MPLS Working Group
Z. Ali
G. Swallow
Internet Draft
Cisco Systems, Inc.
R. Aggarwal
Juniper Networks

Intended status: Standard Track
October 26, 2009
Expires: April 25, 2010

Non PHP Behavior and out-of-band mapping for RSVP-TE LSPs
draft-ietf-mpls-rsvp-te-no-php-oob-mapping-03.txt

Status of this Memo

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

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.

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 RSVP-TE LSP, this document defines a new flag in the Attributes Flags TLV of the LSP_ATTRIBUTES object defined in [RFC5420]:

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

In order to indicate to the Ingress LSR that the Egress LSR recognizes the "non-PHP behavior desired flag", the following new bit is defined in the Flags field of the Record Route object (RRO) Attributes subobject:

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

An Ingress LSR sets the non-PHP behavior desired 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.

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 desired flag", it MUST allocate a non-NULL local label. The egress LSR MUST also set the "Non-PHP behavior acknowledgement flag" in the Flags field of the RRO Attribute subobject.

If the egress LSR 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 flag, then it SHOULD simply ignore the above request.

An ingress LSR requesting non-PHP behavior MAY examine "Non-PHP behavior acknowledgement flag" in the Flags field of the RRO Attribute subobject and MAY send a Path Tear if the Egress has not set the "Non-PHP behavior acknowledgement 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 if the Egress has assigns a Null label value.

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 indication flag.

In order to indicate to the Ingress LSR that the Egress LSR recognizes the "OOB mapping indication flag", the following new bit is defined in the Flags field of the Record Route object (RRO) Attributes subobject:

Bit Number 7 (TBD): OOB mapping acknowledgement flag.

An Ingress LSR sets the OOB mapping indication 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 an egress LSR which supports the "OOB mapping indication flag", receives a Path message with that flag set, the egress LSR MUST set the "OOB mapping acknowledgement flag" in the Flags field of the RRO Attribute 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 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 supports the TLV as well but does not recognize this particular flag, then it SHOULD simply ignore the above request.

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

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 application discussed in this draft, OOB mapping require non-PHP behavior. An Ingress LSR requesting OOB mapping MAY also set non-PHP behavior desired 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 block-holing 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 messages (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 [RFC3209], [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.

- Non-PHP behavior desired flag - Bit Number 6 (Suggested value).

- OOB mapping indication flag - Bit Number 7 (Suggested value).

These flags are only to be used in the Attributes Flags TLV on a Path message. These flags have corresponding new flags to be used in the RRO Attributes subobject. As per RFC5420 [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. Specifically, the numeric values for the corresponding new flags to be used in the RRO Attributes subobject are to be assigned by IANA.

- OOB mapping acknowledgement flag - Bit Number 6 (Suggested value).

- Non-PHP behavior acknowledgement flag - Bit Number 7 (Suggested value).

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

[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

Copyright Notice

documents carefully, as they describe your rights and restrictions with respect to this document.

Legal