

IDR Working Group
INTERNET-DRAFT
Intended status: Experimental
Expires: May 4, 2020

Lo
Juniper
Nov

Color Operation with BGP Label Unicast
draft-chan-idr-bgp-lu2-00.txt

Abstract

This document specifies how to carry colored path advertisement via to the existing protocol BGP Label Unicast. It would allow backward with [RFC8277](#).

The targeted solution is to use stack of labels advertised via BGP L 2.0 for end to end traffic steering across multiple IGP domains. The similar to Segment Routing.

This proposed protocol will convey the necessary reachability information ingress PE node to construct an end to end path

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions 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 May 4, 2020.

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

to IETF Documents (<http://trustee.ietf.org/license-info>) in effect of publication of this document. Please review these documents carefully to describe your rights and restrictions with respect to this document. Components extracted from this document must include Simplified BSD License as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Chan

Expires May 4, 2020

Internet-Draft

[draft-chan-idr-bgp-lu2-00.txt](#)

November 2019

Table of Contents

1.	Introduction.....	
2.	Conventions used in this document.....	
3.	Carrying Label Mapping Information with Color and Label Stack....	
3.1.	Color extended community for BGP Labeled Unicast.....	
3.2.	Color extended community for service prefixes.....	
4.	Uniqueness of path entries.....	
5.	AIGP consideration.....	
6.	Explicit Withdraw of a <color, prefix>.....	
7.	Error Handling Procedure.....	
8.	Security Considerations.....	
9.	IANA Considerations.....	
10.	References.....	
10.1.	Normative References.....	
10.2.	Informative References.....	
11.	Acknowledgments.....	

1. Introduction

The proposed protocol is aimed to solve interdomain traffic steering for different transport services in mind. One application is low latency path across multiple IGP domains, which could scale up to 100k routers network.

BGP is a flexible protocol. With additional of color attribute to BGP Labeled Unicast, a path with specific color would be given a meaning in application. For example, a latency path, a fully protected path, or a path for diversity.

The stack of labels would mean an end to end path across domains through ABR or ASBR. Each ABR or ASBR will take one label from the stack, and he

forwarding path to next ABR, ASBR, or the final destination.

And the label in the stack may be derived from any of the below

- Prefix SID
- Binding SID for RSVP LSP
- Binding SID for SR-TE LSP
- Local assigned label

The enhancement to the original [RFC8277](#) is to add color extended com multiple advertisement allowed. The result is similar to multi-topol different colors.

A new [[BGP-CAP](#)] should be required to enable such slicing.

On the other hand, to enable the service prefixes to be mapped accor L3VPN, L2VPN, EVPN and prefix with BGP signaling, the color extended also added there. In the PE node, the service prefixes with color wi to a transport tunnel with the same color.

Chan

Expires May 4, 2020

Internet-Draft

[draft-chan-idr-bgp-lu2-00.txt](#)

Novemb

The following is an example

PE1----ABR1-----ABR2-----PE2

PE1 will send the following labels with a color 100 path

[2001 13001 801 16], where

2001 - SR label to reach ABR1

13001 - Binding-SID label to reach ABR2. Underlying tunnel type is R

801 - Binding-SID label to reach PE2. Underlying tunnel type is SR-T

16 - a VPN label

If PE1 wants to reach PE2 with another colored path, say color 200, could be different.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",

"SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document interpreted as described in [RFC 2119](#) [[RFC2119](#)].

In this document, these words will appear with that interpretation ALL CAPS. Lower case uses of these words are not to be interpreted with the significance described in [RFC 2119](#).

3. Carrying Label Mapping Information with Color and Label Stack

3.1. Color extended community for BGP Labeled Unicast

The addition of Color Extended Community is an opaque extended community defined in [RFC4360](#) and [RFC5512](#). The draft allows multiple color values advertisement.

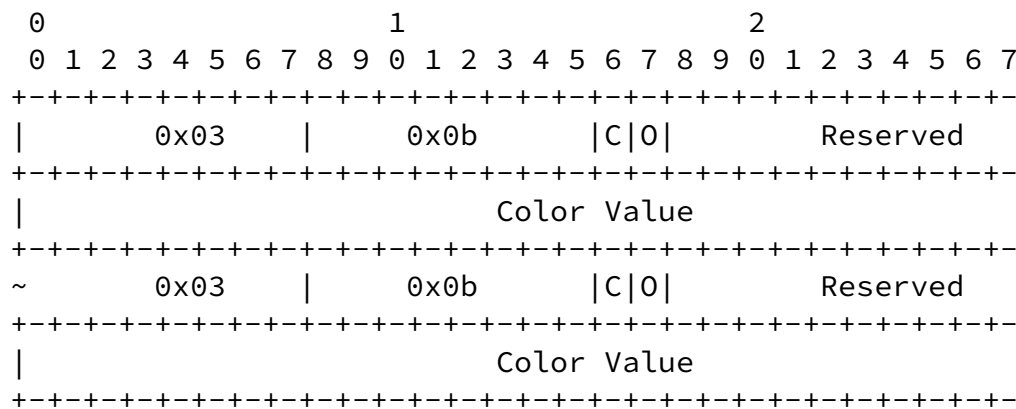


Figure 1: Color value advertisement format
Expires May 4, 2020

Chan

Internet-Draft

[draft-cha-idr-bgp-lu2-00.txt](#)

Novemb

Both in BGP update and MP_UNREACH_NLRI message, multiple color extended communities could be included. It means that multiple colors, indicating different services, could share the same label stack.

If only one color extended community is specified, only prefix with that color value is updated or withdrawn.

If a MP_UNREACH_NLRI message without any color specified is received for a prefix, that prefix with any color should not be affected.

If color extended community is not present in a BGP update message, it is treated as normal BGP-LU without any color.

3 bits of XXX is reserved here for the draft.

Color value 0 is reserved for future interoperability purpose.

3.2. Color extended community for service prefixes

The same format of color extended community is advertised with service prefixes.
The order of the color extended community could be interpreted as

- Order of primary and fallback colors
- Or, ECMP of equal split between color tunnels

The above would be interpreted by the receiving PE upon its local color

4. Uniqueness of path entries

Use of color can be considered to slice into multiple BGP Label Unicast NLRI.
Therefore, it should be treated as unique entries for the <color, prefix>

e.g. <color, prefix>, [labels]

<1, 10.1.1.1/32>, [100 200]

<2, 10.1.1.1/32>, [100 200]

<null, 10.1.1.1/32>, [100 200]

All these 3 NLRI are considered different but valid entries for different instances.

5. AIGP consideration

AIGP ([RFC7311](#)) would be also used in here to embed certain metric across

6. Explicit Withdraw of a <color, prefix>

According to [RFC8277](#), MP_UNREACH_NLRI can be used to remove binding of <color, prefix>.

Chan

Expires May 4, 2020

Internet-Draft

[draft-chan-idr-bgp-lu2-00.txt](#)

November

Compatibility is set to 0xC00000 to specify the use of color. Multiple color extended communities could be applied here.

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Length   |             Compatibility             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```




Figure 2: NLRI for Withdrawal

7. Error Handling Procedure

If BGP receiver could not handle the NLRI, it should silently discard logging.

8. Security Considerations

9. IANA Considerations

TBD. It will require a new BGP capability code to enable such color
New SAFI might be required as well.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirements", [BCP 14](#), [RFC 2119](#), March 1997.

10.2. Informative References

[RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP", [RFC 3107](#), DOI 10.17487/RFC3107, May 2001,

<<https://www.rfc-editor.org/info/rfc3107>>.

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), February 2006

<<https://www.rfc-editor.org/info/rfc4360>>.

Chan

Expires May 4, 2020

[P

Internet-Draft

[draft-chan-idr-bgp-lu2-00.txt](#)

Novemb

[RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation Subsequ

Family Identifier (SAFI) and the BGP Tunnel Encapsulation
[5512](#), April 2009.

<<https://www.rfc-editor.org/info/rfc5512>>.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J.,
McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](#),
DOI 10.17487/RFC5575, August 2009,

<<http://www.rfc-editor.org/info/rfc5575>>.

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Adverti
Multiple Paths in BGP", [RFC 7911](#), DOI 10.17487/RFC7911, Ju

<<https://www.rfc-editor.org/info/rfc7911>>.

[RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefi
DOI 10.17487/RFC8277, October 2017,

<<https://www.rfc-editor.org/info/rfc8277>>.

[BGP-CAP] Chandra, R. and J. Scudder, "Capabilities Advertisement
with BGP-4", [RFC 2842](#), May 2000.

11. Acknowledgments

The following people have contributed to this document:

Jeff Haas, Juniper Networks

Shraddha Hedge, Juniper Networks

Santosh Kolenchery, Juniper Networks

Shihari Sangli, Juniper Networks

Krzysztof Szarkowicz, Juniper Networks

Yimin Shen, Juniper Networks

Chan

Expires May 4, 2020

Internet-Draft

[draft-chan-idr-bgp-lu2-00.txt](#)

Novemb

Authors Addresses

Louis Chan (editor)

Juniper Networks

2604, Cityplaza One, 1111 King's Road

Taikoo Shing

Hong Kong

Phone: +85225876659

Email: louisc@juniper.net

Chan

Expires May 4, 2020

[Pa