

MPLS Working Group

Internet Draft

Intended status: Standard Track
Expires: December 2008

Z. Ali
G. Swallow
Cisco Systems, Inc.
R. Aggarwal
Juniper Networks
June 19, 2008

**Non PHP Behavior and out-of-band mapping for RSVP-TE LSPs
draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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 December 19, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

There are many deployment scenarios which require Egress LSR to receive binding of the RSVP-TE 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

In examples, "C:" and "S:" indicate lines sent by the client and server respectively.

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](#) 0.

Table of Contents

1.	Introduction.....	2
2.	RSVP-TE signaling extensions.....	3
2.1.	Signaling non-PHP behavior.....	3
2.2.	Signaling OOB Mapping Indication.....	4
2.3.	Relationship between OOB and non-PHP bits.....	4
2.4.	Egress Procedure for label binding.....	4
3.	Security Considerations.....	5
4.	IANA Considerations.....	5
4.1.	Attribute Flags for LSP_ATTRIBUTES object.....	5
5.	Acknowledgments.....	6
6.	References.....	6
6.1.	Normative References.....	6
6.2.	Informative References.....	6
	Author's Addresses.....	7
	Intellectual Property Statement.....	7
	Disclaimer of Validity.....	7

[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. Therefore, non Penultimate Hop Popping (non-PHP)

behavior is required at the Egress LSR to apply OOB mapping.

Expires December 2008

[Page 2]

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 [[UPSTREAM](#)] when the context is a MPLS LSP.

This document defines two new bits in the Attributes Flags TLV of the LSP_ATTRIBUTES object defined in [[RFC4420](#)]: one bit for communication of non-PHP behavior, and one bit 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 the 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 RSVP-TE LSP, this document defines a new bit in the Attributes Flags TLV of the LSP_ATTRIBUTES object defined in [[RFC4420](#)]:

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

This bit SHOULD be set by Ingress node in the Attributes Flags TLV of the LSP_ATTRIBUTES object in the Path message for the LSP that desires Non-PHP behavior. This bit MUST NOT be modified by any other nodes in the network. Nodes other than the Egress nodes SHOULD ignore this bit.

If an egress node receiving the Path message, supports the LSP_ATTRIBUTES object and the Attributes Flags TLV, and also recognizes the "non-PHP behavior desired bit", it MUST allocate a non-NULL local label. If the egress node supports the LSP_ATTRIBUTES object but does not recognize the Attributes Flags TLV, or supports the TLV as well but does not recognize this particular bit, then it SHOULD simply ignore the above request.

Expires December 2008

[Page 3]

An ingress node requesting non-PHP behavior MAY examine the label value corresponding to the Egress node(s) in the RRO, and MAY send a Path Tear to the Egress which assigns a Null label value.

2.2. Signaling OOB Mapping Indication

In order to indicate to the Egress LSR that binding of RSVP-TE LSP to an application and payload identification is being communicated by an OOB mechanism, this document defines a new bit in the Attributes Flags TLV of the LSP_ATTRIBUTES object defined in [\[RFC4420\]](#):

Bit Number 7 (TBD): OOB mapping indication bit.

This bit SHOULD be set by Ingress node in the Attributes Flags TLV of the LSP_ATTRIBUTES object in the Path message for the LSP that desires OOB mapping. This bit MUST NOT be modified by any other nodes in the network. Nodes other than the Egress nodes SHOULD ignore this bit.

If an egress node receiving the Path message, supports the LSP_ATTRIBUTES object and the Attributes Flags TLV, and also recognizes the "OOB mapping indication bit", it MUST wait for the OOB mapping before accepting traffic on the P2MP LSP. This implies that the egress node MUST NOT setup forwarding state for the P2MP LSP before it receives the OOB mapping, though it SHOULD proceed with RSVP-TE signaling and send RESV messages as per regular RSVP-TE procedures [\[RFC3209\]](#). It MUST also ignore L3PID in the Label Request Object [\[RFC3209\]](#). If the egress node supports the LSP_ATTRIBUTES object but does not recognize the Attributes Flags TLV, or supports the TLV as well but does not recognize this particular bit, then it SHOULD simply ignore the above request.

2.3. Relationship between OOB and non-PHP bits

Non-PHP behavior desired and OOB mapping indication bit can appear and be processed independently of each other. However, as mentioned earlier, in the context of application discussed in this draft, OOB mapping require non-PHP behavior. An Ingress node requesting OOB mapping MAY also set non-PHP behavior desired bit 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

Expires December 2008

[Page 4]

[Section 2](#), Egress waits for the OOB mapping before accepting traffic on the P2MP LSP.

In order to avoid unnecessary use of the resources and possible block-holding of traffic, if the OOB mapping information is not received within a reasonable time, Egress MAY trigger a Path Error message with the error code/sub-code "Notify Error/ no OOB mapping received" for all affected LSPs. If available, and where notify requests were included when the LSPs were initially setup, Notify message (as defined in [[RFC3473](#)]) MAY also be used for delivery of this information to the Ingress node. Egress node may implement a cleanup timer for this purpose. The time-out value is a local decision at the Egress, with recommended default value is to be added later.

3. Security Considerations

This document does not introduce any new security issues above those identified in [[RFC3209](#)], [[RFC4420](#)] and [[RSVP-TE-P2MP](#)].

4. IANA Considerations

4.1. Attribute Flags for LSP_ATTRIBUTES object

The following new bit is being defined for the Attributes Flags TLV in the LSP_ATTRIBUTES object. The numeric value is to be assigned by IANA.

- o Non-PHP behavior desired bit - Bit Number 6 (Suggested value).
- o OOB mapping indication bit - Bit Number 7 (Suggested value).

These bits are only to be used in the Attributes Flags TLV on a Path message.

The following new error sub-code for Error Code = 25 "Notify Error" (see [[RFC3209](#)]) is needed. The numeric value for this sub-code is to be assigned by IANA.

- o No OOB mapping received.

5. Acknowledgments

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

6. References

6.1. Normative References

- [RFC4420] 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 4420](#), 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.
- [RSVP-TE-P2MP] R. Aggarwal, D. Papadimitriou, S. Yasukawa, et al, "Extensions to RSVP-TE for Point-to-Multipoint TE LSPs", [RFC4875](#).
- [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-06.txt](#).
- [VPLS] R. Aggarwal, et al, "Propagation of VPLS IP Multicast Group Membership Information", [draft-raggarwa-l2vpn-vpls-mcast-ctrl-00.txt](#), work in progress.
- [UPSTREAM] TBA.

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
Email: rahul@juniper.net

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Expires December 2008

[Page 7]

Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.