

MPLS Working Group

Internet Draft

Intended status: Standard Track

Expires: September 7, 2010

Z. Ali
G. Swallow
Cisco Systems, Inc.
R. Aggarwal
Juniper Networks
March 8, 2010

**Non PHP Behavior and out-of-band mapping for RSVP-TE LSPs
draft-ietf-mpls-rsvp-te-no-php-oob-mapping-04.txt**

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). 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.

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/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on September 7, 2010.

Copyright

Copyright (c) 2010 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.

Abstract

There are many deployment scenarios which require Egress Label Switching Router (LSR) to receive binding of the Resource Reservation Protocol Traffic Engineered (RSVP-TE) Label Switched Path (LSP) to an application, and payload identification, using some "out-of-band" (OOB) mechanism. This document proposes 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.

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](#)].

Table of Contents

1.	Introduction.....	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
5.	Acknowledgments.....	7

Expires August 2010

[Page 2]

6. References.....	8
6.1. Normative References.....	8
6.2. Informative References.....	8

1. Introduction

When RSVP-TE is used for applications like MVPN [[MVPN](#)] and VPLS [[VPLS](#)], an Egress LSR receives the binding of the RSVP-TE LSP to an application, and payload identification, using an "out-of-band" (OOB) mechanism (e.g., using BGP). In such cases, the Egress LSR cannot make correct forwarding decision 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.

There are other applications that require non-PHP behavior. When RSVP-TE 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-TP LSPs [[MPLS-TP-Framework](#)].

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 payload identifier (payload-Id) needs to be learned via an out-of-band mapping mechanism. The procedures described in this document are equally applicable for P2P and P2MP LSPs. Specification of the OOB communication mechanism(s) is beyond the scope of this document.

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](#)]:

Expires August 2010

[Page 3]

Bit Number 6 (TBD): non-PHP behavior flag.

An Ingress LSR sets the "non-PHP behavior flag" to signal the egress LSRs SHOULD assign 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.

When signaling a P2MP LSP, a source node may wish to solicit individual response to "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 a source node SHOULD use a separate Path message for each leaf.

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 include the LSP_ATTRIBUTES object with "non-PHP behavior flag" set in the Resv message. For this purpose, as defined in [RFC5420](#), the LSP_ATTRIBUTES object is placed in the flow descriptor and is associated with the FILTER_SPEC object that precedes it.

If the egress LSR

- supports the LSP_ATTRIBUTES object but does not recognize the Attributes Flags TLV; or
- supports the LSP_ATTRIBUTES object and recognize the Attributes Flags TLV, but does not recognize "non-PHP behavior flag";

then it SHOULD silently ignore this request.

An ingress LSR requesting non-PHP behavior SHOULD examine Resv message for presence of "Non-PHP behavior flag" in the LSP_ATTRIBUTES object. An ingress LSR requesting non-PHP behavior MAY send a Path Tear to the Egress which has not included the LSP_ATTRIBUTES object in the Resv or which has included the LSP_ATTRIBUTES object in Resv but has not set the "Non-PHP behavior flag" in it. 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.

Expires August 2010

[Page 4]

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 7 (TBD): OOB mapping flag.

An Ingress LSR sets the OOB mapping flag to signal the Egress LSR that binding of RSVP-TE LSP to an application and payload identification 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 signaling a P2MP LSP, a source node may wish to solicit individual response to "OOB mapping 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 a source node SHOULD use a separate Path message for each leaf.

If an egress LSR receiving the Path message, supports the LSP_ATTRIBUTES object and the Attributes Flags TLV, and also recognizes the "OOB mapping flag", it MUST include the LSP_ATTRIBUTES object with "OOB mapping flag" set in the Resv message. For this purpose, as defined in [RFC5420](#), the LSP_ATTRIBUTES object is placed in the flow descriptor and is associated with the FILTER_SPEC object that precedes it.

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 setup 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 L3PID in the Label Request Object [\[RFC3209\]](#).

If the egress LSR

- supports the LSP_ATTRIBUTES object but does not recognize the Attributes Flags TLV; or

Expires August 2010

[Page 5]

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

then it SHOULD silently ignore this request.

An ingress LSR requesting OOB mapping SHOULD examine Resv message for presence of "OOB mapping flag" in in the LSP_ATTRIBUTES object. An ingress LSR requesting OOB mapping MAY send a Path Tear to the Egress which has not included the LSP_ATTRIBUTES object in the Resv or which has included the LSP_ATTRIBUTES object in Resv but has not set the " OOB mapping flag" in it.

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

2.3. Relationship between OOB and non-PHP flags

"Non-PHP behavior desired" and "OOB mapping indication" flags can appear and be processed independently of each other. However, as mentioned earlier, in the context of applications discussed in this document, OOB mapping require non-PHP behavior. An Ingress LSR requesting OOB mapping MAY also set "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](#), Egress waits for the OOB mapping before accepting traffic on the LSP.

In order to avoid unnecessary use of the resources and possible black-holing of traffic, an Egress LSR MAY send a Path Error

Expires August 2010

[Page 6]

message if the OOB mapping information is not received within a reasonable time. This Path Error message will include the error code/sub-code "Notify Error/ no OOB mapping received" for all affected LSPs. If notify request was included when the LSP was initially setup, 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

This document does not introduce any new security issues above those identified in [\[RFC2205\]](#), [\[RFC3209\]](#), [\[RFC3473\]](#), [\[RFC5420\]](#) and [\[RFC4875\]](#).

4. IANA Considerations

4.1. Attribute Flags for LSP_ATTRIBUTES object

The following new flags are being defined for the Attributes Flags TLV in the LSP_ATTRIBUTES object. The numeric values are to be assigned by IANA.

- o Non-PHP behavior flag - Bit Number 6 (Suggested value).
- o OOB mapping flag - Bit Number 7 (Suggested value).
- o These flags are to be used in the Attributes Flags TLV in both Path and Resv messages.

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

Sub-code	Value
-----	-----
No OOB mapping received	12 (TBD)

5. Acknowledgments

The authors would like to thank Yakov Rekhter for his suggestions on the draft.

6. References

6.1. Normative References

- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5420] A. Farrel, D. Papadimitriou, J. P. Vasseur and A. Ayyangar, "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using RSVP-TE", [RFC 5420](#), February 2006.
- [RFC3209] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [RFC4875] R. Aggarwal, D. Papadimitriou, S. Yasukawa, et al, "Extensions to RSVP-TE for Point-to-Multipoint TE LSPs", [RFC 4875](#).
- [RFC3473] L. Berger, Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.

6.2. Informative References

- [MVPN] E. Rosen, R. Aggarwal et al, "Multicast in MPLS/BGP IP VPNs", [draft-ietf-l3vpn-2547bis-mcast-08.txt](#), work in progress.
- [VPLS] R. Aggarwal, et al, "Propagation of VPLS IP Multicast Group Membership Information", [draft-raggarwa-l2vpn-vpls-mcast-ctrl-00.txt](#), work in progress.
- [MPLS-TP-Framework] M. Bocci, S. Bryant, et al, "A Framework for MPLS in Transport Networks", [draft-ietf-mpls-tp-framework-06](#), work in progress.

Author's Addresses

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

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

Rahul Aggarwal
Juniper Networks
rahul@juniper.net