

INTERNET-DRAFT
Intended status: Proposed Standard

S. Zhuang
Z. Li
Huawei Technologies
D. Eastlake
Futurewei Technologies
L. Yong
Independent
July 9, 2021

Expires: January 8, 2022

BGP Extensions for Enhanced VPN Auto Discovery
draft-zhuang-bess-enhanced-vpn-auto-discovery-08.txt

Abstract

A variety of VPN technologies have been widely deployed to bear different services. As new applications develop, a requirement has been proposed for auto-discovery of Layer 3 Virtual Private Networks (L3VPN) and enhanced auto-discovery requirements for other VPN technologies that already have basic auto-discovery mechanisms.

This document identifies some possible applications of these auto-discovery requirements and defines a new BGP NLRI, called the BGP-VPN-INSTANCE NLRI, to satisfy the requirement for auto-discovery of BGP VPN instances. It also defines a new type of extended community, called the Import Route Target, which can be applied to auto-discovery mechanisms of multiple VPN technologies.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Distribution of this document is unlimited. Comments should be sent to the authors or the BESS working group mailing list: bess@ietf.org.

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 <https://www.ietf.org/1id-abstracts.html>. The list of Internet-Draft

Shadow Directories can be accessed at
<https://www.ietf.org/shadow.html>.

Table of Contents

1. Introduction.....	3
2. Terminologies.....	4
3. Requirements of VPN Auto-Discovery.....	5
3.1 Centralized Traffic Optimization.....	5
3.2 Label/Segment Allocation for VPN Instance.....	5
4. IRT Extended Community.....	6
5. BGP Extensions for L3VPN Auto-Discovery.....	7
5.1 BGP-VPN-INSTANCE SAFI.....	7
5.2 BGP-VPN-INSTANCE NLRI.....	8
5.2.1 VPN Membership A-D Route.....	8
5.3 Procedures.....	9
6. IANA Considerations.....	10
6.1 BGP Extended Communities.....	10
6.2 Subsequent Address Family Identifier.....	10
7. Security Considerations.....	11
Contributors.....	11
Acknowledgements.....	11
Normative References.....	12
Informative References.....	13
Authors' Addresses.....	14

1. Introduction

A variety of VPN technologies have been widely deployed to bear different services. As new applications develop, a requirement has been proposed for auto-discovery of Layer 3 Virtual Private Networks (L3VPN) [[RFC4364](#)] and enhanced auto-discovery requirements for other VPN technologies which already have basic auto-discovery mechanisms.

This document identifies some possible applications of these auto-discovery requirements and defines a new BGP NLRI [[RFC4271](#)], called the BGP-VPN-INSTANCE NLRI, to satisfy the requirement of auto-discovery of BGP VPN instance. It also defines a new type of extended community, called the Import Route Target (IRT), which can be applied to auto-discovery mechanisms of multiple VPN technologies.

2. Terminologies

This document uses the terminologies defined in [[RFC4026](#)]:

A-D: Auto-Discovery

AFI: Address Family Identifier

ERT: Export Route Target

IRT: Import Route Target

LSP: Label Switched Path

NLRI: Network Layer Reachability Information

P2MP: Point to Multi-Point

PE: Provider Edge

RD: Route Distinguisher

VRF: Virtual Routing and Forwarding

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

[3.](#) Requirements of VPN Auto-Discovery

The following subsections are examples of VPN Auto-Discovery requirements.

[3.1](#) Centralized Traffic Optimization

As the development of centrally controlled application such as PCE-initiated LSP [[RFC8281](#)] and PCE-initiated P2MP LSP [[RFC8623](#)], PCE can be used to initiate setup of RSVP-TE LSP or P2MP LSP for the purpose of traffic optimization. In order to support such applications, the controller should learn the relationship of unicast VPN instances or multicast VPN instances distributed on different PEs. According to

the existing VPN auto-discovery mechanism for technologies such as EVPN [[RFC7432](#)] or MVPN [[RFC6514](#)], the A-D routes are always advertised with the Export Route Target (ERT). The ingress PE can use an Import Route Target (IRT) of the local MVPN/EVPN instance to match the route target advertised with the NLRI to determine the relationship of these VPN instances. But the controller, which can be used as the route reflector of VPN routes, cannot learn the relationship of VPN instances since the Import Route Target information is not advertised with these A-D routes. In order to support such applications the IRT can be carried with A-D routes as specified below.

[3.2](#) Label/Segment Allocation for VPN Instance

[I-D.li-mpls-global-label-usecases] proposes use cases of label allocation for unicast VPN or multicast VPN instances.

[[I-D.li-spring-segment-path-programming](#)] proposes use cases of segment allocation for steering traffic. In order to support such applications the PEs need to learn the relationship of VPN instances distributed on other PEs. For L3VPN [[RFC4364](#)] there is no auto-discovery mechanism for BGP VPN instances. In order to support such applications, an auto-discovery mechanism for L3VPN is specified below.

[4.](#) IRT Extended Community

This document defines a new type of transitive extended community, called Import Route Target.

The IANA registry of BGP Extended Communities clearly identifies communities of specific formats: "Two-octet AS Specific Extended

Community" [[RFC4360](#)], "Four-octet AS Specific Extended Community" [[RFC5668](#)], and "IPv4 Address Specific Extended Community" [[RFC4360](#)]. Route Target [[RFC4360](#)] extended communities identify this format in the high-order (Type) octet of the Extended Community. The Import Route Target extended community reuses the same mechanism.

This document defines the following IRT Extended Communities:

Type	Sub-Type	Extended Community	Encoding
0x00	TBD1	AS-2byte IRT	2-octet AS, 4-octet Value
0x01	TBD2	IPv4 IRT	4-octet IPv4 Address, 2-octet Value
0x02	TBD3	AS-4byte IRT	4-octet AS, 2-octet Value

Figure 1. IRT Extended Communities

The IRT Extended Community can be used for MVPN [[RFC6514](#)], L3VPN [[RFC4364](#)], EVPN [[RFC7432](#)], BGP-based VPLS [[RFC4761](#)], and BGP-AD-based VPLS [[RFC6074](#)] and the like. The existing auto-discovery mechanisms of these VPN technologies always carry the ERT extended community. To meet the requirements of applications, they need to carry the IRT extended community with different A-D routes. The local policy, which is out of scope of this document, can be used to control the distribution of IRT information.

[5.](#) BGP Extensions for L3VPN Auto-Discovery

[5.1](#) BGP-VPN-INSTANCE SAFI

The BGP Multiprotocol Extensions [[RFC4760](#)] allow BGP to carry routes from multiple "address families". In this document a new Subsequent Address Family is specified, called "BGP-VPN-INSTANCE Sub Address Family", which uses a specific BGP-VPN-INSTANCE-SAFI (TBD4).

This document also defines a new BGP NLRI, called the BGP-VPN-INSTANCE NLRI to support the BGP VPN instance auto-discovery. BGP-VPN-INSTANCE MP_REACH_NLRI and MP_UNREACH_NLRI (shown in Figures 2 and 3) are formatted as described in [[RFC4760](#)]. The BGP-VPN-INSTANCE NLRI is described in more detail in [Section 5.2](#).

```
+-----+
| Address Family Identifier: 1/2/25      (2 octets)      |
+-----+-----+
| Subsequent AFI:                        |              (1 octet)
| BGP-VPN-INSTANCE-SAFI=TBD4|
+-----+-----+
| Length of Next Hop                    |              (1 octet)
+-----+-----+---...
| Next Hop                             |              (variable)
+-----+-----+---...
| Reserved                             |              (1 octet)
+-----+-----+---...
| BGP-VPN-INSTANCE NLRI                 |              (variable)
+-----+-----+---...
```

Figure 2. BGP-VPN-INSTANCE MP_REACH_NLRI

```
+-----+
| Address Family Identifier: 1/2/25      (2 octets)      |
+-----+-----+
| Subsequent AFI:                        |              (1 octet)
| BGP-VPN-INSTANCE-SAFI=TBD4|
+-----+-----+---...
| BGP-VPN-INSTANCE NLRI                 |              (variable)
+-----+-----+---...
```

Figure 3. BGP-VPN-INSTANCE MP_UNREACH_NLRI

5.2 BGP-VPN-INSTANCE NLRI

The following is the format of the BGP-VPN-INSTANCE NLRI.




Figure 4. BGP-VPN-INSTANCE NLRI

The Route Type field specifies the encoding of the rest of BGP-VPN-INSTANCE NLRI (Route Type specific BGP-VPN-INSTANCE NLRI).

The Length field indicates the length in octets of the Route Type specific field of the BGP-VPN-INSTANCE NLRI.

This document defines the following Route Type for BGP-VPN-INSTANCE routes:

Type 1: VPN Membership A-D Route

5.2.1 VPN Membership A-D Route

The VPN Membership A-D Route is utilized for VPN Membership A-D between PEs. Its format is shown in Figure 5.




Figure 5. VPN Membership A-D Route

- a) Local Router's IP Address: Advertising PE's IPv4/IPv6 address.
- b) RD: RD of one VRF on the advertising PE, encoded as described in [\[RFC4364\]](#).

[5.3](#) Procedures

Every PE needs to process all its VRF configuration and generate one VPN Membership A-D Route for each VRF respectively. The Local Router's IP Address field MUST be filled with the Advertising Router's IP address. The RD field MUST be filled with the VRF's RD value.

All ERTs of the VRF MUST be carried in a BGP Update's RT Extended Community Path Attribute with the Membership A-D Route for the VRF. To meet the requirements of different applications, all IRTs of the VRF SHOULD be able to be carried in BGP Update's IRT Extended Community Path Attribute with the VPN Membership A-D Route for the VRF.

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

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

If IRTs or ERTs of the VRF whose VPN Membership A-D Route has been advertised are changed, then a VPN Membership A-D Route Update with same Prefix and latest IRTs or ERTs MUST be advertised.

When the receiving PE receives a VPN Membership A-D Route, VPN relationship matching MUST be checked with the IRTs carried in the VPN Membership A-D Route and ERTs of each Local VRF.

When the central controller receives a VPN Membership A-D Route, VPN relationship matching MUST be checked with IRTs and ERTs carried in VPN Membership A-D Routes of different VPN instances.

[6.](#) IANA Considerations

[6.1](#) BGP Extended Communities

IANA is requested to assign three BGP Extended Community Sub-Types as shown below.

Transitive Two-Octet AS-Specific Extended Community Sub-Type

Sub-Type	Description	Reference
TBD1	Import Route Target	[this document]

Transitive IPv4-Address-Specific Extended Community Sub-Type

Sub-Type	Description	Reference
TBD2	Import Route Target	[this document]

Transitive Four-Octet AS-Specific Extended Community Sub-Type

Sub-Type	Description	Reference
TBD3	Import Route Target	[this document]

[6.2](#) Subsequent Address Family Identifier

IANA is requested to assign a Subsequent Address Family Identifier (SAFI) from the First Come First Served range as follows:

Value	Description	Reference
-----	-----	-----
TBD4	BGP-VPN-INSTANCE-SAFI	[this document]

Zhuang, et al

Expires January 2022

[Page 10]

INTERNET-DRAFT

BGP Extensions For VPN AD

July 2021

[7.](#) Security Considerations

TBD

Contributors

The following people have substantially contributed to the solution and to the editing of this document:.

Hui Ni
Huawei Technologies
Email: nihui@huawei.com

Acknowledgements

The authors would like to thank Shuanglong Chen and Eric Wu for their contributions to this work.

Zhuang, et al Expires January 2022 [Page 11]

INTERNET-DRAFT BGP Extensions For VPN AD July 2021

Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [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, <<http://www.rfc-editor.org/info/rfc4271>>.

[editor.org/info/rfc4271](http://www.rfc-editor.org/info/rfc4271)>.

- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), DOI 10.17487/RFC4360, February 2006, <<http://www.rfc-editor.org/info/rfc4360>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), DOI 10.17487/RFC4364, February 2006, <<http://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, <<http://www.rfc-editor.org/info/rfc4760>>.
- [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), DOI 10.17487/RFC4761, January 2007, <<http://www.rfc-editor.org/info/rfc4761>>.
- [RFC5668] Rekhter, Y., Sangli, S., and D. Tappan, "4-Octet AS Specific BGP Extended Community", [RFC 5668](#), DOI 10.17487/RFC5668, October 2009, <<http://www.rfc-editor.org/info/rfc5668>>.
- [RFC6074] Rosen, E., Davie, B., Radoaca, V., and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", [RFC 6074](#), DOI 10.17487/RFC6074, January 2011, <<http://www.rfc-editor.org/info/rfc6074>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", [RFC 6514](#), DOI 10.17487/RFC6514, February 2012, <<http://www.rfc-editor.org/info/rfc6514>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", [RFC 7432](#), DOI 10.17487/RFC7432, February 2015, <<http://www.rfc-editor.org/info/rfc7432>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May

2017, <<https://www.rfc-editor.org/info/rfc8174>>.

Informative References

- [I-D.li-mpls-global-label-usecases] Li, Z., Zhao, Q., Yang, T., Raszuk, R., and L. Fang, "Usecases of MPLS Global Label", [draft-li-mpls-global-label-usecases-03](#) (work in progress), October 2015.
- [I-D.li-spring-segment-path-programming] Li, Z., Milojevic, I., Z. Zhuang, "Segment Path Programming (SPP)", [draft-li-spring-segment-path-programming-00](#) (work in progress), October 2015.
- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", [RFC 4026](#), DOI 10.17487/RFC4026, March 2005, <<https://www.rfc-editor.org/info/rfc4026>>.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", [RFC 8281](#), DOI 10.17487/RFC8281, December 2017, <<https://www.rfc-editor.org/info/rfc8281>>.
- [RFC8623] Palle, U., Dhody, D., Tanaka, Y., and V. Beeram, "Stateful Path Computation Element (PCE) Protocol Extensions for Usage with Point-to-Multipoint TE Label Switched Paths (LSPs)", [RFC 8623](#), DOI 10.17487/RFC8623, June 2019, <<https://www.rfc-editor.org/info/rfc8623>>.

Authors' Addresses

Shunwan Zhuang
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing. 100095 China

Email: zhuangshunwan@huawei.com

Zhenbin Li
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing, 100095 China

Email: lizhenbin@huawei.com

Donald Eastlake
Futurewei Technologies
2386 Panoramic Circle
Apopka, FL 32703 USA

Phone: +1-508-333-2270
Email: d3e3e3@gmail.com

Lucy Yong
Independent
USA

Phone: +1-469-227-5837
Email: lucyyong@gmail.com

Copyright, Disclaimer, and Additional IPR Provisions

Copyright (c) 2021 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

