

Carrying Label Information in BGP-4

[draft-ietf-mpls-bgp4-mpls-02.txt](#)

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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>.

2. Abstract

When a pair of Label Switch Routers (LSRs) that maintain BGP peering with each other exchange routes, the LSRs also need to exchange label mapping information for these routes. The exchange is accomplished by piggybacking the label mapping information for a route in the same BGP Update message that is used to exchange the route. This document specifies the way in which this is done.

3. Overview

The Multiprotocol Label Switching (MPLS) architecture [MPLS-ARCH] identifies situations in which the mapping between a label and a route must be distributed between BGP peers. This document specifies how this label mapping information is distributed. The label mapping information is included in the BGP Update message that is used to distribute the route. This is done by utilizing BGP-4 Multiprotocol Extensions attribute [[BGP-MP](#)].

4. Carrying Label Mapping Information

Label mapping information is carried as part of the Network Layer Reachability Information (NLRI) in the Multiprotocol Extensions attributes. Such NLRI is identified by using SAFI TBD.

The Network Layer Reachability information is encoded as one or more triples of the form <label, length, prefix>, whose fields are described below:

```
+-----+
| Length (1 octet) |
+-----+
| Label (3 octets) |
+-----+
.....
+-----+
| Prefix (variable) |
+-----+
```

The use and the meaning of these fields are as follows:

a) Length:

The Length field indicates the length in bits of the address prefix plus the label(s).

b) Label:

The Label field carries one or more labels (that corresponds to the stack of labels [MPLS-ENCAPS]). Each label is encoded as 3 octets, where the high-order bit contains "Bottom of Stack" (as defined in [MPLS-ENCAPS]). The following high-order three bits

must be zero. The remaining 20 bits contain the label value.

c) Prefix:

The Prefix field contains address prefixes followed by enough trailing bits to make the end of the field fall on an octet boundary. Note that the value of trailing bits is irrelevant.

The label(s) specified for a particular route (and associated with it address prefix) must be assigned by the LSR which is identified by the value of the Next Hop attribute of the route.

When a BGP speaker redistributes a route, the label(s) assigned to that route must not be changed (except by omission), unless the speaker changes the value of the Next Hop attribute of the route.

A BGP speaker can withdraw a previously advertised route (as well as the binding between this route and a label) by either (a) advertising a new route (and a label) with the same NLRI as the previously advertised route, or (b) listing the NLRI of the previously advertised route in the Withdrawn Routes field of an Update message. The label information carried (as part of NLRI) in the Withdrawn Routes field should be set to 0x800000.

5. Advertising Multiple Routes to a Destination

A BGP speaker may maintain (and advertise to its peers) more than one route to a given destination, as long as each such route has its own label(s).

The encoding described above allows a single BGP Update message to carry multiple routes, each with its own label(s).

In the case where a BGP speaker advertises multiple routes to a destination, if a route is withdrawn, and a label(s) is specified at the time of withdrawal, only the corresponding route with the corresponding label is withdrawn. If a route is withdrawn, and no label is specified at the time of withdrawal, then only the corresponding unlabeled route is withdrawn; the labeled routes are left in place.

6. Capability Negotiation

A BGP speaker that uses Multiprotocol Extensions to carry label mapping information should use the Capabilities Optional Parameter, as defined in [BGP-CAP], to inform its peers about this capability. The MP_EXT Capability Code, as defined in [BGP-MP], is used to negotiate the (AFI, SAFI) pairs available on a particular connection.

A BGP speaker should not advertise this capability to another BGP speaker unless there is a Label Switched Path (LSP) between the two speakers.

A BGP speaker that is capable of handling multiple routes to a destination (as described above) should use the Capabilities Optional Parameter, as defined in [BGP-CAP], to inform its peers about this capability. The value of this capability is TBD.

7. Security Considerations

Security issues are not discussed in this document.

8. Acknowledgements

To be supplied.

9. References

[BGP-4] "A Border Gateway Protocol 4 (BGP-4), Y. Rekhter, T. Li ([RFC 1771](#)) <http://ds.internic.net/rfc/rfc1771.txt>

[BGP-CAP] "Capabilities Negotiation with BGP-4", R. Chandra, et al, Work in progress, <http://www.internic.net/internet-drafts/draft-ietf-idr-bgp4-cap-neg-03.txt>

[BGP-MP] "Multiprotocol Extensions for BGP-4", T. Bates, et al ([RFC 2283](#)) <http://ds.internic.net/rfc/rfc2283.txt>

[MPLS-ARCH], "A Proposed Architecture for MPLS", E. Rosen, et al, Work in progress, <http://www.internic.net/internet-drafts/draft-ietf-mpls-arch-00.txt>

[MPLS-ENCAPS], "MPLS Label Stack Encoding", E. Rosen, et al, Work in progress, <http://www.internic.net/internet-drafts/draft-ietf-mpls-label-encaps-00.txt>

10. Author Information

Yakov Rekhter
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
email: yakov@cisco.com

Eric Rosen
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824
email: erosen@cisco.com

