

Network Working Group
Internet Draft
Intended status: Informational
Expires April 2014

Kwang-koog Lee
Hosong Lee
KT
Ricard Vilata
CTTC
Victor Lopez
Telefonica

November 10, 2014

ACTN Use-case for On-demand E2E Connectivity Services in Multiple
Vendor Domain Transport Networks

draft-klée-actn-connectivity-multi-vendor-domains-02.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/lid-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 April 10, 2015.

Copyright Notice

Copyright (c) 2014 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.

Abstract

This document provides a use-case that addresses the need for facilitating the application of virtual network abstractions and the control and management of on-demand end-to-end provisioning of connections that traverse multiple vendor domain transport networks.

These abstractions shall help create a virtualized environment supporting operators in viewing and controlling different vendor domains, especially for on-demand network connectivity service for a single operator.

Table of Contents

- 1. Introduction.....2
- 2. On-demand End-to-end Connectivity in Multi-vendor Domain Transport Networks.....3
- 3. Requirements.....5
- 4. References.....8
- 5. Contributors.....8
- Intellectual Property Statement.....9
- Disclaimer of Validity.....9

1. Introduction

Network operators build and operate their network using multiple domains in different dimensions. Domains may be defined by a collection of links and nodes (each of a different technology), administrative zones under the concern of a particular business entity, or vendor-specific "islands" where specific control mechanisms have to be applied. Due to the technology of each vendor, the optical components cannot be interconnected. Therefore each optical domain becomes an isolated island in terms of provisioning. The network operators use vendor-specific NMS implementations along with an operator-tailored umbrella provisioning system, which may include a technology specific Operations Support System (OSS). Thanks to the evolution of vendor specific SDN controllers, the network operators require a network entity, which abstract the details of the optical layer while enabling end-to-end provisioning of services. The establishment of end-to-end connections spanning several of these domains is a perpetual problem for operators, which need to address both interoperability and operational concerns at the control and data planes.

The introduction of new services, often requiring connections that traverse multiple domains, needs significant planning, and several manual operations to interface multiple vendor-specific domains in which specific control/management mechanisms of the vendor equipment have to be applied (e.g., EMS/NMS, OSS/BSS, control plane, SDN controller, etc.). Undoubtedly, establishing an on-demand end-to-end connection which requires provisioning based on dynamic resource information is more difficult in the current network context.

This document provides a use-case that addresses the need for creating a virtualized environment supporting operators in viewing and controlling different vendor domains, especially for on-demand network connectivity service for a single operator. This will accelerate rapid service deployment of new services, including more dynamic and elastic services, and improve overall network operations and scaling of existing services.

This use-case is a part of the overarching work, called Abstraction and Control of Transport Networks (ACTN). Related documents are the ACTN-framework [ACTN-Frame] and the problem statement [ACTN-PS].

2. On-demand End-to-end Connectivity in Multi-vendor Domain Transport Networks

This section provides an architecture example to illustrate the context of the current challenges and issues operators face in delivering on-demand end-to-end connectivity services in operators' multi-vendor domain transport networks.

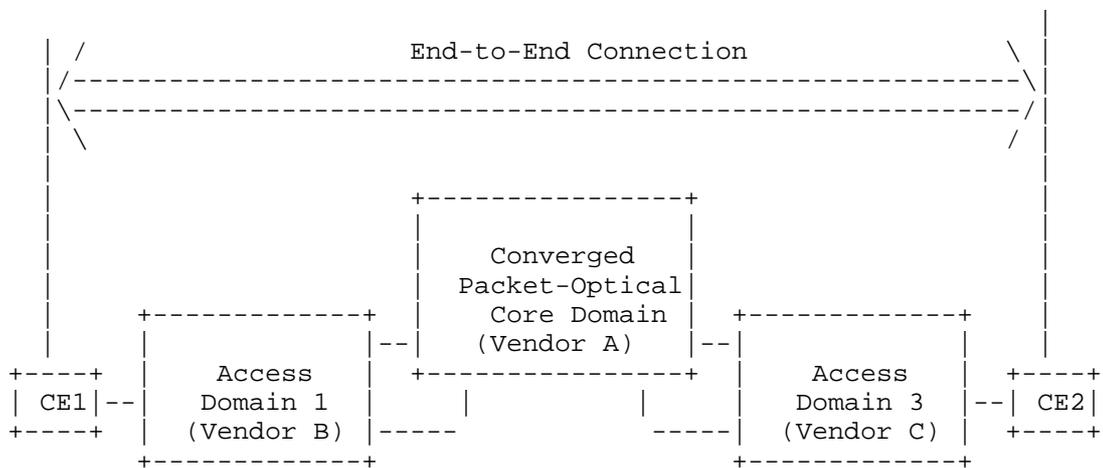


Figure 1. Multi-vendor Domains

As an illustrative example, consider a multi-domain transport network consisting of three domains: one core converged packet-optical domain (Vendor A) and two access domains (Vendors B and C). Each access domain is managed by its domain control/management mechanism which is often a proprietary vendor-specific scheme. The core domain is also managed by Vendor A's proprietary control/management mechanism (e.g., EMS/NMS, OSS/BSS, Control Plane, SDN Controller, or any combination of these entities, etc.) that may not interoperate with access domain control/management mechanisms or at best partially interoperate if Vendor A is same as Vendor B or Vendor C.

Due to these domain boundaries, facilitating on-demand end-to-end connections (e.g., Ethernet Virtual Connections, etc.) that traverse multi-domains is not readily achieved. These domain controls are optimized for its local operation and in most cases not suited for controlling the end-to-end connectivity services. For instance, the discovery of the edge nodes that belong to other domains is hard to achieve partly because of the lack of the common API and its information model and control mechanisms thereof to disseminate the relevant information.

Moreover, the path computation for any on-demand end-to-end connection would need abstraction of dynamic network resources and ways to find an optimal path that meets the connection's service requirements. This would require knowledge of both the domain level dynamic network resource information and the inter-domain connectivity information including domain gateway/peering points and the local domain policy.

From an on-demand connection provisioning perspective, in order to facilitate a fast and reliable end-to-end signaling, each domain operation and management elements should ideally speak the same control protocols to its neighboring domains. However, this is not possible for the current network context unless a folk-lift green field technology deployment with a single vendor solution would be done. Although each domain applies the same protocol for the data plane, an end-to-end connectivity traversing multiple domains might not be provided due to a management and control mechanism focusing only on its own domain.

From a network connectivity management perspective, it would require a mechanism to disseminate any connectivity issues from the local domain to the other domains whenever the local domain cannot resolve a connectivity issues. This is hard to achieve due to the lack of the common API and its agreed-upon information model and control mechanisms thereof to disseminate the relevant information.

From an operation's perspective, the current network environments are not conducive to offering on-demand end-to-end connectivity services in multi-vendor domain transport networks. For instance, when the performance monitoring inquiry is requested, operators manually monitor each domain and aggregate the performance results. However, it may not be precise because of the different measurement timing employed by each domain.

3. Requirements

In the previous section, we discussed the current challenges and issues that prevent operators from offering on-demand end-to-end connectivity services in multi-vendor domain transport networks.

This section provides a high-level requirement for enabling on-demand end-to-end connectivity services in multi-vendor domain transport networks in a single operator environment.

Figure 2 shows information flow requirements of the aforementioned context.

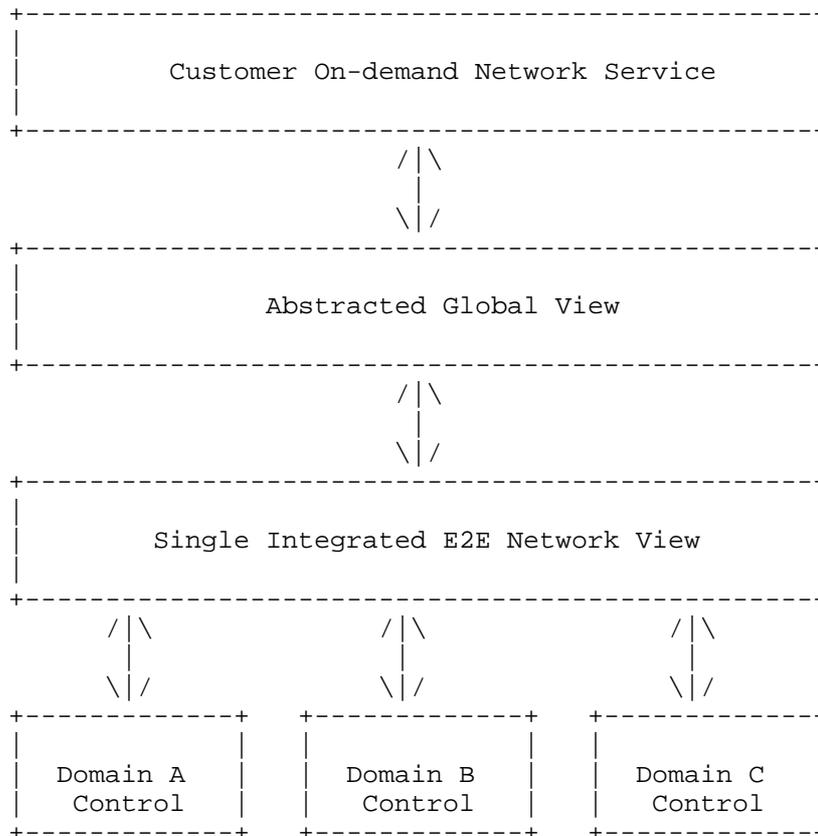


Figure 2. Information Flow Requirements for Enabling On-demand Network Connectivity Service in Multi-vendor Domain Networks

There are a number of key requirements from Figure 2.

- A single integrated end-to-end network view is necessary to be able to provision the end-to-end paths that traverse multiple vendor domains. In this approach the scalability and confidentiality problems are solved, but new considerations must be taken into account:
 - o Limited awareness, by the VNC, of the intra-domain resources availability.
 - o Sub-optimal path selection.

- The path computations shall be performed in two stages: first on the abstracted end-to-end network view (happening at VNC), and on the second stage it shall be expanded by each PNC.
- In order to create a single integrated end-to-end network view, discovery of inter-connection data between domains including the domain border nodes/links is necessary. (The entity to collect domain-level data is responsible for collecting inter-connection links/nodes)
- The entity to collect domain-level data should recognize interoperability method between each domain. (There might be several interoperability mechanisms according to technology being applied.)
- The entity responsible to collect domain-level data and create an integrated end-to-end view should support push/pull model with respect to all its interfaces.
- The same entity should coordinate a signaling flow for end-to-end connections to each domain involved. (This entity to domain control is analogous to an NMS to EMS relationship)
- The entity responsible to create abstract global view should support push/pull model with respect to all its interfaces. (Note that the two entities (an entity to create an integrated end-to-end view and an entity to create an abstracted global view) can be assumed by the same entity, which is an implementation issue.
- Hierarchical composition of integrated network views should be enabled by a common API between NorthBound Interface of the Single Integrated End-to-End view (handled by VNC) and Domain Control (handled by PNC).
- There is a need for a common API between each domain control to the entity that is responsible for creating a single integrated end-to-end network view. At the minimum, the following items are required on the API:
 - o Programmability of the API.
 - o The multiple levels/granularities of the abstraction of network resource (which is subject to policy and service need).
 - o The abstraction of network resource should include customer end points and inter-domain gateway nodes/links.

- o Any physical network constraints (such as SRLG, link distance, etc.) should be reflected in abstraction.
 - o Domain preference and local policy (such as preferred peering point(s), preferred route, etc.)
 - o Domain network capability (e.g., support of push/pull model).
- The entity responsible for abstraction of a global view into a customer view should provide a programmable API to allow the flexibility. Abstraction might be provided by representing each domain as a virtual node (node abstraction) or a set of virtual nodes and links (link abstraction). Node abstraction creates a network topology composed by nodes representing each network domain and the inter-domain links between the border nodes of each domain.
- o Abstraction of a global view into a customer view should be provided to allow customer to dynamically request network on-demand services including connectivity services.
 - o What level of details customer should be allowed to view network is subject to negotiation between the customer and the operator.

4. References

- [ACTN-Frame] D. Ceccarelli, L. Fang, Y. Lee and D. Lopez, "Framework for Abstraction and Control of Transport Networks," draft-ceccarelli-actn-framework, work in progress.
- [ACTN-PS] Y. Lee, D. King, M. Boucadair, R. Jing, and L. Murillo, "Problem Statement for the Abstraction and Control of Transport Networks," draft-leeing-actn-problem-statement, work in progress.

5. Acknowledgement

The authors wish to thank Young Lee for the discussions in the document.

6. Contributors

Authors' Addresses

Kwang-koog Lee

KT

Email: kwangkoog.lee@kt.com

Hosong Lee

KT

Email: hosong.lee@kt.com

Ricard Vilata

CTTC

Email: ricard.vilalta@cttc.es

Victor Lopez

Telefonica

Email: victor.lopezalvarez@telefonica.com