Network Working Group
Internet Draft
Category: Informational

Fatai Zhang
Huawei
O. Gonzalez de Dios
Telefonica Investigacion y Desarrollo
A. Farrel
Old Dog Consulting
Xian Zhang
Huawei
D. Ceccarelli
Ericsson

Expires: January 08, 2014

July 09, 2013

Applicability of Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI)

draft-zhang-ccamp-gmpls-uni-app-04.txt

Status of this Memo

Abstract

Generalized Multiprotocol Label Switching (GMPLS) defines a set of protocols for the creation of Label Switched Paths (LSPs) in various switching technologies. The GMPLS User-Network Interface (UNI) was developed in RFC4208 in order to be applied to an overlay network architectural model.

This document examines a number of GMPLS UNI application scenarios. It shows how techniques developed after the GMPLS UNI can be applied to automate or enable critical processes for these applications. This document also suggests simple extensions that could be made to existing technologies to further enable the UNI and points out some unresolved issues.

Table of Contents

1. Introduction

Generalized Multiprotocol Label Switching (GMPLS) [RFC3945] defines a set of protocols, including Open Shortest Path Fist - Traffic Engineering (OSPF-TE) [RFC4203] and Resource ReserVation Protocol - Traffic Engineering (RSVP-TE) [RFC3473], which can be used to create Label Switched Paths (LSPs) in a number of deployment scenarios with various transport technologies.

The User-Network Interface (UNI) reference point is defined in the Automatically Switched Optical Network (ASON) [G.8080]. According to [G.8080], the UNI may be implemented as a peering between a client-side entity (UNI-C) and a network-side entity (UNI-N). End-to-end connectivity between UNI-C nodes is achieved across the core network by three components: a UNI request from source UNI-C to source UNI-N; a core network connection from source UNI-N to destination UNI-N; and a UNI request from destination UNI-N to destination UNI-C.

The GMPLS overlay model, as per [RFC4208], can be applied at the UNI, as shown in Figure 1.

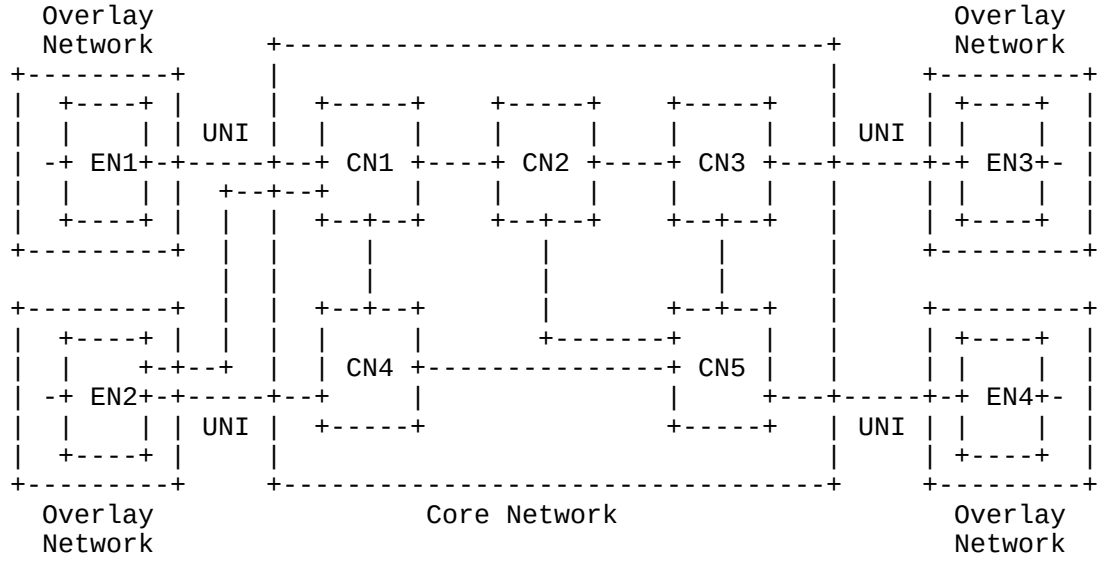

Figure 1 - Applying GMPLS overlay model at UNI

In Figure 1, assume that there is an end-to-end UNI connection passing through EN1-CN1-CN2-CN3-EN3. For convenience, some terms used in this document are defined below:

- "source EN" refers to the edge-node which initiates the connection (i.e., EN1);

"destination EN" refers to the edge-node where the connection is terminated (i.e., EN3);

- "ingress CN" refers to the core-node to which the source EN is attached (i.e., CN1);

- "egress CN" refers to the core-node to which the destination EN is attached (i.e., CN3).

[RFC4208] provides mechanisms for UNI signaling, which are compatible with GMPLS RSVP-TE signaling ([RFC3471] and [RFC3473]). A single end-to-end RSVP session between source EN and destination EN is used for the user connection, just as it would be for connection creation between two core nodes. However, when considering the isolation of topology information between the core network and the overlay network, additional processing of the RSVP-TE Explicit Route Object (ERO) and Record Route Object (RRO) is required. For example, the ingress CN should verify the ERO it receives against its topology database and may enhance it with additional path information before forwarding the PATH message. And the ingress/egress CN may edit or remove the RRO in order to hide the path segment used inside the core network from the EN.

The GMPLS UNI can be used in many application scenarios. For example, in a multi-layer network [RFC6001] the interface between client layer node and server layer node can be seen as a UNI. Or, when deploying VPN services such as Layer One Virtual Private Networks (L1VPNs) [RFC4847], [RFC5253], users can connect to a service provider network via a UNI.

This document examines a number of current and future GMPLS application scenarios. It shows how techniques developed after the GMPLS UNI can be used to automate or enable critical aspects of these application scenarios. It points out some potential technology extensions that could improve UNI operation, and highlights some unresolved issues.

2. UNI Addressing

In [RFC4208], the GMPLS overlay model is applied at the UNI reference point, and it is required that the edge-node and its attached core-node of the overlay network share the same address space that is used by GMPLS to signal between the edge-nodes across the core network. Under this condition, the user connection can be created using a single end-to-end RSVP session, which is consistent with the RSVP model. Therefore, RSVP-TE defined in [RFC3473] can be used for support GMPLS UNI without any extensions.

However, in some deployments of the GMPLS UNI, it is not practical for the EN and its attached CN to share the same address space. This can arise if the core and overlay networks were designed and deployed separately or belong to different carriers. For example, the core network may use IPv6 addresses, while the overlay network uses IPv4 addresses. Or, since the core network is a closed system, the assignment of the IP addresses of the CNs may be independent of other IP addresses outside the core network. This implies that the nodes in the core network may use addresses which could collide with the edge nodes in the overlay network.

[RFC4208] does not state how to ensure that an edge-node and its attached core-node share the same address space. This document analyses the addressing deployment scenarios as follows:

1. Overlay network and core network share a common addressing policy. This might be quite feasible in a multi-layer network operated by a single carrier.

   In this scenario, end-to-end UNI connectivity may use a single RSVP session, and the core routing information (assuming it is shared and not stripped for confidentiality reasons) will be meaningful to the ENs. Note, however, that the overlay model examined by this document assumes that there is some separation between the overlay and core networks, and this might mean that the overlay network is not able to see the topology or routing information of the core network even when they share a common address space.

2. ENs have visibility into the core network, but overlay and core networks have different address spaces. This is the more common model envisaged by [RFC4208] and for basic mode L1VPN deployments [RFC5251]. The previous scenario can be seen to be a special case of this scenario where the two address spaces are complementary. In this deployment the ENs each have two addresses: one in the overlay network and one in the core network. The source EN is

aware of the addresses for itself, the ingress CN, the egress CN, and the destination EN in the address space of the core network. It may also have full visibility into the core network, but this is not a requirement.

In this scenario, the ENs are responsible for performing address mapping between the overlay network's addresses for the ENs, and the core network's addresses for the same nodes and/or its TE links. A typical deployment may assign addresses in the core network address space for the EN and/or its TE links at the EN side, so that EN can use these addresses to communicate with the core network for UNI connection provisioning.

In this deployment, a single end-to-end RSVP-TE session can still be utilized from the source EN to the destination EN using addressing and naming from the core network's address space.

3. ENs do not have any knowledge of the core address space, or do not support the address space the core network uses (e.g., ENs do not support IPv6 that is used by the core network). ENs will have no visibility into the core network.

   In this scenario, the ingress CN is responsible for mapping addresses to the core address space and filling in any additional routing information. A typical deployment is to assign addresses in the overlay address space for the ingress CN and/or its TE links at the CN side, so that the EN can use overlay addresses to reach the ingress CN and to identify the destination EN.

   In this deployment the end-to-end connectivity must be created either using "session stitching" (see Section 6.2) or "session shuffling" (see Section 6.3).

3. UNI Auto Discovery

When the end-to-end connection is set up across the core network, it must be targeted at the destination CN so that it can be extended to the destination EN. This means that either the source EN must know the identity of the destination CN to which the destination EN is attached, or the source CN must know this information. This requires some form of "discovery" (possibly including configuration), and depending on the addressing scheme in use (see Section 2), address mapping needs to be performed by the source EN or the source CN.

The discovery problem may be exacerbated when a variety of services are requested since the source EN will need to know the capabilities and available resources on the link between the destination CN and the destination EN. It could discover this by attempting to set up a connection and by drawing conclusions from connection setup failures, but this is not efficient. Furthermore, in the case of a dual-homed destination EN (such as EN2 in Figure 1), a choice of destination CN must be made, and that choice may be influenced by the capabilities and available resources on the CN-EN links leading to the destination EN.

If the UNI is applied in an L1VPN scenario, two mechanisms for auto discovery have been defined. Auto discovery of UNI using OSPFv2 is provided in [RFC5252] using an L1VPN LSA to advertise the L1VPN information via the L1VPN info TLV and the TE information of the CE-PE link (in the language of UNI, it's the EN-CN link) via the TE link TLV. Auto discovery of UNI using BGP is provided in [RFC5195] by having each edge CN advertise to other edge CN the following information, at a minimum: its own IP address and the list of <private address, provider address> tuples local to that PE. Once that information is received, the remote PEs will identify the list of VPN members they have in common with the advertising PE, and use the information carried within the discovery mechanism to perform address resolution during the signaling phase of Layer-1 VPN connections.

4. UNI Path Computation

End-to-end UNI path computation includes three parts: the selection of the source UNI link, the path computation inside the core network and the selection of the destination UNI link.

The selection of UNI links may not be necessary in all scenarios. One example is in the case of single-homing with only one UNI link between EN and CN. Another example is manual selection of the UNI link when the service is requested (i.e., as a function of the service request such as the port mapping used in a L1VPN). In such cases, the CN to which the source EN is attached, or the path Computation Element (PCE) ([RFC4655]) which is responsible for the core network, can perform the path computation across the core network when the UNI signaling request is sent from the source EN to the source CN.

4.1. UNI Link Selection

This document is specific to the overlay architectural model, and that means that the source EN does not have the topology and TE information of the core network. Therefore, in the case of multi-homing (i.e., the source EN is connected to more than one CN), the source EN does not have enough information to make a correct choice among all the UNI links between itself and the core network for an optimal end-to-end connection.

In this case, a PCE whose computation domain covers both the core network and the ENs attached to it can be used. Note that the GMPLS UNI predates PCE and hence a PCE was not available in early GMPLS UNI deployments. A PCE that has the topology and TE information of the core network can use the UNI discovery mechanism described in Section 3 to learn the EN-CN relationship and the TE information of the UNI links, and therefore has the ability to select the optimal UNI link for the connection.

Figure 2 shows the procedures for UNI path computation using a single PCE with visibility into the core network and information about all of the CN-EN links. When the UNI path computation request is received, the PCE can help the source EN to compute the end-to-end route of the UNI connection based on routing information it has accesses to, so that the source EN can create the UNI connection using the optimal UNI link. As shown in Figure 2, the following steps are carried out:

Step 1: EN1 requests a path from EN1 to EN2 by sending a PCReg message to the PCE;

Step 2: The PCE computes a path based on its view of the core network and knowledge of all the EN-CN links. In this case, it returns the path En1-CN4-CN5-Cn6-EN2 to the EN1 node;

Step 3: EN1 starts the signaling process to set up the LSP by using s standard RSVP signaling process, using the path information as computed.

If confidentiality of the topology within the core network needs to be preserved, the Path Key Subobject (PKS) can be used for either approach outlined here (see [RFC5520] and [RFC5553]). In the PCRep message returned to EN1, the Confidential Path Segment (CPS) (i.e., CN4-CN5-CN6) is encoded as a PKS by the PCE. Therefore, EN1 only learns the selected UNI link from the PCE. When CN4 receives the UNI signaling message from EN1 carrying the PKS, CN4 asks the PCE to decode the PKS and then continues to signal the LSP.

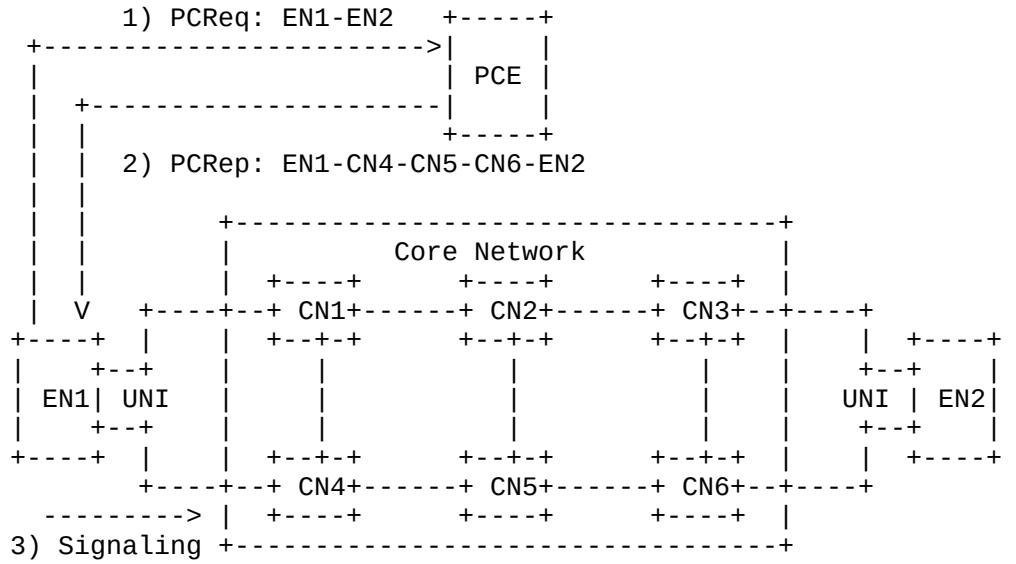

Figure 2 - Procedure using a PCE for UNI path computation

Note that the case described in this section, the PCE needs to be visible to the ENs, and there also needs to be a control channel between the PCE and the ENs for the exchange of PCE Protocol (PCEP) messages. An alternative implementation could be that a PCE is located inside each CN to which the source EN is attached, so that the source EN can use the UNI control channel to send and receive the PCEP messages.

The node requesting for a LSP, crossing UNI, may not be an EN node, as depicted in Figure 3. The procedure described above still applies.

In this case, if an explicit route is desired there is an additional requirement that the PCE needs to have visibility into the overlay networks. Otherwise, the PCE can only provide the route between two EN nodes as illustrated in Figure 3.

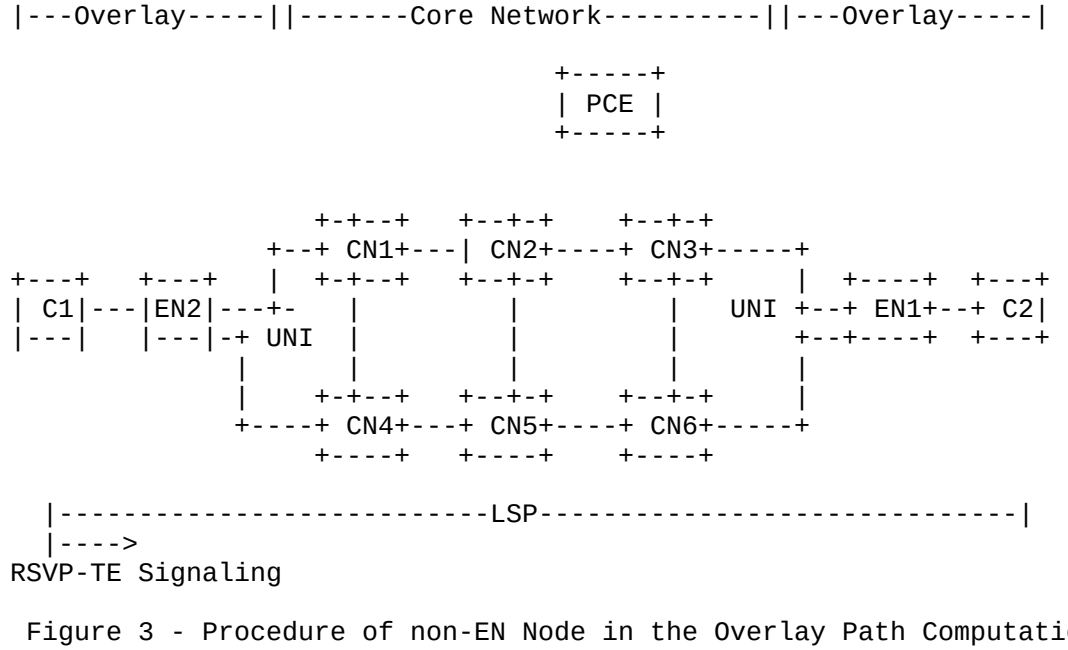

Figure 3 - Procedure of non-EN Node in the Overlay Path Computation

5. Additional Parameters across UNI

With new extensions currently proposed for RSVP-TE protocol, new parameters/functions can also be applicable to UNI.

5.1. Constrained Path Computation

Constraints that can be applied to the path computation in the core network are:

+ 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 can consist of:

-IPv4 prefix, IPv6 prefix, Unnumbered Interface ID, AS Number and SRLG [RFC4874]

-IPv4 P2P subobject and IPv6 P2P subobject [LSP-DIV]

+ Latency, Latency Variation and Cost: max delay/delay variation and cost allowed by the server layer LSP [UNI PLUS]

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. For example, in Figure 2, EN1 can request a path with a constraint: max latency should be 200ms.

If the path computation in the core network is able to provide an LSP meeting the requirements (at least those requirements which must be met) such LSP is established and a RESV message is returned to the Edge node; otherwise an error message (PathErr) is returned.

Use cases described in Section 7 can be viewed as a special use case of diversity.

5.2. Collection Requests over UNI

In addition to the path request with path computation constraints, the overlay nodes can also request for the collection in the core network 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 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-FA]

In the scenario depicted in Figure 2 a request with constraints on max latency might be issued together with the request of collecting e.g. the effective SRLGs of the provided path, in order to set up a SRLG-disjoint recovery path, as explained in Section 7. Collected parameters are returned to the overlay edge node via the Record Route Object (RRO) in the RESV message.

6. UNI Path Provisioning Models

The basic GMPLS UNI application is to provide end-to-end connections between edge-nodes through a core network via the overlay model. This section briefly describes four ways in which the end-to-end LSP can be created and operated across the core network.

6.1. Flat Model

In this model, the edge-nodes have the same switching capability as the nodes in the core network. In this case, one single end-to-end RSVP session through the edge-nodes and a series of core-nodes can be used to create the connection, which forms a flat LSP model, as shown in Figure 4.

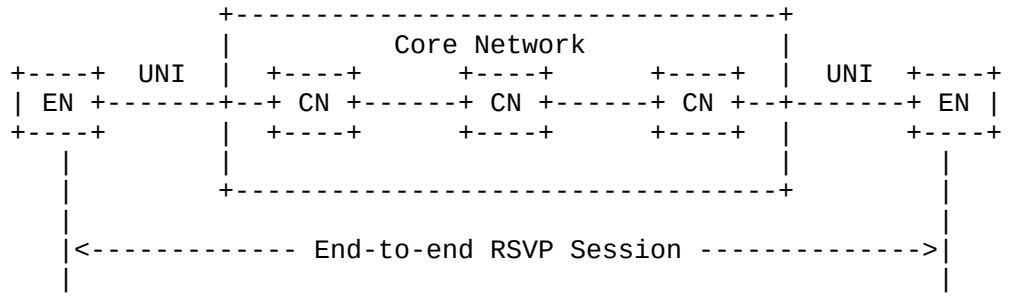

Figure 4 - The Flat Model

If the edge-nodes and their attached core-nodes share the same address space, or the ENs can perform address mapping into the core network address space, the GMPLS signaling described in [RFC3471], [RFC3473] and other related specifications, with special ERO and RRO processing as described in [RFC4208], can be used to create a connection. Note the procedures mentioned still apply in the scenarios where the source node of a connection is not an edge-node but rather nodes within the same domain as EN.

6.2. Stitching Model

The stitching mechanism described in [RFC5150] can be used to create an LSP segment (S-LSP) between the ingress and the egress CN, and to stitch the end-to-end UNI connection to the created S-LSP, as shown in Figure 5.

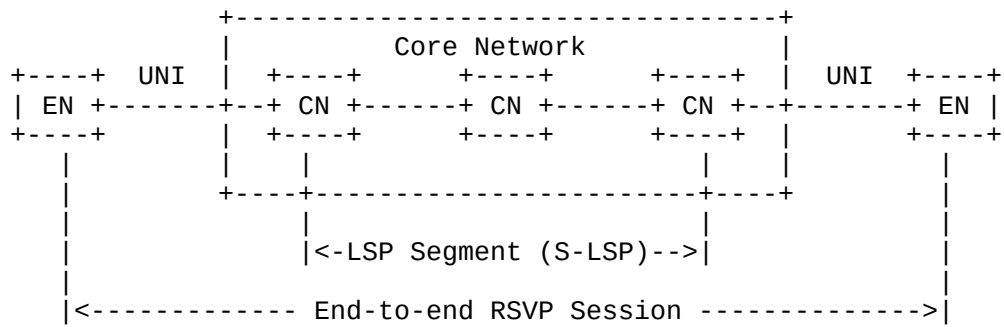

Figure 5 - The Stitching Model

This model allows the core network a degree of independence so that the S-LSP can be set up and modified without the knowledge of the overlay network. Remember that stitching is a data plane function, so that the EN-CN LSP segments are cross-connected to the S-LSP at the edge CNs. This means that, just as in Section 6.1, the overlay and core networks must have the same switching capabilities. However, the control plane for the stitching model operates just as the hierarchical model described in Section 6.4, so the S-LSP appears as a single hop in the overlay network.

6.3. Session Shuffling Model

The session shuffling approach ([RFC5251]) is a modification of the flat model described in Section 6.1. In this approach a single end-to-end session is established, but as the signaling messages pass through the ingress and egress CNs, address mapping is performed on all addresses carried by the messages to replace the addresses with values from the correct address space. The ERO and RRO are stripped from the messages as previously discussed, so there is no need for the CNs to examine those objects to map addresses. However, all other addresses must be mapped including the important session identifiers (the source and destination addresses). Viewed from the outside (perhaps through an NMS) this gives the impression of session stitching because the session has different identifiers as it crosses the core network. An NMS might, therefore, present the shuffling model as the stitching model, or it might operate the same address shuffling/mapping as is used by CNs.

6.4. Hierarchal Model

If the ENs and CNs have the same switching capability, a tunnel between the ingress and egress core-nodes can be provisioned to carry the end-to-end connection. The tunnel may have a larger capacity than

the end-to-end UNI connection, depending on the policies configured at the ingress CN of the core network. The end-to-end connection can be nested into a tunnel, which forms the LSP hierarchy [RFC4206] as shown in Figure 6. If the tunnel has a larger capacity, other LSPs can also be nested within the same tunnel.

Alternatively, if the ENs and CNs have different switching capabilities the LSP hierarchical model can also be used exactly as described in [RFC4206].

In the hierarchal model, the end-to-end connection can be divided into three hops: one for each UNI link and one hop across the core network. The core network tunnel can be pre-provisioned via network planning, or triggered by the UNI signalling. For the latter case, [RFC5212], [RFC6001] and other multi-layer network related specifications can be used to create the hierarchical LSP.

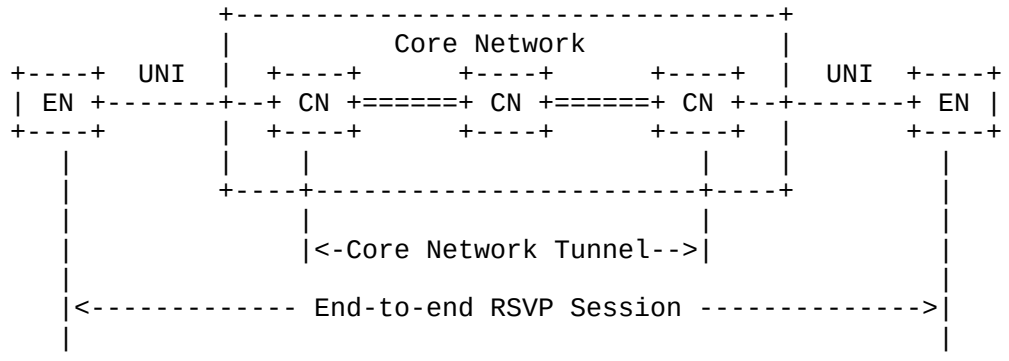

Figure 6 - The Hierarchal Model

7. UNI Recovery

One of the significant uses of GMPLS is to provide recovery mechanisms for connections. Recovery and protection mechanisms are also needed in many UNI scenarios, and the relationship between the overlay and core network provide obvious places at which to operate the recovery techniques.

7.1. End-to-end Recovery

In the case of multi-homing, UNI end-to-end recovery is possible. As shown in Figure 7, the working path (W) and the protection path (P) are disjoint from each other not only inside the core network, but also at both the source and destination sides of the UNI. Mechanisms need to be provided to ensure the selection of disjoint working and backup paths as discussed in the following subsections.

It should be noted that end-to-end recovery can be operated even when the ENs are single-homed. However, obviously, in this case there is no protection against the failure of an EN-CN link, or of the edge CN itself.

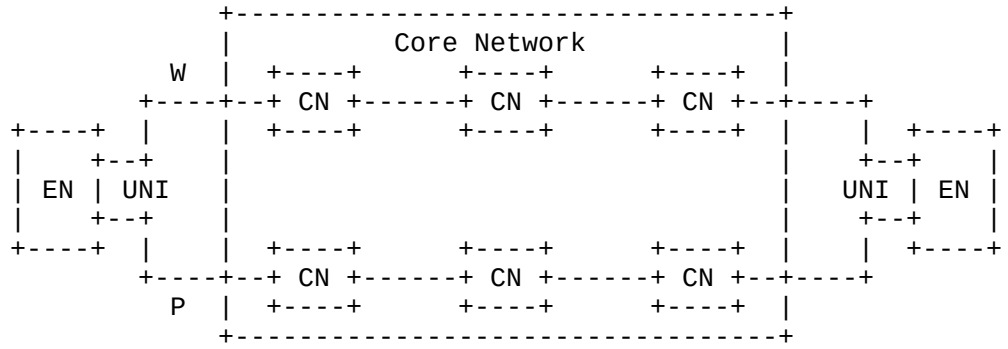

Figure 7 - UNI End-to-End Recovery

7.1.1. Serial Provisioning of Working and Protection Paths

In serial provisioning, one path is computed before another and the associated LSP may even be set up before the second path is computed. In the case where the working path is computed and created before the protection path, path computation for the protection path needs to select a (maximally) disjoint path given this existing working path.

If the EN is allowed to see details of the core network, the EN can use the RRO to collect the route of the working path. It can then use the Exclude Route Object (XRO) to exclude the working path when signaling the protection path, as described in [RFC4874].

But in most cases, in order to preserve the confidentiality of topology within the core network, the route of the working path as it

traverses the core network will be hidden from the EN. In such cases, the RRO and XRO mechanism cannot be used. Alternative includes:

-Only collect the Shared Risk Group (SRG) information, but not the full path information [SLRG-FA]. This is because the SRG information is normally less confidential than the information of node ID and link ID.

-Another possible solution is encrypted the SRG information and provide it to the EN nodes, so that the EN nodes can using this information to convey the diversity constraint, as the method specified in [UNIExt].

-In an application scenario where a PCE is involved inside the core network, then the Path Key mechanism can be used. The confidential path segment, i.e., the route of the working path as it traverses the core network, is encoded as a PKS by the PCE when computing the working path [RFC5520]. This PKS can be used by the EN when it requests the PCE to compute a protection path, to exclude the nodes and links used by the working path. As previously described, the PKS is also used in signaling [RFC5553] so that the EN can indicate to the CN what path to use across the core network.

In order to specify the diversity requirement, it is required that the PKS should be carried in the XRO in both PCEP message and RSVP-TE signaling.

7.1.2. Concurrent Computation of Working and Protection Path

The working and protection path can be computed at the same time (e.g., by PCE or by one of the CNs to which the source EN is attached).

[PCE-GMPLS] adds support for an end node to request a protected service using the protection types defined in [RFC4872].Therefore, it’s possible that the source EN requests the edge CN or PCE to compute both the working and the protection path at the same time. At this time, the disjunction requirement can be resolved inside the path computation server.

Same as described in the previous section, the path segment traversing the core network can be encoded as a PKS if confidentiality is requested.

7.2. Segment Recovery

The UNI connection may request protection only inside the core network, especially in case of single-homing. A UNI segment protection example is shown in Figure 8. In this case, the core network provides a "recovery domain".

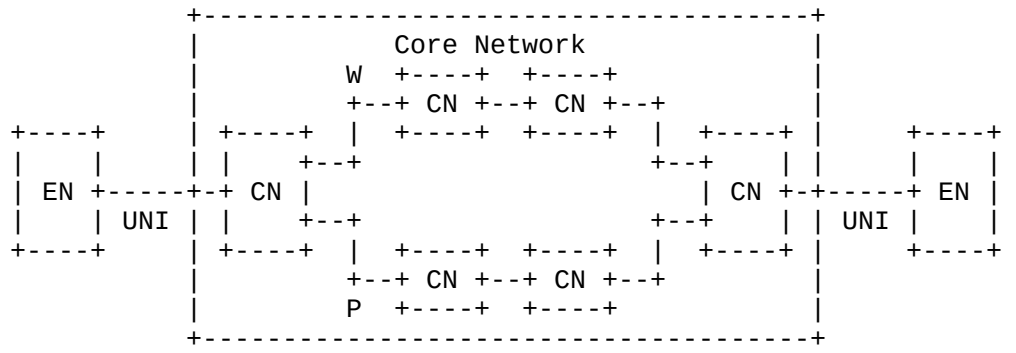

Figure 8 - UNI Segment Recovery

[RFC4873] provides a mechanism for segment recovery, in which the PROTECTION Object is extended to indicate segment recovery, and the Secondary ERO (SERO) is introduced for the explicit control of the protection LSP between the branch node and the merge node.

However, in the overlay model, the mechanisms of segment recovery described in [RFC4873] may not be appropriate. In particular, the source EN might not know the CN to which the destination EN is attached. That means that the source EN knows the branch for the protection segment, but does not know the merge node.

But the model shown in Figure 8 is particularly important because it places the responsibility for service delivery with the edge CNs. This will be a common operational model in overlay networks. Fortunately the stitching model (Section 6.2) and the hierarchical model (Section 6.4) are good at providing the necessary protection within the core network without the ENs having to be aware of the paths in the core network.

8. UNI Call

The Call is a fundamental component of the ASON model [G.8080]. It is used to maintain the association between one or more user

applications and the network, and to control the set-up, release, modification, and maintenance of sets of Connections (LSPs). In simple cases, the Call and Connection can be established at the same time and in a strict one-to-one ratio. In this case, Call signaling requires only minor extensions to connection signaling. However, if Calls are handled separately from Connections, or if more than one Connection can be associated with a single Call, additional Call signaling is required.

The GMPLS Call, defined in [RFC4974], provides a mechanism to negotiate agreement between endpoints possibly in cooperation with the nodes that provide access to the network. Typically the GMPLS Call can be applied in the UNI scenario for access link capability exchange, policy, authorization, security, and so on.

8.1. Exchange of UNI Link Information

It is possible that the TE attributes of the access link (i.e., the UNI link) are not shared across the core network. So the source EN may not have the TE information of the destination access link as well as the capability of the destination EN. For example, in case of TDM network, the Virtual Concatenation (VCAT) and Link Capacity Adjustment Scheme (LCAS) capability of the destination EN may not be known.

In this case, the source EN can raise a Call carrying the LINK_CAPABILITY object to have a capability exchange with the destination EN, as described in [RFC4974].

8.2. Control of Call Route

When applying the Call, it’s possible that there are multiple core network domains between the source EN (Call initiator) and the destination EN (Call terminator), or there is more than one Call manager in the core network (e.g., in the multi-homing scenario where the CNs to which the ENs are attached act as the Call managers).

In the both cases, when establishing the Call, there may be multiple alternative routes for the Call message to reach the destination EN. One can simply use the hop-by-hop manner (i.e., each Call manager determines the next Call manager to which the Call message will be sent by itself) to control the path of the Call.

However, in the practical deployment of UNI Call, commercial and policy motivations normally play an important role in selecting the Call route, especially in the multi-domain scenario. In this case, the hop-by-hop manner is not practical because the route of the Call

needs to be pre-determined in consideration of commercial and policy factors before establishing the Call.

Therefore, it is desirable to allow full control of the Call by the source EN. That is, the source EN can identify the full Call route and signal it explicitly, so that the Call message can be forwarded along the desired route. Moreover, the management plane needs to be able to identify the Call route explicitly as an instruction to the source EN.

9. UNI Multicast

Data plane multicasting is supported in existing Traffic-Engineering networks. GMPLS provides extensions to RSVP-TE to support provisioning of point-to-multipoint (P2MP) TE LSPs via the control plane, as described in [RFC4461] and [RFC4875].

In the scenarios where P2MP is supported using the overlay architectural model, it is a requirement to transport signals from one source EN to multiple destination ENs. One could create a point-to-point (P2P) connection between the source EN and each destination EN, but it will likely be a waste of bandwidth resource both of the UNI link and in the core network.

Therefore, there are some scenarios required to support point-to-multipoint (P2MP) TE LSPs from one source EN to multiple leaf ENs.

9.1. UNI Multicast Connection Model

There are two cases for the UNI multicast. For the first case, only the ingress and egress CNs in the core network support P2MP. The core network has to provide multiple P2P connections between ingress CN and each egress CN for the end-to-end UNI multicast, as shown in Figure 9. This relieves the pressure on the source UNI link, but does not help the over use of the core links such as CN1-CN2.

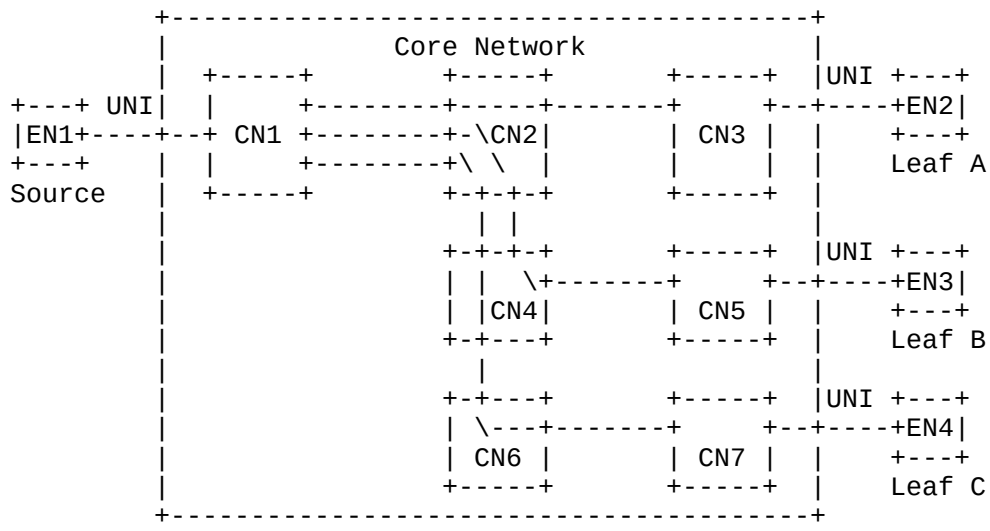

Figure 9 - Only ingress/egress CNs support multicast

For example, in multi-layer scenario of a packet overlay network with a TDM core, the ingress/egress CNs may have packet multicast capabilities and therefore can adapt the packets from EN into multiple TDM connections to transit the core network, but the CNs inside the core network only support point-to-point (P2P) TDM connections.

In another case, all the CNs in the core network can support multicast, so that the core network can create a P2MP LSP to provide the end-to-end UNI multicast, as shown in Figure 10.

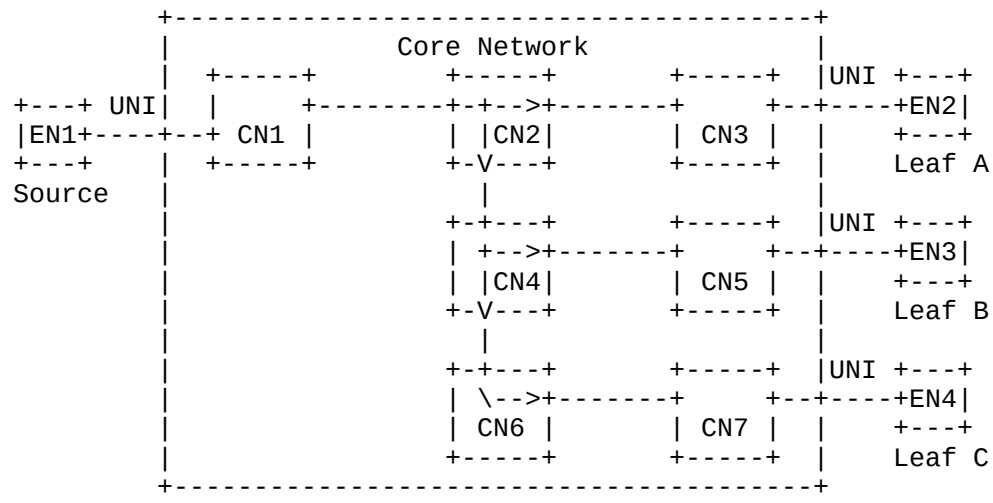

Figure 10 - All CNs support multicast

For example, in the Ethernet over OTN scenario, if the core network can support ODU0 multicast, then an ODU0 P2MP LSP can be created inside the core network to carry the client Gigabit Ethernet (GE) signal for the ENs.

Note that the branching of the P2MP connection could also happen at the source EN if the EN is multi-homed. In this case, each branch from the source EN uses a separate UNI link connecting the source EN to the core network. For each UNI branch, the connection model inside the core network is the same as described in this section.

9.2. UNI Multicast Connection Provisioning

The four UNI connection provisioning models, as described in Section 5, should also be applied in the UNI multicast scenario.

For the flat model, one end-to-end P2MP session as described in [RFC4875] can be used to create the P2MP LSP from source EN to leaf ENs.

For the stitching model, multiple P2P LSP segments or one P2MP LSP segment between the ingress CN and each egress CNs needs to be created and then stitched to the UNI P2MP LSP. GMPLS UNI signaling should have the capability to convey the multicast information by using stitching model.

For the session shuffling model, one end-to-end P2MP session can be used to create the P2MP LSP, with an address mapping performed at both ingress and egress CNs.

For the hierarchical model, multiple P2P LSP tunnels or one P2MP LSP tunnel between the ingress CN and each egress CNs needs be triggered by the UNI signaling for creating the P2MP LSP. GMPLS UNI signaling should have the capability to convey the multicast information by using the hierarchical model.

10. Security Considerations

[RFC5920] provides an overview of security vulnerabilities and protection mechanisms for the GMPLS control plane, which is applicable to this document.

The details of the specific security measures of the overlay network architectural model are provided in [RFC4208], which permits the core network to filter out specific RSVP objects to hide its topology from the EN.

Furthermore, if PCE is used, the security issues described in [RFC4655] should also be considered.

Additionally, when the PKS mechanism is applied, the security issues can be dealt with using [RFC5520] and [RFC5553].

11. IANA Considerations

This informational document does not make any requests for IANA action.

12. Acknowledgments

The authors would like to thank Zafar Ali for his comments.

13. References

13.1. Normative References

[RFC3209] D. Awduche et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC3209, December 2001.

[RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] L. Berger, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

[RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC4203] Kompella, K., and Rekhter, Y., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.

[RFC4206] K. Kompella et al, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC4206, October 2005.

[RFC4208] G. Swallow et al, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC4208, October 2005.

[RFC4655] A. Farrel et al, "A Path Computation Element (PCE)-Based Architecture", RFC4655, August 2006.

[RFC4847] T. Takeda, Ed., "Framework and Requirements for Layer 1 Virtual Private Networks", RFC4847, April 2007.

[RFC4872] J.P. Lang et al, "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC4872, May 2007.

[RFC4873] L. Berger et al, "GMPLS Segment Recovery", RFC4873, May 2007.

[RFC4874] CY. Lee et al, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC4874, April 2007.

[RFC4875] R. Aggarwal et al, "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC4875, May 2007.

[RFC4974] D. Papadimitriou and A. Farrel, Ed., "Generalized MPLS (GMPLS) RSVP-TE Signaling Extensions in Support of Calls", RFC4974, August 2007.

[RFC5150] A. Ayyangar et al, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC5150, February 2008.

[RFC5195] Ould-Brahim, H., Fedyk, D., and Y. Rekhter, "BGP-Based Auto-Discovery for Layer-1 VPNs", RFC 5195, June 2008.

[RFC5251] D. Fedyk and Y. Rekhter, Ed., "Layer 1 VPN Basic Mode", RFC5251, July 2008.

[RFC5252] I. Bryskin and L. Berger Ed., "OSPF-Based Layer 1 VPN Auto-Discovery", RFC5252, July 2008.

[RFC5520] R. Bradford, Ed., "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC5520, April 2009.

[RFC5553] A. Farrel, Ed., "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC5553, May 2009.

[RFC6001] Dimitri Papadimitriou et al, "Generalized Multi-Protocol Label Switching (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC6001, October, 2010.

[RFC6107] K. Shiomoto, A. Farrel, "Procedures for Dynamically Signaled Hierarchical Label Switched Paths", RFC6107, February 2011.

[G.8080] ITU-T Rec. G.8080/Y.1304, "Architecture for the Automatically Switched Optical Network (ASON)," June 2006 (and Amend.2, September 2010).

13.2. Informative References

[RFC4461] S. Yasukawa, Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC4461, April 2006.

[RFC5212] K. Shiomoto et al, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC5212, July 2008.

[RFC5253] T. Takeda, Ed., "Applicability Statement for Layer 1 Virtual Private Network (L1VPN) Basic Mode", RFC 5253, July 2008.

[RFC5339] JL. Le Roux et al, "Evaluation of Existing GMPLS Protocols against Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC5339, September 2008.

[RFC5441] JP. Vasseur et al, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC5441, April 2009.

[RFC5623] Oki, E., Takeda, T., Le Roux, J.L., and Farrel, A., "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", RFC 5623, September 2009.

[RFC5920] L. Fang, Ed., "Security Framework for MPLS and GMPLS Networks", RFC5920, July 2010.

[Call-ext] Fatai Zhang et al, "RSVP-TE extensions to GMPLS Calls", draft-zhang-ccamp-gmpls-call-extensions-01.txt, July 08, 2009.

[PCE-GMPLS] C. Margaria et al, "PCEP extensions for GMPLS", draft-ietf-pce-gmpls-pcep-extensions-07.txt, October 21, 2012.

[SRLG-FA] Fatai Zhang et al, "RSVP-TE Extensions for Configuration SRLG of an FA", draft-ietf-ccamp-rsvp-te-srlg-collect-02.txt, work in progress.

[RFC6344] G. Bernstein et al, "Operating Virtual Concatenation (VCAT) and the Link Capacity Adjustment Scheme (LCAS) with Generalized Multi-Protocol Label Switching (GMPLS)", RFC6344, August 2011.

[UNIExt] D. Fedyk, D. Beller, Lieven Levrau, D. Ceccarelli, F. Zhang, et al, "UNI Extensions for Diversity and Latency Support", draft-fedyk-ccamp-uni-extensions-00.txt, Feb. 2013;

[LSP-DIV] A., Zafar, G., Swallow et al, "Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Path Diversity using Exclude Routes", draft-ietf-ccamp-lsp-diversity-01.txt, work in progress;

[UNI-PLUS] A., Zafar, G., Swallow et al, "Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) extension for signaling Objective Function and Metric Bound", draft-ali-ccamp-rc-objective-function-metric-bound-02.txt, work in progress;

[TE-REC] A., Zafar, G., Swallow et al, "Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)extension for recording TE Metric of a Label Switched Path", draft-ietf-ccamp-te-metric-recording-01.txt, work in progress;

14. Contributors' Address

Yi Lin
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972914
Email: yi.lin@huawei.com

Young Lee
Huawei Technologies
1700 Alma Drive, Suite 100
Plano, TX 75075
USA

Phone: (972) 509-5599 (x2240)
Email: leeyoung@huawei.com

Dan Li
Huawei Technologies

F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28973237
Email: huawei.danli@huawei.com

15. Authors' Addresses

Fatai Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972912
Email: zhangfatai@huawei.com

Oscar Gonzalez de Dios
Telefonica Investigacion y Desarrollo
Emilio Vargas 6
Madrid, 28045
Spain

Phone: +34 913374013
Email: ogondio@tid.es

Adrian Farrel
Old Dog Consulting

EMail: adrian@olddog.co.uk

Xian Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972913
Email: zhang.xian@huawei.com

Daniele Ceccarelli

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

Email: daniele.ceccarelli@ericsson.com

Intellectual Property

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.

The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions.

For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of RFC 5378. No language to the contrary, or terms,

conditions or rights that differ from or are inconsistent with the rights and licenses granted under RFC 5378, shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

Disclaimer of Validity

Copyright Notice