

Network Working Group
Internet Draft
Category: Standards Track
Expiration Date: August 2011

R. Aggarwal
Juniper Networks

J. L. Le Roux
France Telecom

February 02, 2011

MPLS Upstream Label Assignment for LDP

draft-ietf-mpls-ldp-upstream-10.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and 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/lid-abstracts.txt>.

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

Copyright and License Notice

Copyright (c) 2011 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.

Abstract

This document describes procedures for distributing upstream-assigned labels for Label Distribution Protocol (LDP). It also describes how these procedures can be used for avoiding branch Label Switching Router (LSR) traffic replication on a LAN for LDP point-to-multipoint (P2MP) Label Switched Paths (LSPs).

Table of Contents

1	Specification of requirements	3
2	Introduction	3
3	LDP Upstream Label Assignment Capability	4
4	Distributing Upstream-Assigned Labels in LDP	5
4.1	Procedures	5
5	LDP Tunnel Identifier Exchange	6
6	LDP Point-to-Multipoint LSPs on a LAN	10
7	IANA Considerations	12
7.1	LDP TLVs	12
7.2	Interface Type Identifiers	12
8	Security Considerations	12
9	Acknowledgements	13
10	References	13
10.1	Normative References	13
10.2	Informative References	13
11	Author's Address	14

1. 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 [RFC2119].

2. Introduction

This document describes procedures for distributing upstream-assigned labels [RFC5331] for Label Distribution Protocol (LDP) [RFC5036]. These procedures follow the architecture for MPLS Upstream Label Assignment described in [RFC5331].

This document describes extensions to LDP that a Label Switching Router (LSR) can use to advertise to its neighboring LSRs whether the LSR supports upstream label assignment.

This document also describes extensions to LDP to distribute

upstream-assigned labels.

The usage of MPLS upstream label assignment using LDP for avoiding branch LSR traffic replication on a LAN for LDP point-to-multipoint (P2MP) Label Switched Paths (LSPs) [MLDP] is also described.

3. LDP Upstream Label Assignment Capability

According to [RFC5331], upstream-assigned label bindings MUST NOT be used unless it is known that a downstream LSR supports them. This implies that there MUST be a mechanism to enable an LSR to advertise to its LDP neighbor LSR(s) its support of upstream-assigned labels.

A new Capability Parameter, the LDP Upstream Label Assignment Capability, is introduced to allow an LDP peer to exchange with its peers, its support of upstream label assignment. This parameter follows the format and procedures for exchanging Capability Parameters defined in [RFC5561].


Following is the format of the LDP Upstream Label Assignment Capability Parameter:

0																1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9																								
1 0 Upstream Lbl Ass Cap(IANA)																Length (= 1)																																															
1 Reserved																																																															


If an LSR includes the Upstream Label Assignment Capability in LDP Initialization Messages it implies that the LSR is capable of both distributing upstream-assigned label bindings and receiving upstream-assigned label bindings. The reserved bits MUST be set to zero on transmission and ignored on receipt. The Upstream Label Assignment Capability Parameter MUST be carried only in LDP initialization messages and MUST be ignored if received in LDP Capability messages.

4. Distributing Upstream-Assigned Labels in LDP

An optional LDP TLV, Upstream-Assigned Label Request TLV, is introduced. To request an upstream-assigned label an LDP peer MUST include this TLV in a Label Request message.



An optional LDP TLV, Upstream-Assigned Label TLV is introduced to signal an upstream-assigned label. Upstream-Assigned Label TLVs are carried by the messages used to advertise, release and withdraw upstream assigned label mappings.



The Label field is a 20-bit label value as specified in [RFC3032] represented as a 20-bit number in a 4 octet field as specified in section 3.4.2.1 of RFC5036 [RFC5036].

4.1. Procedures

Procedures for Label Mapping, Label Request, Label Abort, Label Withdraw and Label Release follow [RFC5036] other than the modifications pointed out in this section.

A LDP LSR MUST NOT distribute the Upstream Assigned Label TLV to a neighboring LSR if the neighboring LSR had not previously advertised the Upstream Label Assignment Capability in its LDP Initialization messages. A LDP LSR MUST NOT send the Upstream Assigned Label Request TLV to a neighboring LSR if the neighboring LSR had not previously advertised the Upstream Label Assignment Capability in its LDP Initialization messages.

As described in [RFC5331] the distribution of upstream-assigned labels is similar to either ordered LSP control or independent LSP control of the downstream assigned labels.

When the label distributed in a Label Mapping message is an upstream-assigned label, the Upstream Assigned Label TLV MUST be included in the Label Mapping message. When an LSR receives a Label Mapping message with an Upstream Assigned Label TLV and it does not recognize the TLV, it MUST generate a Notification message with a status code of "Unknown TLV" [RFC5036]. If it does recognize the TLV but is unable to process the upstream label, it MUST generate a Notification message with a status code of "No Label Resources". If the Label Mapping message was generated in response to a Label Request message, the Label Request message MUST contain an Upstream Assigned Label Request TLV. A LSR that generates an upstream assigned label request to a neighbor LSR, for a given FEC, MUST NOT send a downstream label mapping to the neighbor LSR for that FEC unless it withdraws the upstream-assigned label binding. Similarly if an LSR generates a downstream assigned label request to a neighbor LSR, for a given FEC, it MUST NOT send an upstream label mapping to that LSR for that FEC, unless it aborts the downstream assigned label request.

The Upstream Assigned Label TLV may be optionally included in Label Withdraw and Label Release messages that withdraw/release a particular upstream assigned label binding.

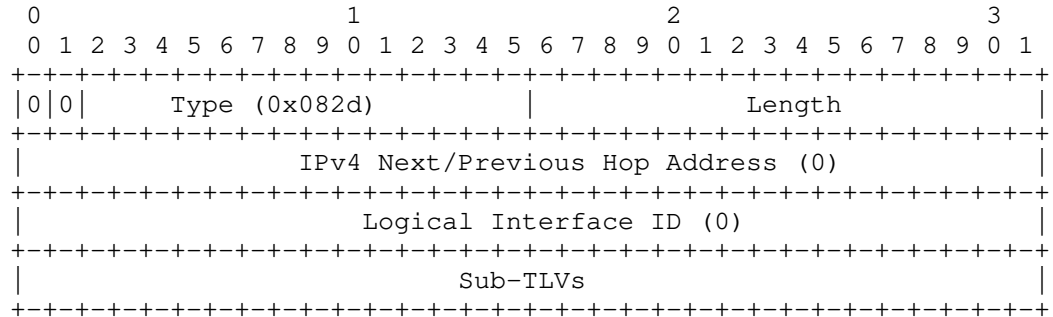
5. LDP Tunnel Identifier Exchange

As described in [RFC5331] an upstream LSR Ru MAY transmit an MPLS packet, the top label of which (L) is upstream-assigned, to a downstream LSR Rd, by encapsulating it in an IP or MPLS tunnel. In this case the fact that L is upstream-assigned is determined by Rd by the tunnel on which the packet is received. There must be a mechanism for Ru to inform Rd that a particular tunnel from Ru to Rd will be used by Ru for transmitting MPLS packets with upstream-assigned MPLS labels.

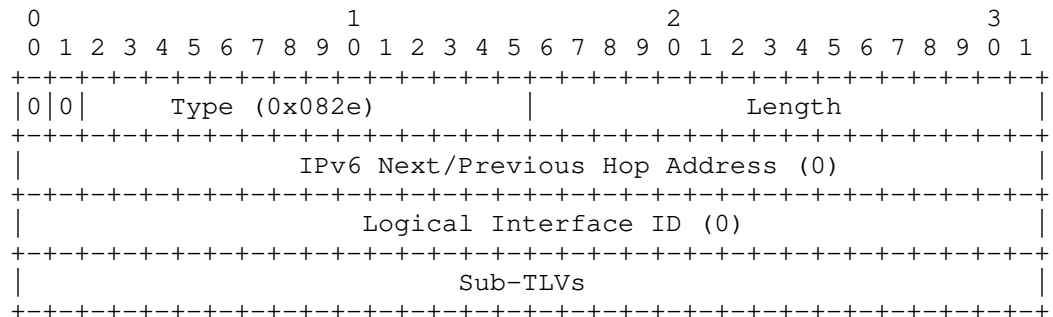
When LDP is used for upstream label assignment, the Interface ID TLV [RFC3472] is used for signaling the Tunnel Identifier. If Ru uses an IP or MPLS tunnel to transmit MPLS packets with upstream assigned labels to Rd, Ru MUST include the Interface ID TLV in the Label Mapping messages along with the Upstream Assigned Label TLV. The IPv4/v6 Next/Previous Hop Address and the Logical Interface ID fields in the Interface ID TLV SHOULD be set to 0 by the sender and ignored by the receiver. The Length field indicates the total length of the TLV, i.e., 4 + the length of the value field in octets. A value field whose length is not a multiple of four MUST be zero-padded so

that the TLV is four- octet aligned.

Hence the IPv4 Interface ID TLV has the following format:



The IPv6 Interface ID TLV has the following format:

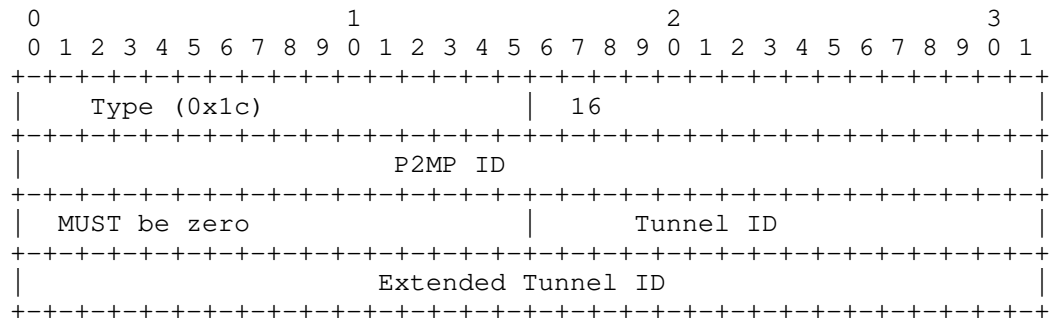


As shown in the above figures the Interface ID TLV carries sub-TLVs. Four new Interface ID sub-TLVs are introduced to support RSVP-TE P2MP LSPs, LDP P2MP LSPs, IP Multicast Tunnels and context labels. The sub-TLV value in the sub-TLV acts as the tunnel identifier.


Following are the sub-TLVs that are introduced:

1. RSVP-TE P2MP LSP TLV. Type = 28 (To be assigned by IANA). Value of the TLV is the RSVP-TE P2MP LSP SESSION Object [RFC4875].

Below is the RSVP-TE P2MP LSP TLV format when carried in the IPv4 Interface ID TLV:




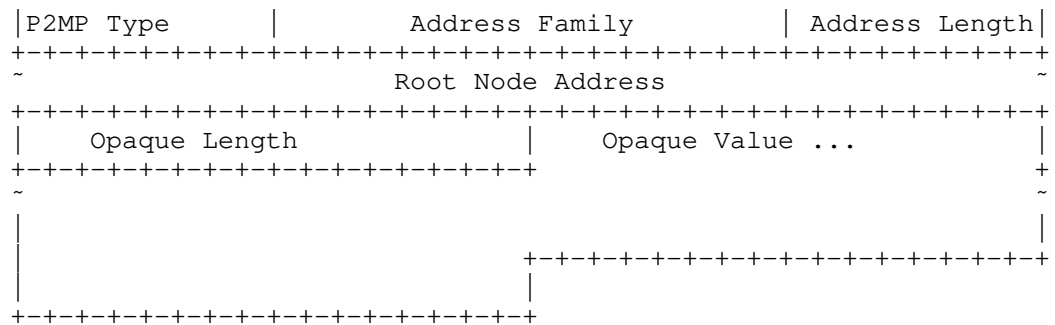
Below is the RSVP-TE P2MP LSP TLV format when carried in the IPv6 Interface ID TLV:



This TLV identifies the RSVP-TE P2MP LSP. It allows Ru to tunnel an "inner" LDP P2MP LSP, the label for which is upstream assigned, over an "outer" RSVP-TE P2MP LSP that has leaves <Rd1...Rdn>. The RSVP-TE P2MP LSP IF_ID TLV allows Ru to signal to <Rd1...Rdn> the binding of the inner LDP P2MP LSP to the outer RSVP-TE P2MP LSP. The control plane signaling between Ru and <Rd1...Rdn> for the inner P2MP LSP uses targeted LDP signaling messages

2. LDP P2MP LSP TLV. Type = 29 (To be assigned by IANA). Value of the TLV is the LDP P2MP FEC as defined in [MLDP] and has to be set as per the procedures in [MLDP]. Here is the format of the LDP P2MP FEC as defined in [MLDP]:





The Address Family MUST be set to IPv4, the Address Length MUST be set to 4 and the Root Node Address MUST be set to an IPv4 address when the LDP P2MP LSP TLV is carried in the IPv4 Interface ID TLV. The Address Family MUST be set to IPv6, the Address Length MUST be set to 16 and the Root Node Address MUST be set to an IPv6 address when the LDP P2MP LSP TLV is carried in the IPv6 Interface ID TLV.

The TLV value identifies the LDP P2MP LSP. It allows Ru to tunnel an "inner" LDP P2MP LSP, the label for which is upstream assigned, over an "outer" LDP P2MP LSP that has leaves <Rd1...Rdn>. The LDP P2MP LSP IF_ID TLV allows Ru to signal to <Rd1...Rdn> the binding of the inner LDP P2MP LSP to the outer LDP- P2MP LSP. The control plane signaling between Ru and <Rd1...Rdn> for the inner P2MP LSP uses targeted LDP signaling messages

3. IP Multicast Tunnel TLV. Type = 30 (To be assigned by IANA) In this case the TLV value is a <Source Address, Multicast Group Address> tuple. Source Address is the IP address of the root of the tunnel i.e. Ru, and Multicast Group Address is the Multicast Group Address used by the tunnel. The addresses MUST be IPv4 addresses when the IP Multicast Tunnel TLV is included in the IPv4 Interface ID TLV. The addresses MUST be IPv6 addresses when the IP Multicast Tunnel TLV is included in the IPv6 Interface ID TLV.

4. MPLS Context Label TLV. Type = 31 (To be assigned by IANA). In this case the TLV value is a <Source Address, MPLS Context Label> tuple. The Source Address belongs to Ru and the MPLS Context Label is an upstream assigned label, assigned by Ru. The Source Address MUST be set to an IPv4 address when the MPLS Context Label TLV is carried in the IPv4 Interface ID TLV. The Source Address MUST be set to an IPv6 address when the MPLS Context Label TLV is carried in the IPv6 Interface ID TLV. This allows Ru to tunnel an "inner" LDP P2MP LSP, the label of which is upstream assigned, over an "outer" one-hop MPLS LSP, where the outer one-hop LSP has the following property:

- + The label pushed by Ru for the outer MPLS LSP is an upstream assigned context label, assigned by Ru. When <Rd1...Rdn> perform an MPLS label lookup on this label a combination of this label and the incoming interface MUST be sufficient for <Rd1...Rdn> to uniquely determine Ru's context specific label space to lookup the next label on the stack in. <Rd1...Rdn> MUST receive the data sent by Ru with the context specific label assigned by Ru being the top label on the label stack.

Currently the usage of the context label TLV is limited only to LDP P2MP LSPs on a LAN as specified in the next section. The context label TLV MUST NOT be used for any other purposes.

Note that when the outer P2MP LSP is signaled with RSVP-TE or MLDP the above procedures assume that Ru has a priori knowledge of all the <Rd1, ... Rdn>. In the scenario where the outer P2MP LSP is signaled using RSVP-TE, Ru can obtain this information from RSVP-TE. However, in the scenario where the outer P2MP LSP is signaled using MLDP, MLDP does not provide this information to Ru. In this scenario the procedures by which Ru could acquire this information are outside the scope of this document.

6. LDP Point-to-Multipoint LSPs on a LAN

This section describes one application of upstream label assignment using LDP. Further applications are to be described in separate documents.

[MLDP] describes how to setup P2MP LSPs using LDP. On a LAN the solution relies on "ingress replication". A LSR on a LAN, that is a branch LSR for a P2MP LSP, (say Ru) sends a separate copy of a packet that it receives on the P2MP LSP to each of the downstream LSRs on the LAN (say <Rd1...Rdn> that are adjacent to it in the P2MP LSP.

It is desirable for Ru to send a single copy of the packet for the LDP P2MP LSP on the LAN, when there are multiple downstream routers on the LAN that are adjacent to Ru in that LDP P2MP LSP. This requires that each of <Rd1...Rdn> must be able to associate the label L, used by Ru to transmit packets for the P2MP LSP on the LAN, with that P2MP LSP. It is possible to achieve this using LDP upstream-assigned labels with the following procedures.

Consider an LSR Rd that receives the LDP P2MP FEC [MLDP] from its downstream LDP peer. Further the upstream interface to reach LSR Ru which is the next-hop to the P2MP LSP root address, Pr, in the LDP P2MP FEC, is a LAN interface, Li. Further Rd and Ru support upstream-assigned labels. In this case Rd instead of sending a Label Mapping

message as described in [MLDP] sends a Label Request message to Ru. This Label Request message MUST contain an Upstream Assigned Label Request TLV.

On receiving this message, Ru sends back a Label Mapping message to Rd with an upstream-assigned label. This message also contains an Interface ID TLV with a MPLS Context Label sub-TLV, as described in the previous section, with the value of the MPLS label set to a value assigned by Ru on interface Li as specified in [RFC5331]. Processing of the Label Request and Label Mapping messages for LDP upstream-assigned labels is as described in section 4.1. If Ru receives a Label Request for an upstream assigned label for the same P2MP FEC from multiple downstream LSRs on the LAN, <Rd1...Rdn>, it MUST send the same upstream-assigned label to each of <Rd1...Rdn>.

Ru transmits the MPLS packet using the procedures defined in [RFC5331] and [RFC5332]. The MPLS packet transmitted by Ru contains as the top label the context label assigned by Ru on the LAN interface, Li. The bottom label is the upstream label assigned by Ru to the LDP P2MP LSP. The top label is looked up in the context of the LAN interface, Li, [RFC5331] by a downstream LSR on the LAN. This lookup enables the downstream LSR to determine the context specific label space to lookup the inner label in.

Note that <Rd1...Rdn> may have more than one equal cost next-hop on the LAN to reach Pr. It MAY be desirable for all of them to send the label request to the same upstream LSR and they MAY select one upstream LSR using the following procedure:

1. The candidate upstream LSRs are numbered from lower to higher IP address
2. The following hash is performed: $H = (\text{Sum Opaque value}) \bmod N$, where N is the number of candidate upstream LSRs. Opaque value is defined in [MLDP] and comprises the P2MP LSP identifier.
3. The selected upstream LSR U is the LSR that has the number H.

This allows for load balancing of a set of LSPs among a set of candidate upstream LSRs, while ensuring that on a LAN interface a single upstream LSR is selected. It is also to be noted that the procedures in this section can still be used by Rd and Ru if other LSRs on the LAN do not support upstream label assignment. Ingress replication and downstream label assignment will continue to be used for LSRs that do not support upstream label assignment.

7. IANA Considerations

7.1. LDP TLVs

IANA maintains a registry of LDP TLVs at the registry "Label Distribution Protocol" in the sub-registry called "TLV Type Name Space".

This document defines a new LDP Upstream Label Assignment Capability TLV (Section 3). IANA is requested to assign the value 0x0507 to this TLV.

This document defines a new LDP Upstream-Assigned Label TLV (Section 4). IANA is requested to assign the type value of 0x204 to this TLV.

This document defines a new LDP Upstream-Assigned Label Request TLV (Section 4). IANA is requested to assign the type value of 0x205 to this TLV.

7.2. Interface Type Identifiers

[RFC3472] defines the LDP Interface ID IPv4 and IPv6 TLV. These top-level TLVs can carry sub-TLVs dependent on the interface type. These sub-TLVs are assigned "Interface ID Types". IANA maintains a registry of Interface ID Types for use in GMPLS in the registry "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters" and sub-registry "Interface_ID Types". IANA is requested to make corresponding allocations from this registry as follows:

- RSVP-TE P2MP LSP TLV (requested value 28)
- LDP P2MP LSP TLV (requested value 29)
- IP Multicast Tunnel TLV (requested value 30)
- MPLS Context Label TLV (requested value 31)

8. Security Considerations

The security considerations discussed in RFC 5036, RFC 5331 and RFC 5332 apply to this document.

More detailed discussion of security issues that are relevant in the context of MPLS and GMPLS, including security threats, related defensive techniques, and the mechanisms for detection and reporting, are discussed in "Security Framework for MPLS and GMPLS Networks

[MPLS-SEC].

9. Acknowledgements

Thanks to Yakov Rekhter for his contribution. Thanks to Ina Minei and Thomas Morin for their comments. The hashing algorithm used on LAN interfaces is taken from [MLDP]. Thanks to Loa Andersson, Adrian Farrel and Eric Rosen for their comments and review.

10. References

10.1. Normative References

[RFC5331] R. Aggarwal, Y. Rekhter, E. Rosen, "MPLS Upstream Label Assignment and Context Specific Label Space", RFC5331

[RFC5332] T. Eckert, E. Rosen, R. Aggarwal, Y. Rekhter, RFC5332

[RFC2119] "Key words for use in RFCs to Indicate Requirement Levels.", Bradner, March 1997

[RFC5036] L. Andersson, et. al., "LDP Specification", RFC5036.

[RFC4875] R. Aggarwal, D. Papadimitriou, S. Yasukawa [Editors], "Extensions to RSVP-TE for Point to Multipoint TE LSPs", RFC 4875

[MLDP] I. Minei et. al, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", draft-ietf-mpls-ldp-p2mp-08.txt

10.2. Informative References

[RFC5561] B. Thomas, K. Raza, S. Aggarwal, R. Aggarwal, JL. Le Roux, "LDP Capabilities", RFC5561

[MPLS-SEC] L. fang, ed, "Security Framework for MPLS and GMPLS Networks", draft-ietf-mpls-mpls-and-gmpls-security-framework-07.txt

[RFC3032] E. Rosen et. al, "MPLS Label Stack Encoding", RFC 3032

[RFC3472] Ashwood-Smith, P. and L. Berger, Editors, " Generalized Multi-Protocol Label Switching (GMPLS) Signaling - Constraint-based Routed Label Distribution Protocol (CR-LDP) Extensions", RFC 3472, January 2003.

11. Author's Address

Rahul Aggarwal
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089
Phone: +1-408-936-2720
Email: rahul@juniper.net

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
France
E-mail: jeanlouis.leroux@orange-ftgroup.com

MPLS Working Group

Internet Draft

Intended status: Standard Track

Expires: February 16, 2012

Z. Ali

G. Swallow

Cisco Systems, Inc.

R. Aggarwal

Juniper Networks

August 17, 2011

Non Penultimate Hop Popping Behavior and out-of-band mapping for
RSVP-TE Label Switched Paths
draft-ietf-mpls-rsvp-te-no-php-oob-mapping-09.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on February 16, 2012.

Copyright Notice

Copyright (c) 2011 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.

Expires February 2012

[Page 1]

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.

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 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.

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

Copyright Notice	1
1. Introduction	3
2. RSVP-TE signaling extensions	4
2.1. Signaling non-PHP behavior	4
2.2. Signaling OOB Mapping Indication	5
2.3. Relationship between OOB and non-PHP flags	7
2.4. Egress Procedure for label binding	7
3. Security Considerations	8
4. IANA Considerations	8
4.1. Attribute Flags for LSP_ATTRIBUTES object	8
4.2. New RSVP error sub-code	9

5. Acknowledgments	9
6. References	9
6.1. Normative References	9
6.2. Informative References	10

1. Introduction

When Resource ReserVation Protocol Traffic Engineered (RSVP-TE) is used for applications like Multicast Virtual Private Network (MVPN) [MVPN] 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 payload identification, using an "out-of-band" (OOB) mechanism (e.g., using Border Gateway Protocol (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. 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-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 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 RRO Attributes subobject, corresponding flags to be carried in RRO Attributes subobject are also defined.

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]:

Bit Number (to be assigned by IANA): non-PHP behavior requested flag.

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

Bit Number (same as bit number assigned for non-PHP behavior requested flag): Non-PHP behavior acknowledgement flag.

An Ingress LSR sets the "non-PHP behavior requested 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 requested 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 LSP_ATTRIBUTES object and recognize the Attributes Flags TLV, but does not recognize the "non-PHP behavior requested flag";

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

An ingress LSR requesting non-PHP behavior SHOULD examine "Non-PHP behavior 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 "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 to the Egress which assigns a Null label value.

When signaling a P2MP LSP, a source node may wish to solicit individual response to the "non-PHP behavior requested 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 [ATTRIBUTE-BNF] is not supported. In networks with [ATTRIBUTE-BNF] deployed either separate Path message for each leaf or multiple leafs 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 (to be assigned by IANA): 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 (same as bit number assigned for OOB mapping indication flag): 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 LSP_ATTRIBUTES object and recognizes the Attributes Flags TLV, but does not recognize the "OOB mapping indication flag";

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

An ingress LSR requesting OOB mapping SHOULD 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".

When signaling a P2MP LSP, a source node may wish to solicit individual response to the "OOB mapping indication 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 [ATTRIBUTE-BNF] is not supported. In

networks with [ATTRIBUTE-BNF] deployed either separate Path message for each leaf or multiple leafs per Path message MAY be used by the source node.

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 the OOB mapping is sent asynchronously with respect to the 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, an Ingress LSR may need to know when the Egress is properly setup for forwarding (i.e., has received the OOB mapping). How the Ingress LSR determines that the LSR is properly setup 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

"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 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 requested 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 OAM mechanisms, e.g., LSP Ping and Trace route as defined in [RFC4379], [P2MP-OAM], are expected to work independent 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 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

Addition of "non-PHP behavior" adds a variable 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 label 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. The numeric values are to be assigned by IANA.

o Non-PHP behavior flag:

This flag is used in the Attributes Flags TLV in a Path message. The flag has 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: To be assigned by IANA.
- 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 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: To be assigned by IANA.
- 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	to be assigned by IANA.

5. Acknowledgments

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

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.

- [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] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003..
- [RFC2205] R. Braden, Ed., "Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification", RFC 2205, September 1997.
- [ATTRIBUTE-BNF] Berger, L. and Swallow, G., "LSP Attributes Related Routing Backus-Naur Form", draft-ietf-ccamp-attribute-bnf, work in progress.

6.2. Informative References

- [MVPN] E. Rosen, R. Aggarwal et al, "Multicast in MPLS/BGP IP VPNs", draft-ietf-l3vpn-2547bis-mcast-10.txt, work in progress.
- [RFC4761] Kompella, K., Ed., and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007.
- [RFC5921] M. Bocci, S. Bryant, et al, "A Framework for MPLS in Transport Networks", RFC 5921, January 2007.
- [RFC5920] L. Fang, Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

- [RFC4379] K. Kompella, and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [P2MP-OAM] S. Saxena, Ed., G. Swallow, Z. Ali, A. Farrel, S. Yasukawa, T. Nadeau, "Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping", draft-ietf-mpls-p2mp-lsp-ping-17.txt, 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

Expires February 2012

[Page 11]