

Workgroup: Network Working Group
Internet-Draft: draft-fan-bess-pm-bgp-02
Published: 3 March 2024
Intended Status: Standards Track
Expires: 4 September 2024
Authors: J. Fan Z. Li S. Zhuang
 Huawei Huawei Huawei

BGP Extension For L3VPN Performance Monitoring

Abstract

This document describes a new VT address family in BGP to exchange information required for apply performance monitoring in MPLS/BGP VPN, as described in [[I-D.dong-l3vpn-pm-framework](#)].

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 <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 4 September 2024.

Copyright Notice

Copyright (c) 2024 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. Terminologies](#)
- [3. The New VT Sub Address Family](#)
 - [3.1. VT Sub Address Family](#)
 - [3.2. VT NLRI](#)
 - [3.2.1. VT VPN-membership A-D Route](#)
 - [3.2.2. VT Labeled Route](#)
- [4. Operations](#)
 - [4.1. VRF-to-VRF VPN Membership Auto Discovery](#)
 - [4.2. VRF-to-VRF Labeled Route Exchange](#)
 - [4.2.1. VT Labeled Route Update](#)
 - [4.2.2. VT Labeled Route Withdrawal](#)
 - [4.3. VRF-to-VRF Labeled Route Application](#)
- [5. VT Route Selection Consideration](#)
- [6. Deployment Consideration](#)
- [7. IANA Considerations](#)
- [8. Security Considerations](#)
- [9. Acknowledgements](#)
- [10. References](#)
 - [10.1. Normative References](#)
 - [10.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

This document describes the BGP encodings and procedures for exchanging the information elements required by applying traffic performance monitoring in MPLS/BGP VPN, as specified in [\[I-D.dong-l3vpn-pm-framework\]](#).

Current BGP Labeled VPN Route exchange procedure combines VRF VPN-membership Auto-Discovery and L3VPN Label allocation together. While applying PM for L3VPN needs BGP extended to support VPN membership Auto-Discovery and L3VPN Label allocation in a VRF-to-VRF manner. To achieve this, a new Sub address family, called VRF-to-VRF Tunnel(VT) Subsequent Address Family, is introduced.

This document defines two kinds of routes for VT NLRI:

VPN-Membership A-D Route: for the use of doing VRF VPN membership auto-discovery in VRF-to-VRF manner

VT Labeled Route: for the use of allocating VT Label from Local VRF to Remote VRF to setup VRF-to-VRF Tunnel between the pair of VRFs.

2. Terminologies

This document uses the terminologies defined in [RFC4026]:

*ERT: Export Route Target

*IRT: Import Route Target

*PE: Provider Edge

*RD: Route Distinguisher

*VRF: Virtual Routing and Forwarding

*VT: VRF-to-VRF Tunnel

3. The New VT Sub Address Family

The BGP Multiprotocol Extensions [[RFC4760](#)] allow BGP to carry routes from multiple "address families". In this document a new Subsequent Address Family is introduced, called "VT Sub Address Family".

3.1. VT Sub Address Family

VT Address Family uses AFI 1/2 to present IPv4/IPv6 Address Family and a specific VT_SAFI(TBD) to present VT Subsequent Address Family.

VT MP_REACH_NLRI and MP_UNREACH_NLRI are formatted as described in [RFC4760]

```

+-----+
| Address Family Identifier (2 octets): 1/2 |
+-----+
| Subsequent AFI (1 octet): VT_SAFI (TBD) |
+-----+
| Length of Next Hop (1 octet): 4 |
+-----+
| Next Hop: IPv4 Address |
+-----+
| Reserved (1 octet) |
+-----+
| BGP VT NLRI (Variable) |
+-----+

```

Figure 1 VT MP_REACH_NLRI

```

+-----+
| Address Family Identifier (2 octets): 1/2 |
+-----+
| Subsequent AFI (1 octet): VT_SAFI (TBD) |
+-----+
| BGP VT NLRI (Variable) |
+-----+

```

Figure 2 VT MP_UNREACH_NLRI

3.2. VT NLRI

BGP VT NLRI has format as depicted in following diagram:

```

+-----+
| Route Type (1 octet) |
+-----+
| Length (1 octet) |
+-----+
| Route Type Specific (Variable) |
+-----+

```

Figure 3 IPv4 VT-Family NLRI

Route Type indicates type of route under VT SAFI.

Type 1: VT VPN membership A-D Route

Type 2: VT Labeled Route

Length defines Route Type specific routes length in octets

Route Type specific route information field, encoded according to Route Type definition.

3.2.1. VT VPN-membership A-D Route

VT VPN membership A-D Route, concisely named as VT A-D Route hereafter, is utilized for VRF-to-VRF VPN Membership Auto-Discovery between PEs.

Its format is defined as following diagram:




Figure 4 VT VPN Membership A-D Route

a) RD: RD of one VRF on advertising PE, encoded as described in [RFC4364].

b) Local Router's IP Address: Advertising PE's IPv4/IPv6 address
<RD, Local Router's IP Address> is defined as Prefix of VT A-D Route.

3.2.2. VT Labeled Route

VT Labeled Route is utilized for VRF-To-VRF Label(s) allocation and advertisement, its format is defined as following diagram.

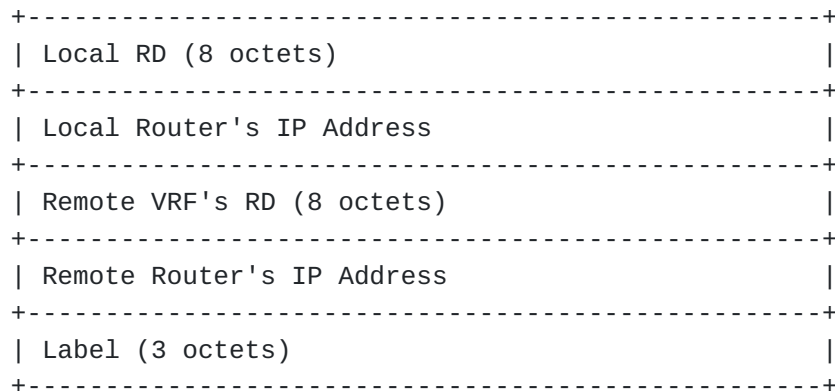


Figure 5 VT Labeled Route Format

a) Local RD: Route Distinguisher value of one VRF on advertising PE, encoded as described in [RFC4364].

b) Local Router's IP Address: Advertising PE's IPv4/IPv6 address.

c) Remote VRF's RD: Route Distinguisher value of Remote VRF encoded as described in [RFC4364].

d) Remote Router's IP Address: Remote PE's IPv4/IPv6 address.

e) Label: The Label field carries one or more labels that corresponds to the stack of labels [[RFC3032](#)]. Each label is encoded as 3 octets, where the high-order 20 bits contain the label value, and the low-order bit contains "Bottom of Stack" as defined in [[RFC3032](#)].

<Local RD, Local Router's IP Address, Remote VRF's RD, Remote Router's IP Address> which indicates a pair of VRFs is defined as the Prefix of VT Labeled Route.

4. Operations

4.1. VRF-to-VRF VPN Membership Auto Discovery

For every PE, it needs to process all its VRF configured and generate one VT A-D Route for each VRF respectively.

RD field MUST be filled with the VRF's RD value.

Local Router's IP Address field MUST filled with the Advertising Router's IP address.

The VT A-D Route MUST carry all IRTs of the VRF in BGP Update's Ext-Community Path Attribute, route importing request of one VRF is described by its corresponding VT A-D route. In contrast VPN Labeled Routes carry ERTs in BGP Update's Ext-Community Path Attribute.

If a VRF is created, then its corresponding VT A-D Route MUST be generated and advertised.

If the VRF whose VT A-D Route has been advertised is deleted, then the VT A-D Route Withdrawal message MUST be generated and advertised.

If IRT of the VRF whose VT A-D Route has been advertised is changed, then a VT A-D Route Update with same Prefix and latest IRTs MUST be advertised.

When receiving PE receives VT A-D Route, VPN relationship matching MUST be checked between IRTs in VT A-D Route and ERTs of each Local VRF, this process is called VRF-to-VRF VPN membership Auto Discovery.

Either finding one VRF-to-VRF VPN membership newly formed or released, receiving PE MUST proceed to the VT Labeled Route processing described in next section.

4.2. VRF-to-VRF Labeled Route Exchange

4.2.1. VT Labeled Route Update

If Receiving PE finds one new VRF-to-VRF VPN membership formed, it MUST allocate one VT MPLS Label for the VRF-to-VRF VPN membership and the label is advertised to the Remote VRF by VT Labeled Route.

Local RD MUST filled with RD value of the Local VRF which is found belong to the same VPN with Remote VRF.

Local Router's IP Address Must filled with the advertising PE's IPv4/ IPv6 address.

Remote VRF's RD MUST filled with RD value of the Remote VRF which belongs to a same VPN with the Local VRF.

Remote Router's IP Address: Remote PE's IPv4/IPv6 address.

Label: MUST be filled with one or more MPLS Labels allocated by advertising PE for the pair of VRFs.

Only both sides of a pair of VRFs learnt each other's VT Labeled Route advertisement, the VRF-to-VRF tunnel between the pair of VRFs is considered setup.

4.2.2. VT Labeled Route Withdrawal

If receiving PE finds one existing VRF-to-VRF VPN membership released then it MUST send out the VT Labeled Route Withdrawal message, then release the MPLS Label(s) allocated.

Local RD MUST be filled with RD value of the Local VRF.

Local Router's IP Address MUST be filled with the advertising PE's IPv4/IPv6 address.

Remote VRF's RD MUST be filled with RD value of the Remote VRF.

Remote Router's IP Address: MUST be filled with Remote PE's IPv4/ IPv6 address.

Label: MUST be filled with ZERO or the MPLS Labels value allocated for the VT Labeled Route.

4.3. VRF-to-VRF Labeled Route Application

To achieve the goal of converting normal L3VPN MP2P forwarding model into P2P model which is required in [[I-D.zheng-l3vpn-pm-analysis](#)], after VPNv4 routes received, Receiving PE MUST apply VT Labels when downloading VPNv4 Route into Data Plan which is in detail described in [[I-D.dong-l3vpn-pm-framework](#)].

Between a pair of PEs both support VT capability, It COULD be an implementation option that VPNv4 Routes from a remote VRF WOULD NOT be downloaded into a Local VRF's Forwarding Plan until a VT Labeled route received from same Remote VRF for the Local VRF.

If VT Labeled Route withdrawal message is received, receiving PE MUST delete VT Labels from Forwarding Plane and VPNv4 Routes MUST be kept on Forwarding Plane with original VPNv4 Label as inner Label.

5. VT Route Selection Consideration

When receiving and processing VT A-D Route, the BGP best route selection procedure described in [[RFC4271](#)] MUST be followed.

When receiving and processing VT Labeled Route, the BGP best route selection procedure described in [[RFC4271](#)] COULD be followed.

Especially VT Labeled Route MUST be advertised ONLY to the BGP peer from which the best VT A-D route is received, the VT A-D route contains the Remote VRF's RD and Remote PE's IP address.

If a Peer receives VT A-D or VT Labeled Route originated from itself, the route MUST be ignored.

6. Deployment Consideration

This document currently supports deploying VT SAFI in following two manners:

- a) Inner-AS L3VPN with Full-mesh IBGP sessions or Router Reflectors.
- b) Inter-AS L3VPN with Option A(VRF-to-VRF)[[RFC4364](#)].

How to support Inter-AS L3VPN Option B(MP-EBGP) and Option-C [[RFC4364](#)] will be described in this draft's future version.

7. IANA Considerations

A new SAFI value to present VT Subsequent Address Family is required and to be allocated by IANA.

8. Security Considerations

This extension to BGP does not change the underlying security issues.

9. Acknowledgements

TBD.

10. References

10.1. Normative References

- [I-D.dong-l3vpn-pm-framework] Dong, J., Li, Z., and B. Parise, "A Framework for L3VPN Performance Monitoring", Work in Progress, Internet-Draft, draft-dong-l3vpn-pm-framework-03, 27 October 2014, <<https://datatracker.ietf.org/doc/html/draft-dong-l3vpn-pm-framework-03>>.
- [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>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.

10.2. Informative References

- [I-D.zheng-l3vpn-pm-analysis] Zheng, L., Li, Z., Aldrin, S., and B. Parise, "Performance Monitoring Analysis for L3VPN", Work in Progress, Internet-Draft, draft-zheng-l3vpn-pm-analysis-03, 3 July 2014, <<https://datatracker.ietf.org/doc/html/draft-zheng-l3vpn-pm-analysis-03>>.

Authors' Addresses

Jiajia Fan
Huawei
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095

China

Email: fanjiajia@huawei.com

Zhenbin Li

Huawei

Huawei Bld., No.156 Beiqing Rd.

Beijing

100095

China

Email: lizhenbin@huawei.com

Shunwan Zhuang

Huawei

Huawei Bld., No.156 Beiqing Rd.

Beijing

100095

China

Email: zhuangshunwan@huawei.com