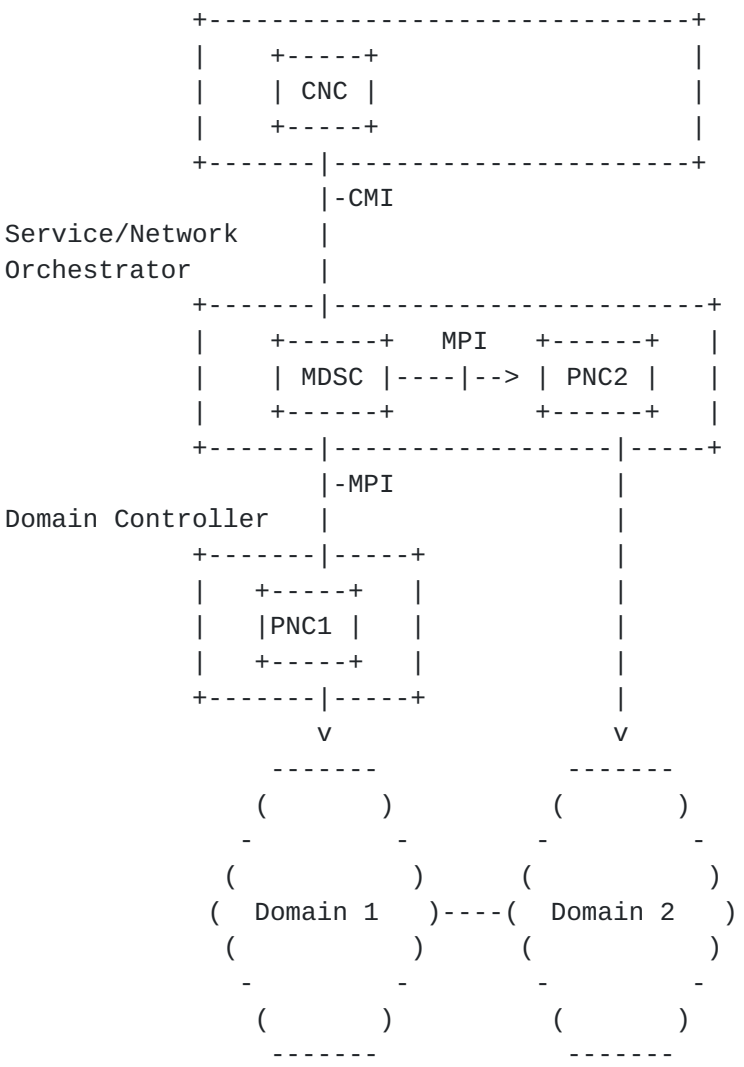

APPENDIX B - Example of IP + Optical network with L3VPN service

This section provides a more complex deployment scenario in which ACTN hierarchy is deployed to control a multi-layer network via an IP/MPLS PNC and an Optical PNC. The scenario is further enhanced by the introduction of an upper layer service configuration (e.g.
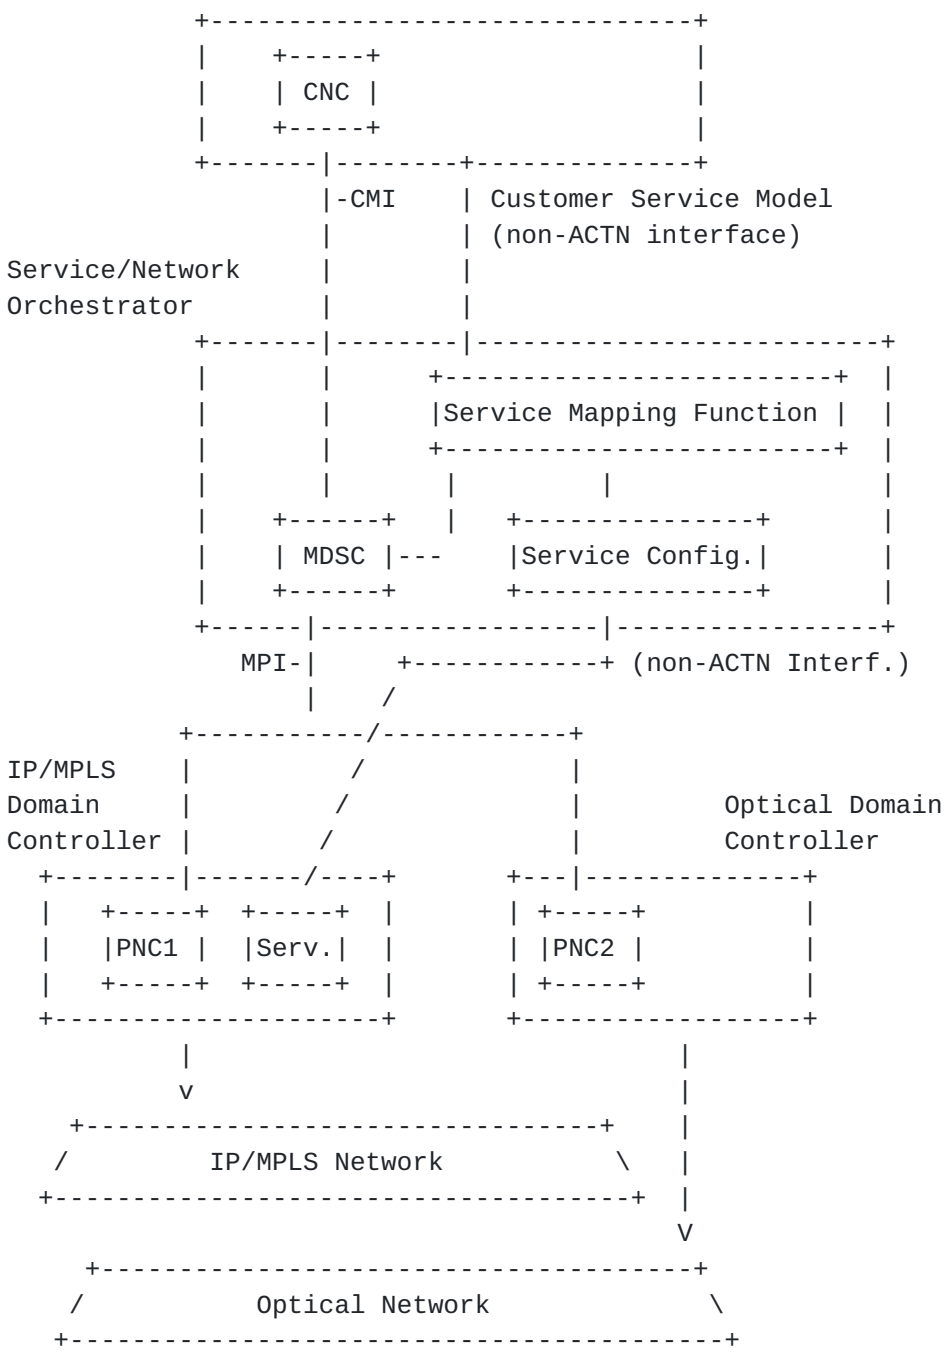
L3VPN). The provisioning of the L3VPN service is outside ACTN scope but it is worth showing how the two parts are integrated for the end to end service fulfilment. An example of service configuration function in the Service/Network Orchestrator is discussed in [I-D.dhjain-bess-bgp-l3vpn-yang].