

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 22, 2015

O. Gonzalez de Dios, Ed.
Telefonica GCTO
J. Meuric, Ed.
Orange
D. Ceccarelli
Ericsson
July 21, 2014

Terminology and Models for Control of Traffic Engineered Networks with
Client-Server Relationship
draft-dios-ccamp-control-models-customer-provider-01

Abstract

Different kinds of relationships can be established among interconnected Traffic Engineered Networks. In particular, this document focuses on the case where there is a client-server relation between the network domains. The domain interconnection is a policy and administrative boundary. This informational document collects current terminology and provides a taxonomy for the possible control plane based operation models.

Each control model defines, on the one hand, the level of information that the domain acting as client receives by control plane means from the domain acting as server and, on the other hand, the control model will determine what can be requested from the client domain to the server domain.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 22, 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Examples of Client-Server TE Network Domain Scenarios . .	3
2. Terminology	4
2.1. Routing domain	4
2.2. Overlay of routing domains	4
2.3. Multilayer	4
2.4. Policy	5
2.5. Client Domain - Server Domain Interface	5
2.5.1. UNI in IP over Optical Networks	5
2.5.2. ITU-T Definition of UNI	5
2.5.3. OIF Definition of UNI	6
2.5.4. Proposed Vocabulary	6
2.6. Reachability	7
2.6.1. Unqualified Reachability	7
2.6.2. Qualified Reachability	7
2.6.3. Qualified Reachability with associated potential TE path	8
3. Control Models	8
3.1. Signaling Only	8
3.1.1. Signaling with Requirements	9
3.1.2. Signaling with Collection	9
3.2. Signaling and Reachability Model	9
3.2.1. Signalling + Basic Reachability	10
3.2.2. Signalling + Qualified Reachability	10
3.2.3. Signalling + Qualified Reachability + Potential Services	10
3.3. Service Attributes vs service constraints	10
3.4. Other Models	11
3.4.1. Multi-Layer Networks / Multi-Region Networks	11
3.4.2. Management Model	11
4. Abstraction	11

5. Security Considerations	11
6. Contributing Authors	11
7. Acknowledgments	12
8. References	12
8.1. Normative References	12
8.2. Informative References	12
Authors' Addresses	12

1. Introduction

Traffic Engineered Networks can be interconnected, establishing different types of relationships among them. For example, both network can have a peering relation, where connections starting in one domain and end in the other domain. This document is focused on the case where the interconnected network domains have a client-server relationship among them. Such client-server relation comes from the two main points. On the one hand, end-to-end services in the client network can be set up using services of a network acting as server. On the other hand, the client-server relation comes from the fact that their interconnection is a policy and administrative boundary, limiting the amount of information allowed to be exchanged between networks. In the case of interconnected TE domains where there is no administrative nor strict policy boundary between client and server (typically, just a technology change), the MLN/MRN model can be applied.

The interface between the client and the server domain is typically called "User-to-Network Interface" (UNI), and regarded as signaling-only [RFC4208]. Due to the strict association of functionality to the UNI term, its exact scope has become highly controversial. This document compiles different definitions of the term used so far and propose some terminology to serve as a foundation to move the work forward.

What is more, the document compiles the possible operation models of client-server network from the control plane perspective. Each control model defines, on the one hand, the level of information of the domain acting as client provides through the control plane to the domain acting as server. On the other hand, the control model will determine what can be requested from the client domain to the server domain.

1.1. Examples of Client-Server TE Network Domain Scenarios

The most typical example of interconnected TE domains that follow a client-server relation is an IP/MPLS domain using the services of an optical OTN/WDM network. Note that the interconnected domain can be part of the same organization, but with different administration.

A particular network scenario that has attracted lot of attention from the industry is the IP/MPLS/OTN/WDM over WDM. The client network is based on multi-layer routers able to set up packet-based TE connections over wavelengths. The server network is a WDM optical network that provides the switching for the wavelengths as well as restoration capabilities of the optical channels.

Another example is MPLS over MPLS, where both client and server networks are able to set up packet based TE connections. This is the case, for example, of carrier-over-carrier scenarios.

Summing up, there number of applicable scenarios is wide.

2. Terminology

2.1. Routing domain

A routing domain is made of GMPLS enabled nodes (i.e., a network device including a GMPLS entity). These nodes can be either edge nodes (i.e., hosts, ingress LSRs or egress LSRs), or internal LSRs. An example of non-PSC host is an SONET/SDH Terminal Multiplexer (TM). Another example is an SONET/SDH interface card within an IP router or ATM switch.

A routing domain is characterized by being under the control of the same administration and by running a common set of protocols to exchange routing information

2.2. Overlay of routing domains

In an overlay environment we have a client routing domain and a server routing domain, each of which running its own routing protocol instance. Connectivity in the client routing domain can be made by connectivity services of the server domain.

2.3. Multilayer

As per RFC 5212 "UA data plane layer is a collection of network resources capable of terminating and/or switching data traffic of a particular format [RFC4397]. These resources can be used for establishing LSPs for traffic delivery. For example, VC-11 and VC4-64c represent two different layers."U

In a Multilayer network, each layer can be or not a routing domain. In fact, a multi-layer network can be controled with a single control plane instance in which all resources are adverstised in the same IGP instance

2.4. Policy

In an overlay network, policy is the set of rules that apply in the interface between two routing domains, and that restrict the level of information exchanged and the operations allowed. The policy decisions obey to confidentiality reasons (typically, the routing domains operate under the control of different administrations) and scalability (to avoid excessive flow of information that collapse the processing capacity of the nodes)

An example of policy example, visibility of the server domain could be restricted to the client domain.

2.5. Client Domain - Server Domain Interface

The interface between the client and the server domain is typically called "User-to-Network Interface" (UNI). However, the term "UNI" has been used in different contexts and SDOs. As a consequence, the exact definition of UNI and the functionalities included depend on the application. Bellow, as a reference, it is shown a set of the different definitions of UNI.

2.5.1. UNI in IP over Optical Networks

[RFC3717] says: "The client-optical internetwork interface (UNI) represents a service boundary between the client (e.g., IP router) and the optical network. The client and server (optical network) are essentially two different roles: the client role requests a service connection from a server; the server role establishes the connection to fulfill the service request -- provided all relevant admission control conditions are satisfied."

In other words, this definition refers to a signaling protocol between two administrative domains with a client-server relationship. It is agnostic to the existence of a data plane client-server relationship and to the side(s) of the boundary where it may happen, if any.

2.5.2. ITU-T Definition of UNI

ITU-T has defined the term UNI in the context of control plane. [G.807] [G.8081] (ITU-T): "User-Network Interface for the control plane (UNI): A bidirectional signaling interface between service requester and service server control plane entities."

The terms "requester/provider" are used to refer to the relationship.

2.5.3. OIF Definition of UNI

UNI: "The service control interface between a client device and the transport network."

UNI-C: "The logical entity that terminates UNI signalling on the client device side."

UNI-N: "The logical entity that terminates UNI signalling on the transport network side."

The terms "client/transport" and "client/network" are used to refer to the relationship.

2.5.4. Proposed Vocabulary

As listed above, the existing terminology is far from unique. To avoid overloaded concepts, this document proposes to use the "client/server" terms.

Unless stated, this document focuses on control protocol exchanges and their uses across administrative boundaries for client-server interconnection. Data plane transition and/or client-server relationship may not be aligned with the boundary.

2.5.4.1. Client network

A Client network is defined as a network domain able to request a connectivity service to a server network domain across an administrative boundary.

2.5.4.2. Server network

A Server network is defined as a network domain able to deliver connectivity services to a client network domain across an administrative boundary.

2.5.4.3. Client-Server Control Plane Interface

The control plane interface between the client network domain and the server network domain convey a set of control functionalities that help to operate such kind of networks. The exact functionalities of this Interface (and then the level of information exchanged) depend on the chosen control model. This document presents a taxonomy with the possible control models.

2.6. Reachability

In graph theory, reachability refers to the ability to get from one vertex to another within a graph. Thus, a vertex can reach another vertex if there exists a sequence of adjacent vertices which starts with the source vertex and ends with the destination vertex.

The document [draft-farrel-interconnected-te-info-exchange-04] provides the definition of what is reachability for client-server networks. [EDITOR's note: Text from draft-farrel-interconnected-te-info-exchange has been borrowed for this first version. Duplicated text will be deleted at later stages]

In an IP network, reachability is the ability to deliver a packet to a specific address or prefix. That is, the existence of an IP path to that address or prefix. TE reachability is the ability to reach a specific address along a TE path.

In the context of Traffic Engineered networks with client and server relationships, we can define several types of reachability:
[draft-farrel-interconnected-te-info-exchange-04]

2.6.1. Unqualified Reachability

Two client domain nodes are said to be reachable if, either there exists at least one path through the client domain that connects both nodes, or, in the case that there is no path exclusively through the client domain network, there exists at least one path connecting nodes of client and server domain by which both client nodes can be connected.

In the case of basic reachability, it is only known that it is possible to connect the nodes, but there is no notion of the details of such possible connections, such as, for example, bandwidth available or performance metrics. Also, the exact path to connect both nodes is not known to the client network. Note that, even if two nodes are reachable, there may not be enough resources for a desired TE connection with specific TE constraints.

2.6.2. Qualified Reachability

In this case, on top of the basic reachability, it is known some TE attributes of the possible connection (or connections). Examples of such attributes are: TE metrics, hop count, available bandwidth, delay, SRLG list. Note that this information is specific per connection. Thus, if there are several possible TE paths, there are a set of attributes.

2.6.3. Qualified Reachability with associated potential TE path

In this particular case, on top of the qualified reachability, there exists an associated potential TE path that satisfies the TE connection between two client nodes. Thus, in this case, the client Network has the information that there exists a TE path that can be set up at any time.

3. Control Models

The control of the networks formed by interconnected domains with a client-server relations between them can be done following different models. Each control model defines, on the one hand, the level of information that the domain acting as client receives by control plane means about the services given by the domain acting as server. This information, for example, can vary from a complete lack of information, so the client domain only knows that it could be possible to reach another point of its domain via the server network, to a detailed view on the possibilities offered by the server network. The level of detail of this information will determine which information is exchanged between both networks. On the other hand, the control model will determine what can be requested from the client domain to the server domain. As an example, the most basic use is specifying just the end-points to connect. Other cases may include the possibility to request a service specifying a set of constraints, like bandwidth, diversity, an optimization criteria, etc.

Which control model to choose depends on several factors. For the network operators, the main concern will be related to the level of trustness and relationship between client and server domains. Also, one key factor to take into account is the protocol interoperability. Note that, equipment in the interconnected domains may be from different technologies (but not necessarily) and are likely to use different implementations. The higher the level of functionality included in the control plane, the higher the protocol interoperability requirements, as it will force all implementations to support many functionality. Finally, scalability, that is, the ability of the control plane to provide the same functionality regarding the number of equipment, needs to be taken into account: the amount of information in each option will have different limits in terms on number of interconnected nodes.

3.1. Signaling Only

This first model considers that the sole functionality allowed in the control plane is signaling, that is the ability to request services from client to server domain.

In this model, the control plane does not provide a priori hints to the client domain about the state of the server domain (e.g., resource availability). This model does not preclude that, by other means like the management plane, the client domain knows what is possible or not. Such management actions are out of the scope of the control plane. Thus, it is perfectly feasible that the reachability information is provided either statically or by some management platform.

The most basic case relies on sending a loose ERO from the client, specifying the edges of the connection.

In a trusted interconnection mode, the signaling allows the client domain to provide a full ERO, given to the client network by external tools.

3.1.1. Signaling with Requirements

The control plane may allow to express complex requests to the server domain. That is, through the signaling protocol, it is allowed to not only request a connection between two points of the client domain, but also to include some constraints: e.g., minimum bandwidth, maximum delay, optimization criteria, or request diversity from another service. The policy at the edges of the server network will determine which constraints are accepted. Note the many of the requirements that can be expressed in the request are similar to what would be asked to a path computation function.

3.1.2. Signaling with Collection

Even though the only protocol enabled is signaling, it may be beneficial for the client domain to be able to know some updated information of the services that it has requested to the server. Thus, this case considers the possibility that, through the signaling protocol, the client domain can receive some information. What information it is allowed to collect will be determined by the policy of the server domain.

3.2. Signaling and Reachability Model

This second model considers that, in addition to signaling, the client domain receives some reachability information through a control plane mechanism.

3.2.1. Signalling + Basic Reachability

In this particular case, through control plane mechanisms, the client domain knows whether it is possible to reach a remote end point. The client domain should also remain aware of this information if there are failures in the server domain or if the associated capacity has been filled.

3.2.2. Signalling + Qualified Reachability

The control plane will provide information not only about the possibility to reach a remote end point, but also some TE information of possible connections. For example, the client domain will know that it is possible to reach another point with some bandwidth or delay. Note that, in this case, such information is sent by control plane mechanisms (not statically configured by management plane).

3.2.3. Signalling + Qualified Reachability + Potential Services

In addition to the TE information of the possible connections between two points, the control plane will also provide to the client domain information about potential server's services which could satisfy given requirements. By control plane procedures, the client domain can request, with respect to its needs, a service using such potential service and make high level path selection within the server domain.

3.3. Service Attributes vs service constraints

When asking for the setup of a service in the server domain, the client domain can put constraints on such request. Constraints can consist on the utilization of a path that minimizes a given metric (e.g. TE metric or end to end delay) or on a set of lower/upper bounds that must be followed (e.g. maximum number of hops or maximum end to end delay). Once the service has been provisioned (or just its paths computed), it is possible to identify (e.g. measure or collect) the attributes that characterize such service. For example the path has been computed so to meet the constraint of maximum end to end delay of 20ms, while one of its attributes is the effective end to end delay that is experimented along its path, which could be of e.g. 14 ms. Other examples of constraints and attributes can be found in path diversity. A typical constraint in LSP provisioning is diversity, which is a constraint, but then attributes of the two diverse LSPs like e.g. SRLGs can be collected. Both constraints and attributes need to be exchanged between a client and a server domain.

3.4. Other Models

3.4.1. Multi-Layer Networks / Multi-Region Networks

MLN/MRN extensions to control protocols have been defined. They are well scoped for client and server data plane domains without administrative boundary between them. This allows MLN nodes to participate in common control protocol instances. There is a full set of mechanisms to operate such networks [Editor's note: add refs to MLN/MRN)]. Typical use cases are switches combining both low- and high-order Sonet/SDH, or both ODUk and wavelengths.

However, MLN/MRN assumes no policy boundary between client and server domains. Thus, the level of information exchanged is not restricted, and full interoperability of both the signaling and routing protocols is required.

3.4.2. Management Model

In this particular case, the role of the control plane is limited to operate independently in each of the domains. [Editor's note: Common Control... WG => do we leave it?]

4. Abstraction

Abstraction:

- a physical topology is made of actual nodes interconnected by existing links, i.e. without abstraction;
- a virtual topology is made of nodes and/or links which may (or may not) exist or be instantiated to look the same as the advertised abstraction;
- a potential topology is made of nodes and/or links which are not existing at advertising time but could be instantiated on demand, i.e. a virtual topology which can be actually provided by a network.

5. Security Considerations

TBD

6. Contributing Authors

7. Acknowledgments

The authors would like to thank Lou Berger for pointing out the direction of the document and Dieter Beler for his review. The authors would like to specially thank all the authors of draft-farrel-interconnected-te-info-exchange-02

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, May 2001.
- [RFC3717] Rajagopalan, B., Luciani, J., and D. Awduche, "IP over Optical Networks: A Framework", RFC 3717, March 2004.
- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.

8.2. Informative References

- [draft-farrel-interconnected-te-info-exchange-04]
"Farrel, A., Drake, J., Bitar, N., Swallow, G.,
Ceccarelli, D. draft-farrel-interconnected-te-info-
exchange-04 Problem Statement and Architecture for
Information Exchange Between Interconnected Traffic
Engineered Networks", 2014.

Authors' Addresses

Oscar Gonzalez de Dios (editor)
Telefonica GCTO
Dis
Madrid 28045
Spain

Phone: +34913128832
Email: oscar.gonzalezdedios@telefonica.com

Julien Meuric (editor)
Orange
2 avenue Pierre Marzin
Lannion 22300
France

Email: julien.meuric@orange.com

Daniele Ceccarelli
Ericsson
Via Calda 5
Genova
Italy

Phone: +39 010 600 2512
Email: daniele.ceccarelli@ericsson.com