

IDR
Internet-Draft
Intended status: Standards Track
Expires: 8 September 2022

R. Chen
Ch. Dai
Sh. Peng
ZTE Corporation
7 March 2022

Inter-domain Network Slicing via BGP-LU
draft-zhou-idr-inter-domain-lcu-04

Abstract

This document aims to solve inter-domain network slicing problems using existing technologies. It attempts to establish multiple BGP-LU LSPs of different colors for a/multiple prefix to stitch multiple network segments.

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 <https://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 8 September 2022.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. Color	3
3. Advertising multiple paths	3
4. Assigning Label(s)	4
5. Inter-domain Network Slicing via BGP-LU	4
5.1. Colored BGP-LU Capability Advertisement	4
5.2. Colored BGP-LU realized	5
6. SRv6 support	7
7. Deploy Considerations	7
8. Acknowledgements	8
9. IANA Considerations	8
10. Security Considerations	8
11. Normative References	8
Authors' Addresses	9

1. Introduction

In the traditional end to end inter-domain network slicing, BGP-LU is used to build inter-domain MPLS LSP, and overlay service will be directly over BGP-LU LSP. For an E2E BGP-LU LSP, if overlay service has TE requirements that defined by a color, the BGP-LU LSP need also have a sense of color, i.e., BGP-LU label could be allocated per color.

[RFC8277] specifies a set of procedures for using BGP to advertise that a specified router has bound a specified MPLS label to a specified address prefix. It's an effective way for inter-domain labels, but it does not have the ability to select the underlying network resources.

This document describes the colored BGP-LU LSP, which contains two options:

- * One is to define the multiple paths for the same destination prefix and advertise in BGP UPDATE message, and each UPDATE message can contain the color Extended Community [RFC9012] with different color value, which helps to select the underlying resources. This mode require additional path function defined in [RFC7911].
- * The other is that multiple prefixes and multiple colors are configured on PE. One prefixes corresponds to one color. This mode does not require to additional path function.

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. Color

[RFC9012] introduces the concept of color, which is used as one of the KEY of SR policy [I-D.ietf-spring-segment-routing-policy]. The color of SR policy defines a TE purpose, which includes a set of constraints such as bandwidth, delay, TE metric, etc.

TO help routing decisions , each UPDATE may contain a Color Extended Community with a specific color value, the Color Sub-TLV is only an opaque extended community.

3. Advertising multiple paths

A BGP speaker can advertise multiple paths for a particular address prefix by a Path identifier in the Extended NLRI Encoding as defined in [RFC7911].




Figure 1

The Path Identifier only identifies a path, not carrying any particular semantics. In this document, it can be generated by the <Prefix,Color> tuple. The assignment to the Path Identifier for a path by a BGP speaker is purely a local matter.

Therefore, if a BGP speaker has two colors for the prefix P, which correspond to two different paths, it may advertise two UPDATE NLRIs, <prefix, pathid1> with color1 extended community and <prefix, pathid2> with color2 extended community. Pathid1 and pathid2 in two UPDATE NLRIs MUST be different.

Note that in this document, BGP speakers acting as border routers that interact with external neighbors need to support advertising multiple paths corresponding to the same prefix. Although multiple

paths have different path ids, they have the same next hop. As for the procedures of mutual backup paths with the same prefix and the different next hops, refer to [RFC7911].

4. Assigning Label(s)

[RFC8277] describes how to use BGP to bind MPLS label(s) to the address prefixes. The specific format of the UPDATE message is detailed in Section 2 of [RFC8277].

[RFC8277] Section 3.2 details the process of modifying the Label field during propagation. When propagating a SAFI-4 or SAFI-128 route, if the Network Address of Next Hop field has never changed, the label field must remain unchanged. Otherwise, if the Network Address of Next Hop field is changed, the label field(s) of the propagating route must contain the label(s) that is (are) bound to the prefix at the new next hop. What the label changes to depends on the local policy. However, LSPs with different color paths need to have different label(s).

5. Inter-domain Network Slicing via BGP-LU

5.1. Colored BGP-LU Capability Advertisement

A BGP speaker that uses Colored BGP-LU Extensions SHOULD use the Capability Advertisement procedures [RFC3392] to determine whether the speaker could use Colored BGP-LU Extensions with a particular peer.

The fields in the Capabilities Optional Parameter are set as follows:

- * The Capability Code field TBD1 (which indicates Colored BGP-LU Extensions capabilities).
- * The Capability Length field is set to 4.
- * The Capability Value field is defined as:

```

+-----+
| Address Family Identifier (2 octets) |
+-----+
| Subsequent Address Family Identifier (1 octet) |
+-----+
| reserve (1 octet) |
+-----+

```

Figure 2

where:

AFI-Address Family Identifier (16 bit), The values is 1 "IPv4" or 2 "IPv6".

SAFI-Subsequent Address Family Identifier (8 bit), The values is 1 "Unicast" or 4 (BGP LU).


Res.-Reserved (8 bit) field. SHOULD be set to 0 by the sender and ignored by the receiver.

Note that not setting the field value to 0 may create issues for a receiver not ignoring the field. In addition, this definition is problematic if it is ever attempted to redefine the field.

5.2. Colored BGP-LU realized

[RFC7911] defined that multiple paths for a particular address prefix by a Path identifier can be advertised in BGP UPDATE message, and each UPDATE message can contain the Color Extended Community [RFC9012] with different color value. That is a simple existing way to realize BGP-LU color function, and only an extension of Colored BGP-LU capability advertisement is required.

Consider the following example of establishing multiple BGP-LU LSPs per different colors in a cross-domain scenario.



Label Exchange Tables:

ASBR1:			ASBR2:		
inLabel	outLabel	nextHop	inLabel	outLabel	nextHop
201	200	SR-TE1	201	201	sub-if1
202	200	SR-TE2	202	202	sub-if2
PE2:					
prefix	color	outLabel	nextHop		
1.1.1.1	1	201	SR-TE1		
1.1.1.1	2	202	SR-TE2		

Figure 3

In figure 1, PE1 advertises two paths: <1.1.1.1, path-id1> carries the color1 attribute and <1.1.1.1, path-id2> carries the color2 attribute to ASBR1. PE1 advertises the binding between the prefix 1.1.1.1 and label 200. Because of the end node, both paths have the same label value 200.

ASBR1 receives these two paths from PE1, and when sending to ASBR2, it modifies the next hop to itself. And allocate two new labels based on <prefix, path-id, color>. As shown in Figure 1, ASBR1 sends two paths to ASBR2, <1.1.1.1, path-id1> carries color1+label201, and <1.1.1.1, path-id2> carries color2+label202.

Similarly, ASBR2 also generates two different labels based on the <prefix, path-id, color>. As shown in Figure 1, multiple end to end BGP-LU LSPs are established. Different BGP-LU LSPs select the underlay SR-BE/TE tunnels according to their colors.

6. SRv6 support

Colored BGP-LU can be also used to setup end-to-end color-aware connectivity using Segment Routing over IPv6 (SRv6) [RFC8402].

As defined in [I-D.ietf-bess-srv6-services], to provide SRv6 service with underlay SRv6 policy connectivity, the egress PE signals the BGP overlay service route with SRv6 Service SID and color extended community . The ingress PE encapsulates the payload in an outer IPv6 header which contains the underlay SRv6 policy segment list and the overlay Service SID.

In addition, another solution is to provide SRv6 service with underlay SRv6 best-effort connectivity that is created by global IPv6 (AFI/SAFI 2/1) with color extended community. The underlay SRv6 SID is allocated based on <global IPv6, path-id, color>. The ingress PE encapsulates the payload in an outer IPv6 header which contains the underlay SRv6 SID and the Service SID.

7. Deploy Considerations

All BGP routers (PE1--ASBR1, ASBR1---ASBR2, ASBR2---PE2) SHOULD be Colored BGP-LU neighbors in advance. There may be multiple border routers to ensure multipath backup. All routers require the Colored BGP-LU Capability Advertisement. If transit network domains that do not support Colored BGP-LU are processed as follows;

- * When the Colored BGP-LU neighbor receives the BGP-LU routes, if it continues to advertise the BGP-LU routes to the upstream neighbor that supports the Colored BGP-LU, the BGP-LU routes shouldn't be changed to the Colored BGP-LU routes.
- * When receiving the Colored BGP-LU advertisement from the neighbor that supports Colored BGP-LU, if the advertisement continues to be advertised to the upstream neighbor that does not support Colored BGP-LU, the advertisement should be changed to BGP-LU advertisement, that is, advertise one out of multiple path.

This document not only supports interprovider VPNs while the customer sites belong to different ASs, but also supports the Carrier-of-Carriers VPNs while the customer site belong to the same AS. Multiple operators are involved, so AS border routers may involve color mapping, color namespaces, or color service chains. These services can be delivered by the controller configurations or the local configurations.

8. Acknowledgements

TBD.

9. IANA Considerations

TBD.

10. Security Considerations

TBD.

11. Normative References

[I-D.ietf-bess-srv6-services]

Dawra, G., Filsfils, C., Talaulikar, K., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "SRv6 BGP based Overlay Services", Work in Progress, Internet-Draft, draft-ietf-bess-srv6-services-12, 5 March 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-srv6-services-12>>.

[I-D.ietf-spring-segment-routing-policy]

Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-19, 5 March 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-routing-policy-19>>.

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

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.

[RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder,
"The BGP Tunnel Encapsulation Attribute", RFC 9012,
DOI 10.17487/RFC9012, April 2021,
<<https://www.rfc-editor.org/info/rfc9012>>.

Authors' Addresses

Ran Chen
ZTE Corporation
Nanjing
China
Email: chen.ran@zte.com.cn

Chunning Dai
ZTE Corporation
Nanjing
China
Email: dai.chunning@zte.com.cn

Shaofu Peng
ZTE Corporation
Nanjing
China
Email: peng.shaofu@zte.com.cn