

**Advertising MPLS labels in IGP
draft-gredler-rtgwg-igp-label-advertisement-01**

Abstract

Historically MPLS label distribution was driven by session oriented protocols. In order to obtain a particular routers label binding for a given destination FEC one needs to have first an established session with that node.

This document describes a mechanism to distribute FEC/label mappings through flooding protocols. Flooding protocols publish their objects for an unknown set of receivers, therefore one can efficiently scale label distribution for use cases where the receiver of label information is not directly connected.

Application of this technique are found in the field of backup (LFA) computation, Label switched path stitching, traffic engineering and egress link load balancing.

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](#)].

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 August 22, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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	3
2.	Motivation and Applicability	3
2.1.	Explicit One hop tunnels	4
2.2.	Egress ASBR Link Selection	5
2.3.	Explicit Path Routing through Label Stacking	6
2.4.	Stitching MPLS Label Switched Path Segments	7
3.	Acknowledgements	8
4.	IANA Considerations	8
5.	Security Considerations	8
6.	References	9
6.1.	Normative References	9
6.2.	Informative References	9
6.3.	References	9
	Author's Address	9

1. Introduction

MPLS label allocations are predominantly distributed by using the LDP [[RFC5036](#)], RSVP [[RFC5151](#)] or labeled BGP [[RFC3107](#)] protocol. All of those protocols have in common that they are session oriented, which means that in order to learn the Label Information database of a particular router one needs to have a direct control-plane session using the given protocol.

There are a couple of interesting use cases where the consumer of a MPLS label allocation may not be adjacent to the router having allocated the label. Bringing up an explicit session using existing label distribution protocols between the non-adjacent label allocator and the label consumer is the existing remedy for this dilemma.

For LDP protection routing LDP next next hop labels [[NNHOP](#)] have been proposed to provide the 2 hop neighborhood labels. While the 2 hop neighborhood provides good backup coverage for the typical network operator topology it is inadequate for some sparse for example ring like topologies.

Depending on the application, retrieval and setup of forwarding state of such >1 hop label allocations may only be transient. As such configuring and un-configuring the explicit session is an operational burden and therefore should be avoided.

2. Motivation and Applicability

It may not be immediate obvious, however introduction of Remote LFA [[I-D.ietf-rtgwg-remote-lfa](#)] technology has implied important changes for an IGP implementation. Previously the IGP had a one-way communication path with the LDP module. The IGP supplies tracking routes and LDP selects the best neighbor based upon FEC to tracking routes exact matching results. Remote LFA changes that relationship such that there is a bi-directional communication path between the IGP and LDP. Now the IGP needs to learn about if a label switched path to a given destination prefix has been established and what the ingress label for getting there is. The IGP needs to push that label for the tracking routes of destinations beyond a remote LFA neighbor.

Since the IGP now creates forwarding state based on label information it may make sense to distribute label by the IGP as well. This section lists example applications of IGP distribution of MPLS labels.

2.1. Explicit One hop tunnels

Deployment of Loop free alternate backup technology [RFC 5286](#) [[RFC5286](#)] results in backup graphs whose coverage is highly dependent on the underlying Layer-3 topology. Typical network deployments provide backup coverage less than 100 percent (see [RFC 6571 Section 4.3](#) for Results [[RFC6571](#)]) for IGP destination prefixes.

By closer examining the coverage gaps from the referenced production network topologies, it becomes obvious that most topologies lacking backup coverage are close to ring shaped topologies (Figure 1).

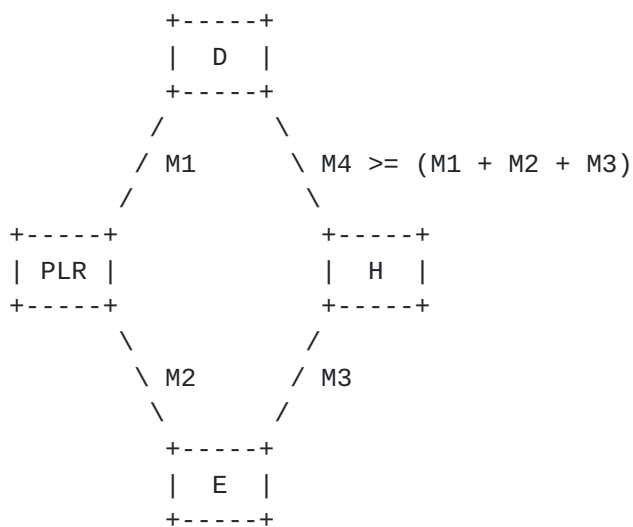


Figure 1: Coverage gap analysis

The protection router (PLR) evaluates for a primary path to destination 'D' if {E -> H -> D} is a viable backup path. Because the metric M4 {H -> D} is higher than the sum of the original primary path and the path from router 'H' to the PLR, this particular path would result in a loop and therefore is rejected.

Remote LFA [[I-D.ietf-rtgwg-remote-lfa](#)] has introduced the notion of "remote" LFA neighbor. A helper Router which is both in P and Q space could forward the traffic to the final destination. Router 'H' is in P space, however due to the actual metric allocation router 'H' is not in Q space.

Now consider that router 'H' would advertise a label for FEC 'D', which has the semantics that H will POP the label and forward to the destination node 'D'. This is done irrespective of the underlying IGP metric 'M4' it is a 'strict forwarding' label. The PLR router can now construct a label stack where the outermost label provides transport to router 'H'. The next label on the MPLS stack is the IGP

learned 'strict forwarding label' label. Note that the label 'strict forwarding' semantics are similar to a 1-hop ERO (Explicit route object). The Remote 'LFA' calculation would need to get changed, such that even if a node is not in PQ space, but rather in P space, it may get used as a backup neighbor if it advertises a strict forwarding label to the final destination. A recursive version of the algorithm is applicable as well as long as a node in P space has some non looping LSP path to the final destination. The PLR router can now program a backup path irrespective of the undesirable underlying layer-3 topology.

2.2. Egress ASBR Link Selection

In the ingress router 'S' is facing a dilemma (Figure 2). Router S receives a BGP route from all of its 4 upstream routers. Using existing mechanism the provider owning AS1 can control the loading of its direct links *to* its ASBR3 and ASBR4, however it cannot control the load of the links beyond the ASBRs, except manually tweaking the BGP import policy and filtering out a certain prefix. It would be more desirable to have visibility of all four BGP paths and be able to control the loading of those four paths using Weighted ECMP. Note that the computation of the 'Weight' percentage and the component doing this computation (Network embedded/SDN) is outside the scope of this document.

If all the ASes would be under one common administrative control then the network operator could deploy a forwarding hierarchy by using [\[RFC3107\]](#) to learn about the remote-AS BGP nexthop addresses and associated labels. An ingress router 'S' would then stack the transport label to its local egress ASBR and the remote ASBR supplied label. In reality it is hard to convince a peering AS to deploy another protocol just in order to easier control the egress load on the WAN links for the ingress AS.

A 'strict forwarding' paradigm would solve this problem: An Egress ASBR (e.g. ASBR 3 and 4) allocates a strict forwarding label toward all of its peering ASes and advertises it into its local IGP. The forwarding state of all those labels is to POP off the label and forward to the respective interface. The ingress router 'S' then builds a MPLS label stack by combining its local transport label to ASBR3 or ASBR4 with the IGP learned label pointing to the remote-AS ASBR.

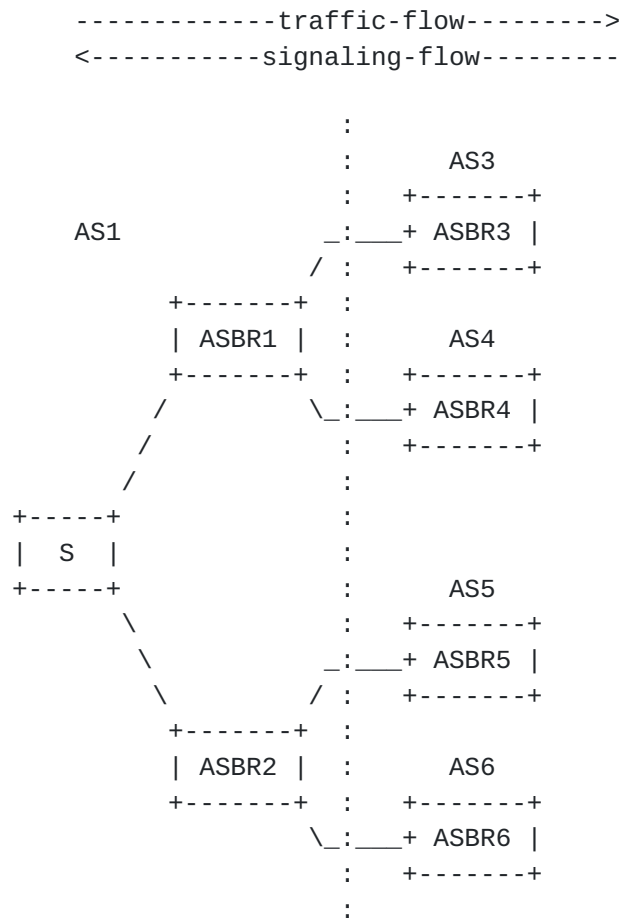


Figure 2: Egress ASBR Link selection

2.3. Explicit Path Routing through Label Stacking

IGP advertised strict forwarding labels can be utilized for constructing simple EROs via virtue of the MPLS label stack. In a classical traffic engineering problem (Figure 3) is illustrated. The best IGP path between {S,D} is {S, R3, R4, D}. Unfortunately this path is congested. It turns out that the links {S, R1}, {R1, R4} and {R2, R4} do have some spare capacity. In the past a C-SPF calculation would have passed the ERO {S, R1, R4, R2, D} down to RSVP for signaling. The conceptional problem with RSVP signaled paths is that they cannot be shared with other nodes in the network. Hence all potential ingress routers need to setup their "private" ERO path and allocate network signaling resources and forwarding state.

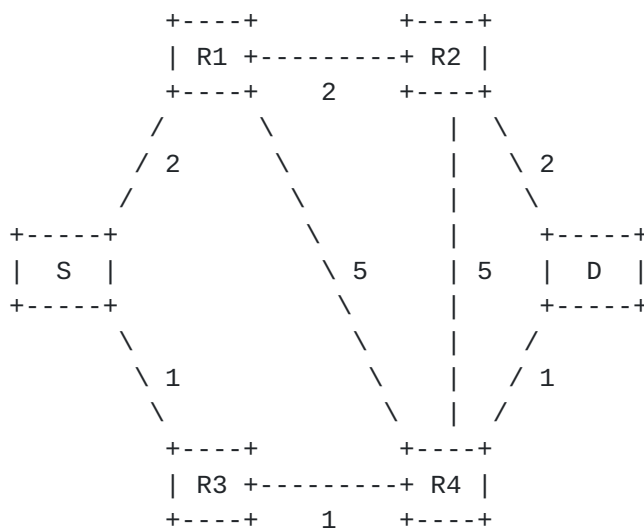


Figure 3: Explicit Routing using Label stacking

Consider now every router along the path does advertise a strict forwarding label for its direct neighbor. Router S could now construct a couple of paths for avoiding the hot links without explicitly signaling them.

- o {S, R1, R2, D}
- o {S, R1, R4, D}
- o {S, R1, R4, R2, D}

Note that not every hop in the ERO needs to be unique label in the label stack. This is undesired as existing forwarding hardware technology has got upper limits how much labels can get pushed on the label stack. In fact an existing tunnel (for example LDP tunnel {S, R1, R2}) can be reused for certain path segments.

2.4. Stitching MPLS Label Switched Path Segments

One of the shortcomings of existing traffic-engineering solutions is that existing label switched paths cannot get advertised and shared by many ingress routers in the network. In the example network (Figure 4) a LSP with an ERO of {R4, R2, R6} has been established in order to utilize two unused north / south links. The only way to attract traffic to that LSP is to advertise the LSP as a forwarding adjacency. This causes loss of the original path information which might be interesting for a potential router which might wants to use this LSP for backup purposes. A computing router would need to have all underlying fate-sharing and bandwidth utilization information.

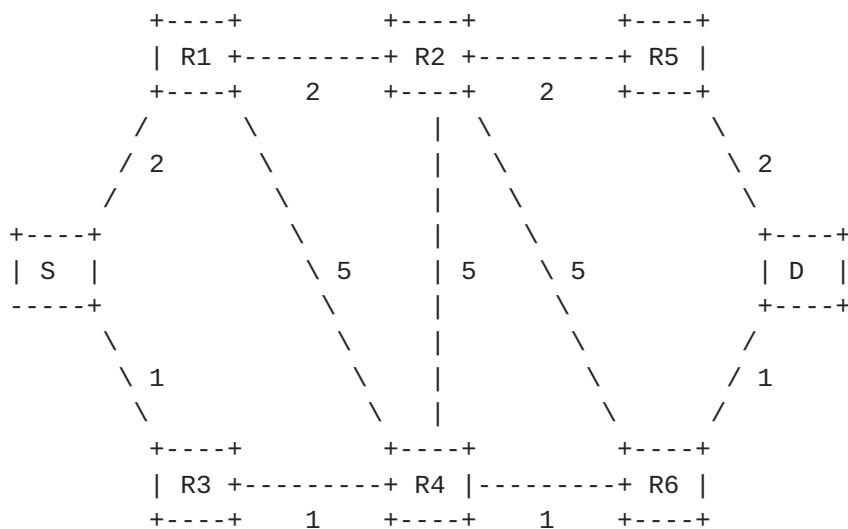


Figure 4: Advertising path segments

The IGP on R4 can now advertise the LSP segment by advertising its ingress label and optionally pass the original ERO, such that any upstream router can do their fate-sharing computations. Potential ingress routers now can use this LSP as a segment of the overall LSP. Furthermore ingress routers can combine label advertisements from different routers along the path. For example router S could stack its LDP path to R2 {S, R1, R2} plus the IGP learned RSVP LSP {R4, R5, R6} plus a strict forwarding label {R6, D}.

3. Acknowledgements

Many thanks to Yakov Rehkter, Stephane Likowski and Bruno Decraene for their useful comments.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

This document does not introduce any change in terms of IGP security. It simply proposes to flood existing information gathered from other protocols via the IGP.

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.
- [RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", [RFC 3107](#), May 2001.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.
- [RFC5151] Farrel, A., Ayyangar, A., and JP. Vasseur, "Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 5151](#), February 2008.
- [RFC5286] Atlas, A. and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates", [RFC 5286](#), September 2008.
- [RFC6571] Filsfils, C., Francois, P., Shand, M., Decraene, B., Uttaro, J., Leymann, N., and M. Horneffer, "Loop-Free Alternate (LFA) Applicability in Service Provider (SP) Networks", [RFC 6571](#), June 2012.

6.2. Informative References

- [I-D.ietf-rtgwg-remote-lfa]
Bryant, S., Filsfils, C., Previdi, S., Shand, M., and S. Ning, "Remote LFA FRR", [draft-ietf-rtgwg-remote-lfa-01](#) (work in progress), December 2012.

6.3. References

- [NNHOP] Chen, H., Shen, H., and H. Tian, "Discovering LDP Next-Nexthop Labels", November 2005, <<http://tools.ietf.org/html/draft-shen-mpls-ldp-nnhop-label-02>>.

Author's Address

Hannes Gredler
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: hannes@juniper.net

