

Contribution Number: oif2000.171.0

Working Group: Signaling Working Group

TITLE: RSVP Extensions for Optical UNI Signaling

SOURCE: Eric Gray, Fong Liaw, John Yu Zaffire, Inc. Emails: {egray, fliwa, jzyu}@zaffire.com	Jonathan Lang, John Drake Calient Networks Emails: {jplang, jdrake}@calient.net
George Swallow, Yakov Rekhter Cisco Systems, Inc. Email: {swallow, yakov}@cisco.com	Kireeti Kompella Juniper Networks kireeti@juniper.com
Bala Rajagopalan Tellium, Inc. Email: Bala@tellium.com	Yangguang Xu Lucent Technologies, Inc. xuyg@lucent.com
Raj Jain Nayna Networks, Inc. Email: raj@nayan.com	Osama Aboul-Maged Nortel Networks osama@nortelnetworks.com
Greg Bernstein Ciena Email: GregB@ciena.com	Zhensheng Zhang Sorrento Networks Email: zzhang@sorrento.com

DATE: August 7, 2000

Notice: This Technical Document has been created by the Optical Internetworking Forum (OIF). This document is offered to the OIF Membership solely as a basis for agreement and is not a binding proposal on the companies listed as resources above. The OIF reserves the rights to at any time to add, amend, or withdraw statements contained herein. Nothing in this document is in any way binding on the OIF or any of its members.

The user's attention is called to the possibility that implementation of the OIF implementation agreement contained herein may require the use of inventions covered by the patent rights held by third parties. By publication of this OIF implementation agreement, the OIF makes no representation or warranty whatsoever, whether expressed or implied, that implementation of the specification will not infringe any third party rights, nor does the OIF make any representation or warranty whatsoever, whether expressed or implied, with respect to any claim that has been or may be asserted by any third party, the validity of any patent rights related to any such claim, or the extent to which a license to use any such rights may or may not be available or the terms thereof.

For additional information contact:
The Optical Internetworking Forum, 39355 California Street,
Suite 307, Fremont, CA 94538
510-608-5990 phone ♦ info@oiforum.com

© 2000 Optical Internetworking Forum

Abstract

This draft defines a signaling mechanism based on RSVP-TE ([2]) to support an Optical User Network Interface (UNI). This effort is in part driven by work in the OIF as well as the recent draft on signaling requirements for the optical UNI ([3]), and is consistent with recent work on Generalized MPLS (see [4], [5], [6], and [7]). The main function of this draft is to identify the relevant mechanisms in RSVP-TE (including further extensions) to satisfy functional requirements for an Optical UNI. This draft reflects ongoing work at the Optical Interworking Forum (OIF), however, not all of the concepts/requirements have been approved by the OIF.

1. Introduction

There has recently been a significant effort amongst carriers, service providers, and vendors in the optical networking space to eliminate proprietary control protocols and develop a common control plane. The Optical Internetworking Forum (OIF) has initiated work on an Optical User-Network Interface (Optical UNI) as a step in this direction. Recently, a draft [3] was submitted to the IETF defining proposed requirements and abstract messages for the Optical UNI.

This document describes how a signaling mechanism based on RSVP-TE [2] may be used for an Optical UNI. In particular, we identify the mechanisms already defined for RSVP-TE that can be used to satisfy the proposed requirements of [3]. This work is also based on recent Internet Drafts defining Generalized MPLS signaling (e.g. - [4]).

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 RFC-2119 [8].

2. Overview

This document is compatible to the general RSVP-TE baseline document, and intends to highlight the new objects/procedures in particular relevant to UNI signaling. Therefore, this document is not a self-contained document as it currently stands. Required objects/procedures for UNI signaling are included by reference, but not all listed, except the new ones. A new version of this document may be expanded to make it self-contained.

RSVP-TE is one of the candidate protocols described in [3] for Optical UNI signaling implementation. As part of this Optical UNI, the signaling protocol must have the capability to create, delete, and modify lightpaths across a network, as well as query the status of an existing lightpath. Most of these capabilities may be directly supported by re-using existing procedures, messages, and objects defined in RSVP-TE [2] and in Generalized MPLS Signaling [4].

This document further extends [2] and [4] to support Optical UNI signaling requirements as following:

- Use of DREQ and DREP message and procedure as defined in [9] for Lightpath status enquiry and response.
- Use of Message_ID and Message_Ack objects, Ack_desire flag, and Ack message as defined in [10] to support confirmation of Lightpath deletion and reliable messaging.
- PathTear and ResvTear MUST be used to delete a lightpath due to explicit deletion or RSVP state timeout.

This document does not specify procedures to support the following proposed requirements listed in [3]:

- Concept of "indirect interface" as defined in [3]. It is envisioned that such a requirement can be better serviced via a network management station.
- Different source and destination user groups. The use and procedure of these two attributes need further clarification, see below.
 - Procedures for client registration. This shall be supported by a link management protocol or by manual configuration. Both are out of scope of a signaling specification.
 - Light path ID. This is a NNI attributes per OIF contribution oif2000-0611 [13].

2.1 Use of RSVP-TE and Generalized MPLS signaling for Optical UNI

2.1.1 In-band and out-of-band signaling channels

Optical UNI requires support of in-band and out-of-band signaling channels. The in-band signaling channel is supported by the use of a regular link; and the concept of out-of-band signaling channel is supported by the use of non-packet link [7] where a non-packet link is associated (by configuration or other protocols) with the control channel.

When an out-of-band signaling channel is used, the procedure of using the "non-packet link" in [7] can be applied for advertising the control channel.

2.1.2 Reliable messaging

To support reliable messaging across the UNI, the Message_ID and Message_Ack objects, Ack message, and Ack desired flag of [10] MUST be used in UNI RSVP messages. Acknowledgements apply to hop-by-hop, as opposed to end-to-end, message delivery.

2.1.2 Lightpath creation

Lightpath creation uses the procedure described in section 2.2 "Operation of LSP tunnels" of [2]. Additional objects and their use and procedure are defined in Sections 3 and 4 of [4].

2.1.3 Lightpath modification

Lightpath modification uses procedure as described in section 2.5 "Rerouting of Traffic Engineered Tunnels" of [2].

2.1.4 Lightpath deletion

Lightpath deletion can be done by either the client that originated or terminated the lightpath. The lightpath originator will use the PathTear message while the lightpath terminator will use the ResvTear message. Acknowledgement mechanisms of [10] MUST be used to provide confirmation. The network side of the originator shall use ResvTear to delete a lightpath toward the originator. The network side of the terminator shall use PathTear to delete a lightpath toward the terminator.

2.1.5 Lightpath status enquiry and response

Any client interested in the lightpath status shall send a Diagnostic Request (DREQ) message toward the termination end point of the lightpath. The DREQ message specifies the Session Object, Sender Template Object, and an ending node. Starting at the last hop of the lightpath, the DREQ message is sent along the lightpath toward the sender and start collecting information hop-by-hop in the DREQ. When the DREQ message reaches the ending node, the message type is changed to Diagnostic Reply (DREP) and forwarded to the original requestor node. The DREQ originator can select specific RSVP objects to be collected by including a DIAG_SELECT object in the DREQ message. The full procedure of DREQ and DREP messages is described in [9].

3. RSVP Message Extensions for OPTICAL UNI signaling

In addition to objects defined in [2] and [4], new objects may need to be defined to address additional requirements. Additional objects defined in this specification are OPTIONAL with respect to RSVP and RSVP-TE.

3.1 Propagation Delay

If a maximum propagation delay is desired, the lightpath originator shall include the Tspec object for Propagation_Delay_with a new c-type, and an ADSPEC object in its Path or Resv Messages for the lightpath, which may facilitate the clients understanding the actual latency the traffic would experience.

3.2 Labels

Generalized MPLS signaling [4] defines several types of labels that may be represented in a Generalized Label object. For Optical applications a label may be arbitrarily associated with all or part of a component link, or may be a superset of multiple component links. A common understanding of the meaning of a Generalized Label - in particular the meaning of the link ID portion of the Generalized Label - must exist between the user and the network across any Optical UNI. This common understanding may be dynamically derived (e.g. using LMP [5]), or it may be configured.

This specification uses the Generalized Label, Generalized Label Request, Upstream Label, Suggested Label and Label Set objects defined in [4].

4. RSVP objects related to OPTICAL signaling requirements

Optical UNI signaling requirements [3] specify a set of attributes to be signaled during lightpath creation and modification. The following table summarizes the RSVP objects that are used to signaling a particular OPTICAL signaling attribute.

OPTICAL signaling attribute	RSVP object
Light_Path ID	Not supported
Source termination point	Sender Template (or Session) [2]
Destination termination point	Session [2]
Source Termination Point Port	Label Set/Suggest Label [4]
Destination Termination Label	Egress Label [4]
Contract ID	Session Attribute/Name [2]
Framing	Generalized Label Request [4]
Bandwidth	Generalized Label Request [4]
Transparency	Generalized Label Request [4]
Directionality	Upstream Label [4]
Service Level	Session Attribute [2]
Diversity	Session Attribute [2] and Generalized Label Request Object [4]
Return code	Error Spec
Source user group ID	Session Attributes/LSP_TUNNEL_RA [2]
Destination user group ID	Session Attributes/LSP_TUNNEL_RA [2]
Propagation Delay	Propagation Delay (Tspec)

Note that

(1) There is an attribute called Light Path ID in [3]. However, according to the light path attributes requirements in [13], light path ID is declared as a NNI feature, which hence should not be included.

(2) The support of Transparency is defined in OIF as a requirement, yet not captured in [4]; they should be incorporated in new version of [4].

5. RSVP messages related to OPTICAL UNI signaling

5.1 Path Message

As described in [4], RSVP-TE signaling for support of lightpath creation allows for labels to be suggested by the upstream LSR that is sending a Path message. In addition, the upstream node may provide a label for use in bi-directional setup. The format for the Path message to be used for the OPTICAL UNI is given below.

```
<Path Message> ::= <Common Header> [ <INTEGRITY> ]
                   <SESSION> <RSVP_HOP>
```

```

[ <TIME_VALUES> ]
[ <EXPLICIT_ROUTE> ]
<GENERALIZED_LABEL_REQUEST>
[ <LABEL_SET> | <SUGGESTED_LABEL> ]
[ <UPSTREAM_LABEL> ]
[ <SESSION_ATTRIBUTE> ]
[ <POLICY_DATA> ... ]
<sender_descriptor>

<sender_descriptor> ::= <SENDER_TEMPLATE> <SENDER_TSPEC>
[ <ADSPEC> <PROPAGATION_DELAY> ]
[ <RECORD_ROUTE> ]

```

Note that

(1) The EXPLICIT_ROUTE object may be included to provide an egress label ([4]). Additional explicit hops included in the EXPLICIT_ROUTE object may be ignored by the network in the Optical UNI. Similarly, the RECORD_ROUTE object may be ignored by the Network side of the Optical UNI.

(2) User may need to indicate bi-directional path by including an upstream label object [4]. ERO/egress label [4] is used to indicate termination port.

5.2 Resv Message

RSVP-TE signaling support of the lightpath creation response is via the OPTICAL UNI Resv message. The format for the OPTICAL UNI Resv Message is as shown below.

```

<Resv Message> ::= <Common Header> [ <INTEGRITY> ]
<SESSION> <RSVP_HOP>
[ <TIME_VALUES> ]
[ <RESV_CONFIRM> ] [ <SCOPE> ]
[ <POLICY_DATA> ... ]
<STYLE> <flow_descriptor_list>

<flow_descriptor_list> ::= <FF_flow_descriptor_list>
| <SE_flow_descriptor>

<FF_flow_descriptor_list> ::= <FLOWSPEC> <FF_flow_descriptor>
| <FF_flow_descriptor_list> <FF_flow_descriptor>

<FF_flow_descriptor> ::= [ <FLOWSPEC> ]
<FILTER_SPEC>
<GENERALIZED_LABEL>
[ <RECORD_ROUTE> ]

<SE_flow_descriptor> ::= <FLOWSPEC> <SE_filter_spec_list>
<SE_filter_spec_list> ::= <SE_filter_spec>
| <SE_filter_spec_list> <SE_filter_spec>
<SE_filter_spec> ::= <FILTER_SPEC> <GENERALIZED_LABEL>
[ <RECORD_ROUTE> ]

```

Note that

(1) Filter spec is used to identify the LSP, per RSVP-TE, required to identify the LSP in Resv. The flow spec is used to convey the delay back to head-end of lightpath, similar to Tspec+ADSPEC in Path. This would require a new c-type.

(2) Network may ignore the RRO. The reason to list it here (as an option) is that some UNIs may be owned by network providers, named as private-UNIs, in which RRO may be meaningful.

3.4 PathTear Message

The lightpath origination node uses a PathTear message to explicitly delete a lightpath. The acknowledgement procedures of [10] are used to provide confirmation. Receiving of a MESSAGE_ID_ACK object of a PathTear message SHALL be interpreted to mean that the downstream node has cleared its state associated with the specified lightpath, but not to mean a confirmation of an end-to-end lightpath deletion. If a MESSAGE_ID_ACK object of a PathTear Message cannot be piggybacked in an pending RSVP message immediately, a downstream node SHALL generate a Notify message including the MESSAGE_ID_ACK object to its upstream node. This allows a faster confirmation for the PathTear message.

5.4 ResvTear Message

A lightpath termination node uses ResvTear Message to explicitly delete a lightpath. The acknowledgement procedures of [10] are used to provide confirmation. Receiving of a MESSAGE_ID_ACK object of a ResvTear message SHALL be interpreted to mean that the upstream node has cleared its state associated with the specified lightpath, but not to mean a confirmation of an end-to-end lightpath deletion. If a MESSAGE_ID_ACK object of a ResvTear Message cannot be piggybacked in a pending RSVP message immediately, an upstream node SHALL generate a Notify message including the MESSAGE_ID_ACK object to its upstream node. This allows a faster confirmation for the ResvTear message.

5.5 Notify Message

The Notify message can be generated at any time to allow expedited notification of change in the status of a lightpath. Consequently, both the user and network sides of the Optical UNI MUST be prepared to receive a Notify message. The format of the Notify message is given in [4].

5.6 Diagnostic (DREQ) Message

The DREQ Message is used to initiate Lightpath Status Enquiry required by [3]. The format of the Diagnostic Request (DREQ) message is given in [9]. Refer to [9] for full procedures on how to handle DREQ messages.

5.7 Diagnostic Reply (DREP) Message

The DREP Message is used to send the collected lightpath status to originator of the DREQ message. It has the same format as the DREQ message. Refer to [9] for procedures on how to forward a DREP message.

Motion: Adopt the method proposed in this contribution to the technical document for optical UNI signaling in the Signaling WG.

6. References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [2] Awduche, D. O., Berger, L., Gan, D.-H., Li, T., Swallow, G., Srinivasan, V., "Extensions to RSVP for LSP Tunnels", Work in Progress (Internet Draft), draft-ietf-mpls-rsvp-lsp-tunnel-06.txt, July 2000.
- [3] Aboul-Magd, O., Aparicio, O., Barry, R., Bernstein, G., Jain, R., Jia, L., Dulepet, R., Lazer, M., Yates, J., Pendarakis, D., Rajagopalan, B., Rennison, R., Xu, Y., Xue, Y., Yu, J. and Zhang, Z., "Signaling Requirements at the Optical UNI", Work in Progress (Internet Draft), July 2000.
- [4] Ashwood-Smith, P., Fan, Y., Banerjee, A., Drake, J., Lang, J., Berger, L., Bernstein, G., Kompella, K., Mannie, E., Rajagopalan, B., Saha, S., Tang, Z., Rekhter, Y. and Sharma, V., "Generalized MPLS - Signaling Functional Description", Work in Progress (Internet Draft), draft-ashwood-generalized-mpls-signaling-00.txt, June 2000.
- [5] Lang, J.P., Mitra, K., Drake, J., Kompella, K., Rekhter, Y., Saha, D., Berger, L., Basak, D., and Sandick, H., "Link Management Protocol (LMP)", Work in Progress (Internet Draft), draft-lang-mpls-lmp-01.txt, July 2000.
- [6] Kompella, K., Rekhter, Y., Banerjee, A., Drake, J., Bernstein, G., Fedyk, D., Mannie, E., Saha, D., and Sharma, V., "OSPF Extensions in Support of MPL(ambda)S", Work in Progress (Internet Draft), draft-kompella-ospf-ompls-extensions-00.txt, July 2000.
- [7] Kompella, K., Rekhter, Y., Banerjee, A., Drake, J., Bernstein, G., Fedyk, D., Mannie, E., Saha, D., and Sharma, V., "IS-IS Extensions in Support of MPL(ambda)S", Work in Progress (Internet Draft), draft-kompella-ospf-ompls-extensions-00.txt, July 2000.

- [8] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119 (BCP 14), March 1997.
- [9] Terzis, A., etal., "RSVP Diagnostic Messages", RFC 2745, January 2000.
- [10] Berger, L. etal., "RSVP Refresh Overhead Reduction Extensions", work in progress, June 2000.
- [11] Kompella, K. etal., "Link Bundling in MPLS Traffic Engineering", work in progress, July 2000.
- [12] Kompella, K. etal., "Traffic Engineering with Unnumbered Links", work in progress, July 2000.
- [13] Larry McAdams, Jennifer Yates, "User to Network Interface (UNI) Service Definition and Lightpath Attributes", OIF2000.0611, February 2000.