

Internet Engineering Task Force (IETF)
Request for Comments: 6511
Category: Standards Track
ISSN: 2070-1721

Z. Ali
G. Swallow
Cisco Systems
R. Aggarwal
Juniper Networks
February 2012

Non-Penultimate Hop Popping Behavior and Out-of-Band Mapping for RSVP-TE Label Switched Paths

Abstract

There are many deployment scenarios that require an egress Label Switching Router (LSR) to receive binding of the Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Label Switched Path (LSP) to an application and a payload identifier using some "out-of-band" (OOB) mechanism. This document defines protocol mechanisms to address this requirement. The procedures described in this document are equally applicable for point-to-point (P2P) and point-to-multipoint (P2MP) LSPs.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6511>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	2
1.1. Conventions Used in This Document	3
2. RSVP-TE Signaling Extensions	3
2.1. Signaling Non-PHP Behavior	3
2.2. Signaling OOB Mapping Indication	5
2.3. Relationship between OOB and Non-PHP Flags	6
2.4. Egress Procedure for Label Binding	6
3. Security Considerations	7
4. IANA Considerations	7
4.1. Attribute Flags for LSP Attributes Object	7
4.2. New RSVP Error Sub-Code	8
5. Acknowledgements	8
6. References	8
6.1. Normative References	8
6.2. Informative References	9

1. Introduction

When Resource Reservation Protocol - Traffic Engineering (RSVP-TE) is used for applications like Multicast Virtual Private Network (MVPN) [RFC6513] and Virtual Private LAN Service (VPLS) [RFC4761], an egress Label Switching Router (LSR) receives the binding of the RSVP-TE Label Switched Path (LSP) to an application and a payload identifier using an "out-of-band" (OOB) mechanism (e.g., Border Gateway Protocol (BGP)). In such cases, the egress LSR cannot make correct forwarding decisions until such OOB mapping information is received. Furthermore, in order to apply the binding information, the egress LSR needs to identify the incoming LSP on which traffic is coming.

Therefore, non-Penultimate Hop Popping (non-PHP) behavior is required to apply OOB mapping. Non-PHP behavior requires the egress LSRs to assign a non-NULL label for the LSP being signaled.

There are other applications that require non-PHP behavior. When RSVP-TE point-to-multipoint (P2MP) LSPs are used to carry IP multicast traffic non-PHP behavior enables a leaf LSR to identify the P2MP TE LSP on which traffic is received. Hence, the egress LSR can determine whether traffic is received on the expected P2MP LSP and discard traffic that is not received on the expected P2MP LSP. Non-PHP behavior is also required to determine the context of upstream assigned labels when the context is a MPLS LSP. Non-PHP behavior may also be required for MPLS Transport Profile (MPLS-TP) LSPs [RFC5921].

This document defines two new flags in the Attributes Flags TLV of the LSP Attributes object defined in [RFC5420]: one flag for communication of non-PHP behavior and one flag to indicate that the binding of the LSP to an application and a payload identifier (Payload ID) needs to be learned via an out-of-band mapping mechanism. As there is one-to-one correspondence between bits in the Attribute Flags TLV and the Record Route Object (RRO) Attributes subobject, corresponding flags to be carried in the RRO Attributes subobject are also defined.

The procedures described in this document are equally applicable for point-to-point (P2P) and P2MP LSPs. Specification of the OOB communication mechanism(s) is beyond the scope of this document.

1.1. Conventions Used in This Document

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 [RFC2119].

2. RSVP-TE Signaling Extensions

This section describes the signaling extensions required to address the above-mentioned requirements.

2.1. Signaling Non-PHP Behavior

In order to request non-PHP behavior for an RSVP-TE LSP, this document defines a new flag in the Attributes Flags TLV of the LSP Attributes object defined in [RFC5420]:

Bit Number 7: Non-PHP behavior flag

In order to indicate to the ingress LSR that the egress LSR recognizes the "Non-PHP behavior flag", the same bit is used in the Flags field of the Record Route Object (RRO) Attributes subobject.

An ingress LSR sets the "Non-PHP behavior flag" to signal that the egress LSRs SHOULD assign a non-NULL label for the LSP being signaled. This flag MUST NOT be modified by any other LSRs in the network. LSRs other than the egress LSRs SHOULD ignore this flag.

If an egress LSR receiving the Path message supports the LSP Attributes object and the Attributes Flags TLV and also recognizes the "Non-PHP behavior flag", it MUST allocate a non-NULL local label. The egress LSR MUST also set the "Non-PHP behavior flag" in the Flags field of the RRO Attributes subobject.

If the egress LSR

- supports the LSP Attributes object but does not recognize the Attributes Flags TLV; or
- supports the LSP Attributes object and recognizes the Attributes Flags TLV, but does not recognize the "Non-PHP behavior flag";

then it silently ignores the request according to the processing rules of [RFC5420].

An ingress LSR requesting non-PHP behavior SHOULD examine the "Non-PHP behavior flag" in the Flags field of the RRO Attributes subobject and MAY send a Path Tear to the egress, which has not set the "Non-PHP behavior flag". An ingress LSR requesting non-PHP behavior MAY also examine the label value corresponding to the egress LSR(s) in the RRO and MAY send a Path Tear to the egress, which assigns a NULL label value.

When signaling a P2MP LSP, a source node may wish to solicit an individual response to the "Non-PHP behavior flag" from the leaf nodes. Given the constraints on how the LSP Attributes may be carried in Path and Resv messages according to [RFC5420], in this situation, the source node MUST use a separate Path message for each leaf in networks where [RFC6510] is not supported. In networks with [RFC6510] deployed, either a single leaf per Path message or multiple leaves per Path message MAY be used by the source node.

2.2. Signaling OOB Mapping Indication

This document defines a single flag to indicate that the normal binding mechanism of an RSVP session is overridden. The actual out-of-band mappings are beyond the scope of this document. The flag is carried in the Attributes Flags TLV of the LSP Attributes object defined in [RFC5420] and is defined as follows:

Bit Number 8: OOB mapping flag

In order to indicate to the ingress LSR that the egress LSR recognizes the "OOB mapping flag", the following same bit is used in the Flags field of the Record Route object (RRO) Attributes subobject.

An ingress LSR sets the "OOB mapping flag" to signal the egress LSR that the binding of RSVP-TE LSP to an application and a payload identifier is being signaled out-of-band. This flag MUST NOT be modified by any other LSRs in the network. LSRs other than the egress LSRs SHOULD ignore this flag.

When an egress LSR that supports the "OOB mapping flag" receives a Path message with that flag set, the egress LSR MUST set the "OOB mapping flag" in the Flags field of the RRO Attributes subobject. The rest of the RSVP signaling proceeds as normal. However, the LSR MUST have received the OOB mapping before accepting traffic on the LSP. This implies that the egress LSR MUST NOT set up forwarding state for the LSP before it receives the OOB mapping.

Note that the payload information SHOULD be supplied by the OOB mapping. If the egress LSR receives the payload information from OOB mapping, then the LSR MUST ignore the L3PID (Layer 3 Protocol ID) in the Label Request Object [RFC3209].

If the egress LSR

- supports the LSP Attributes object but does not recognize the Attributes Flags TLV; or
- supports the LSP Attributes object and recognizes the Attributes Flags TLV, but does not recognize the "OOB mapping flag";

then it silently ignores the request according to the processing rules of [RFC5420].

An ingress LSR requesting OOB mapping SHOULD examine the "OOB mapping flag" in the Flags field of the RRO Attributes subobject and MAY send a Path Tear to the egress, which has not set the "OOB mapping flag".

When signaling a P2MP LSP, a source node may wish to solicit an individual response to the "OOB mapping flag" from the leaf nodes. Given the constraints on how the LSP Attributes object may be carried in Path and Resv messages according to [RFC5420], in this situation, the source node MUST use a separate Path message for each leaf in networks where [RFC6510] is not supported. In networks with [RFC6510] deployed, either a single leaf per Path message or multiple leaves per Path message MAY be used by the source node.

In deploying applications where the egress LSR receives the binding of the RSVP-TE LSP to an application and a payload identifier using an OOB mechanism, it is important to recognize that the OOB mapping is sent asynchronously with respect to the signaling of RSVP-TE LSP. The egress LSR only installs forwarding state for the LSP after it receives the OOB mapping. In deploying applications using an OOB mechanism, an ingress LSR may need to know when the egress is properly set up for forwarding (i.e., has received the OOB mapping). How the ingress LSR determines that the LSR is properly set up for forwarding at the egress LSR is beyond the scope of this document. Nonetheless, if the OOB mapping is not received by the egress LSR within a reasonable time, the procedure defined in Section 2.4 to tear down the LSP is followed.

2.3. Relationship between OOB and Non-PHP Flags

The "Non-PHP behavior flag" and "OOB mapping flag" can appear and be processed independently of each other. However, as mentioned earlier, in the context of the applications discussed in this document, OOB mapping requires non-PHP behavior. An ingress LSR requesting the OOB mapping MAY also set the "Non-PHP behavior flag" in the LSP Attributes object in the Path message.

2.4. Egress Procedure for Label Binding

RSVP-TE signaling completion and the OOB mapping information reception happen asynchronously at the egress. As mentioned in Section 2.2, egress waits for the OOB mapping before accepting traffic on the LSP. Nonetheless, MPLS Operations, Administration, and Maintenance (OAM) mechanisms, e.g., LSP ping and traceroute, as defined in [RFC4379] and [RFC6425], are expected to work independently of OOB mapping learning process.

In order to avoid unnecessary use of the resources and possible black-holing of traffic, an egress LSR MAY send a Path Error message if the OOB mapping information is not received within a reasonable time. This Path Error message SHOULD include the error code/sub-code "Notify Error / no OOB mapping received" for all affected LSPs. If a notify request was included when the LSP was initially set up, a

Notify message (as defined in [RFC3473]) MAY also be used for delivery of this information to the ingress LSR. An egress LSR MAY implement a cleanup timer for this purpose. The time-out value is a local decision at the egress, with a RECOMMENDED default value of 60 seconds.

3. Security Considerations

The addition of non-PHP behavior adds a variety of attacks on the label assigned by the egress node. As change in the value of the egress label reported in the RRO can cause the LSP to be torn down, additional security considerations for protecting labels assigned by the egress node are required. Security mechanisms as identified in [RFC5920], [RFC2205], [RFC3209], [RFC3473], [RFC5420], and [RFC4875] can be used for this purpose. This document does not introduce any additional security issues above those identified in [RFC5920], [RFC2205], [RFC3209], [RFC3473], [RFC5420], and [RFC4875].

4. IANA Considerations

The following changes to the Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Parameters registry are required.

4.1. Attribute Flags for LSP Attributes Object

The following new flags are defined for the Attributes Flags TLV in the LSP Attributes object.

- o Non-PHP behavior flag:

This flag is used in the Attributes Flags TLV in a Path message. The flag has a corresponding new flag to be used in the RRO Attributes subobject. As per [RFC5420], the bit numbering in the Attribute Flags TLV and the RRO Attributes subobject is identical. That is, the same attribute is indicated by the same bit in both places. This flag is not allowed in the Attributes Flags TLV in a Resv message. Specifically, attributes of this flag are as follows:

- Bit Number: 7
- Attribute flag carried in Path message: Yes
- Attribute flag carried in Resv message: No
- Attribute flag carried in RRO message: Yes

- o OOB mapping flag:

This flag is used in the Attributes Flags TLV in a Path message. The flag has a corresponding new flag to be used in the RRO Attributes subobject. As per [RFC5420], the bit numbering in the Attribute Flags TLV and the RRO Attributes subobject is identical. That is, the same attribute is indicated by the same bit in both places. This flag is not allowed in the Attributes Flags TLV in a Resv message. Specifically, attributes of this flag are as follows:

- Bit Number: 8
- Attribute flag carried in Path message: Yes
- Attribute flag carried in Resv message: No
- Attribute flag carried in RRO message: Yes

4.2. New RSVP Error Sub-Code

For Error Code = 25 "Notify Error" (see [RFC3209]), the following sub-code is defined.

Sub-code	Value
-----	-----
No OOB mapping received	12

5. Acknowledgements

The authors would like to thank Yakov Rekhter for his suggestions on this document.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 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.

- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [RFC5420] Farrel, A., Ed., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.
- [RFC6510] Berger, L. and G. Swallow, "Resource Reservation Protocol (RSVP) Message Formats for Label Switched Path (LSP) Attributes Objects", RFC 6510, February 2012.

6.2. Informative References

- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [RFC4761] Kompella, K., Ed., and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010.
- [RFC6425] Saxena, S., Ed., Swallow, G., Ali, Z., Farrel, A., Yasukawa, S., and T. Nadeau, "Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping", RFC 6425, November 2011.
- [RFC6513] Rosen, E., Ed., and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, February 2012.

Authors' Addresses

Zafar Ali
Cisco Systems, Inc.
EMail: zali@cisco.com

George Swallow
Cisco Systems, Inc.
EMail: swallow@cisco.com

Rahul Aggarwal
Juniper Networks
EMail: raggarwa_1@yahoo.com

