

Network Working Group
Internet Draft
Intended status: Informational
Expires: March 29, 2015

Luyuan Fang
Microsoft
September 29, 2014

ACTN Use Case for Multi-domain Data Center Interconnect

draft-fang-actn-multidomain-dci-01.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 March 29, 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 discusses a use case for data center operators that need to interface multi-domain transport networks to offer their global data center applications and services. As data center operators face multi-domain and diverse transport technology, interoperability based on standard-based abstraction is required to support dynamic and flexible applications and services.

Table of Contents

- 1. Introduction.....2
- 2. Multi-domain Data Center Interconnection Applications.....3
 - 2.1. VM Migration.....3
 - 2.2. Global Load Balancing.....4
 - 2.3. Disaster Recovery.....4
 - 2.4. On-demand Virtual Connection/Circuit Services.....5
- 3. Issues and Challenges for Multi-domain Data Center Interconnection Operations.....5
- 4. Control Hierarchy.....7
- 5. Requirements.....9
- 6. References.....10
- 7. Contributors.....11
- Authors' Addresses.....11
- Intellectual Property Statement.....11
- Disclaimer of Validity.....11

1. Introduction

This document discusses a use case for data center operators that need to interface multi-domain transport networks to offer their global data center applications and services. As data center providers face multi-domain and diverse transport technology, interoperability based on standard-based abstraction is required to support dynamic and flexible applications and services.

This use case is a part of the overarching work, called Abstraction and Control of Transport Networks (ACTN). The goal of ACTN is to facilitate virtual network operation by:

- . The creation of a virtualized environment allowing operators to view the abstraction of the underlying multi-admin, multi-vendor, multi-technology networks and

- . The operation and control/management of these multiple networks as a single virtualized network.

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.

Related documents are the ACTN-framework [ACTN-Frame] and the problem statement [ACTN-PS].

Multi-domain transport networks herein are referred to physical WAN infrastructure whose operation may or may not belong to the same administrative domain as the data center operation. Some data center operators may wholly own the entire physical WAN infrastructure while others may own partially or even not at all. In all cases, data center operation needs to establish multi-domain relationships with one or more physical network infrastructure operations.

Data center based applications are used to provide a wide variety of services such as video gaming, cloud storage and computing, grid application, data base tools, and mobile applications, and others. High-bandwidth video applications such as remote medical surgery, video streaming for live concerts and sporting events are also emerging. This document is mainly concerned with data center applications that in aggregate or individually make substantial bandwidth demands that traverse multi-domain transport networks, some of which may belong to different administrative domains. In addition, these applications may require specific bounds on QoS related parameters such as guaranteed bandwidth, latency and jitter and others.

The organization of this document is as follows: Section 2 will discuss multi-domain Data Center interconnection and its various application scenarios. Section 3 will discuss the issues and challenges for Multi-domain Data Center Interconnection Operations Architecture. Section 4 will provide high-level requirements.

2. Multi-domain Data Center Interconnection Applications

2.1. VM Migration

A key enabler for data center cost savings, consolidation, flexibility and application scalability has been the technology of compute virtualization or Virtual Machines (VMs). A VM to the software application looks like a dedicated processor with dedicated memory and dedicated operating system. In modern data centers or "computing clouds", the smallest unit of computing resource is the VM. In public data centers one can buy computing capacity in terms of VMs for a particular amount of time. Though different VM

configurations may be offered that are optimized for different types of processing (e.g., memory intensive, throughput intensive).

VMs offer not only a unit of compute power but also as an "application environment" that can be replicated, backed up and moved. Although VM migration started in the LAN, the need for inter-DC VM migration for workload burst/overflow management on the WAN has been a real need for Data Center Operators.

Virtual machine migration has a variety of modes: (i) scheduled vs. dynamic; (ii) bulk vs. sequential; (iii) point-to-point vs. point-to-multi-point. Transport network capability can impact virtual machine migration strategy. For certain mission critical applications, dynamic bandwidth guarantee as well as performance guarantee must be provided by the network. Make-before-break capability is also critical to support seamless migration.

2.2. Global Load Balancing

As the many data center applications are distributed geographically across many data centers and over multi-domain networks, load balancing is no longer a local decision. As such, the decision as to selecting a server for an application request from the users or selecting data centers for migrating or instantiating VMs needs to be done globally. This refers to global load balancing.

There are many factors that can negatively affect the quality of experience (QoE) for the application. Among them are: the utilization of the servers, the underlying network loading conditions within a data center (LAN), the underlying network loading conditions between data centers (MAN/WAN), the underlying network conditions between the end-user and data center (Access Network). To allow data center operators to facilitate global load balancing over heterogeneous multi-domain transports from access networks to metro/core transport networks, on-line network resource information needs to be abstracted and represented from each involving network domain.

2.3. Disaster Recovery

For certain applications, disaster recovery in real-time is required. This requires transport of extremely large amount of data from various data center locations to other locations and a quick feedback mechanism between data center operator and infrastructure network providers to facilitate the complexity associated with real-time disaster recovery.

As this operation requires real-time concurrent connections with a large amount of bandwidth, a strict guarantee of bandwidth and a

very low latency between a set of data centers, the underlying physical network infrastructure is required to support these network capability. Moreover, as the data center operator interfaces multiple network infrastructure providers, standard-based interfaces and a common ways to abstract network resources and connections are necessary to facilitate its operations.

2.4. On-demand Virtual Connection/Circuit Services

Related to the real-time operations discussed in other applications in the previous sections, many applications require on-demand virtual connection/circuit services with an assured quality of service across multiple domain transport networks.

The on-demand aspect of this service applies not only in setting up the initial virtual connections/circuits but also in increasing bandwidth, changing the QoS/SLA, adding a new protection scheme to an existing service.

The on-demand network query to estimate available SLA/QoS (e.g., BW availability, latency range, etc.) between a few data center locations is also part of this application.

3. Issues and Challenges for Multi-domain Data Center Interconnection Operations

This section discusses operational issues and challenges for multi-domain data center interconnection. Figure 1 shows a typical multi-domain data center interconnection operations architecture.

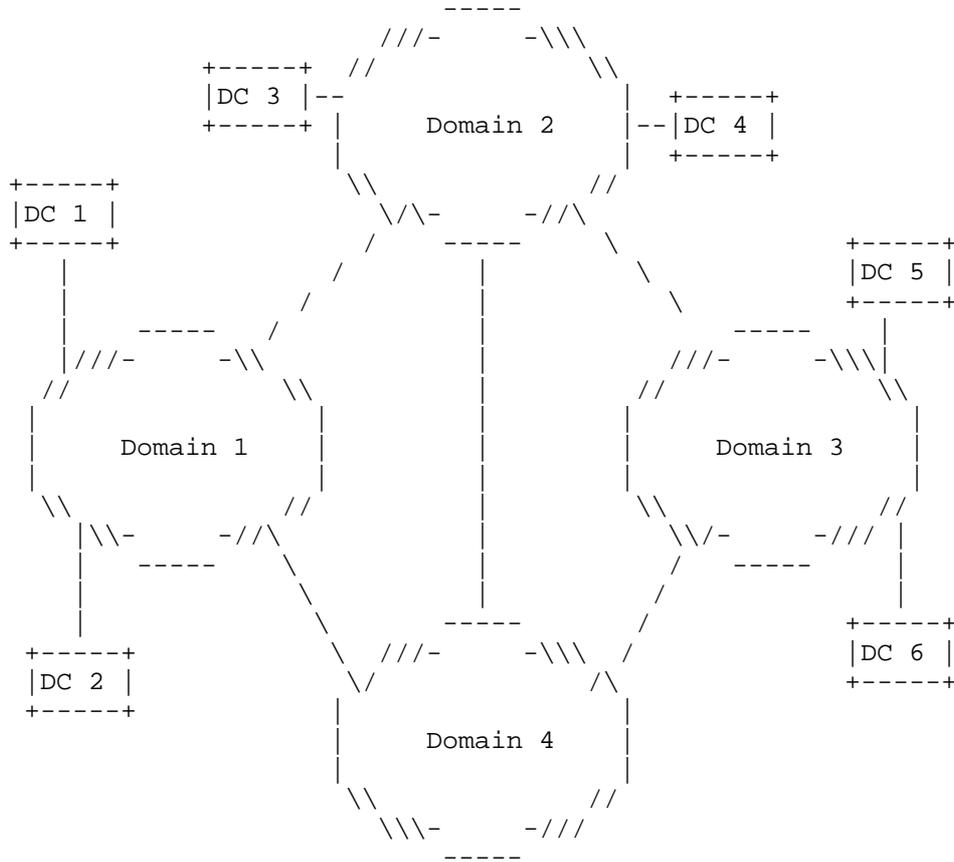


Figure 1. Multi-domain Data Center Interconnect Operations Architecture

Figure 1 shows several characteristics pertaining to current multi-domain data center operations.

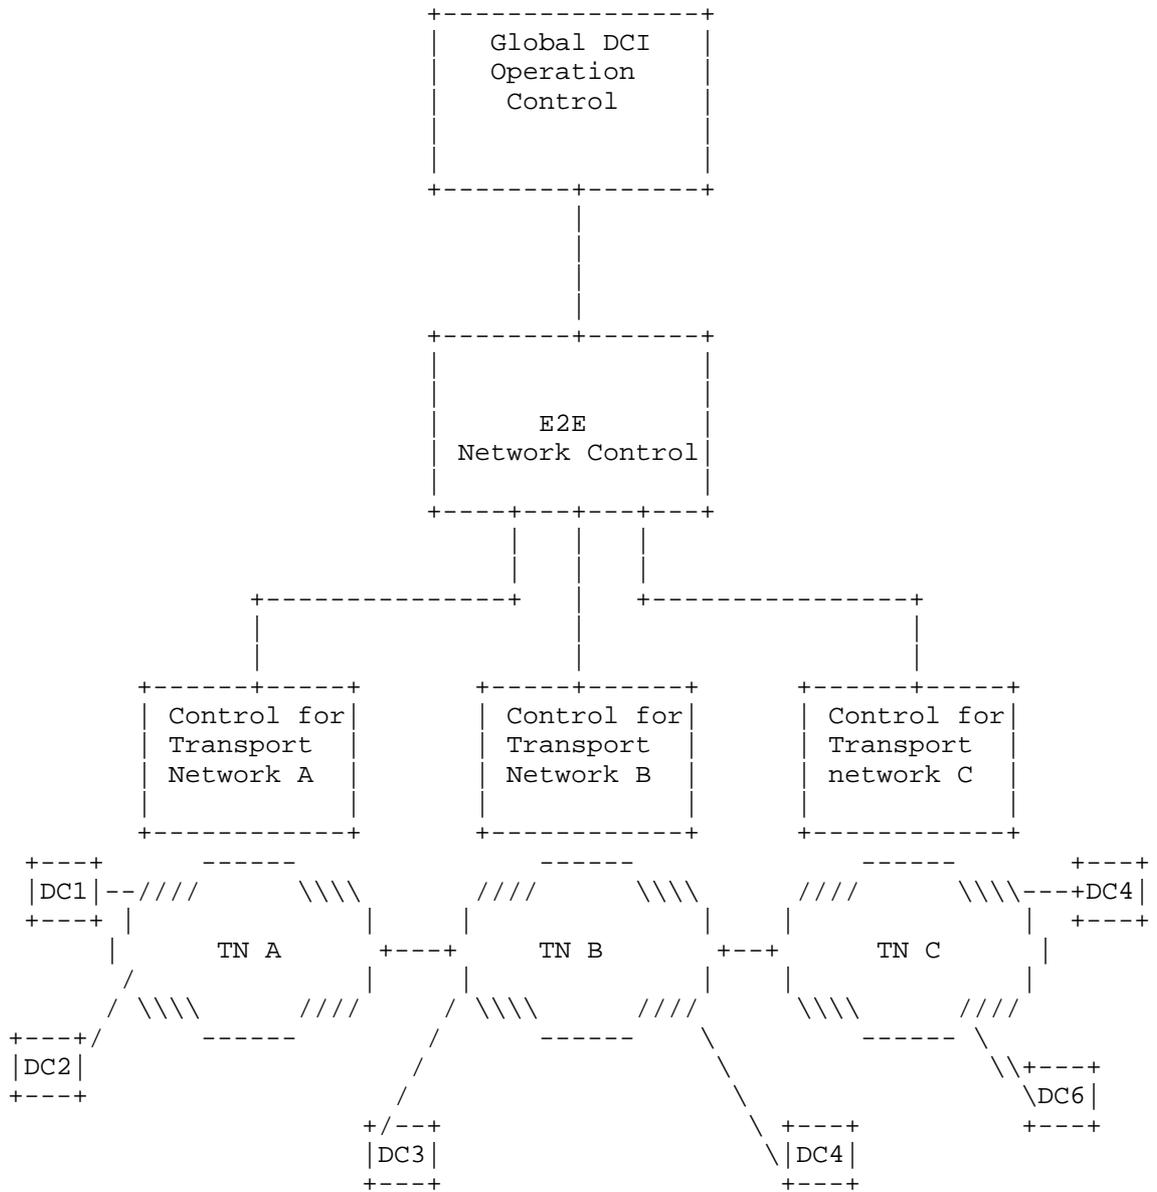
1. Data centers are geographically spread and homed on possibly a number of mutually independent physical network infrastructure provider domains.
2. Between the data center operator domain and each of mutually independent physical network provider domains must establish trusted relationships amongst the involved entities. In some cases where data center operator owns the whole or partial physical

network infrastructure domains, a trusted relationship is still required between the data center operation and the network operations due to organizational boundaries although it is less strict than a pure multi-domain case.

3. Data center operator may lease facility from physical network infrastructure providers for intra-domain connectivity or own the facility. For instance, there may be an intra-domain leased facility for connectivity between DC 1 to DC 2. It is also possible that the data center provider may own this intra-domain facility such as dark fibers for connectivity between DC 1 and DC 2.
 4. There may be need for connectivity that may traverse multi-domain networks. For instance, Data Center 1 may have VMs that need to be transported to Data Center 6. Typically, multi-domain connectivity is arranged statically such that the routes are pre-negotiated with the involved operators. For instance, if Data Center 1 were to send its VMs to Data Center 6, the route may take on Domain 1 - Domain 4 - Domain 3 based on a pre-negotiated agreement prior to connectivity request. In such case, the inter-domain facilities between Domains 1 & 4 Domains 4 & 3 are a part of this pre-negotiated agreement. There could be alternative route choices. Whether there may be alternate routing or not is subject to policy. Alternate routing may be static or dynamic depending on policy.
 5. These transport network domains may be diverse in terms of local policy, transport technology and its capability and vendor equipment. Due to this diversity, new service introduction, requiring connections that traverse multiple domains, need significant planning, and several manual operations to interface different vendor equipment and technology. New applications requiring dynamic and elastic services and real-time mobility may be hampered by these manual operational factors.
4. Control Hierarchy

This section provides a control hierarchy for multi-domain DC operations.

Figure 2 shows a control hierarchy for multi-domain Data Center Interconnection operation.



There are a number of important considerations to support a global multi-domain data center interconnection operation.

1. Need a hierarchical operation/control.

2. Build on top of existing network control technologies/domains to be able to E2E network control to help global DCI operation/control.
3. Need standard-based abstraction/APIs and protocols between E2E network control and global DCI operation control and between E2E network control and domain transport network controls.

5. Requirements

This section provides high-level requirements to fulfill multi-domain data center interconnection to support various applications discussed in the previous sections.

1. The interfaces between the Data Center Operation and each transport network domain SHOULD support standards-based abstraction with a common information/data model.
2. The Data Center Operation should be able to create a single virtual network view.
3. The following capability should be supported:
 - a. Network Query (Pull Model) from the Data Center Operation to each transport network domain to collect potential resource availability (e.g., BW availability, latency range, etc.) between a few data center locations.
 - i. The level of abstracted topology (e.g., tunnel-level, graph-form, etc.)
 - b. Network Path Computation Request from the Data Center Operation to each transport network domain to estimate the path availability.
 - c. Network Virtual Connections/Circuits Request from the Data Center Operation to each transport domain to establish an end-to-end virtual connections/circuits.
 - i. The type of the connection: P2P, P2MP, etc.
 - ii. Concurrency of the request (this indicates if the connections must be simultaneously available or not in case of multiple connection requests).
 - iii. The duration of the connections

- iv. SLA/QoS parameters: minimum guaranteed bandwidth, latency range, etc.
 - v. Protection/Reroute Options (e.g., SRLG requirement, etc.)
 - vi. Policy Constraints (e.g., peering preferences, etc.)
- d. Network Virtual Connections/Circuits Modification Request from the Data Center Operation to each transport domain to change QoS/SLA, protection schemes of the existing connections/circuits.
- e. Network Abnormality Report (Push Model) from each transport domain to the Data Center Operation indicating the service impacting network conditions or the potential degradation indications of the existing virtual connections/circuits.

6. 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-leeking-actn-problem-statement, work in progress.

7. Authors' Addresses

Luyuan Fang
Microsoft
Email : lufang@microsoft.com

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.