

Network Working Group
Internet Draft

Daniele Ceccarelli
Ericsson

Intended status: Informational
Expires: July 2014

Luyuan Fang
Microsoft

Young Lee
Huawei

Diego Lopez
Telefonica

January 3, 2014

Framework for Abstraction and Control of Transport Networks

[draft-ceccarelli-actn-framework-00.txt](#)

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 June 3, 2014.

Copyright Notice

Copyright (c) 2013 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.

Abstract

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

Table of Contents

1.	Terminology	3
2.	Introduction	3
3.	Business Model of ACTN	5
3.1.	Consumers	6
3.2.	Service Providers	6
3.3.	Network Providers	7
4.	Computation Model of ACTN	7
4.1.	Request Processing	7
4.2.	Types of Network Resources	8
4.3.	Accuracy of Network Resource Representation	8
4.4.	Resource Efficiency	8
4.5.	Guarantee of Client Isolation	8
4.6.	Computing Time	8
4.7.	Admission Control	8
4.8.	Path Constraints	9
5.	Control and Interface Model for ACTN	9
5.1.	A High-level ACTN Control Architecture	9
5.2.	Consumer Controller	11
5.3.	Abstracted Topology	12
5.4.	Workflows of ACTN Control Modules	15
5.5.	Programmability of the ACTN Interfaces	17
6.	Design Principles of ACTN	17
6.1.	Network Security	17
6.2.	Privacy and Isolation	17

6.3. Scalability	18
6.4. Manageability and Orchestration	18
6.5. Programmability	18
6.6. Network Stability	18
7. References	19
7.1. Informative References	19
8. Contributors	19
Authors' Addresses	20
Intellectual Property Statement	20
Disclaimer of Validity	21

1. Terminology

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

CVI	Consumer-VNC Interface
PNC	Physical Network Controller
VL	Virtual Link
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 WSON), 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 network control plane from the data plane has been already achieved for 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 wavelength) and a packet switched control plane. The physical separation of the control plane and the data plane is a major step toward 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 resource 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 consumers 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.

Network virtualization, in general, refers to allowing the consumers 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 consumer control facilitates introduction of new services and applications as the consumers are given to create, modify, and delete their virtual network services. The level of virtual control given to the consumers 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 consumer control capabilities are added to the traditional VPN along with a consumer specific virtual network view. Consumers 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 the particular consumer. As the consumer controller is envisioned to support a plethora of distinct applications, there would be another level of virtualization from the consumer 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 (consumers);
- Sharing of network resources, to meet specific application and users requirements;
- A computation scheme, via an information model, to serve various consumers that request network connectivity and properties associated with it;
- A virtual network controller that adapts consumer 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 consumers 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 consumers. 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:

- Consumers

- Network Providers
- Service Providers

3.1. Consumers

Consumers are the users of virtual network services. As consumers are geographically spread over multiple network provider domains, the necessary control and data interfaces to support such consumer needs is no longer a single interface between the consumer and one single network provider. With this premise, consumers have to interface multiple providers to get their end-to-end network connectivity service and the associated topology information. Consumers may have to support multiple virtual network services with differing service objectives and QoS requirements. For flexible and dynamic applications, consumers may want to control their allocated virtual network resources in a dynamic fashion. To allow that, consumers should be given an abstracted view of topology on which they can perform the necessary control decisions and take the corresponding actions.

3.2. Service Providers

Service providers are the providers of virtual network services to their consumers. 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 consumer maintains a single interface with a single provider. When consumer 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 consumers. In this case, service providers themselves are consumers 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.

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. There may be other types of service providers such as mobile virtual network operators (MVNO) that 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 consumers. The layered model proposed by this draft separates the concerns of network providers and consumers, with service providers acting as aggregators of consumer requests.

4. Computation Model of ACTN

This section discusses ACTN from a computational point of view. As multiple consumers 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). (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.

4.1. Request Processing

This is concerned about whether a set of consumer requests for VN creation can be dealt with simultaneously or not. This depends on the nature of applications the consumer support. If the consumer 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.

4.2. Types of Network Resources

When a consumer makes a VN creation request to the substrate network, what kind of network resources is consumed is of concern of both the consumer and network providers. 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 consumer.

4.3. 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 consumer requests into the physical network resource and topology.

4.4. Resource Efficiency

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

4.5. Guarantee of Client Isolation

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

4.6. 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, the computation model/algorithm should support this constraint.

4.7. Admission Control

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

4.8. 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.

5. Control and Interface Model for ACTN

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

5.1. A High-level ACTN Control Architecture

To allow virtualization, the network has to provide open, programmable interfaces in which consumer applications can create, replace and modify virtual network resources in an interactive, flexible and dynamic fashion while having no impact on other network consumers. Direct consumer 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 consumer to directly interface with transport network elements.

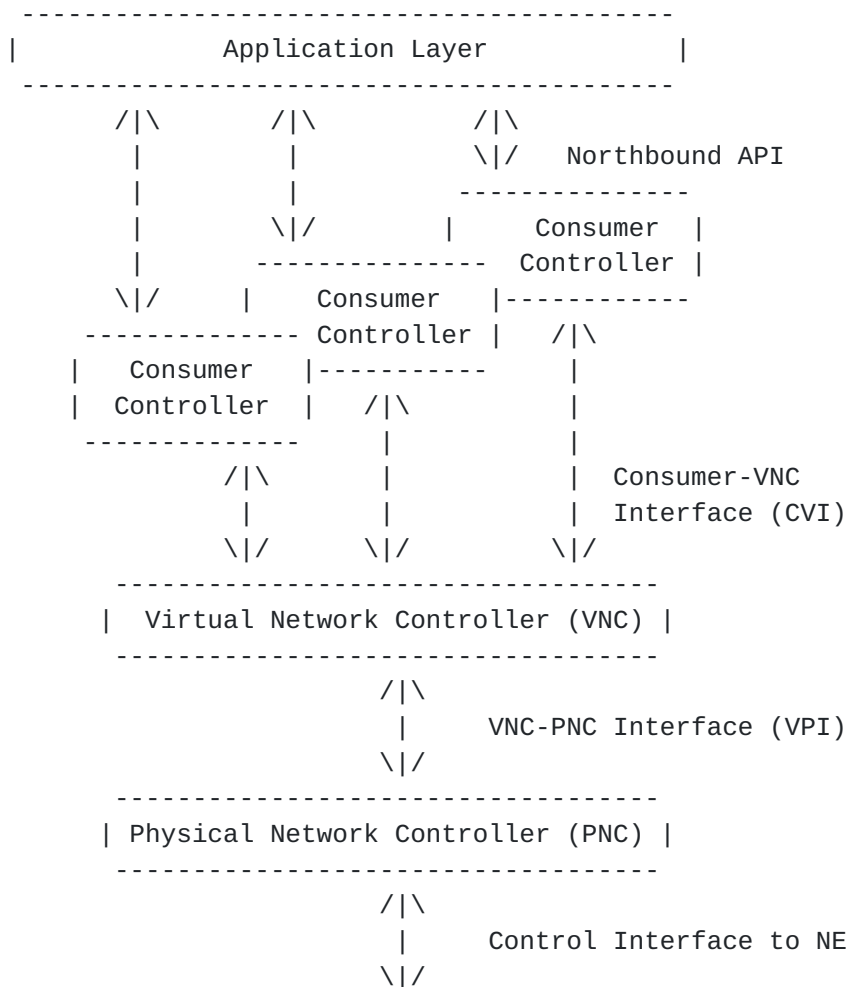
While the current network control plane is well suited for control of physical network resources via dynamic provisioning, path computation, etc., a virtual network controller needs to be built on top of physical network controller to support network virtualization. On a high-level, virtual network control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

- Computation of consumer resource requests into virtual network paths based on the global network-wide abstracted topology;
- Mapping and translation of consumer virtual network slices into physical network resources;
- Creation of an abstracted view of network slices allocated to each consumer, according to consumer-specific objective functions, and to the consumer 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 assumed by MPLS/GMPLS and PCE, which is termed as the VNC-PNC Interface (VPI);
- Another interface with the consumer controller for the virtual network, which is termed as Client-VNC Interface (CVI).

Figure 1 depicts a high-level control and interface architecture for ACTN.



Physical Network Infrastructure

Figure 1: Control and Interface Architecture for ACTN.

Figure 1 shows that there are multiple consumer controllers, which are independent to one another, and that each consumer supports various business applications over its NB API. There are layered client-server relationships in this architecture. As various applications are clients to the consumer 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.

5.2. Consumer Controller


A Virtual Network Service is instantiated by the consumer controller via the CVI. As the consumer controller directly interfaces the application stratum, it understands multiple application requirements and their service needs. It is assumed that the consumer 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 consumer-network physical interfaces that connect consumer premise equipment to network provider equipment. Figure 2 shows an example physical network topology that supports multiple consumers. In this example, consumer A has three end-points A.1, A.2 and A.3. The interfaces between consumers 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. Consumer 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.

5.3. Abstracted Topology

There are two levels of abstracted topology that needs to be maintained and supported for ACTN. Consumer-specific Abstracted Topology refers to the abstracted view of network resources allocated to the consumer. The granularity of this abstraction

varies depending on the nature of consumer applications. Figure 2 illustrates this.

Figure 2 shows how three independent consumers A, B and C provide its respective traffic demand matrices 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 consumers. 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 Consumer A			Traffic Matrix for Consumer B			Traffic Matrix for Consumer C					
A.1	A.2	A.3	B.1	B.2	B.3	C.1	C.2	C.3			
A.1	-	20G	20G	B.1	-	40G	40G	C.1	-	20G	20G
A.2	20G	-	10G	B.2	40G	-	20G	C.2	20G	-	10G
A.3	20G	10G	-	B.3	40G	20G	-	C.3	20G	10G	-

Figure 2: Physical network topology shared with multiple consumers


Figure 3 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 consumer 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 3.

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



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





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




Figure 3: Topology Abstraction Examples for Consumers

As different consumers have different control/application needs, abstracted topologies for consumers 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 consumer 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 consumer controller. If the consumer controller has applications that require more granular control of virtual network resources, then the abstracted topology shown for consumer A may be the right abstraction level for such controller. For instance, if the consumer 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 consumer were only to support simple tunnel services to its applications, then the abstracted topology shown for consumer C (one VNE and three VLs) would suffice.

5.4. Workflows of ACTN Control Modules

Figure 4 shows workflows across the consumer controller, VNC and PNC for the VNS instantiation, topology exchange, and VNS setup.

The consumer 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 Consumer Controller triggers a path computation request by the virtual PCE (vPCE) agent in the VNC after VNC's proxy's interlay of this request to the vPCE. vPCE sends a concurrent path computation request that is converted according to the traffic demand matrix as part of the VNS instantiation request from the Consumer Controller. Upon receipt of this path computation request, the PCE in the PNC block computes paths and updates network topology DB and informs the vPCE agent of the VNC of the paths and topology updates.

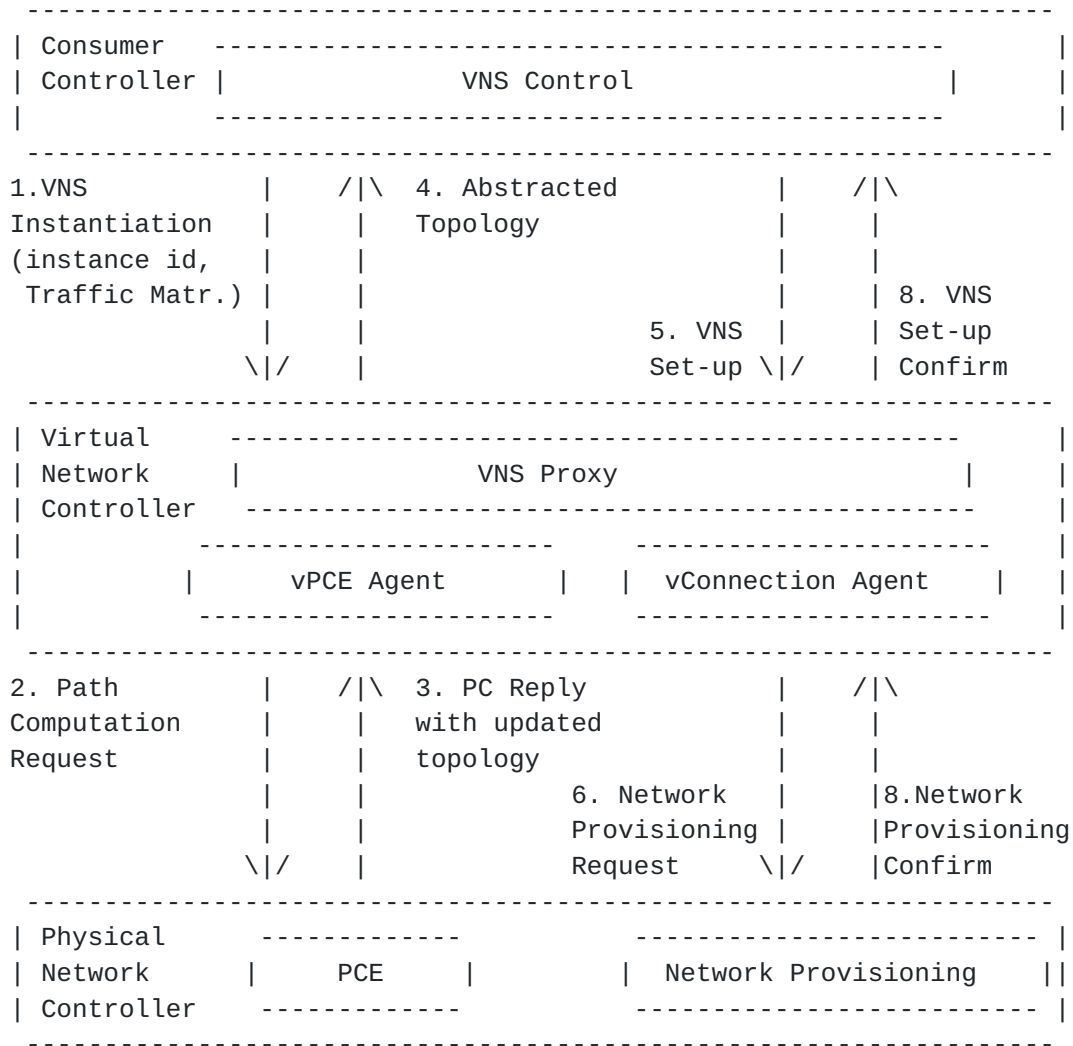


Figure 4. Workflows across Consumer Controller, VNC and PNC

It is assumed that the PCE in PNC is a stateful PCE [PCE-S]. vPCE agent abstracts the network topology into an abstracted topology for the consumer based on the agreed-upon granularity level. The abstracted topology is then passed to the VNS control of the Consumer 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 using control plane messages such as GMPLS, etc.

5.5. Programmability of the ACTN Interfaces

From Figures 1 and 4, we have identified several interfaces that are of interest of the ACTN model. More precisely, ACTN concerns the following interfaces:

- Consumer-VNC Interface (CVI): an interface between a consumer 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 consumer so they can create, modify and delete virtual 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.

6. Design Principles of ACTN

6.1. 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.

6.2. Privacy and Isolation

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

Policy should be applied per client.

6.3. 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 consumers. 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 consumer application request.

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

6.4. 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 consumers and service/network providers.

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

6.5. 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.

6.6. 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.

7. References

7.1. 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.

8. Contributors

Authors' Addresses

Daniel Ceccarelli
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

The IETF Trust takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in any IETF Document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights.

Copies of Intellectual Property disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

All IETF Documents and the information contained therein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgment

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