

CCAMP Working Group
Internet Draft
Intended status: Standards Track

Vishnu Pavan Beeram (Ed)
Juniper Networks
Igor Bryskin (Ed)
ADVA Optical Networking

Expires: September 05, 2015

March 05, 2015

Network Assigned Upstream-Label
draft-ietf-teas-network-assigned-upstream-label-01

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), 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 05, 2015.

Copyright Notice

Copyright (c) 2014 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

This document discusses a GMPLS RSVP-TE protocol mechanism that enables the network to assign an upstream-label for a given 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 network to pick an appropriate label.

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.....	2
2. Use-Case: Alien Wavelength Setup.....	3
3. The "crank-back" approach.....	3
4. Symmetric Labels.....	5
5. Unassigned Upstream Label.....	5
5.1. Processing Rules.....	5
5.2. Backwards Compatibility.....	6
6. Applicability.....	6
6.1. Initial Setup.....	7
6.2. Wavelength Change.....	8
7. Security Considerations.....	8
8. IANA Considerations.....	8
9. Normative References.....	8
10. Acknowledgments.....	8

[1. Introduction](#)

The GMPLS RSVP-TE extensions for setting up a Bidirectional LSP are discussed in [[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 LSP onto the network.

2. Use-Case: Alien Wavelength Setup

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.




Figure 1: Sample topology

3. The "crank-back" 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 following setup sequence is an attempt to address the above use-case using existing protocol mechanisms:




Figure 2: Setup Sequence - Crank-back Approach

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

- Since "Router-A" does not know which upstream-label to use, it picks some random label and signals it without programming its data-plane. As a result, the outgoing PATH message has no indication of whether the upstream-label has been installed along the data-path or not.
- If "Router-A" somehow correctly guesses (by sheer luck) 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)
- The "Path-Error 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).
- There is an awkward window between the time the network sends out the Path-Error (with the ACCEPTABLE_LABEL_SET) and receives the corresponding Path (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.

- The above solution assumes the use of "symmetric labels" by default.

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

4. Symmetric Labels

As per [[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" - to indicate an Unassigned Label. 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 can also be interpreted as a request to mandate "symmetric labels" for the LSP at the given hop.

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 has the option to specify a list of valid labels via the LABEL_SET object.

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 symmetric label if and when it needs to change the RESV label at a later point in time.




Figure 3: Unassigned UPSTREAM_LABEL

5.2. Backwards Compatibility

If the downstream-node is running an older implementation (which may be using the "crank-back" approach discussed in [Section 3](#)) and doesn't understand the semantics of an Unassigned UPSTREAM LABEL, it will either (a) reject the special label value and generate an error 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. Ingress implementations may even choose to adopt the "crank-back" approach in such cases. If there is some existing implementation that exhibits the behavior in (b), then there could be some potential issues. The use-case discussed in this draft pertains to LSC LSPs and it is safe to assume that the behavior in (b) will not be exhibited for such LSPs.

6. Applicability

Let us revisit the "alien wavelength" use-case discussed in [Section 2](#) and 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




Figure 4: Alien Wavelength - Initial Setup

Steps:

- "Router A" does not have enough information to pick an appropriate client wavelength. It sends a PATH downstream requesting the network to assign an appropriate symmetric label for it to use. Since the client wavelength is unknown, the laser is off at the ingress client.
- The network receives the PATH, chooses the appropriate wavelength values and forwards them in appropriate label fields to the egress client ("Router B")
- "Router B" receives the PATH, 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 upstream.
- The RESV 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:

- "RESV-CONF" setup procedure (or)
- 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 SHOULD 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 SHOULD retune the laser at the egress to the new wavelength.

7. Security Considerations

This document introduces no new security considerations.

8. IANA Considerations

This document makes no requests for IANA action.

9. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching Signaling Functional Description", [RFC 3471](#), January 2003
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching Signaling Resource Reservation Protocol-Traffic Engineering Extensions", [RFC 3473](#), January 2003.

10. Acknowledgments

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

Authors' Addresses

Vishnu Pavan Beeram
Juniper Networks
Email: vbeeram@juniper.net

John Drake
Juniper Networks
Email: jdrake@juniper.net

Gert Grammel
Juniper Networks
Email: ggrammel@juniper.net

Igor Bryskin
ADVA Optical Networking
Email: ibryskin@advaoptical.com

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

Daniele Ceccarelli
Ericsson
Email: daniele.ceccarelli@ericsson.com

Oscar Gonzalez de Dios
Telefonica
Email: ogondio@tid.es

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