

CCAMP Working Group  
Internet-Draft  
Intended status: Informational  
Expires: January 5, 2014

D. Ceccarelli  
Ericsson  
O. Gonzalez de Dios  
Telefonica I+D  
July 4, 2013

**Use cases for operating networks in the overlay model context**  
**draft-ceccadedios-ccamp-overlay-use-cases-00**

Abstract

This document defines a set of use cases for operating networks in the overlay model context through the Generalized Multiprotocol Label Switching (GMPLS) overlay interfaces.

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 5, 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.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Terminology . . . . .</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Variants of UNI interface . . . . .</a>	<a href="#">5</a>
<a href="#">3.1.</a>	<a href="#">UNI . . . . .</a>	<a href="#">5</a>
<a href="#">3.1.1.</a>	<a href="#">UNI example . . . . .</a>	<a href="#">6</a>
<a href="#">3.2.</a>	<a href="#">ONI . . . . .</a>	<a href="#">7</a>
<a href="#">3.2.1.</a>	<a href="#">Routing info over the ONI . . . . .</a>	<a href="#">8</a>
<a href="#">4.</a>	<a href="#">Hybrid Nodes . . . . .</a>	<a href="#">9</a>
<a href="#">5.</a>	<a href="#">Adding the PCEP to the UNI . . . . .</a>	<a href="#">10</a>
<a href="#">5.1.</a>	<a href="#">Edge-node to core-node PCEP interface . . . . .</a>	<a href="#">10</a>
<a href="#">5.2.</a>	<a href="#">PCEP and colored UNI . . . . .</a>	<a href="#">12</a>
<a href="#">5.3.</a>	<a href="#">Using PCEP to obtain TE reachability info . . . . .</a>	<a href="#">12</a>
<a href="#">6.</a>	<a href="#">Provisioning . . . . .</a>	<a href="#">12</a>
<a href="#">6.1.</a>	<a href="#">Use case P1 - Local Trigger . . . . .</a>	<a href="#">13</a>
<a href="#">6.2.</a>	<a href="#">Use case P2 - Remote Signaling . . . . .</a>	<a href="#">13</a>
6.3.	<a href="#">Use case P3 - Provisioning with requirements over the UNI . . . . .</a>	<a href="#">14</a>
6.4.	<a href="#">Use case P4 - Provisioning with requirements and collection request over the UNI . . . . .</a>	<a href="#">16</a>
6.5.	<a href="#">Use case P5 - Advertisement of collected metrics in the client layer . . . . .</a>	<a href="#">16</a>
6.6.	<a href="#">Use case P6 - Provisioning with requirements over the ONI . . . . .</a>	<a href="#">16</a>
<a href="#">6.7.</a>	<a href="#">Use case P7 - Dual homing . . . . .</a>	<a href="#">17</a>
<a href="#">7.</a>	<a href="#">Recovery . . . . .</a>	<a href="#">18</a>
<a href="#">7.1.</a>	<a href="#">Use case R1 - Segment Recovery - Single homing . . . . .</a>	<a href="#">18</a>
7.2.	<a href="#">Use case R2 - Local recovery - Dual Homing - Single overlay node . . . . .</a>	<a href="#">19</a>
7.3.	<a href="#">Use case R3 -End to end Recovery - Dual homing - Double overlay node . . . . .</a>	<a href="#">19</a>
7.4.	<a href="#">Use case R4 - Combination of client protection and server restoration . . . . .</a>	<a href="#">20</a>
<a href="#">8.</a>	<a href="#">Optimization . . . . .</a>	<a href="#">22</a>
<a href="#">8.1.</a>	<a href="#">Use case O1 - . . . . .</a>	<a href="#">22</a>
<a href="#">9.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">22</a>
<a href="#">10.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">22</a>
<a href="#">11.</a>	<a href="#">Contributors . . . . .</a>	<a href="#">22</a>
<a href="#">12.</a>	<a href="#">Acknowledgements . . . . .</a>	<a href="#">23</a>
<a href="#">13.</a>	<a href="#">References to be added . . . . .</a>	<a href="#">23</a>
<a href="#">14.</a>	<a href="#">References . . . . .</a>	<a href="#">23</a>
<a href="#">14.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">23</a>
<a href="#">14.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">23</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">23</a>



## **1. Introduction**

The GMPLS overlay model [[RFC 4208](#)] specifies a client-server relationship between networks where client and server layers are managed as separate domains because of trustiness, scalability and operational issue. By means of procedures from the GMPLS protocol suite it is possible to build a topology in the client (overlay) network from Traffic Engineering paths in the server network. In this context, the UNI (User to Network Interface) is the demarcation point between networks. It is a boundary where policies, administrative and confidentiality issues apply that limit the exchange of information.

This GMPLS overlay model supports a wide variety of network scenarios. The packet over optical scenario is probably the most popular example where the overlay model applies.

In order to exploit the full potential of client/server network interworking in the overlay model, it may be desirable to know in advance whether is it feasible or not to connect two client network nodes [INTERCON-TE]. This requires to have a certain amount of TE information of the server network in the client network. This need not be the full set of TE information available within each network, but does need to express the potential of providing TE connectivity. This subset of TE information is called TE reachability information.

Reachability can be classified according to the available information:

- Non-qualified reachability: The most basic form of TE reachability is just the information of whether is it physically feasible to connect two nodes. Thus, Non-qualified reachability is the ability to be able to connect one end point to another.
- Qualified reachability: In addition to knowing of whether is it physically feasible to connect two nodes, the qualified reachability information indicates metrics of the potential connection. This metric can be a cost, a service metric (delay), bandwidth, etc.
- Reachability with associated Potential Virtual Link: In this case, the client node, in addition to have the information of the feasibility of reaching a given destination, it has the information of a possible path, that can be established at any time through UNI signalling. This possible path may have been pre-computed in advance and may contain the explicit path (e.g. ERO). Thus, client layer topology could be richened with this potential virtual TE link information.



Current standard GMPLS UNI [[RFC 4208](#)] is focused on signaling and extensions to the GMPLS UNI [[RFC4208](#)] are being proposed to gather a tighter integration between the client packet layer and the server optical one. Also, new elements, like the Path Computation element (PCE), have become more popular and may also play a role in the overlay context. In order to understand what are the requirements that need to be fulfilled, it is necessary to understand current use cases, that is, what are the network operators expecting the UNI to do.

The first group of uses cases of the overlay model is related to the automatic provisioning, by which the client (overlay) network is able to set-up a connection through the server network. The second group of use cases is related to the enhancement of the recovery mechanism by a higher coordination of overlay and server networks.

Summing up, the goal of this document is to overview the network scenarios where the overlay model applies and analyze the most interesting aspects of the use cases that fall under the above definitions, with particular focus on signaling, exchange of info, operations, path computation and L1VPN aspects.

## 2. Terminology

The following terms are used within the document:

- Edge node [[RFC4208](#)]: node of the client domain belonging to the overlay network, i.e. nodes with at least one interface connected to the server domain.
- Core node [[RFC4208](#)]: node of the server domain.
- Access link: link between core node and edge node. It is the link where the UNI is usually implemented.
- Remote node: node in the client domain which has no direct access to the server domain but can reach it through an edge node in its same administrative domain.
- Local trigger: LSP setup request issued to an edge node. It triggers the setup of a client layer FA through the server domain via a UNI interface.
- Remote trigger: LSP setup request issued to a remote node. It triggers the setup of a client layer LSP which, upon reaching an edge node, will use an FA through the server domain dynamically provided via an UNI interface.



### **3. Variants of UNI interface**

[RFC4208] defines the GMPLS UNI as an overlay interface allowing for the exchange of both routing and signaling messages between a client and a server domain. At that time only signaling extensions and procedures for the UNI interface were defined but since then a lot of RSVP-TE extensions have been defined (e.g. [LSP-DIV][TE-REC]).

In this section a taxonomy of different variants of UNI interfaces is provided.

#### **3.1. UNI**

GMPLS UNI defined in [RFC4208] allows the exchange of RSVP-TE Path messages between edge and core nodes including the Explicit Route Object (ERO) or Record Route Object (RRO). This allows for the explicit indication of the path (or part of it) to choose in the server domain or the recording of the nodes and links chosen respectively.

Subsequently new requirements have been added to the exchange of messages between edge and core nodes. In the message flow from edge to core nodes it is desirable to define a number of path computation constraints that the server domain needs to consider when providing a path that will be used as connection between two edge nodes. An example of path computation constraints consists of: SRLG diversity, max latency, max number of hops, etc. Encoding extensions for adding constraints to the server domain path computation are defined in [UNI-PLUS] and [LSP-DIV].

Similarly it is desirable that the core nodes are able to provide the edge nodes with the parameters qualifying the path provided. Such set of parameters is the same identified above and encoding extensions for its collection are defined in [TE-REC] and [SRLG-COLL].

RSVP-TE extensions are obviously inherited by the UNI interface, which has significantly been augmented with respect to the [RFC4208](#) version.

In the overlay model applied to the packet-opto environment the UNI can be implemented over grey or colored interfaces. In the former case the transponder is placed on the RROADM and the traffic from the router to the RROADM has no G.709 framing, while in the latter the transponder is moved on the router and the traffic injected into the WDM domain is G.709 framed and managed as an alien lambda. From a procedure point of view there is no difference between the grey and the colored UNI, what is different is the type of information that





the WDM PCE need to be provided with in order to compute a path in the WDM domain but check its feasibility from router to router.

For applicability considerations regarding the GMPLS UNI please refer to [UNI-APP]

### 3.1.1. UNI example

This section provides an example of how SRLGs, latency and other types of parameters can be used in the edge node to core node direction as path computation constraints and in the reverse direction from core node to edge node as path qualifiers. The example is applied to a packet-opto environment.

In the reference network below, suppose that router 1 wants RROADM A to provide a path between A and G with max unidirectional latency 20ms and SRLG different from (37;52). Such parameters are passed to A via the RSVP-TE Path message.

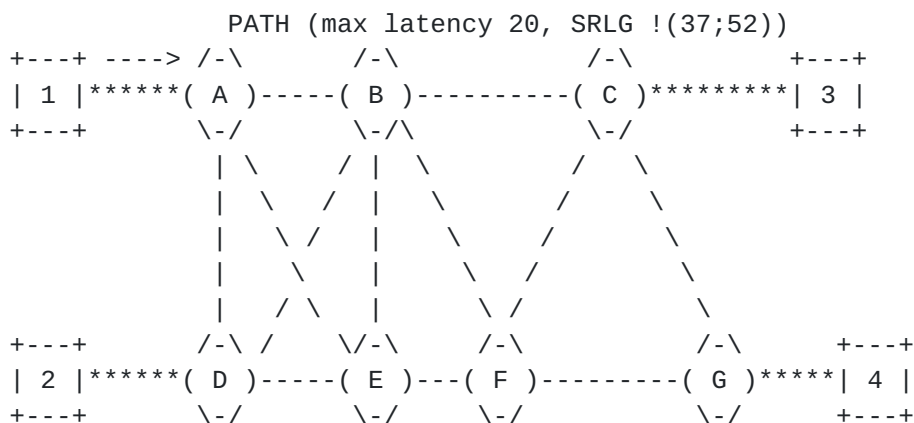


Figure 1: Path computation constraints

Node A performs a constrained path computation in the optical domain (A-B-C-G) and provides 1 (via the RSVP-TE RESV message) with the parameters qualifying the provided path e.g. max latency 12ms and SRLG (11,55).



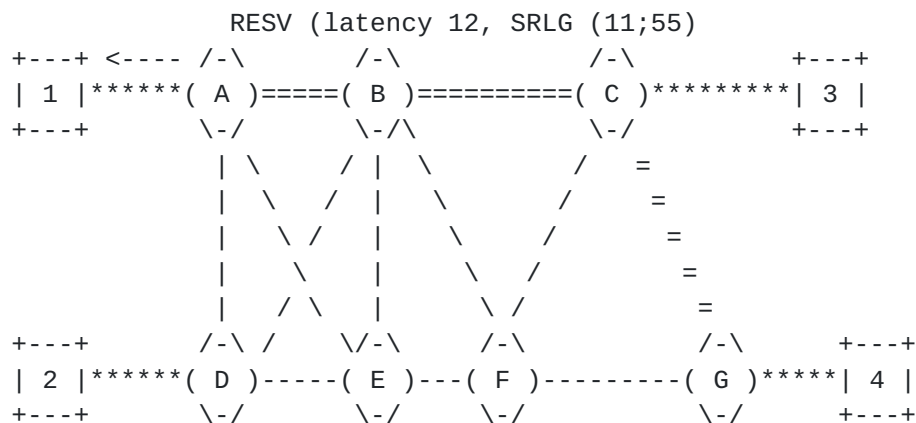


Figure 2: Real network topology

### 3.2. ONI

A different variant of GMPLS UNI interface can be implemented within the overlay model context. It foresees adding some TE (Traffic Engineering) topological information exchange between edge and core nodes. We use the term ONI (Overlay Network Interface) to identify this variant of UNI interface with such capabilities.

The ONI foresees the definition of a virtual topology in the server layer (arbitrarily configured by the network operator) which is exported by each core border node to its peering edge node by means of a routing protocol (e.g. OSPF-TE, BGP-LS). Such virtual topology comprises a mix of potential virtual TE-links [[RFC4847](#)] and virtual nodes.

A virtual TE-link, as defined in [RFC4847](#), is a provider network Traffic Engineering (TE) link advertised to customers in routing information for purposes that include path computation. We introduce the term potential virtual TE-link to indicate those virtual TE-links whose resources are not necessarily reserved or committed in the server layer network.

A potential virtual TE link is advertised via the ONI as a real TE link and hence contributes to augmenting the client layer topology. Moreover it does not require the allocation of resources in the server layer until it is really used thus allowing the mutually exclusive sharing of server layer network resources with other potential Virtual TE Links.

On the other side a virtual node is a collection of zero or more provider network domain nodes that are collectively represented to the clients as a single node that exists in the customer layer



network and is capable of terminating of access, inter-domain and virtual links. (a virtual node can correspond 1:1 to a physical node, to a part of it or to a set of nodes).

As per the UNI, we define two variations of ONI for the packet-opto scenario, namely the ONI when we refer to grey interfaces (i.e. transponder on the RROADM) and ONI++ when referring to colored interfaces (i.e. transponder on the router).

### 3.2.1. Routing info over the ONI

The following reference network in an IPoWDM environment is used to depict what kind of information needs to be advertised to the edge nodes in order to augment the client layer topology.




Figure 3: Real network topology

Where:

```

+---+
| 1 | = Edge node      ***** = Access Link (UNI link)
+---+

/-\
( A ) = Core node      ----- = Core domain real link
\-/

```

Suppose that the network operator decides to let client network see two potential virtual TE-links, one between A and C and the other between D and G. Figure below show the potential virtual TE-links on



the left and corresponding real topology on the right.

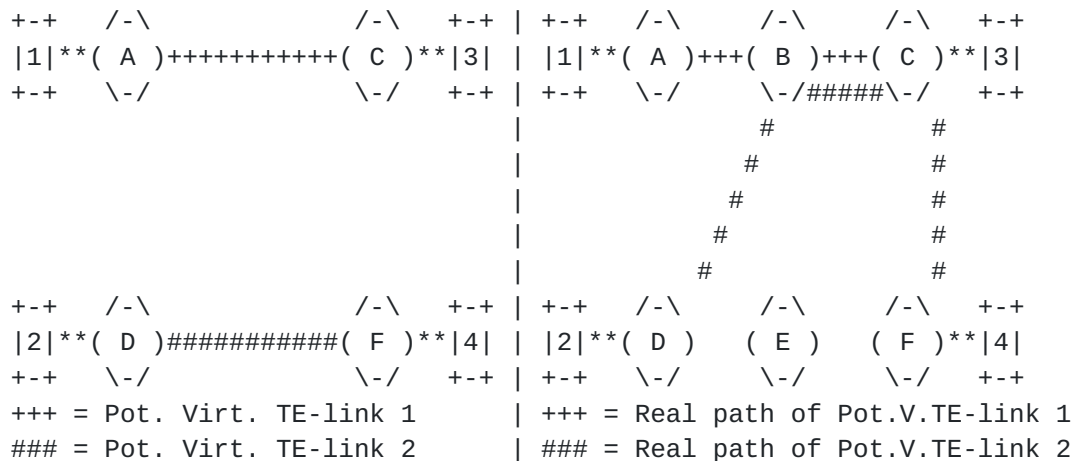


Figure 4: Virtual network topology

From the above figure it is possible to see that the edge nodes, in order to be able to compute a path, need to know the following:

- Access links: availability and TE parameters
- Potential virtual TE-links: availability, TE parameters, diversity, constraints (e.g. latency, packet loss)
- Virtual nodes: (in figure above a 1 to 1 relationship between real and virtual nodes is assumed) availability, connectivity matrix, TE parameters, diversity, constraints (in those cases where a virtual node is composed of 2 or more real nodes)

More details on potential virtual TE links diversity can be found at [MELG].

#### 4. Hybrid Nodes

The overlay model and interfaces can be applied also to hybrid nodes [RFC5212], where it is possible to have different control plane instances in the client and server domain and keep on exploiting the advantages of managing the domains separately without losing the capability of making them tightly interwork.





## **5. Adding the PCEP to the UNI**

In the GMPLS overlay model, the edge nodes do not participate in the routing instance of the core network. This means that the overlay network is not aware of the details (topology and TE information) of the core network. Thus, it is generally assumed that, in the case of a UNI request, the routing is performed in the core network.

The operator is able to include a set of constraints that have impact on the routing that will be performed in the core network in two different ways. One is to enrich RSVP-TE to include such constraints (please see above), and the other one is use the well-defined PCEP protocol.

The main advantage of using the PCEP interface is that the semantics for the different queries are well defined. Moreover, in case confidentiality and trustiness is an issue, the Path Key mechanism can be used in the answers to hide the details of the path. Different possible scenarios are shown bellow.

### **5.1. Edge-node to core-node PCEP interface**

The first scenario foresees the provisioning of a PCEP interface from the core nodes to the edge nodes. Such interface is reachable from the same interfaces used for the exchange of UNI signaling messages. The path computation in the server domain can be centralized or distributed. In the former scenario the core node acts as a proxy between the PCC on the edge node and the PCE as shown in figure below.






Figure 5: PCEP proxy on core node

On the other hand, when the path computation in the server domain is performed in a distributed way, the PCEP interface can be exported by the core nodes so to allow the edge node issue path computation queries. See figure below.

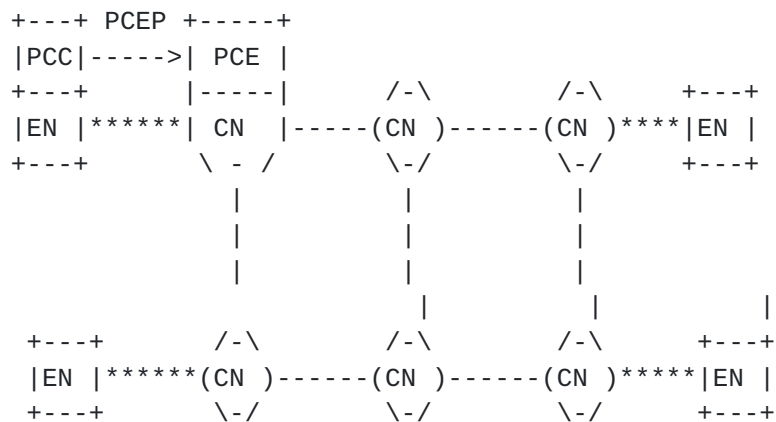


Figure 6: PCEP interface between EN and CN

In those cases where the addressing space is common between the client and the server layer, and hence with the PCE, it is also possible to have the PCC on the client layer directly speaking with the server domain PCE as shown in figure below.






Figure 7: PCE and PCC direct connection

## 5.2. PCEP and colored UNI

TBD

## 5.3. Using PCEP to obtain TE reachability info

Although the overlay network does not participate in the routing instance of the server network, [RFC 4208](#) allows TE reachability to be exchanged between both overlay and server networks. TE reachability, as defined in [INTERCON-TE], is the ability to reach a specific address along a TE path. In the overlay context, TE reachability indicates if it is possible to reach from one node of the overlay network to another one through the server network. Also, TE reachability can indicate some TE metrics of the possible paths metrics, hop count, available bandwidth, delay, shared risk.

The PCEP interface can be used at the UNI border to query from overlay network to server network about the possible paths between nodes. The answers do not need to include a full detailed ERO, but just a collection of the desired metrics and, either a path key, or an ERO with just the end-points.

## 6. Provisioning

Two different use cases can be identified for path provisioning depending on where the setup trigger is issued. They will be called Local and Remote Trigger in the following. Further details regarding



different provisioning models over the UNI can be found in [UNI-APP]. What described below is in addition to the simple provisioning use case without constraints defined in [RFC4208].

### 6.1. Use case P1 - Local Trigger

Local Trigger is the use case defined in [RFC4208], where a setup trigger is issued to one of the edge nodes belonging to the overlay network, i.e. directly connected to the UNI.




Figure 8: Local trigger

As it is possible to see in the figure above, a trigger is issued on R3 (edge node) for starting the setup request procedure over the UNI interface (R3-A). Such trigger might include an arbitrary set of constraints for the path computation in the optical domain. Once the optical LSP is setup and an adjacency in the packet layer between R3 and R5 is created, it is advertised in the rest of the packet layer and used by the signaling protocol (e.g. LDP) for setting up end-to-end (e.g. from R1 to R7) packet LSPs.

### 6.2. Use case P2 - Remote Signaling

A second use case consists on the utilization of a connection oriented (e.g. RSVP-TE) signaling protocol in the packet domain that allows issuing the setup trigger directly on the end nodes of the packet domain and using the RSVP-TE message for triggering the setup of the LSP in the optical domain.





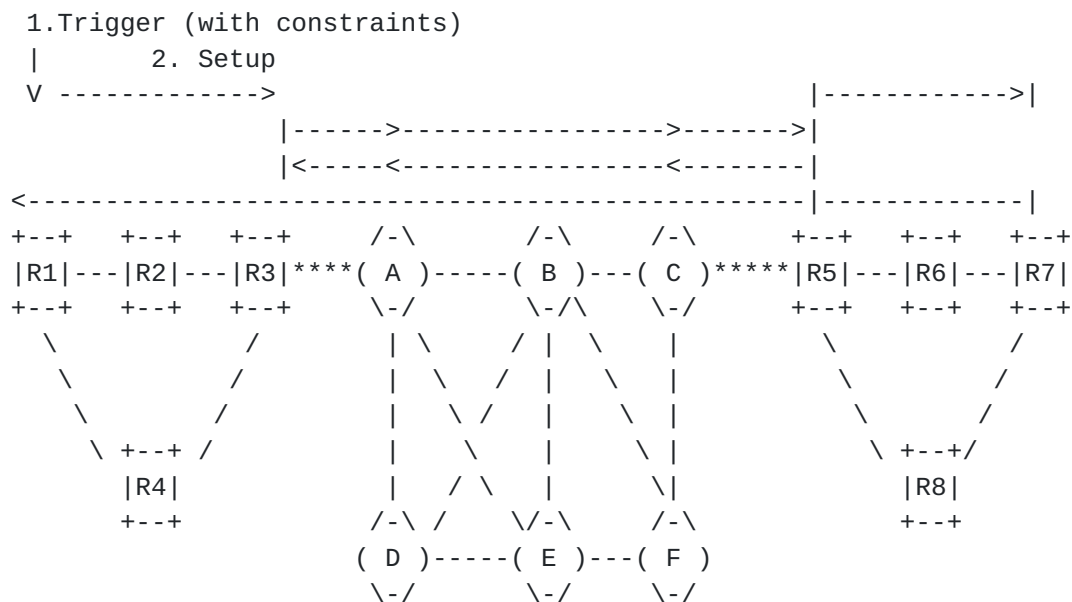


Figure 9: Local trigger

The utilization of the remote trigger allows for a strict control of the resources that will be used for the setup of the end to end path. In particular, two further approaches are possible when using the remote trigger: RSVP-TE or BGP approach. In the case of RSVP-TE it is possible to exploit the capabilities of the protocol of signaling LSPs across multiple domains and indicate in the ERO (explicit route object) the resources that need to be included in the path. It is possible to indicate a strict ERO (i.e. all the indicated resources and only those ones must be used) or a loose ERO (i.e. at least the indicated resources must be used but the PCEs of the different domains decide autonomously how to fill the gaps in the ERO).

### 6.3. Use case P3 - Provisioning with requirements over the UNI

This use case applies to both the local and remote trigger scenarios and describes how it is possible for the client layer to ask for a connection in the server layer with some constraints. An example of the constraints that can be applied to the path computation in the server layer consists of the following criteria:

- + Diversity: it is possible to indicate the resources must or should be avoided during the path computation by means of the Exclude Router Object (XRO) [RFC4874], the Explicit Exclusion Route Subobject (EXRS) [RFC4874] and the LSP subobject [LSP-DIV]. Such resources consists on:



- IPv4 prefix [[RFC4874](#)]
- IPv6 prefix [[RFC4874](#)]
- Unnumbered Interface ID [[RFC4874](#)]
- Autonomous system number [[RFC4874](#)]
- SRLG [[RFC4874](#)]
- IPv4 P2P subobject [LSP-DIV]
- IPv6 P2P subobject [LSP-DIV]
- + Latency: max delay introduced by the server layer LSP [OF-TEMB]
- + Latency variation: max latency variation allowed for the server layer LSP [OF-TEMB]
- + Cost: max cost allowed for the server layer LSP [OF-TEMB]

The overlay Edge Node can include into the RSVP-TE Path message an arbitrary number of path computation constraints for the provisioning of the LSP in the server domain. In figure below and example is shown using as path computation constraints: (i) max latency should be 20ms and (ii) SRLG must be different from (A;B).

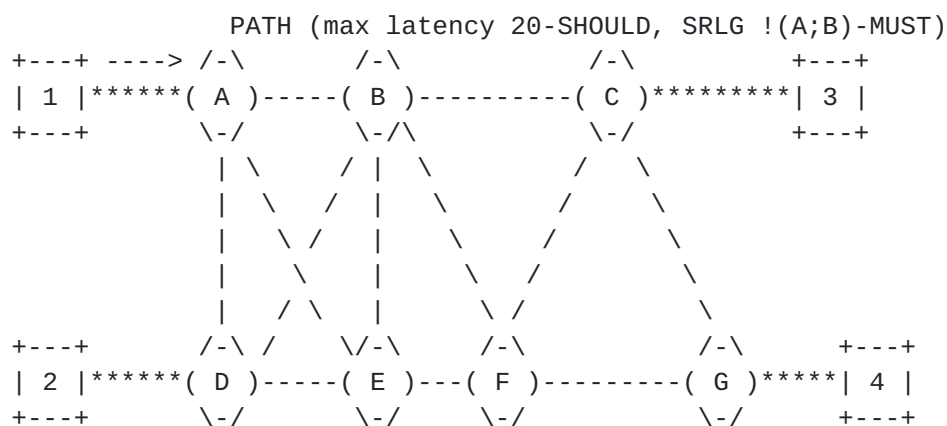


Figure 10: Use Case P3 - provisioning with requirements

If the path computation in the server domain is able to provide an LSP meeting the requirements (at least the must ones) such LSP is established and a RESV message is returned to the Edge node, otherwise an error message (PattErr) is returned.



#### **6.4. Use case P4 - Provisioning with requirements and collection request over the UNI**

This use case is an extension of Use case P3. In addition to the path request with path computation constraints, the client layer is also allowed to request for the collection in the server domain of the effective values of the parameters indicated as path computation constraints. The collection of such parameters is indicated via dedicated flags in the LSP\_ATTRIBUTES and LSP\_REQUIRED\_ATTRIBUTES in the Path Message. Flags defined are:

- Cost collection flag [TE-REC]
- Latency collection flag [TE-REC]
- Latency Variation collection flag [TE-REC]
- SRLG collection flag [SRLG-COLL]

In the scenario depicted in figure Figure 10 a request with constraints on max latency and SRLG diversity might be issued together with the request of collecting e.g. the effective SRLGs of the provided path. Collected parameters are returned to the overlay edge node via the Record Route Object (RRO) in the RESV message.

#### **6.5. Use case P5 - Advertisement of collected metrics in the client layer**

In the case of local trigger, that is when a forwarding adjacency (FA) is created between two nodes in the client domain by means of an LPS in the server domain, it is desirable to advertise the characteristics of such FA in the client domain so that an informed path computation in the client layer can be performed.

Interior Gateway Protocol (IGP) extensions for the advertisement of the TE metrics characterizing such forwarding adjacencies have been defined for IS-IS [ISIS-TEMET] and OSPF-TE [OSPF-TEMET].

It is also possible to perform such advertisement by means of BGP-LS, where the collected data are sent to the route reflector and then forwarded to all the BGP speakers in the different client domains.

#### **6.6. Use case P6 - Provisioning with requirements over the ONI**

In the case of ONI/ONI++ each overlay edge node is provided with the overlay topology of the server domain as indicated in section [Section 3.2](#). When a trigger is received by the edge node (local or remote), it is able to choose among the different available paths



depending on which one meets the constraints included into the trigger. When a path is selected, the signaling procedure occurs accordingly with a standard RSVP-TE signaling procedure including a strict (or loose) ERO. For more details please see [ENNI].

### 6.7. Use case P7 - Dual homing

A quite common case is the need for provisioning two (or more) optical paths with different ingress and egress nodes (i.e. different CNs) with diversity constraints. This is needed to provide two totally disjoint paths in the client layer without any kind of single point of failure (e.g. a single UNI link or a single edge node) in order to create protection schemes in the packet layer like e.g. 1+1 protection.

In this case some sort of info exchange in the client domain is needed and can be provided in two ways:

- Coordination between 1 and 3
- Coordination between 2-1 and 2-3 [UNI EXT]

In both cases node 1 is supposed to ask the optical domain to provide a path (e.g. restorable [[RFC4873](#)]) and collect its SRLGs, then pass such SRLGs info to node 3 (directly or via 2). Hence node 3 will ask the optical domain to provide a path (e.g. restorable) avoiding the given SRLGs.




Figure 11: Use Case P7 - Dual homing

As per previous use cases it is possible to either provide the path between 1 and 4, collect the SRLGs, provide the path between 3 and 6





with SRLG diversity constraints and then let the packet domain use those two packet layer adjacencies or have 2 asking 1 and 3 to respectively set up 1-4 and 3-6.

Whenever one of the paths is restored in the optical domain (e.g. B-C fails and the path is restored using A-B-E-C), an SRLG collection procedure (by 1) and exchange (to F) is needed so that a possible restoration of the protection path would avoid the new SRLGs of working path. Please see use case R4.

## 7. Recovery

### 7.1. Use case R1 - Segment Recovery - Single homing

This is one of the simplest recovery scenarios, where each overlay node has a single UNI interface available on the overlay nodes and hence only the server layer is in charge of restoring or protecting the traffic within its domains boundaries. As it is possible to see in figure below the client layer is able to recover from failures within the server domain.

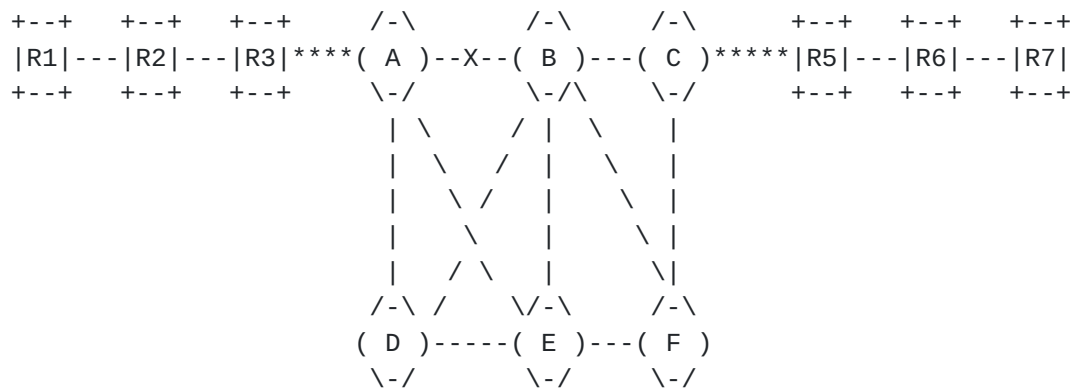


Figure 12: Use case R1 - Segment Recovery - Single homing

This kind of scenario is extremely cheap in terms of resource utilization but has some limitations due to a high number of single points of failures which, in case of unavailability, prevent any type of recovery attempt. The single points of failure are: the overlay nodes (R3 and R5), the UNI links (R3-A and C-R5) and the server domain border nodes (A and C). In the following sections higher degrees of reliability are provided.



### 7.2. Use case R2 - Local recovery - Dual Homing - Single overlay node

In this use case it is possible to combine local recovery in the overlay nodes (e.g. Fast Rerouting) with restoration or protection in the optical domain so to be able to react to failures not only in the client domain but also affecting the server layer border nodes (A and C) and the UNI links (R3-A and R5-C). The number of single points of failure is hence significantly reduced to just the overlay nodes (R3 and R5).

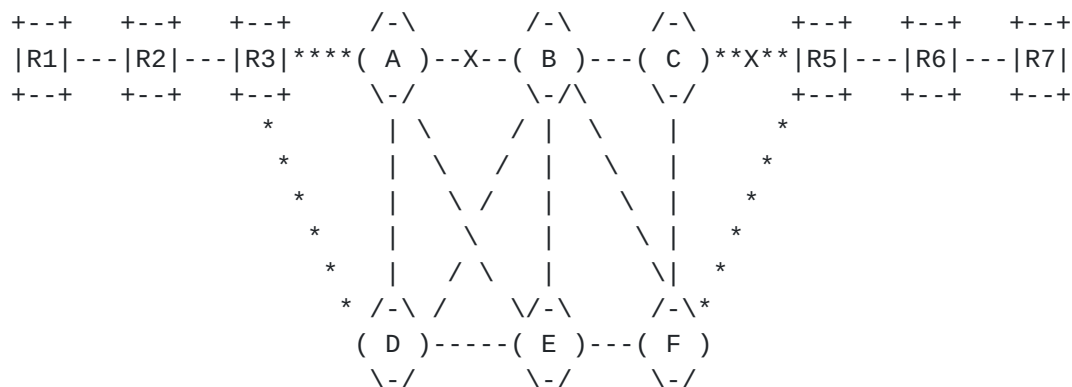


Figure 13: Use case R2 - Local recovery - Dual Homing - Single overlay node

### 7.3. Use case R3 -End to end Recovery - Dual homing - Double overlay node

This use case focuses on recovery just in the client layer, no recovery is requested to the server layer but just the provisioning of LSPs with requirements (in this case diversity).

The provisioning of the working path occurs like the provisioning of LSP over the UNI with remote trigger and with the collection of SRLGs enabled. Once that the optical LSP supporting the working client LSP is setup (A-B-C), the SRLGs in the optical domain are collected and passed to the overlay node (R3) which forwards them to the ingress node (R1) in the reverse path of the signaling process of the working LSP (i.e. RSVP-TE Resv Message).

Once that R1 is provided with the SRLGs in the optical domains, it is able to issue the signaling of the protection LSP indicating in a loose ERO R5 and R7 as loose hops and the SRLG diversity in the ERO expansions between them. This will lead in the provisioning of a path in the photonic domain diverse from the previous one (S2-S4-S6) and hence two total diverse end to end paths. In this case the



network will be able to recover to a single failure in any part of the network without any single point of failure.




Figure 14: Use case R3 -End to end Recovery - Dual homing - Double overlay node

#### **7.4. Use case R4 - Combination of client protection and server restoration**

This is the most interesting use case when the overlay model is applied to the packet/opto scenario. The packet protection and optical restoration allows for a fast recovery of the traffic upon a failure in the network, while the restoration in the optical domain allows for always providing the packet layer with two available path against an arbitrarily high number of failure in the optical network (the number of failures depend on the meshing degree of the photonic network).

Moreover the exchange of SRLGs over the UNI interface (please see Provisioning use cases) makes sure that the optical path used for the working branch of the packet protection and the optical path used for the protection branch of the packet protection do not share any resource and hence that a single failure in the optical domain does not cause any traffic hit; i.e. traffic can be recovered with packet protection time (tens of ms) instead of photonics restoration time (tens of s).

The SRLG collection procedure is pretty simple during the provisioning phase, but more complicated when one of the paths fail and is restored.

In figure below an example is considered where a protection is provided in the client layer where the working path is setup along path (2-1-A-C-4-5) and its protection path along (2-3-F-H-G-6-5).



After the setup of the optical path between A and C (A-B-C), the SRLGs of such path are collected and the overlay edge node 3 can ask F for the setup of a path between F and H avoiding such SRLGs. This ensures that a single failure in the optical domain will not affect both the W and P.




Figure 15: Use Case R4 - Setup of paths in full diversity

In order to guarantee that different SRLGs will always be used for W and P it is necessary to perform two different steps after the provisioning of the two LSPs:

1. Each ingress node of the server domain paths (namely A and F) must be informed that in case of restoration the path computation must avoid the *actual* SRLG of the other path
2. Every time that either of the server domain paths is restored, the collection of SRLGs must be performed and communicated to the ingress node of the other LSP.








Figure 16: Use Case R4 - Restoration of paths keeping full diversity

In the example above a failure affecting W (link B-C) and a failure affecting P (link F-G) occur, but the server network is able to provide alternate paths for both of them and still provide the same level of fault resiliency to the client network.

## 8. Optimization

### 8.1. Use case 01 -

TBD

## 9. Security Considerations

TBD

## 10. IANA Considerations

TBD

## 11. Contributors

Diego Caviglia

Email: [diego.caviglia@ericsson.com](mailto:diego.caviglia@ericsson.com)



Francesco Fondelli

Email: francesco.fondelli@ericsson.com

## **12. Acknowledgements**

The authors would like to thank Manuela Scarella and Enrico Dutti for the comments and review.

## **13. References to be added**

[EDITOR NOTE] : section to be fixed

[LSP-DIV] - <http://tools.ietf.org/html/draft-ietf-ccamp-lsp-diversity-01>

[TE-REC] - <http://tools.ietf.org/html/draft-ietf-ccamp-te-metric-recording-01>

[OF-TEMB] - <http://tools.ietf.org/html/draft-ali-ccamp-rc-objective-function-metric-bound-02>

[SRLG-COLL] - <http://tools.ietf.org/html/draft-ietf-ccamp-rsvp-te-srlg-collect-02>

[MELG] - <http://tools.ietf.org/html/draft-beeram-ccamp-melg-00>

[UNI-APP] - <http://tools.ietf.org/html/draft-zhang-ccamp-gmpls-uni-app-03>

[ENNI] - <http://datatracker.ietf.org/doc/draft-beeram-ccamp-gmpls-enni/>

[ISIS-TEMET] - <http://tools.ietf.org/html/draft-previdi-isis-te-metric-extensions-03>

[OSPF-TEMET] - <http://tools.ietf.org/html/draft-ietf-ospf-te-metric-extensions-04>

[INTERCON-TE] - <http://tools.ietf.org/html/draft-farrel-interconnected-te-info-exchange-00>

[UNI EXT] - <http://tools.ietf.org/html/draft-fedyk-ccamp-uni-extensions-01>

## **14. References**

### **14.1. Normative References**

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

### **14.2. Informative References**

#### Authors' Addresses

Daniele Ceccarelli  
Ericsson  
Via A. Negrone 1/A  
Genova - Sestri Ponente  
Italy

Email: [daniele.ceccarelli@ericsson.com](mailto:daniele.ceccarelli@ericsson.com)

Ceccarelli & Gonzalez de Dios Expires January 5, 2014

[Page 23]

Oscar Gonzalez de Dios  
Telefonica I+D  
Don Ramon de la Cruz 82-84  
Madrid 28045  
Spain  
  
Email: ogondio@tid.es