Network Working Group
Internet Draft

Daniele Ceccarelli
Ericsson

Intended status: Informational
Expires: November 2014

Luyuan Fang
Microsoft

Young Lee
Huawei

Diego Lopez
Telefonica

May 30, 2014

Framework for Abstraction and Control of Transport Networks

draft-ceccarelli-actn-framework-02.txt

Abstract

This draft provides a framework for abstraction and control of transport networks.

Table of Contents

1. Terminology

This document uses the terminology defined in [RFC4655], and [RFC5440].

CVI Customer-VNC Interface

PCA Path Computation Agent

PNC Physical Network Controller

VL Virtual Link

VN Virtual Network

VNM Virtual Network Mapping

VNC Virtual Network Controller

VNE Virtual Network Element

VNS Virtual Network Service

VPI VNC-PNC Interface

2. Introduction

Transport 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.

Transport networks in this draft refer to a set of different type of connection-oriented networks, primarily Connection-Oriented Circuit Switched (CO-CS) networks and Connection-Oriented Packet Switched (CO-PS) networks. This implies that at least the following transport networks are in scope of the discussion of this draft: L1 optical networks (e.g., OTN and WDM), MPLS-TP, MPLS-TE, as well as other emerging connection-oriented networks such as Segment Routing (SR). One of the characteristics of these network types is the ability of dynamic provisioning and traffic engineering such that resource guarantee can be provided to their clients.

One of the main drivers for Software Defined Networking (SDN) is a physical separation 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 PCE [PCE] for TE-based transport networks. In fact, in transport networks such separation of data and control plane was dictated at the onset due to the very different natures of the data plane (circuit switched TDM or WDM) and a packet switched control plane. The physical separation of the control plane and the data plane is a major step towards allowing operators to gain the full control for optimized network design and operation. Moreover, another advantage of SDN is its logically centralized control regime that allows a global view of the underlying network under its control. Centralized control in SDN helps improve network resources utilization from a distributed network control. For TE-based transport network control, PCE is essentially equivalent to a logically centralized control for path computation function.

As transport networks evolve, the need to provide network abstraction has emerged as a key requirement for operators; this implies in effect the virtualization of network resources so that the network is "sliced" for different uses, applications, services, and customers each being given a different partial view of the total topology and each considering that it is operating with or on a single, stand-alone and consistent network. Moreover, particular attention needs to be paid to the multi-domain case, where ACTN can facilitate virtual network operation via the creation of a single virtualized network. 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.

Network virtualization, in general, refers to allowing the customers to utilize a certain amount of network resources as if they own them and thus control their allocated resources in a way most optimal with higher layer or application processes. This empowerment of customer control facilitates introduction of new services and applications as the customers are permitted to create, modify, and delete their virtual network services. The level of virtual control given to the customers can vary from a tunnel connecting two end-points to virtual network elements that consist of a set of virtual nodes and virtual links in a mesh network topology. More flexible, dynamic customer control capabilities are added to the traditional VPN along with a customer specific virtual network view. Customers control a view of virtual network resources, specifically allocated to each one of them. This view is called an abstracted network topology. Such a view may be specific to the set of consumed services as well as to a particular customer. As the customer controller is envisioned to support a plethora of distinct applications, there would be another level of virtualization from the customer to individual applications.

The virtualization framework described in this draft is named Abstraction and Control of Transport Network (ACTN) and facilitates:

- Abstraction of the underlying network resources to higher-layer applications and users (customers);

- Slicing infrastructure to connect multiple customers to meet specific application and users requirements;

- Creation of a virtualized environment allowing operators to view and control multi-subnet multi-technology networks into a single virtualized network;

- A computation scheme, via an information model, to serve various customers that request network connectivity and properties associated with it;

- A virtual network controller that adapts customer requests to the virtual resources (allocated to them) to the supporting physical network control and performs the necessary mapping, translation, isolation and security/policy enforcement, etc.;

- The coordination of the underlying transport topology, presenting it as an abstracted topology to the customers via open and programmable interfaces.

The organization of this draft is as follows. Section 3 provides a discussion for a Business Model, Section 4 a Computation Model, Section 5 a Control and Interface model and Section 6 Design Principles.

3. Business Model of ACTN

The traditional Virtual Private Network (VPN) and Overlay Network (ON) models are built on the premise that one single network provider provides all virtual private or overlay networks to its customers. This model is simple to operate but has some disadvantages in accommodating the increasing need for flexible and dynamic network virtualization capabilities.

The ACTN model is built upon entities that reflect the current landscape of network virtualization environments. There are three key entities in the ACTN model [REF probl stat]:

- Customers
- Service Providers
- Network Providers

3.1. Customers

Within the ACTN framework, different types of customers may be taken into account depending on the type of their resource needs, on their number and type of access. As 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 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 service and internet access.

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 demand significantly varying in time and from customer to customer. This is one of reasons why a bundled services offer is not enough but it is desirable to provide each of them with customized virtual network services. As customers are geographically spread over multiple network provider domains, the necessary control and data interfaces to support such customer needs is no longer a single interface between the customer and one single network provider. With this premise, customers have to interface multiple providers to get their

end-to-end network connectivity service and the associated topology information. Customers may have to support multiple virtual network services with differing service objectives and QoS requirements. For flexible and dynamic applications, customers may want to control their allocated virtual network resources in a dynamic fashion. To allow that, customers should be given an abstracted view of topology on which they can perform the necessary control decisions and take the corresponding actions.

Customers of a given service provider can in turn offer a service to other customers in a recursive way. An example of recursiveness with 2 service providers is shown below.

- Customer (of service B)
- Customer (of service A) & Service Provider (of service B)
- Service Provider (of service A)
- Network Provider

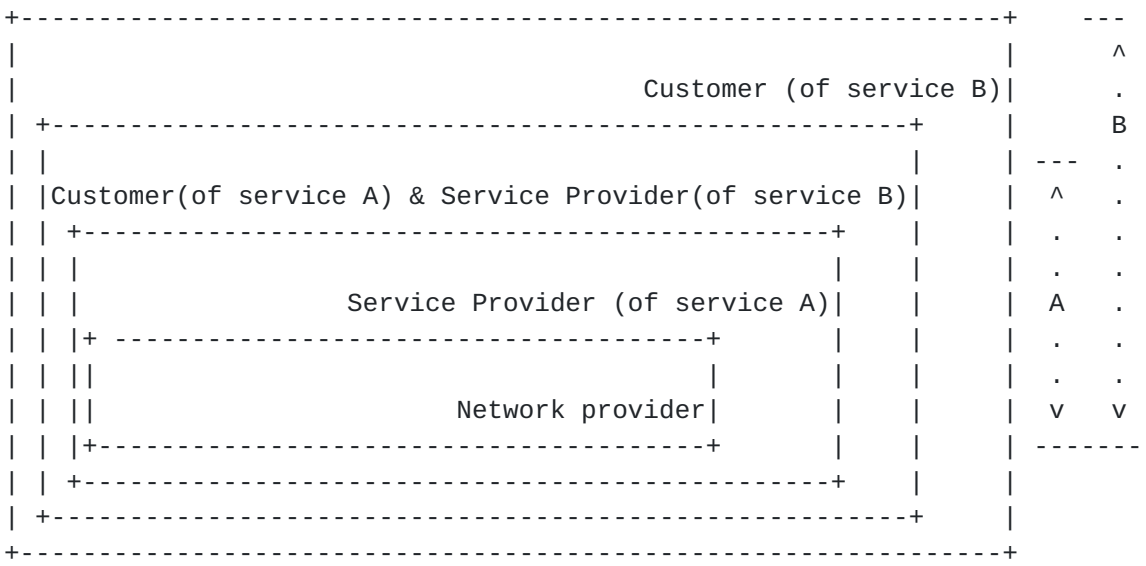

3.2. Service Providers

Service providers are the providers of virtual network services to their customers. Service providers may or may not own physical network resources. When a service provider is the same as the network provider, this is similar to traditional VPN models. This model works well when the customer maintains a single interface with a single provider. When customer location spans across 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 service providers directly interface their customers. In this case, service providers themselves are 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.

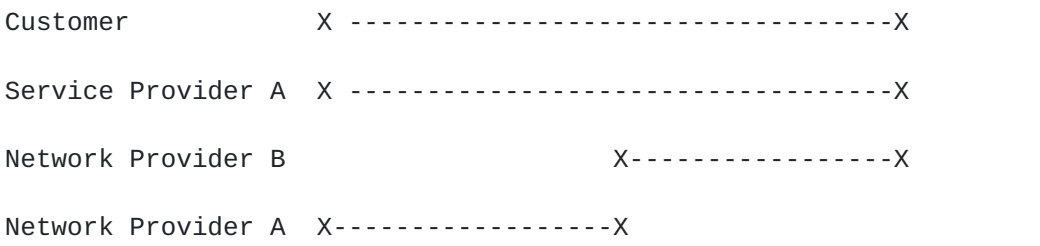

The ACTN network model is predicated upon this three tier model and is summarized in figure below:

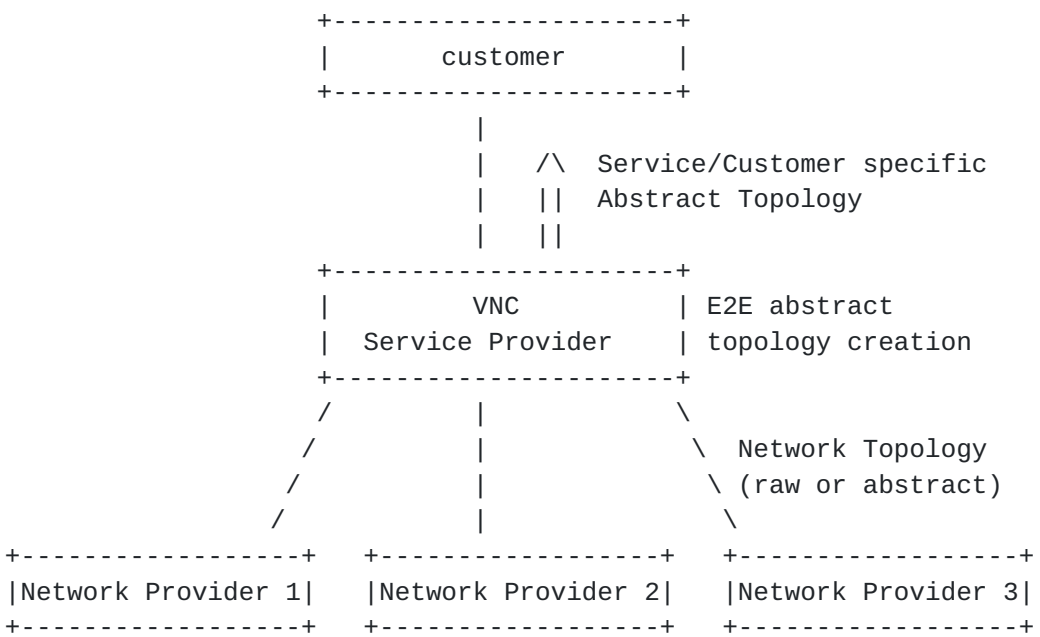

Figure 1: Three tier model.

There can be multiple types of service providers.

- Data Center providers: can be viewed as a service provider type as they own and operate data center resources to various WAN clients, they can lease physical network resources from network providers.
- Internet Service Providers (ISP): can be a service provider 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.

3.3. Network Providers

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

4. Multi domain management

Network operators build and operate multi-domain networks and these domains may be technology, administrative or vendor specific (vendor islands). Interoperability for dealing with different domains is a perpetual problem for operators. Due to these issues, new service introduction, often requiring connections that traverse multiple domains, need significant planning, and several manual operations to interface different vendor equipment and technology.

The creation of a virtualized environment allowing operators to view and control multi-subnet multi-technology networks into a single virtualized network highly facilitates network operators and will accelerate rapid service deployment of new services, including more dynamic and elastic services, and improve overall network operations and scaling of existing services.

Figure 2 depict a common scenario in which two different domains can be managed by a single VNC, which is in charge of acting as orchestrator between them and presenting them as a single entity to its clients.

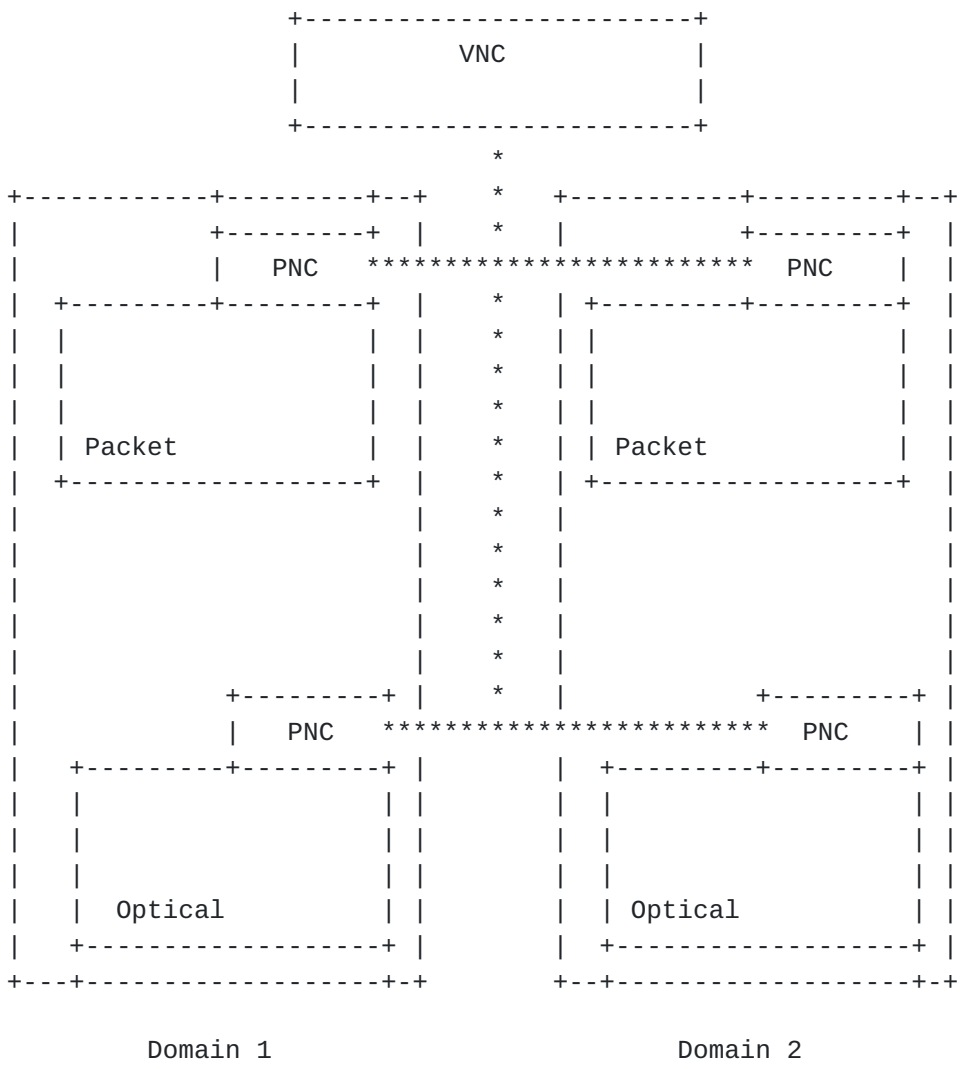

Figure 2: Multi domain management

In this figure the case of packet and optical domains controlled by different PNCs is shown but any combination can be considered, like e.g. a single PNC controlling the packet+optical domain 1 and different PNCs for domains 2.

5. Computation Model of ACTN

This section discusses ACTN framework from a computational point of view. As multiple customers run their virtualized network on a

shared infrastructure, making efficient use of the underlying resources requires effective computational models and algorithms. This general problem space is known as Virtual Network Mapping or Embedding (VNM or VNE). [Editors's note(Put some reference)].

As VNM/VNE issues impose some additional compute models and algorithms for virtual network path computation, this section discusses key issues and constraints for virtual network path computation.

5.1. Request Processing

This is concerned about whether a set of customer requests for VN creation can be dealt with in real-time or off line, and in the latter case, simultaneously or not. This depends on the nature of applications the customer support. There are applications and use cases, like e.g. management of catastrophic events or real time SLA negotiation, that require a real-time VN creation. If the customer does not require real-time instantiation of VN creation, the computation engine can process a set of VN creation requests simultaneously to improve network efficiency.

5.2. Types of Network Resources

When a customer makes a VN creation request to the substrate network, what kind of network resources is consumed is of concern of both the customer and service/network providers. The customer needs to put constraints (e.g. TE parameters, resiliency) for the provisioning of the VN, while the service and network providers need to choose which resources meet such constraints and possibly have fewest impact on the capability of serving other customers. For transport network virtualization, the network resource consumed is primarily network bandwidth that the required paths would occupy on the physical link(s). However, there may be other resource types such as CPU and memory that need to be considered for certain applications. These resource types shall be part of the VN request made by the customer.

5.3. Accuracy of Network Resource Representation

As the underlying transport network in itself may consist of a layered structure, it is a challenge how to represent these underlying physical network resources and topology into a form that can be reliably used by the computation engine that assigns customer requests into the physical network resource and topology.

5.4. Resource Sharing and Efficiency

Related to the accuracy of network resource representation is resource efficiency. As a set of independent customer VN is created and mapped onto physical network resources, the overall network resource utilization is the primary concern of the network provider.

In order to provide an efficient utilization of the resources of the provider network, it should be possible to share given physical resources among a number of different VNs. Whether a virtual resource is sharable among a set of VNs (and hence of customers) is something the service provider needs to agree with each customer. Preemption and priority management are tools that could help provide an efficient sharing of physical resources among different VNs.

5.5. Guarantee of Client Isolation

While network resource sharing across a set of customers for efficient utilization is an important aspect of network virtualization, customer isolation has to be guaranteed. Admissions of new customer requests or any changes of other existing customer VNs must not affect any particular customer in terms of resource guarantee, security constraints, and other performance constraints.

5.6. Computing Time

Depending on the nature of applications, how quickly a VN is instantiated from the time of request is an important factor. For dynamic applications that require instantaneous VN creation or VN changes from the existing one, the computation model/algorithm should support this constraint.

5.7. Admission Control

To coordinate the request process of multiple customers, an admission control will help maximize an overall efficiency.

5.8. Path Constraints

There may be some factors of path constraints that can affect the overall efficiency. Path Split can lower VN request blocking if the underlying network can support such capability. A packet-based TE network can support path split while circuit-based transport may have limitations.

Path migration is a technique that allows changes of nodes or link assignments of the established paths in an effort to accommodate new

requests that would not be accepted without such path migration(s). This can improve overall efficiency, yet additional care needs to be applied to avoid any adverse impacts associated with changing the existing paths.

Re-optimization is a global process to re-shuffle all existing path assignments to minimize network resource fragmentation. Again, an extra care needs to be applied for re-optimization.

6. Control and Interface Model for ACTN

This section provides a high-level control and interface model of ACTN.

6.1. A High-level ACTN Control Architecture

To allow virtualization, the network has to provide open, programmable interfaces, in which customer applications can create, replace and modify virtual network resources in an interactive, flexible and dynamic fashion while having no impact on other customers. Direct customer control of transport network elements over existing interfaces (control or management plane) is not perceived as a viable proposition for transport network providers due to security and policy concerns among other reasons. In addition, as discussed in the previous section, the network control plane for transport networks has been separated from data plane and as such it is not viable for the customer to directly interface with transport network elements.

While the current network control plane is well suited for control of physical network resources via dynamic provisioning, path computation, etc., a virtual network controller needs to be built on top of physical network controller to support network virtualization. On a high-level, virtual network control refers to a mediation layer that performs several functions:

- Computation of customer resource requests into virtual network paths based on the global network-wide abstracted topology;

- Mapping and translation of customer virtual network slices into physical network resources;

- Creation of an abstracted view of network slices allocated to each customer, according to customer-specific objective functions, and to the customer traffic profile.

In order to facilitate the above-mentioned virtual control functions, the virtual network controller (aka., "virtualizer") needs to maintain two interfaces:

- One interface with the physical network controller functions which is termed as the VNC-PNC Interface (VPI).

- Another interface with the customer controller for the virtual network, which is termed as Client-VNC Interface (CVI).

Figure 2 depicts a high-level control and interface architecture for ACTN.

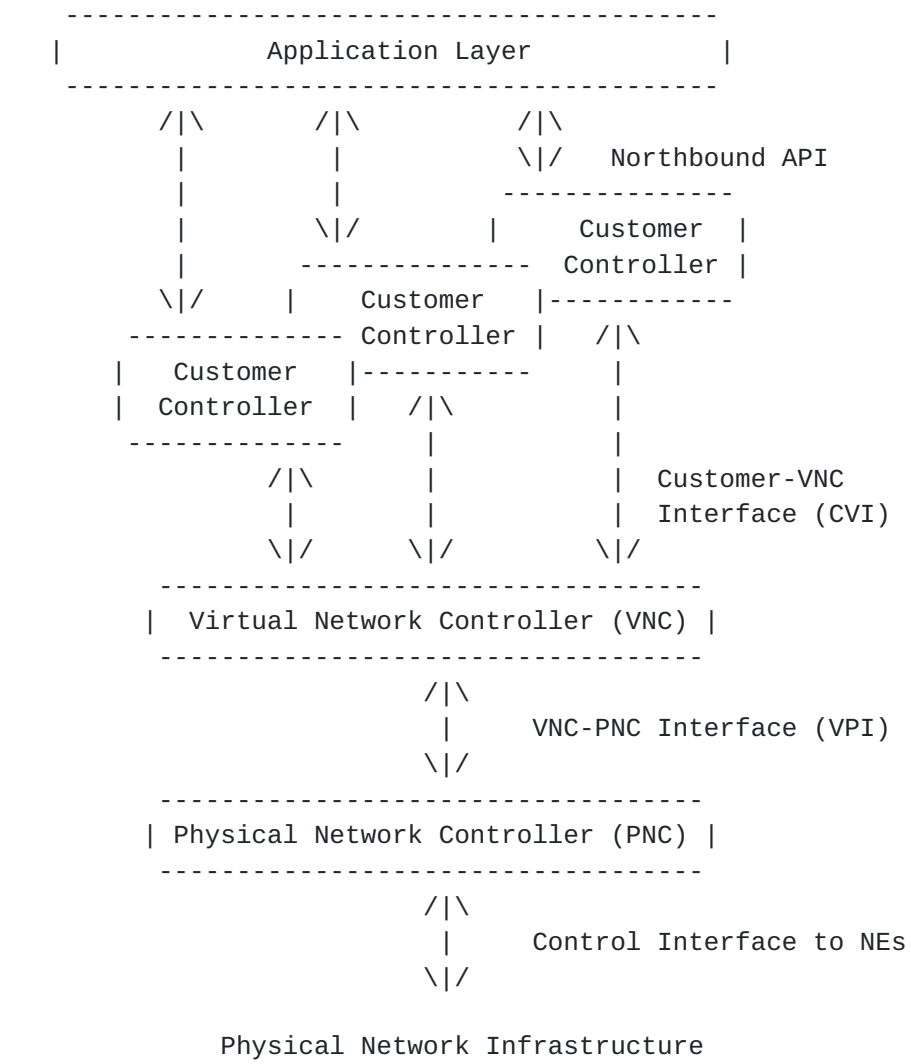

Figure 3: Control and Interface Architecture for ACTN.

Figure 2 shows that there are multiple customer controllers, which are independent to one another, and that each customer supports various business applications over its NB API. There are layered client-server relationships in this architecture. As various applications are clients to the customer controller, it also becomes itself a client to the virtual network controller. Likewise, the virtual network controller is also a client to the physical network controller. This layered relationship is important in the protocol definition work on the NB API, the CVI and VPI interfaces as this allows third-party software developers to program client controllers and virtual network controllers independently.

There are several ways in which the Physical Network Controller manages the network elements, e.g. via management protocols, PCEP+GMPLS, or any other type of protocol. In other words the ACTN architecture both applies to physical networks controlled by control plane protocols (e.g. PCEP+GMPLS) or management plane protocols (e.g. SNMP).

6.2. Customer Controller

A Virtual Network Service is instantiated by the customer controller via the CVI. As the customer controller directly interfaces the application stratum, it understands multiple application requirements and their service needs. It is assumed that the customer controller and the VNC have a common knowledge on the end-point interfaces based on their business negotiation prior to service instantiation. End-point interfaces refer to customer-network physical interfaces that connect customer premise equipment to network provider equipment. Figure 5 shows an example physical network topology that supports multiple customers. In this example, customer A has three end-points A.1, A.2 and A.3. The interfaces between customers and transport networks are assumed to be 40G OTU links. For simplicity's sake, all network interfaces are assumed to be 40G OTU links and all network ports support ODU switching and grooming on the level of ODU1 and ODU2. Customer controller for A provides its traffic demand matrix that describes bandwidth requirements and other optional QoS parameters (e.g., latency, diversity requirement, etc.) for each pair of end-point connections.

6.3. Virtual Network Controller

The virtual network controller sits between the consumer controller (the one issuing connectivity requests) and the physical network controller (the one managing the resources). The Virtual Network controller can be collocated with the physical network controller, especially in those cases where the service provider and the network provider are the same entity.

The virtual network controller is composed by the following functional components:

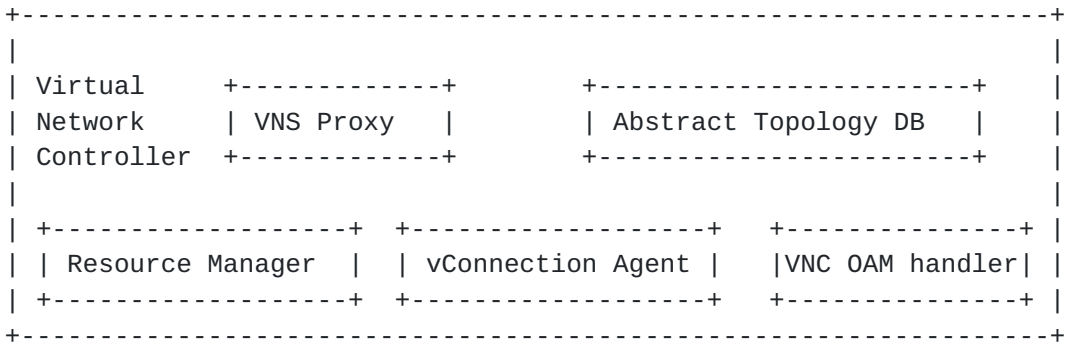

- VNS proxy: The VNS proxy is the functional module in charge of performing policy management and AAA (Authentication, authorization, and accounting) functions. It is the one that receives that VN instantiation and resource allocation requests from the Customer controllers.
- Abstract Topology DB: This is the database where the abstract topology, generated by the VNC or received from the PNC, is stored. A different VN instance is kept for every different customer.
- Resource Manager: The resource manager is in charge of receiving VNS instantiation requests from the customer controller and, as a consequence, triggering a concurrent path computation request to the PCE in the PNC based on the traffic matrix. The Resource manager is also in charge of generating the abstract topology for the customer.
- vConnection Agent: This module is in charge of mapping VN setup commands into network provisioning requests to the PNC.
- VNC OAM handler: The VNC OAM handler is the module that is in charge of understanding how the network is operating, detecting faults and reacting to problems related to the abstract topology.

6.4. Physical Network Controller

The physical network controller is the one in charge of configuring the network elements, monitoring the physical topology of the network and passing it, either raw or abstracted, to the VNC.

It is composed by the following functional components:

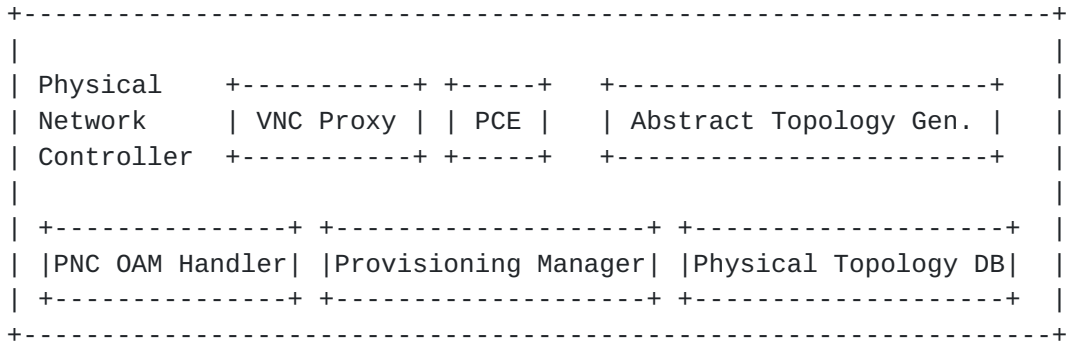

- VNC proxy: The VNC proxy is the functional module in charge of performing policy management and AAA (Authentication, authorization, and accounting) functions on requests coming from the VNC.
- PCE: This is the stateful PCE performing the path computation over the physical topology and that provides the vConnection agent with the network topology.
- Abstract topology generator: the network topology can be passed to the VNC as raw or abstract. In case the topology is passed as abstract topology, this module is in charge of generating it from the physical topology DB. The module is optional.
- ONC OAM handler: it verifies that connections exists, implements monitoring functions to see if failures occurs. It is the proxy to an OSS/NMS system but does not duplicate any of OSS/NMS functionalities.
- Physical topology database: The physical topology database is mainly composed by two databases: the Traffic Engineering Database (TED) and the LSP Database (LSP-DB).
- Provisioning manager: The Provisioning Manager is responsible for making or channeling requests for the establishment of LSPs. This may be instructions to the control plane running in the networks, or may involve the programming of individual network devices. In the latter case, the Provisioning Manager may act as an OpenFlow Controller [ONF].

6.5. Abstracted Topology

There are two levels of abstracted topology that needs to be maintained and supported for ACTN. Customer-specific Abstracted Topology refers to the abstracted view of network resources allocated (shared or dedicated) to the customer. The granularity of this abstraction varies depending on the nature of customer applications. Figure 3 illustrates this.

Figure 3 shows how three independent customers A, B and C provide its respective traffic demand matrix to the VNC. The physical network topology shown in Figure 2 is the provider's network topology generated by the PNC topology creation engine such as the link state database (LSDB) and Traffic Engineering DB (TEDB) based on control plane discovery function. This topology is internal to PNC and not available to customers. What is available to them is an abstracted network topology (a virtual network topology) based on the negotiated level of abstraction. This is a part of VNS instantiation between a client control and VNC.

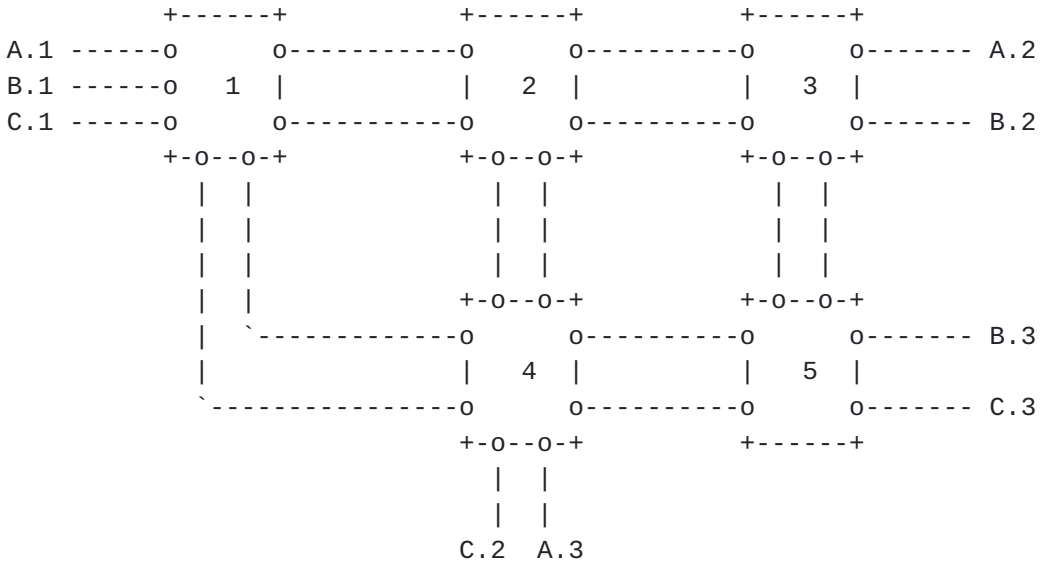

Traffic Matrix for Customer A

	A.1	A.2	A.3
A.1	-	20G	20G
A.2	20G	-	10G
A.3	20G	10G	-

Traffic Matrix for Customer B

	B.1	B.2	B.3
B.1	-	40G	40G
B.2	40G	-	20G
B.3	40G	20G	-

Traffic Matrix for Customer C

	C.1	C.2	C.3
C.1	-	20G	20G
C.2	20G	-	10G
C.3	20G	10G	-

Figure 4: Physical network topology shared with multiple customers

Figure 5 depicts illustrative examples of different level of topology abstractions that can be provided by the VNC topology abstraction engine based on the physical topology base maintained by the PNC. The level of topology abstraction is expressed in terms of the number of virtual network elements (VNEs) and virtual links (VLs). For example, the abstracted topology for customer A shows there are 5 VNEs and 10 VLs. This is by far the most detailed topology abstraction with a minimal link hiding compared to other abstracted topologies in Figure 4.

(a) Abstracted Topology for Customer A (5 VNEs and 10 VLs)

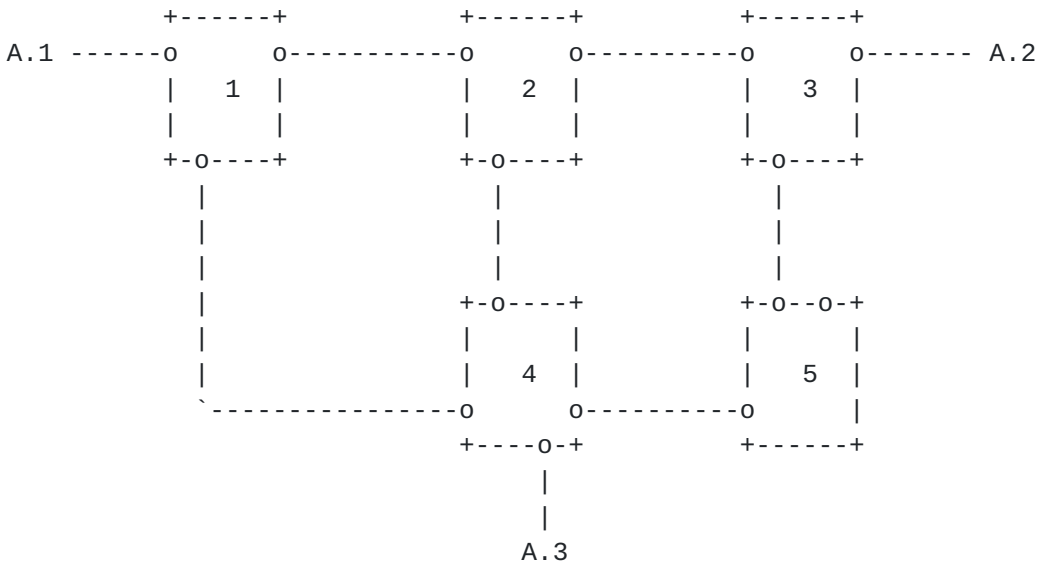

(b) Abstracted Topology for Customer B (3 VNEs and 6 VLs)

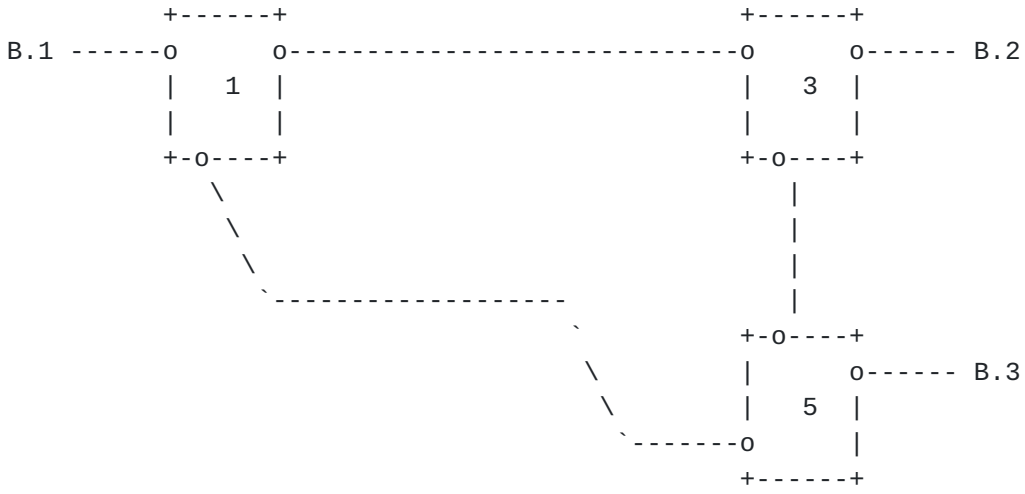

(c) Abstracted Topology for Customer C (1 VNE and 3 VLs)

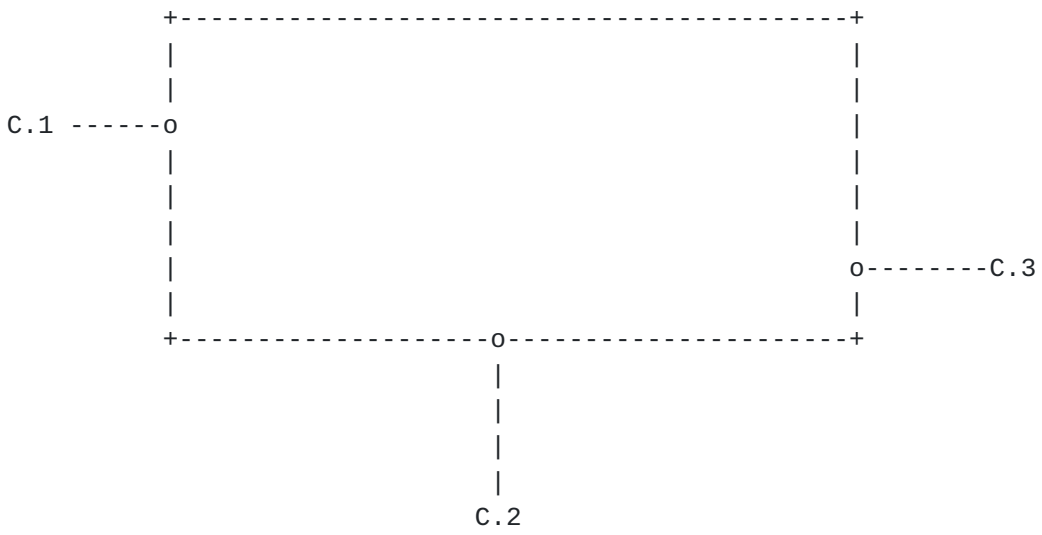

Figure 5: Topology Abstraction Examples for Customers

As different customers have different control/application needs, abstracted topologies for customers B and C, respectively show a much higher degree of abstraction. The level of abstraction is determined by the policy (e.g., the granularity level) placed for the customer and/or the path computation results by the PCE operated by the PNC. The more granular the abstraction topology is, the more control is given to the customer controller. If the customer controller has applications that require more granular control of virtual network resources, then the abstracted topology shown for customer A may be the right abstraction level for such controller. For instance, if the customer is a third-party virtual service broker/provider, then it would desire much more sophisticated control of virtual network resources to support different application needs. On the other hand, if the customer were only to support simple tunnel services to its applications, then the abstracted topology shown for customer C (one VNE and three VLs) would suffice.

6.6. Workflows of ACTN Control Modules

Figure 5 shows workflows across the customer controller, VNC and PNC for the VNS instantiation, topology exchange, and VNS setup.

The customer controller "owns" a VNS and initiates it by providing the instantiation identifier with a traffic demand matrix that includes path selection constraints for that instance. This VNS instantiation request from the Customer Controller triggers a path computation request by the Resource Manager in the VNC after VNC's proxy's interlay of this request to the Resource Manager. PCA sends a concurrent path computation request that is converted according to the traffic demand matrix as part of the VNS instantiation request from the Customer Controller. Upon receipt of this path computation request, the PCE in the PNC block computes paths and updates network topology DB and informs the Resource Manager of the VNC of the paths and topology updates.

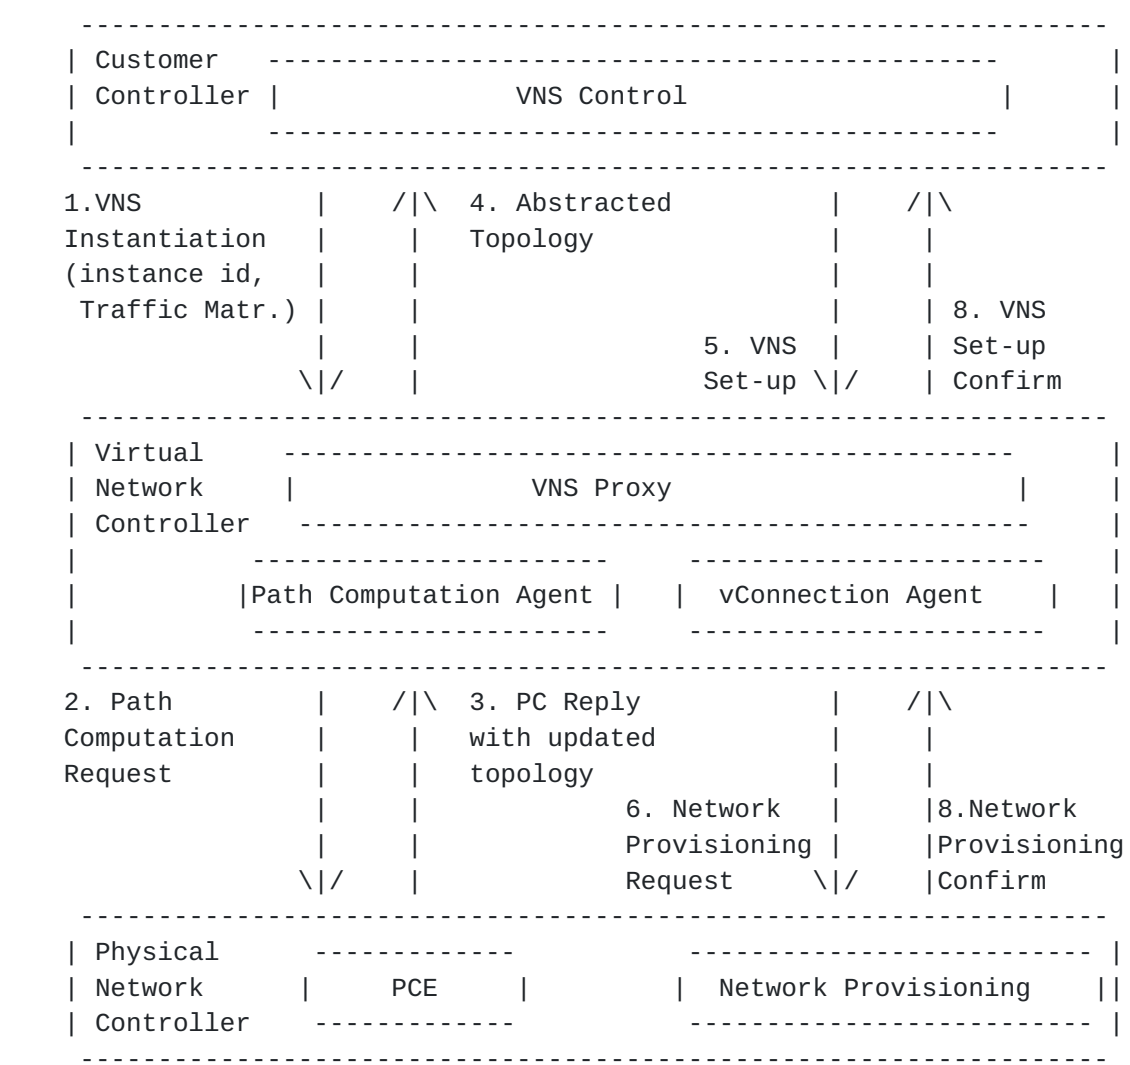

Figure 6. Workflows across Customer Controller, VNC and PNC

It is assumed that the PCE in PNC is a stateful PCE [PCE-S]. PCA abstracts the physical network topology into an abstracted topology for the customer based on the agreed-upon granularity level. The abstracted topology is then passed to the VNS control of the Customer Controller. This controller computes and assigns virtual network resources for its applications based on the abstracted topology and creates VNS setup command to the VNC. The VNC vConnection module turns this VN setup command into network provisioning requests over the network elements.

6.7. Programmability of the ACTN Interfaces

From Figures 2 and 5, we have identified several interfaces that are of interest of the ACTN model. More precisely, ACTN concerns the following interfaces:

- Customer-VNC Interface (CVI): an interface between a customer controller and a virtual network controller.

- VNC-PNC Interface (VPI): an interface between a virtual network controller and a physical network controller.

The NBI interfaces and direct control interfaces to NEs are outside of the scope of ACTN.

The CVI interface should allow programmability, first of all, to the customer so they can create, modify and delete virtual network service instances. This interface should also support open standard information and data models that can transport abstracted topology.

The VPI interface should allow programmability to service provider(s) (through VNCs) in such ways that control functions such as path computation, provisioning, and restoration can be facilitated. Seamless mapping and translation between physical resources and virtual resources should also be facilitated via this interface.

7. Design Principles of ACTN

7.1. Network Security

Network security concerns are always one of the primary principles of any network design. ACTN is no exception. Due to the nature of heterogeneous VNs that are to be created, maintained and deleted flexibly and dynamically and the anticipated interaction with

physical network control components, secure programming models and interfaces have to be available beyond secured tunnels, encryption and other network security tools.

7.2. Privacy and Isolation

As physical network resources are shared with and controlled by multiple independent customers, isolation and privacy for each customer has to be guaranteed.

Policy should be applied per client.

7.3. Scalability

As multiple VNs need to be supported seamlessly, there are potentially several scaling issues associated with ACTN. The VN Controller system should be scalable in supporting multiple parallel computation requests from multiple customers. New VN request should not affect the control and maintenance of the existing VNs. Any VN request should also be satisfied within a time-bound of the customer application request.

Interfaces should also be scalable as a large amount of data needs to be transported across customers to virtual network controllers and across virtual network controllers and physical network controllers.

7.4. Manageability and Orchestration

As there are multiple entities participating in network virtualization, seamless manageability has to be provided across every layer of network virtualization. Orchestration is an important aspect of manageability as the ACTN design should allow orchestration capability.

ACTN orchestration should encompass network provider multi-domains, relationships between service provider(s) and network provider(s), and relationships between customers and service/network providers.

Ease of deploying end-to-end virtual network services across heterogeneous network environments is a challenge.

7.5. Programmability

As discussed earlier in Section 5.5, the ACTN interfaces should support open standard interfaces to allow flexible and dynamic virtual service creation environments.

7.6. Network Stability

As multiple VNs are envisioned to share the same physical network resources, combining many resources into one should not cause any network instability. Provider network oscillation can affect readily both on virtual networks and the end-users.

Part of network instability can be caused when virtual network mapping is done on an inaccurate or unreliable resource data. Data base synchronization is one of the key issues that need to be ensured in ACTN design.

8. References

8.1. Informative References

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

[PCE-S] Crabbe, E, et. al., "PCEP extension for stateful PCE",draft-ietf-pce-stateful-pce, work in progress.

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

[NFV-AF] "Network Functions Virtualization (NFV); Architectural Framework", ETSI GS NFV 002 v1.1.1, October 2013.

9. Contributors

Authors' Addresses

Daniele Ceccarelli
Ericsson
Via Melen, 77
Genova, Italy
Email: daniele.ceccarelli@ericsson.com

Luyuan Fang
Email: luyuanf@gmail.com

Young Lee
Huawei Technologies
5340 Legacy Drive
Plano, TX 75023, USA
Phone: (469)277-5838
Email: leeyoung@huawei.com

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

Intellectual Property Statement

any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.