Network Working Group
Internet Draft
Expiration Date: March 2001

Kireeti Kompella
Juniper Networks
Yakov Rekhter
Cisco Systems

Signalling Unnumbered Links in RSVP-TE

draft-kompella-mpls-rsvp-unnum-00.txt

1. Status of this Memo

2. Abstract

Current signalling used by MPLS TE doesn't provide support for unnumbered links. This document defines procedures and extensions to RSVP-TE, one of the MPLS TE signalling protocols, that are needed in order to support unnumbered links.

3. Overview

Supporting MPLS TE over unnumbered links (i.e., links that do not have IP addresses) involves two components: (a) the ability to carry (TE) information about unnumbered links in IGP TE extensions (ISIS or OSPF), and (b) the ability to specify unnumbered links in MPLS TE signalling. The former is covered in [ISIS-TE, OSPF-TE]. The focus of this document is on the latter.

Current signalling used by MPLS TE doesn't provide support for unnumbered links because the current signalling doesn't provide a way to indicate an unnumbered link in its Explicit Route and Record Route Objects. This document proposes simple procedures and extensions that allow RSVP-TE signalling [RSVP-TE] to be used with unnumbered links.

4. Interface Identifiers

Since unnumbered links are not identified by an IP address, then for the purpose of MPLS TE they need some other identifier. We assume that each unnumbered link on a Label Switched Router (LSR) is given a unique 16-bit identifier. The scope of this identifier is the LSR to which the link belongs; moreover, the IS-IS and/or OSPF and RSVP modules on an LSR must agree on interface identifiers.

Note that links are directed, i.e., a link l is from some LSR A to some other LSR B. LSR A chooses the interface identifier for link l. To be completely clear, we call this the "outgoing interface identifier from LSR A's point of view". If there is a reverse link from LSR B to LSR A (for example, a point-to-point SONET interface connecting LSRs A and B would be represented as two links, one from A to B, and another from B to A), B chooses the outgoing interface identifier for the reverse link. There is no a priori relationship between the two interface identifiers.

5. Unnumbered Forwarding Adjacencies

If an LSR that originates an LSP advertises this LSP as an unnumbered Forwarding Adjacency in IS-IS or OSPF [LSP-HIER], the LSR MUST allocate an interface ID to that Forwarding Adjacency. Moreover, the Tunnel ID in the Session Object of the Path Message for the LSP MUST be set to that interface ID, and the Extended Tunnel ID in the Session Object of the LSP MUST be set to the Router ID of the LSR that originates the LSP.

If the LSP is bidirectional, and the tail-end LSR (of the forward LSP) advertises the reverse LSP as an unnumbered Forwarding Adjacency, the tail-end LSR MUST allocate an interface ID to the reverse Forwarding Adjacency. Furthermore, it MUST set the "Reverse Interface ID" field in the Filter Specification object in the flow descriptor list for this LSP to the reverse FA's interface ID (note that while in general there can be multiple Filter Specifications, it is expected in the case of point-to-point LSPs that there is only one). To accommodate this, the LSP_TUNNEL_IPv4 Filter Specification Object's format is modified per Figure 1:

Figure 1: LSP_TUNNEL_IPv4 Filter Specification Object

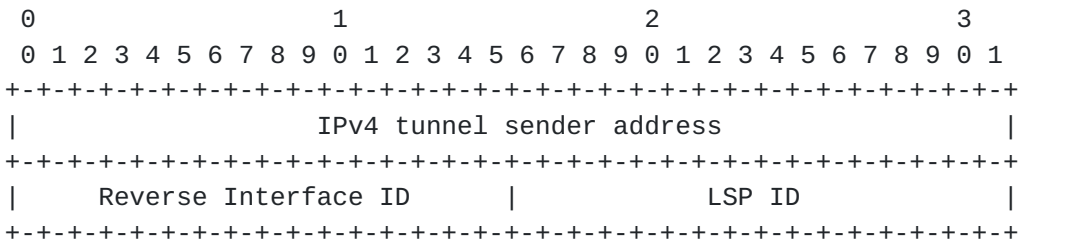

6. Signalling Unnumbered Links in EROs

A new subobject of the Explicit Route Object (ERO) is used to specify unnumbered links. This subobject has the following format:

Figure 2: Unnumbered Interface ID Subobject

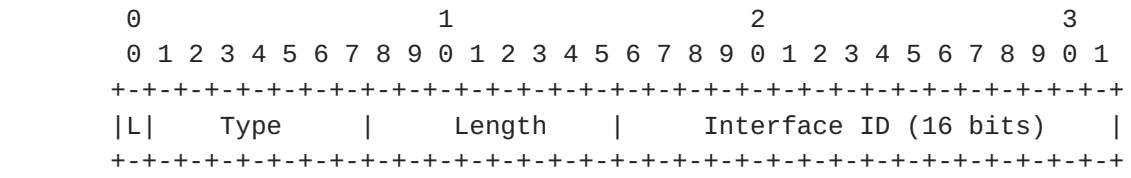

This subobject MUST be strict (i.e., the L bit MUST be 0). The Type is 4 (Unnumbered Interface ID). The Length is 4.

6.1. Interpreting the Unnumbered Interface ID Subobject

The Interface ID is the outgoing interface identifier with respect to the previous node in the path (i.e., the PHOP). If the Path message contains an Unnumbered Interface ID subobject as the first subobject in the ERO, then the PHOP object in the message must contain the router ID of the previous node.

6.2. Processing the Unnumbered Interface ID Subobject

A node that receives a Path message with an Unnumbered Interface ID as the first subobject in the ERO carried by the message MUST check whether the tuple <PHOP, Interface ID> matches the tuple <Extended Tunnel ID, Tunnel ID> of any of the LSPs for which the node is a tail-end. If a match is found, the match identifies the Forwarding Adjacency for which the node has to perform label allocation.

Otherwise, the node MUST check whether the tuple <PHOP, Interface ID> matches the tuple <Extended Tunnel ID, Reverse Interface ID> of any of the bidirectional LSPs for which the node is the head-end. If a match is found, the match identifies the Forwarding Adjacency for which the node has to perform label allocation, namely, the reverse Forwarding Adjacency for the LSP identified by the match.

Otherwise, it is assumed that the node has to perform label allocation for the link over which the Path message was received. In this case the receiving node MAY validate that it received the Path message correctly. To do so, the node must maintain a database of Traffic Engineering information distributed by IS-IS and/or OSPF.

To validate that it received the Path message correctly, the node looks up in its Traffic Engineering database for the node corresponding to the router ID in the PHOP object in the Path. It then checks that there is a link from the previous node to itself that carries the same Interface ID as the one in the ERO subobject. If this is not the case, the receiving node has received the message in error and SHOULD return a "Bad initial subobject" error. Otherwise, the receiving node removes the first subobject, and continues processing the ERO.

6.3. Selecting the Next Hop

If, after processing and removing all initial subobjects in the ERO that refer to itself, the receiving node finds a subobject of type Unnumbered Interface ID, it determines the next hop as follows. The Interface ID MUST refer to an outgoing interface identifier that this

node allocated; if not, the node SHOULD return a "Bad EXPLICIT_ROUTE object" error. The next hop is the node at the other end of the link that the Interface ID refers to.

Furthermore, when sending a Path message to the next hop, the ERO to be used is the current ERO (starting with the Unnumbered Interface ID subobject); the PHOP object is the sending node's router ID.

7. Record Route Object

A new subobject of the Record Route Object (RRO) is used to record that the LSP path traversed an unnumbered link. This subobject has the following format:

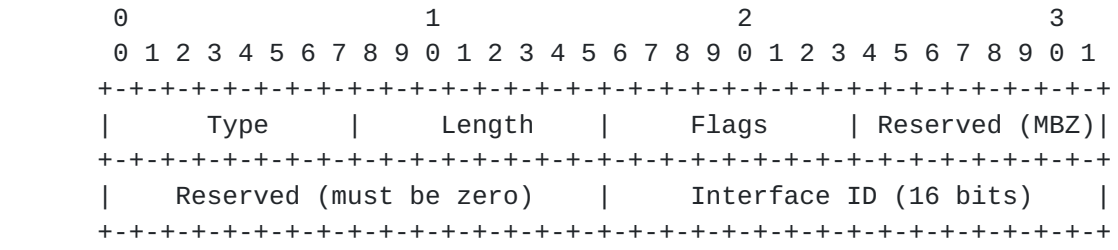

The Type is 4 (Unnumbered Interface ID); the Length is 8. Flags are defined below.

0x01 Local protection available

Indicates that the link downstream of this node is protected via a local repair mechanism. This flag can only be set if the Local protection flag was set in the SESSION_ATTRIBUITE object of the cooresponding Path message.

0x02 Local protection in use

Indicates that a local repair mechanism is in use to maintain this tunnel (usually in the face a an outage of the link it was previously routed over).

7.1. Handling RRO

If at an intermediate node (or at the head-end), the ERO subobject that was used to determine the next hop is of type Unnumbered Interface ID, and a RRO object was received in the Path message (or is desired in the original Path message), an RRO subobject of type Unnumbered Interface ID MUST be appended to the received RRO when sending a Path message downstream.

If the ERO subobject that was used to determine the next hop is of any other type, the handling procedures of [RSVP-TE] apply. Also, if Label Recording is desired, the procedures of [RSVP-TE] apply.

8. Security Considerations

This document raises no new security concerns for RSVP.

9. IANA Considerations

The responsible Internet authority (presently called the IANA) assigns values to RSVP protocol parameters. The current document defines a new subobject for the EXPLICIT_ROUTE object and for the ROUTE_RECORD object. The rules for the assignment of subobject numbers have been defined in [RSVP-TE], using the terminology of BCP 26 "Guidelines for Writing an IANA Considerations Section in RFCs". Those rules apply to the assignment of subobject numbers for the new subobject of the EXPLICIT_ROUTE and ROUTE_RECORD objects.

10. Acknowledgments

Thanks to Lou Berger and Markus Jork for pointing out that the RRO should be extended in like fashion to the ERO. Thanks also to Rahul Aggarwal and Alan Kullberg for their comments on the text.

11. References

[ISIS-TE] Smit, H., and Li, T., "IS-IS extensions for Traffic Engineering", draft-ietf-isis-traffic-02.txt (work in progress)

[LSP-HIER] Kompella, K., and Rekhter, Y., "LSP Hierarchy with MPLS TE", draft-ietf-mpls-lsp-hierarchy-01.txt (work in progress)

[OSPF-TE] Katz, D., and Yeung, D., "Traffic Engineering Extensions to OSPF", draft-katz-yeung-ospf-traffic-02.txt (work in progress)

[RSVP-TE] Awduche, D., Berger, L., Gan, D. H., Li, T., Srinivasan, V., and Swallow, G., "RSVP-TE: Extensions to RSVP for LSP Tunnels", draft-ietf-mpls-rsvp-lsp-tunnel-07.txt (work in progress)

12. Author Information

Kireeti Kompella
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
e-mail: kireeti@juniper.net

Yakov Rekhter
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
e-mail: yakov@cisco.com