

TEAS Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 12, 2017

X. Zhang, Ed.
Huawei Technologies
V. Beeram, Ed.
Juniper Networks
I. Bryskin
Huawei Technologies
D. Ceccarelli
Ericsson
O. Gonzalez de Dios
Telefonica
March 11, 2017

Network Assigned Upstream-Label
draft-ietf-teas-network-assigned-upstream-label-04

Abstract

This document discusses a Generalized Multi-Protocol Label Switching (GMPLS) Resource reSerVation Protocol with Traffic Engineering (RSVP-TE) mechanism that enables the network to assign an upstream label for a bidirectional LSP. This is useful in scenarios where a given node does not have sufficient information to assign the correct upstream label on its own and needs to rely on the downstream node to pick an appropriate label.

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 September 12, 2017.

Copyright Notice

Copyright (c) 2017 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.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	Use-Case: Wavelength Setup for IP over Optical Networks	3
3.	The "Crankback Signaling" Approach	3
4.	Symmetric Labels	5
5.	Unassigned Upstream Label	5
5.1.	Processing Rules	6
5.2.	Backwards Compatibility	6
6.	Applicability	7
6.1.	Initial Setup	7
6.2.	Wavelength Change	8
7.	Acknowledgements	8
8.	Contributors	9
9.	IANA Considerations	9
10.	Security Considerations	9
11.	References	9
11.1.	Normative References	9
11.2.	Informative References	10
	Authors' Addresses	10

[1.](#) Introduction

The Generalized Multi-Protocol Label Switching (GMPLS) Resource reSerVation Protocol with Traffic Engineering (RSVP-TE) extensions for setting up a bidirectional LSP are specified in [RFC 3473](#) [[RFC3473](#)]. The bidirectional LSP setup is indicated by the presence of an UPSTREAM_LABEL Object in the PATH message. As per the existing setup procedure outlined for a bidirectional LSP, each upstream node must allocate a valid upstream label on the outgoing interface before sending the initial PATH message downstream. However, there are certain scenarios where it is not desirable or possible for a given node to pick the upstream label on its own. This document defines the protocol mechanism to be used in such scenarios. This mechanism enables a given node to offload the task of assigning the upstream label for a given bidirectional LSP onto the network.


1.1. Requirements Language

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. Use-Case: Wavelength Setup for IP over Optical Networks

Consider the network topology depicted in Figure 1. Nodes A and B are client IP routers that are connected to an optical WDM transport network. F, H and I represent WDM nodes. The transponder sits on the router and is directly connected to the add-drop port on a WDM node.

The optical signal originating on "Router A" is tuned to a particular wavelength. On "WDM-Node F", it gets multiplexed with optical signals at other wavelengths. Depending on the implementation of this multiplexing function, it may not be acceptable to have the router send signal into the optical network unless it is at the appropriate wavelength. In other words, having the router send signal with a wrong wavelength may adversely impact existing optical trails. If the clients do not have full visibility into the optical network, they are not in a position to pick the correct wavelength up-front.



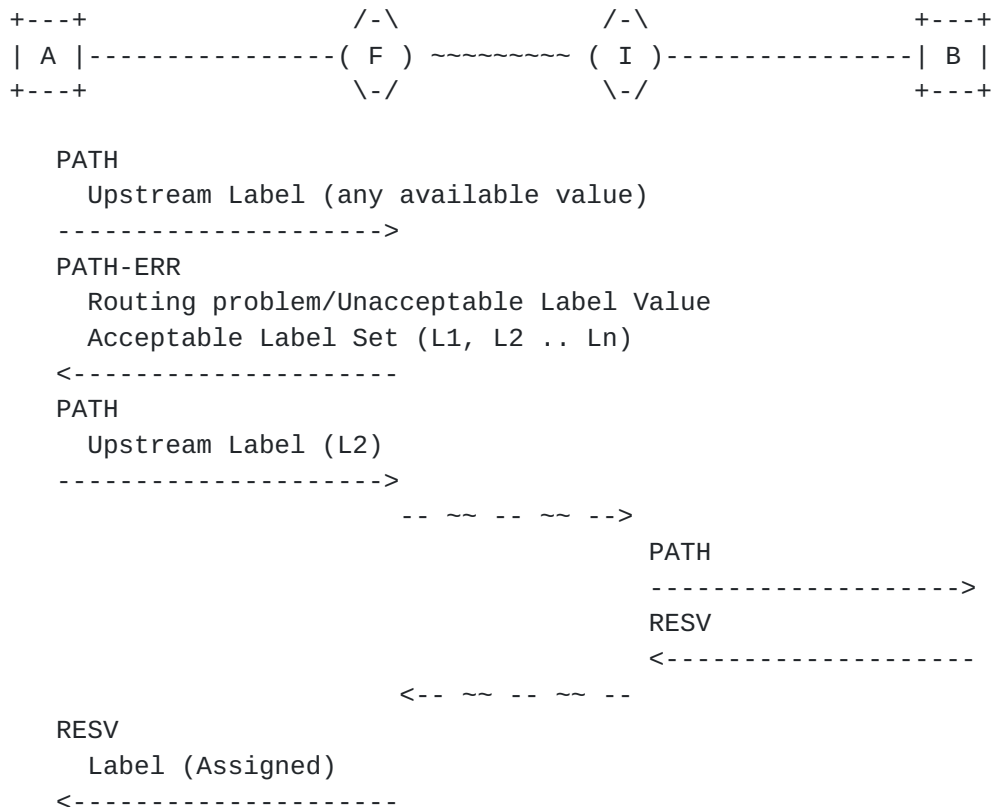
Sample Topology

Figure 1

3. The "Crankback Signaling" Approach

There are currently no GMPLS RSVP-TE protocol mechanisms that an upstream node can use for indicating that it does not know what upstream label to use and that it needs the downstream node to pick the label on its behalf.

The "Crankback Signaling" [RFC 4920](#) [[RFC4920](#)] approach can be applied to address the above use-case as shown in the following setup sequence:



Setup Sequence - Crank-back Approach

Figure 2

The above approach does work, but there are a few obvious concerns:

- o Since "Router-A" does not know which upstream label to use, it picks some random label and signals it without programming its data-plane (this is a deviation from the UPSTREAM_LABEL processing procedures outlined in [RFC 3473](#) [[RFC3473](#)]). As a result, the outgoing PATH message has no indication of whether the upstream label has been installed along the data-path or not.
- o Even if "Router-A" somehow correctly guesses an acceptable upstream label upfront, the network may end up finding a path which is suboptimal (there could be a different acceptable upstream label which corresponds to a better path in the network)

- o The "PATH-ERR with Acceptable Label Set" retry approach is usually used for exception handling. The above solution uses it for almost every single setup request (except in the rare scenario where the appropriate upstream label is guessed correctly).
- o There is an awkward window between the time the network sends out the PATH-ERR message (with the ACCEPTABLE_LABEL_SET) and receives the corresponding PATH message (with the selected UPSTREAM_LABEL); this window opens up the possibility of the selected UPSTREAM_LABEL to be stale by the time the network receives the retry PATH.
- o The above solution assumes the use of "symmetric labels" by default.

The rest of the sections in this draft present a solution proposal that is devoid of any of the above concerns.

4. Symmetric Labels

As per [RFC 3471](#) [[RFC3471](#)], the upstream label and the downstream label for an LSP at a given hop need not be the same. The use-case discussed in this document pertains to Lambda Switch Capable (LSC) LSPs and it is an undocumented fact that in practice, LSC LSPs always have symmetric labels at each hop along the path of the LSP.

The use of the protocol mechanism discussed in this document mandates "Label Symmetry". This mechanism is meant to be used only for bidirectional LSPs that assign symmetric labels at each hop along the path of the LSP.

5. Unassigned Upstream Label

This document proposes the use of a special label value - "0xFFFFFFFF" (for a 4-byte label) - to indicate an Unassigned Upstream Label. Similar "all-ones" patterns are expected to be used for labels of other sizes. The presence of this value in the UPSTREAM_LABEL object of a PATH message indicates that the upstream node has not assigned an upstream label on its own and has requested the downstream node to provide a label that it can use in both forward and reverse directions. The presence of this value in the UPSTREAM_LABEL object of a PATH message MUST also be interpreted by the receiving node as a request to mandate "symmetric labels" for the LSP.

5.1. Processing Rules

The Unassigned Upstream Label is used by an upstream node when it is not in a position to pick the upstream label on its own. In such a scenario, the upstream node sends a PATH message downstream with an Unassigned Upstream Label and requests the downstream node to provide a symmetric label. If the upstream node desires to make the downstream node aware of its limitations with respect to label selection, it MUST specify a list of valid labels via the LABEL_SET object as specified in [RFC 3473](#) [[RFC3473](#)].

In response, the downstream node picks an appropriate symmetric label and sends it via the LABEL object in the RESV message. The upstream node would then start using this symmetric label for both directions of the LSP. If the downstream node cannot pick the symmetric label, it MUST issue a PATH-ERR message with a "Routing Problem/Unacceptable Label Value" indication.

The upstream node will continue to signal the Unassigned Upstream Label in the PATH message even after it receives an appropriate symmetric label in the RESV message. This is done to make sure that the downstream node would pick a different symmetric label if and when it needs to change the label at a later point in time.




Figure 3

5.2. Backwards Compatibility

If the downstream node is running an older implementation and doesn't understand the semantics of an Unassigned UPSTREAM LABEL, it will either (a) reject the special label value and generate an error as

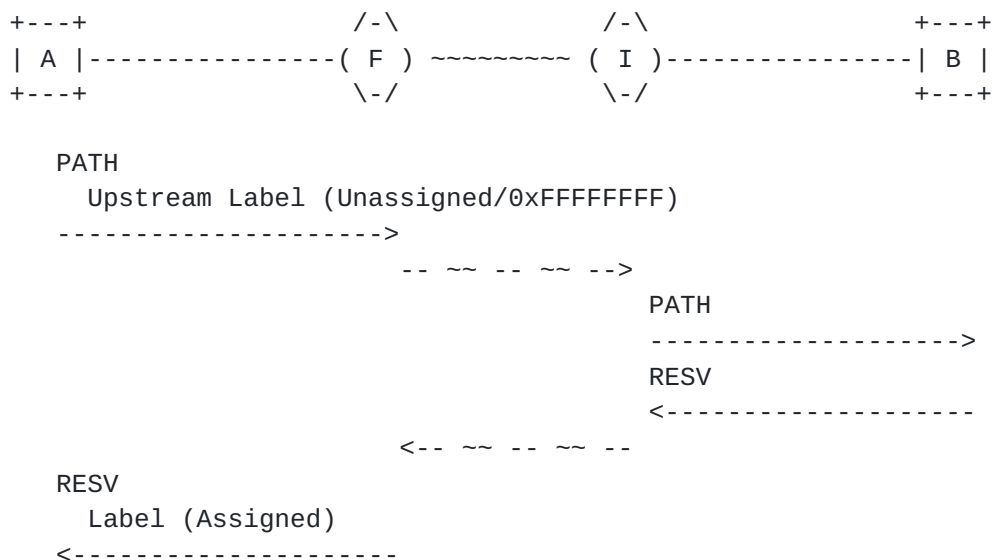
specified in [Section 3.1 of RFC 3473](#) [RFC3473] or (b) accept it and treat it as a valid label.

If the behavior that is exhibited is (a), then there are obviously no backwards compatibility concerns. If there is some existing implementation that exhibits the behavior in (b), then there could be some potential issues. However, at the time of publication, there is no documented evidence of any existing implementation that uses the "all-ones" bit pattern as a valid label. Thus, it is safe to assume that the behavior in (b) will never be exhibited.

6. Applicability

The use-case discussed in [Section 2](#) is revisited to examine how the mechanism proposed in this document allows the optical network to select and communicate the correct wavelength to its clients.

6.1. Initial Setup



Initial Setup Sequence

Figure 4

Steps:

- o "Router A" does not have enough information to pick an appropriate client wavelength. It sends a PATH message downstream requesting the network to assign an appropriate symmetric label for its use.

Since the client wavelength is unknown, the laser is off at the ingress client.

- o The downstream node (Node F) receives the PATH message, chooses the appropriate wavelength values and forwards them in appropriate label fields to the egress client ("Router B")
- o "Router B" receives the PATH message, turns the laser ON and tunes it to the appropriate wavelength (received in the UPSTREAM_LABEL/LABEL_SET of the PATH) and sends out a RESV message upstream.
- o The RESV message received by the ingress client carries a valid symmetric label in the LABEL object. "Router A" turns on the laser and tunes it to the wavelength specified in the network assigned symmetric LABEL.

For cases where the egress-node relies on RSVP signaling to determine exactly when to start using the LSP, this draft recommends integrating the above sequence with any of the existing graceful setup procedures:

- o "RESV-CONF" setup procedure (or)
- o 2-step "ADMIN STATUS" based setup procedure ("A" bit set in the first step; "A" bit cleared when the LSP is ready for use).

6.2. Wavelength Change

After the LSP is set up, the network MAY decide to change the wavelength for the given LSP. This could be for a variety of reasons - policy reasons, restoration within the core, preemption etc.

In such a scenario, if the ingress client receives a changed label via the LABEL object in a RESV modify, it MUST retune the laser at the ingress to the new wavelength. Similarly, if the egress client receives a changed label via UPSTREAM_LABEL/LABEL_SET in a PATH modify, it MUST retune the laser at the egress to the new wavelength. If the node receiving the changed label in a PATH/RESV message does not find the label acceptable, then the corresponding error procedures defined in [RFC 3473](#) [[RFC3473](#)] MUST be followed.

7. Acknowledgements

The authors would like to thank Adrian Farrel and Chris Bowers for their inputs.

8. Contributors

John Drake
Juniper Networks
Email: jdrake@juniper.net

Gert Grammel
Juniper Networks
Email: ggrammel@juniper.net

Pawel Brzozowski
ADVA Optical Networking
Email: pbrzozowski@advaoptical.com

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

9. IANA Considerations

This document makes no requests for IANA action.

10. Security Considerations

This document defines a special label value to be carried in the UPSTREAM_LABEL object of a PATH message. This special label value is used to enable the function of requesting network assignment of an upstream label. The changes proposed in this document pertain to the semantics of a specific field in an existing RSVP object and the corresponding procedures. Thus, there are no new security implications raised by this document and the security considerations put together by [RFC 3473](#) [[RFC3473](#)] still applies.

For a general discussion on MPLS and GMPLS related security issues, see the MPLS/GMPLS security framework [RFC 5920](#) [[RFC5920](#)].

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#), DOI 10.17487/RFC3471, January 2003, <<http://www.rfc-editor.org/info/rfc3471>>.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), DOI 10.17487/RFC3473, January 2003, <<http://www.rfc-editor.org/info/rfc3473>>.
- [RFC4920] Farrel, A., Ed., Satyanarayana, A., Iwata, A., Fujita, N., and G. Ash, "Crankback Signaling Extensions for MPLS and GMPLS RSVP-TE", [RFC 4920](#), DOI 10.17487/RFC4920, July 2007, <<http://www.rfc-editor.org/info/rfc4920>>.

11.2. Informative References

- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", [RFC 5920](#), DOI 10.17487/RFC5920, July 2010, <<http://www.rfc-editor.org/info/rfc5920>>.

Authors' Addresses

Xian Zhang (editor)
Huawei Technologies

Email: zhang.xian@huawei.com

Vishnu Pavan Beeram (editor)
Juniper Networks

Email: vbeeram@juniper.net

Igor Bryskin
Huawei Technologies

Email: igor.bryskin@huawei.com

Daniele Ceccarelli
Ericsson

Email: daniele.ceccarelli@ericsson.com

Oscar Gonzalez de Dios
Telefonica

Email: ogondio@tid.es