

Network Working Group
Request for Comments: 3477
Category: Standards Track

K. Kompella
Y. Rekhter
Juniper Networks
January 2003

Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

Current signalling used by Multi-Protocol Label Switching Traffic Engineering (MPLS TE) does not provide support for unnumbered links. This document defines procedures and extensions to Resource ReSerVation Protocol (RSVP) for Label Switched Path (LSP) Tunnels (RSVP-TE), one of the MPLS TE signalling protocols, that are needed in order to support unnumbered links.

Specification of Requirements

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 BCP 14, RFC 2119 [RFC2119].

1. 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 [GMPLS-ISIS, GMPLS-OSPF]. The focus of this document is on the latter.

Current signalling used by MPLS TE does not provide support for unnumbered links because the current signalling does not 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 [RFC3473] to be used with unnumbered links.

2. Link Identifiers

An unnumbered link has to be a point-to-point link. An LSR at each end of an unnumbered link assigns an identifier to that link. This identifier is a non-zero 32-bit number that is unique within the scope of the LSR that assigns it. If one is using OSPF or ISIS as the IGP in support of traffic engineering, then the IS-IS and/or OSPF and RSVP modules on an LSR must agree on the identifiers.

There is no a priori relationship between the identifiers assigned to a link by the LSRs at each end of that link.

LSRs at the two end points of an unnumbered link exchange with each other the identifiers they assign to the link. Exchanging the identifiers may be accomplished by configuration, by means of a protocol such as LMP ([LMP]), by means of RSVP/CR-LDP (especially in the case where a link is a Forwarding Adjacency, see below), or by means of IS-IS or OSPF extensions ([ISIS-GMPLS], [OSPF-GMPLS]).

Consider an (unnumbered) link between LSRs A and B. LSR A chooses an identifier for that link. So does LSR B. From A's perspective, we refer to the identifier that A assigned to the link as the "link local identifier" (or just "local identifier"), and to the identifier that B assigned to the link as the "link remote identifier" (or just "remote identifier"). Likewise, from B's perspective, the identifier that B assigned to the link is the local identifier, and the identifier that A assigned to the link is the remote identifier.

In the context of this document the term "Router ID" means a stable IP address of an LSR that is always reachable if there is any connectivity to the LSR. This is typically implemented as a "loopback address"; the key attribute is that the address does not become unusable if an interface on the LSR is down. In some cases this value will need to be configured. If one is using the OSPF or ISIS as the IGP in support of traffic engineering, then it is RECOMMENDED for the Router ID to be set to the "Router Address" as defined in [OSPF-TE], or "Traffic Engineering Router ID" as defined in [ISIS-TE].

This section is equally applicable to the case of unnumbered component links (see [LINK-BUNDLE]).

3. Unnumbered Forwarding Adjacencies

If an LSR that originates an LSP advertises this LSP as an unnumbered Forwarding Adjacency in IS-IS or OSPF (see [LSP-HIER]), or the LSR uses the Forwarding Adjacency formed by this LSP as an unnumbered component link of a bundled link (see [LINK-BUNDLE]), the LSR MUST allocate an identifier to that Forwarding Adjacency (just like for any other unnumbered link). Moreover, the Path message used for establishing the LSP that forms the Forwarding Adjacency MUST contain the LSP_TUNNEL_INTERFACE_ID object (described below), with the LSR's Router ID set to the head end's Router ID, and the Interface ID set to the identifier that the LSR allocated to the Forwarding Adjacency.


If the Path message contains the LSP_TUNNEL_INTERFACE_ID object, then the tail-end LSR MUST allocate an identifier to that Forwarding Adjacency (just like for any other unnumbered link). Furthermore, the Resv message for the LSP MUST contain an LSP_TUNNEL_INTERFACE_ID object, with the LSR's Router ID set to the tail-end's Router ID, and the Interface ID set to the identifier allocated by the tail-end LSR.

For the purpose of processing the ERO and the IF_ID RSVP_HOP objects, an unnumbered Forwarding Adjacency is treated as an unnumbered (TE) link or an unnumbered component link as follows. The LSR that originates the Adjacency sets the link local identifier for that link to the value that the LSR allocates to that Forwarding Adjacency, and the link remote identifier to the value carried in the Interface ID field of the Reverse Interface ID object. The LSR that is a tail-end of that Forwarding Adjacency sets the link local identifier for that link to the value that the LSR allocates to that Forwarding Adjacency, and the link remote identifier to the value carried in the Interface ID field of the Forward Interface ID object.

3.1. LSP_TUNNEL_INTERFACE_ID Object

The LSP_TUNNEL_INTERFACE_ID object has a class number of of 193, C-Type of 1 and length of 12. The format is given below.

Figure 1: LSP_TUNNEL_INTERFACE_ID Object




This object can optionally appear in either a Path message or a Resv message. In the former case, we call it the "Forward Interface ID" for that LSP; in the latter case, we call it the "Reverse Interface ID" for the LSP.

4. 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



The Type is 4 (Unnumbered Interface ID). The Length is 12.

The Interface ID is the identifier assigned to the link by the LSR specified by the router ID.

4.1. Processing the IF_ID RSVP_HOP object

When an LSR receives a Path message containing the IF_ID RSVP_HOP object (see [RFC3473], [RFC3471]) with the IF_INDEX TLV, the LSR processes this TLV as follows. The LSR must have information about the identifiers assigned by its neighbors to the unnumbered links between the neighbors and the LSR. The LSR uses this information to find a link with tuple <Router ID, local identifier> matching the tuple <IP Address, Interface ID> carried in the IF_INDEX TLV. If the matching tuple is found, the match identifies the link for which the LSR has to perform label allocation.

Otherwise, the LSR SHOULD return an error using the IF_ID ERROR_SPEC object (see [RFC3473], [RFC3471]). The Error code in the object is set to 24. The Error value in the object is set to 16.

4.2. Processing the ERO

The Unnumbered Interface ID subobject is defined to be a part of a particular abstract node if that node has the Router ID that is equal to the Router ID field in the subobject, and if the node has an


(unnumbered) link or an (unnumbered) Forwarding Adjacency whose local identifier (from that node's point of view) is equal to the value carried in the Interface ID field of the subobject.

With this in mind, the ERO processing in the presence of the Unnumbered Interface ID subobject follows the rules specified in section 4.3.4.1 of [RFC3209].

As part of the ERO processing, or to be more precise, as part of the next hop selection, if the outgoing link is unnumbered, the Path message that the node sends to the next hop MUST include the IF_ID RSVP_HOP object, with the IP address field of that object set to the Router ID of the node, and the Interface ID field of that object set to the identifier assigned to the link by the node.

5. 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 12. 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_ATTRIBUTE object of the corresponding Path message.

0x02 Local protection in use

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

5.1. 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 [RFC3209] apply. Also, if Label Recording is desired, the procedures of [RFC3209] apply.

6. Security Considerations

This document makes a small extension to RFC 3209 [RFC3209] to refine and explicate the use of unnumbered links. As such it poses no new security concerns. In fact, one might argue that use of the extra interface identify could make an RSVP message harder to spoof.

7. IANA Considerations

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 [RFC3209], using the terminology of BCP 26, RFC 2434, "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.

Furthermore, the same Internet authority needs to assign a class number to the LSP_TUNNEL_INTERFACE_ID object. This must be of the form 11bbbbbb (i.e., RSVP silently ignores this unknown object but forwards it).

8. Intellectual Property Considerations

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

9. 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. Finally, thanks to Bora Akyol, Vach Kompella, and George Swallow.

10. References

10.1. Normative references

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3209] Awduche, D., Berger, L., Gan, D. Li, T., Srinivasan, V. and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC3471] Berger, L., Editor, "Generalized Multi-Protocol Label Switching (MPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., Editor, "Generalized Multi-Protocol Label Switching (MPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

10.2. Non-normative references

- [GMPLS-ISIS] Kompella, K., Rekhter, Y., Banerjee, A. et al., "IS-ISIS Extensions in Support of Generalized MPLS", Work in Progress.
- [GMPLS-OSPF] Kompella, K., Rekhter, Y., Banerjee, A. et al., "OSPF Extensions in Support of Generalized MPLS", Work in Progress.

- [ISIS-TE] Li, T. and H. Smit, "IS-IS extensions for Traffic Engineering", Work in Progress.
- [LINK-BUNDLE] Kompella, K., Rekhter, Y. and L. Berger, "Link Bundling in MPLS Traffic Engineering", Work in Progress.
- [LSP-HIER] Kompella, K. and Y. Rekhter, "LSP Hierarchy with MPLS TE", Work in Progress.
- [LMP] Lang, J., Mitra, K., et al., "Link Management Protocol (LMP)", Work in Progress.
- [OSPF-TE] Katz, D., Yeung, D., Kompella, K., "Traffic Engineering Extensions to OSPF Version 2", Work in Progress.

11. Authors' Addresses

Kireeti Kompella
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089

EMail: kireeti@juniper.net

Yakov Rekhter
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089

EMail: yakov@juniper.net

12. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

