IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 3, 2020

J. Zhou
Chunning. Dai
Shaofu. Peng
ZTE Corp.

Dec 31, 2019

Inter-domain Network Slicing via BGP-LU
draft-zhou-idr-inter-domain-lcu-00

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 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 July 3, 2020.

Table of Contents

1. Introduction

As described in [I-D.peng-teas-network-slicing], in the traditional end to end inter-domain network slicing, BGP-LU is used as a distributed slicing scheme. [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, in which the routing prefix not only carries a label(or a sequence of labels), but also carries an unique color attribute which helps to select the underlying logic slice.

[RFC7911] defines a BGP extension that allows the advertisement of multiple paths of the same address prefix. It can help with optimal routing and in a network by providing potential alternate or backup paths. In this document, it is not used for the link backup. It is only used to advertise multiple paths, whether intra-domain or inter-domain.

2. Requirements notation

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 [RFC2119].

3. Color

[I-D.ietf-idr-tunnel-encaps] 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.

Color is used a policy keyword to help routing decisions and can be carried through Community attribute [I-D.ietf-idr-tunnel-encaps] in BGP. Each UPDATE SHOULD contain a Color Extended Community with a specific color value, the Color Sub-TLV is only an opaque extended community.

4. Advertising multiple paths

Consider the following LSR topology in Figure 1: PE1---ASBR1---ASBR2 ---PE2. Packet transfers from PE2 to PE1. The overlay service of PE1 to PE2 has two colors, color1 and color2.

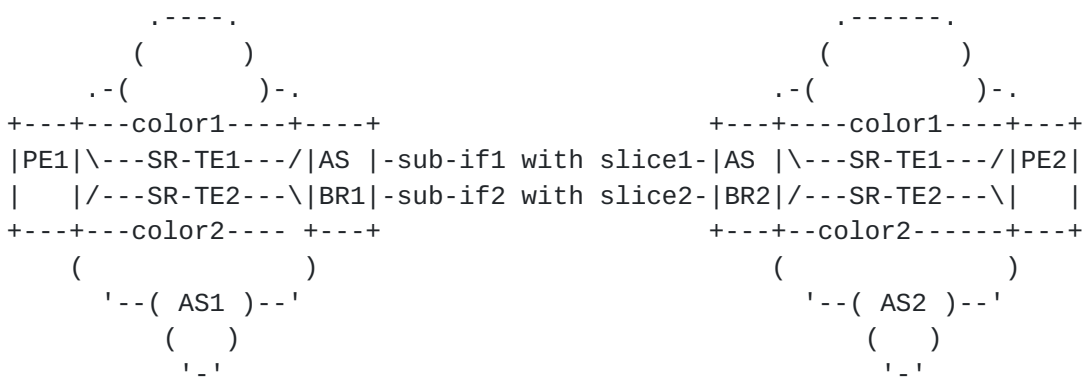

Figure 1: Network Slicing Example

According to the extended NLRI Encodings in [RFC7911], in order to advertise multiple paths for the same address prefix, PE1 MUST encode two different NLRIs to advertise the two paths. Every NLRI has a unique Path Identifier.

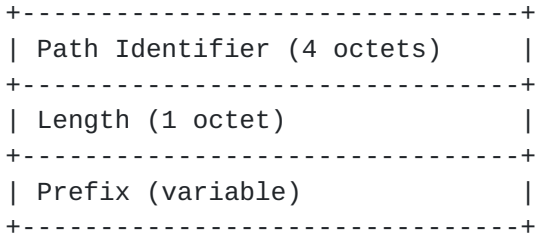

The Path Identifier only identifies a path, not carrying any particular semantics. It MAY be generated from the two-tuple <Prefix,Color>. Also, it MAY be generated from the information containing Prefix and Color.

Finally, There are two BGP UPDATE NLRIs from PE1 to ASBR1, <prefix, pathid1> with color1 extended community and <prefix, pathid2> with color2 extended community. Pathid1 and pathid2 in two UPDATE NLRI MUST be different. As described in [RFC8277], they have the same label value, however, different label value is also possible.

5. Generating Labels by Path

To realize the inter-domain multiple paths, the border routers of per domain SHOULD consider the diversity of paths when generating the labels.

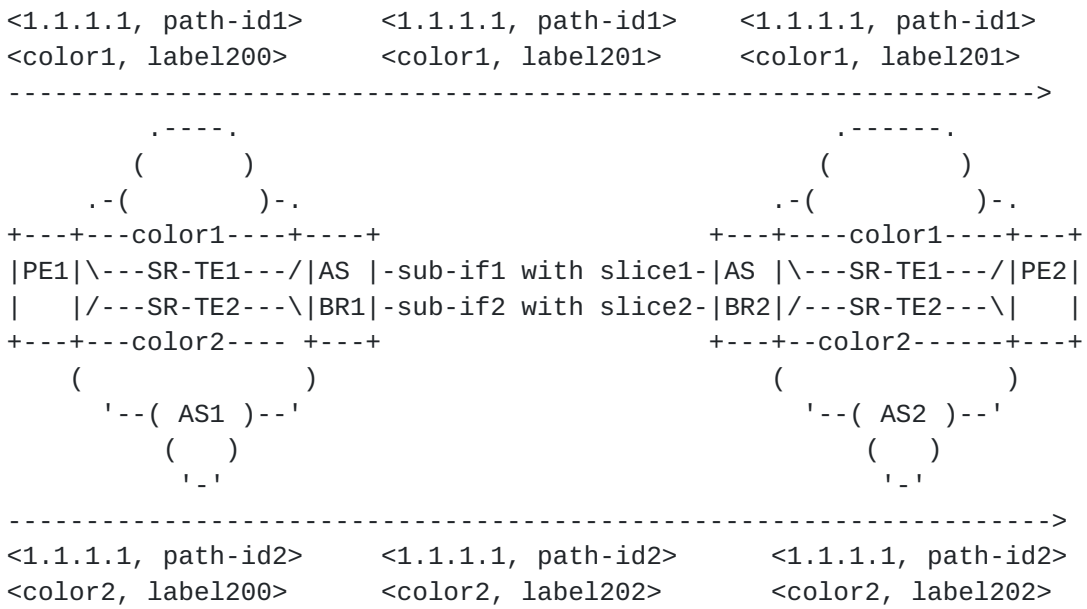

Label Exchange Tables:

ASBR1:

inLabel	outLabel	nextHop
201	200	SR-TE1
202	200	SR-TE2

ASBR2:

inLabel	outLabel	nextHop
201	201	sub-if1
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 2: Details of Network Slicing Example

Suppose ASBR1 receives two paths from PE1, with next hop 1.1.1.1, path identifier 1 and 2, ASBR1 MUST modify the next hop to itself, and generate two new labels based on two paths. As depicted in Figure 2, <1.1.1.1, color1> generates label 201, <1.1.1.1, color2>

generates label 202. Both of them are the local label exchange table items. Similarly, ASBR2 also generates two different labels based on the two paths. As shown in Figure 2, multiple end to end BGP-LU LSP are established.

So, at the entry node of per domain, BGP-LU LSP generated for specific color is over intra-domain SR-TE or SR Best-effort path generated for the color again. At exit node of per domain, BGP-LU LSP generated for specific color selects inter-domain forwarding resource per color.

6. Deploy Considerations

- The premise of this document is that all border routers SHOULD have the same understanding of a color value. For example, PE1 understands color1 representing a low packet loss rate, and ASBR1 must think so and apply it to the local forwarding policy.

- All BGP routers(PE1--ASBR1, ASBR1---ASBR2, ASBR2---PE2) SHOULD be BGP-LU neighbors in advance.

- All routers require the ADD-PATH Capability which is described in chapter 4 of [RFC7911].

- If any Router Reflector existed in the network, it SHOULD support this function. Later will discuss Confederation.

7. Security Considerations

TBD

8. References

8.1. Normative References

[I-D.ietf-idr-tunnel-encaps]
Patel, K., Velde, G., and S. Ramachandra, "The BGP Tunnel Encapsulation Attribute", draft-ietf-idr-tunnel-encaps-15 (work in progress), December 2019.

[I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Sivabalan, S., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy-06 (work in progress), December 2019.

[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>.

8.2. Informative References

[I-D.peng-teas-network-slicing] Peng, S., Chen, R., Mirsky, G., and F. Qin, "Packet Network Slicing using Segment Routing", draft-peng-teas-network-slicing-02 (work in progress), December 2019.

Authors' Addresses

Jin Zhou
ZTE Corp.

Email: zhou.jin6@zte.com.cn

Chunning Dai
ZTE Corp.

Email: dai.chunning@zte.com.cn

Shaofu Peng
ZTE Corp.

Email: peng.shaofu@zte.com.cn