

Network Working Group
Internet Draft
Category: Standards Track
Expiration Date: October 2004

G. Swallow
Cisco Systems, Inc

J. Drake
Calient Networks, Inc

H. Ishimatsu
Japan Telecom

Y. Rekhter
Juniper Networks, Inc

April 2004

GMPLS UNI: RSVP Support for the Overlay Model

[draft-ietf-ccamp-gmpls-overlay-04.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

Generalized Multiprotocol Label Switching (GMPLS) defines both routing and signaling protocols for the creation of Label Switched Paths (LSPs) in various transport technologies. These protocols can be used to support a number of deployment scenarios. This memo addresses the application of GMPLS to the overlay model.

Contents

1	Introduction	3
1.1	GMPLS User-Network Interface (GMPLS UNI)	4
2	Addressing	5
3	ERO Processing	6
3.1	Path Message without ERO	6
3.2	Path Message with ERO	6
3.3	Explicit Label Control	6
4	RRO Processing	7
5	Notification	7
6	Connection Deletion	7
6.1	Alarm-Free Connection Deletion	7
6.2	Connection Deletion with PathErr	8
7	VPN Connections	8
8	Security Considerations	9
9	Acknowledgments	9
10	References	9
10.1	Normative References	9
10.2	Informative References	10
11	Authors' Addresses	10
12	Intellectual Property Notice	11
12.1	IPR Disclosure Acknowledgement	11
13	Full Copyright Statement	12

1. Introduction

Generalized Multiprotocol Label Switching (GMPLS) defines both routing and signaling protocols for the creation of Label Switched Paths (LSPs) in various transport technologies. These protocols can be used to support a number of deployment scenarios. In a peer model, edge-nodes support both a routing and a signaling protocol. The protocol interactions between an edge-node and a core node are the same as between two core-nodes. In the overlay model, the core-nodes act more as a closed system. The edge nodes do not participate in the routing protocol instance that runs among the core nodes; in particular, the edge nodes are unaware of the topology of the core nodes. There may, however, be a routing protocol interaction between a core node and an edge node for the exchange of reachability information to other edge nodes.

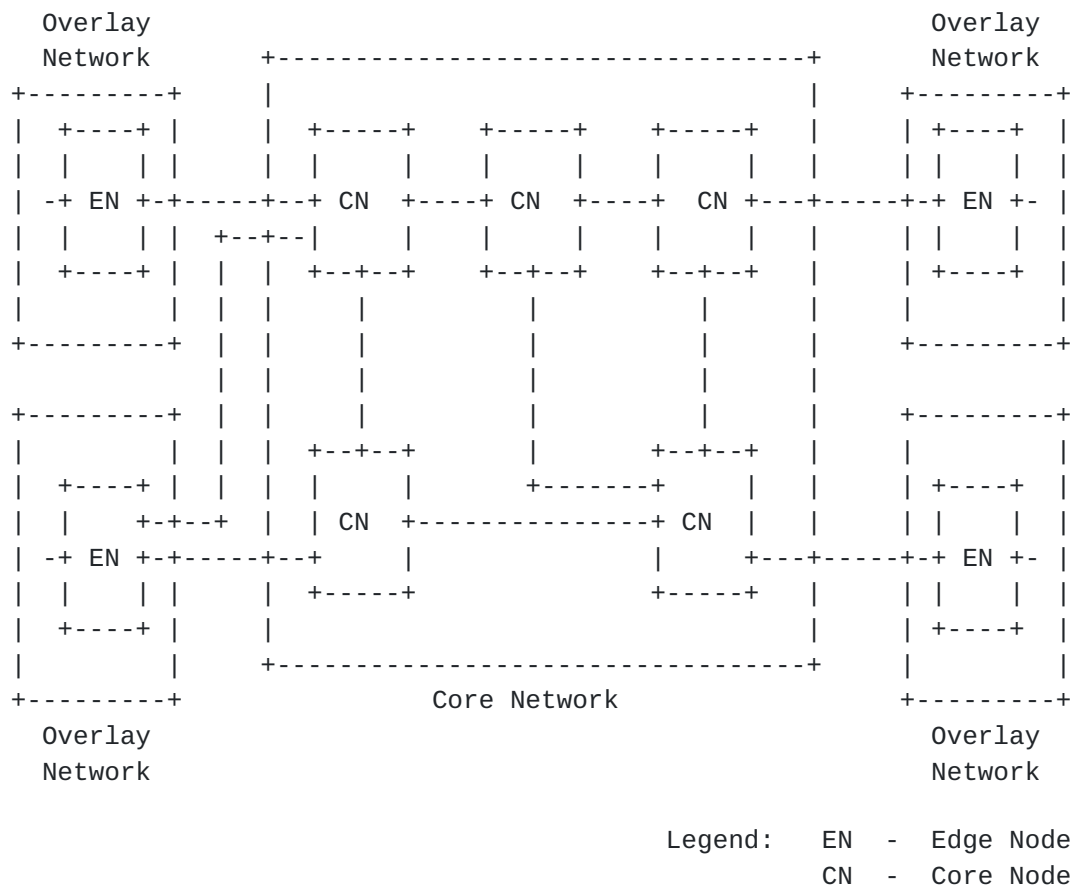


Figure 1: Overlay Reference Model

Figure 1 shows a reference network. The core network is represented by the large box in the center. It contains five core-nodes marked 'CN'. The four boxes around the edge marked "Overlay Network" represent four islands of a single network overlaid network. Only

the nodes of this network with TE links into the core network are shown. These nodes are called edge-nodes; the terminology is in

respect to the core network, not the overlay network. Note that each box marked "Overlay Network" could contain many other nodes. Such nodes are not shown; they do not participate in the signaling described in this document. Only the edge-nodes can signal to set up links across the core to other edge-nodes. Once a link is established, presumably the edge-node will inform the other nodes of the overlay network of its existence. How this is accomplished is beyond the scope of this document. However, one possible means is using a forwarding adjacency as described in [[MPLS-HIER](#)].

In the overlay model there may be restrictions on what may be signaled between an edge-node and a core-node. This memo addresses the application of GMPLS to the overlay model. Specifically it addresses RSVP procedures between an edge-node and a core-node in the overlay model. All RSVP procedures are assumed to be identical to [[RFC3473](#)] except as noted in this document.

This document primarily addresses interactions between an edge-node and it's adjacent (at the data plane) core-node. Except where noted, the term core-node refers to the node which is immediately adjacent to an edge-node across a particular data plane interface. The term core-nodes, however, refers to all nodes in the core.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

1.1. GMPLS User-Network Interface (GMPLS UNI)

One can apply the GMPLS Overlay model at the User-Network Interface (UNI) reference point defined in the Automatically Switched Optical Network (ASON) [[G.8080](#)]. Consider the case where the 'Core Network' in Figure 1 is a Service Provider network, and the Edge Nodes are 'user' devices. The interface between EN and CN is the UNI reference point, and to support the ASON model, one must define signaling across the UNI.

The extensions described in this memo provide mechanisms for UNI signaling that are compatible with GMPLS signaling [[RFC3471](#), [RFC3473](#)]. Moreover, these mechanisms for UNI signaling are in line with the RSVP model, namely, there is a single end-to-end RSVP session for the user connection. The first and last hops constitute the UNI, and the RSVP session carries the user parameters end-to-end. This obviates the need to map (or carry) user parameters to (in) the format expected by the network-to-network interface (NNI) used within the Service Provider network. This in turn means that the UNI and NNI can be independent of one another, which is a requirement of the ASON architecture. However, in the case that the UNI and NNI are both GMPLS RSVP-based, the methodology specified in this memo allows

for a single RSVP session to instantiate both UNI and NNI signaling,
if so desired, and if allowed by Service Provider policy.

2. Addressing

Addresses for edge-nodes in the overlay model are drawn from the same address space as the edge-nodes use to address their adjacent core-nodes. This may be the same address space as used by the core-nodes to communicate among themselves or it may be a VPN space supported by the core-nodes as an overlay.

To be more specific, an edge-node and its attached core-node must share the same address space which is used by GMPLS. A set of <edge-node, core-node> tuples share the same address space if the edge-nodes in the set could establish LSPs (through the core-nodes) among themselves (note that edge-nodes in the set may be a subset of all the edge-nodes). The address space used by the core-nodes to communicate among themselves may, but need not be shared with the address space used by any of the <edge-node, core-node> tuples.

An edge-node is identified by either a single IP address representing its Node-ID, or by one or more numbered TE links that connect the edge-node to the core-nodes. Core-nodes are assumed to be ignorant of any other addresses associated with an edge-node (i.e. addresses which are not used in signaling connections through the GMPLS core).

An edge-node need only know its own address, an address of the adjacent core-node, and know (or be able to resolve) the address of any other edge-node to which it wishes to connect.

A core-node need only know (and track) the addresses on interfaces between that core-node and its attached edge-nodes, as well as the Node IDs of those edge-nodes. In addition, a core-node needs to know the interface addresses and Node IDs of other edge-nodes to which an attached edge-node is permitted to connect.

When forming a SENDER_TEMPLATE the ingress edge-node includes either its Node-ID or the address of one of its numbered TE links. In the latter case the connection will only be made over this interface.

When forming a SESSION_OBJECT, the ingress edge-node includes either the Node-ID of the egress edge-device or the address of one of the egress' numbered TE links. In the latter case the connection will only be made over this interface. The Extended_Tunnel_ID of the SESSION Object is set to either zero or to an address of the ingress edge-device.

Links may be either numbered or unnumbered. Further, links may be bundled or unbundled. See [[GMPLS-ARCH](#)], [[RFC3471](#)], [[BUNDLE](#)], and [[RFC3477](#)].

3. ERO Processing

An edge-node MAY include an ERO. A core-node MAY reject a Path message that contains an ERO. Such behavior is controlled by (hopefully consistent) configuration. If a core-node rejects a Path message due to the presence of an ERO it SHOULD return a PathErr message with an error code of "Unknown object class" toward the sender. This causes the path setup to fail.

Further a core-node MAY accept EROs which only include the ingress edge-node, the ingress core-node, the egress core-node and the egress edge-node. This is to support explicit label control on the edge-node interface, see below. If a core-node rejects a Path message due to the presence of an ERO not of the permitted format it SHOULD return a PathErr message with an error code of Bad Explicit Route Object as defined in [[RFC3209](#)].

3.1. Path Message without ERO

When a core-node receives a Path message from an edge-node that contains no ERO, it MUST calculate a route to the destination and include that route in a ERO, before forwarding the PATH message. One exception would be if the egress edge-node were also adjacent to this core-node. If no route can be found, the core-node SHOULD return a PathErr message with an error code and value of 24,5 - "No route available toward destination".

3.2. Path Message with ERO

When a core-node receives a Path message from an edge-node that contains an ERO, it SHOULD verify the route against its topology database before forwarding the PATH message. If the route is not viable, then a PathErr message with an error code and value of 24,5 - "No route available toward destination" should be returned.

3.3. Explicit Label Control

In order to support explicit label control and full identification of the egress link an ingress edge-node may include an ERO that consists of only the last hop. This is signaled by setting the first subobject of the ERO to the node-ID of the egress core-node with the L-bit set. Following this subobject are all other subobjects necessary to identify the link and labels as they would normally appear.

This process is described in [[RFC3473](#)] and [[EXPLICIT](#)].

4. RRO Processing

An edge-node MAY include an RRO. A core-node MAY remove the RRO from the Path message before forwarding it. Further the core-node may remove the RRO from a Resv message before forwarding it to the edge-node. Such behavior is controlled by (hopefully consistent) configuration.

Further a core-node MAY edit the RRO in a Resv message such that it includes only the subobjects from the egress core-node through the egress edge-node. This is to allow the ingress node to be aware of the selected link and labels on at the far end of the connection.

5. Notification

An edge-node MAY include a NOTIFY_REQUEST object in both the Path and Resv messages it generates. Core-nodes may send Notification messages to edge-nodes which have included the NOTIFY_REQUEST object.

A core-node MAY remove a NOTIFY_REQUEST object from a Path or Resv message received from an edge-node before forwarding it.

If no NOTIFY_REQUEST object is present in the Path or Resv received from an edge-node, the core-node adjacent to the edge-node may include a NOTIFY_REQUEST object and set its value to its own address.

In either of the above cases, the a core-node SHOULD NOT send Notify messages to the edge-node.

When a core-node receives a NOTIFY_REQUEST object from an edge-node, it MAY update the Notify Node Address with its own address before forwarding it. In this case, when NOTIFY messages are received, they MAY be selectively (based on local policy) forwarded to the edge-node.

6. Connection Deletion

6.1. Alarm-Free Connection Deletion

RSVP currently deletes connections using either a single pass PathTear message, or a ResvTear and PathTear message combination. Upon receipt of the PathTear message, a node deletes the connection state and forwards the message. In optical networks, however, it is possible that the deletion of a connection (e.g., removal of the cross-connect) in a node may cause the connection to be perceived as failed in downstream nodes (e.g., loss of frame, loss of light, etc.). This may in turn lead to management alarms and perhaps the triggering of restoration/protection for the connection.

To address this issue, the graceful connection deletion procedure SHOULD be followed. Under this procedure, an ADMIN_STATUS object MUST be sent in Path or Resv message along the connection's path to inform all nodes enroute of the intended deletion, prior to the actual deletion of the connection. The procedure is described in [\[RFC3473\]](#).

An ingress core-node receiving a PathTear without having first seen an ADMIN_STATUS object informing it that the connection is about to be deleted MAY pause the PathTear and first send a Path message with an ADMIN_STATUS object to inform all downstream LSRs that the connection is about to be deleted. When the Resv is received echoing the ADMIN_STATUS or using a timer as described in [\[RFC3473\]](#), the ingress core-node MUST forward the PathTear.

6.2. Connection Deletion with PathErr

[\[RFC3473\]](#) introduces the Path_State_Removed flag to a PathErr message to indicate that the sender has removed all state associated with the LSP and does not need to see a PathTear. A core-node next to an edge-node MAY map between teardown using ResvTear/PathTear and PathErr with Path_state_Removed.

- A core-node next to an edge-node receiving a ResvTear from its downstream neighbor MAY respond with a PathTear and send a PathErr with Path_State_Removed further upstream.
- A core-node next to an edge-node receiving a PathErr with Path_State_Removed from its downstream neighbor MAY retain Path state and send a ResvTear further upstream.

7. VPN Connections

As stated in the addressing section above, the extensions in this document are designed to be compatible with the support of VPNs. Since the core network may be some technology other than GMPLS, no mandatory means of mapping core connections to access connections is specified. However, when GMPLS is used for the core network, we RECOMMEND the following procedure which is based on [\[LSP-HIER\]](#).

The VPN connection is modeled as being three hops. One for each access link and one hop across the core network.

The VPN connection is established using a two step procedure. When a Path message is received at a core node on an interface which is part of a VPN, the Path message is held until a core connection is established.

The connection across the core is setup as a separate signaling exchange between the core nodes, using the address space of the core nodes. While this exchange is in progress, the original Path message is held at the ingress core node. Once the exchange for the core connection is complete, this connection is used in the VPN connection as if it were a single link. This is signaled by including an IF_ID RSVP_HOP object (defined in [[RFC3473](#)]) using the procedures defined in [LSP-HIER].

The original Path message is then forwarded within the VPN addressing realm to the core node attached to the destination edge node. Many ways of accomplishing this are available, including IP and GRE tunnels and BGP/MPLS VPNs. Specifying a particular means is beyond the scope of this document.

8. Security Considerations

This draft imposes no new security risks over [[RFC3473](#)]. In fact it represents a lower trust model between the core and edge-nodes as the core is permitted to hide its topology from the edge-nodes. The core is also permitted to restrict the actions of edge-nodes by filtering out specific RSVP objects.

9. Acknowledgments

The authors would like to thank Kireeti Kompella, Jonathan Lang, Dimitri Papadimitriou, Dimitrios Pendarakis, Bala Rajagopalan, and Adrian Farrel for their comments and input.

10. References

10.1. Normative References

- [RFC3471] Berger, L., "Generalized MPLS Signaling Functional Description", [RFC 3471](#), January 2003.
- [RFC3473] Berger, L., "Generalized MPLS Signaling RSVP-TE Extensions", [RFC 3473](#), January 2003.
- [RFC3209] Awduche, et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [RFC3667] Bradner, S., "IETF Rights in Contributions", [BCP 78](#), [RFC 3667](#), February 2004.
- [RFC3668] Bradner, S., Ed., "Intellectual Property Rights in IETF Technology", [BCP 79](#), [RFC 3668](#), February 2004.

10.2. Informative References

- [GMPLS-ARCH] Mannie, E, et al., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", [draft-ietf-ccamp-gmpls-architecture-04.txt](#), February 2003, work in progress.
- [MPLS-HIER] Kompella, K., and Rekhter, Y., "LSP Hierarchy with MPLS TE", [draft-ietf-mpls-lsp-hierarchy-08.txt](#), September 2001, work in progress.
- [BUNDLE] Kompella, K., Rekhter, Y., and Berger, L., "Link Bundling in MPLS Traffic Engineering", [draft-ietf-mpls-bundle-04.txt](#), July 2002, work in progress.
- [RFC3477] Kompella, K., and Rekhter, Y., "Signaling Unnumbered Links in RSVP-TE", [RFC 3477](#), January 2003.
- [G.8080] ITU-T Rec. G.8080/Y.1304, "Architecture for the Automatically Switched Optical Network (ASON)," November 2001 (and Revision, January 2003).
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," [RFC 2119](#).
- [EXPLICIT] Berger, L., "GMPLS Signaling Procedure For Egress Control", [draft-ccamp-gmpls-egress-control-00.txt](#), March 2004, work in progress.

11. Authors' Addresses

John Drake
Calient Networks
5853 Rue Ferrari
San Jose, CA 95138
Phone: +1 408 972 3720
Email: jdrake@calient.net

Hirokazu Ishimatsu
Japan Telecom Co., Ltd.
2-9-1 Hatchobori, Chuo-ku, Tokyo, 104-0032, Japan
Tel: +81 3 5540 8493
Email: hirokazu.ishimatsu@japan-telecom.co.jp

Yakov Rekhter
Juniper Networks, Inc.
Email: yakov@juniper.net

George Swallow
Cisco Systems, Inc.
1414 Massachusetts Ave,
Boxborough, MA 01719
Phone: +1 978 936 1398
Email: swallow@cisco.com

12. Intellectual Property Notice

The IETF 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 this 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. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR 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 this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

12.1 IPR Disclosure Acknowledgement

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

13. Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

