

INTERNET-DRAFT
Intended status: Proposed Standard

Chunxia Sun
Yaokun Zhang
Donald Eastlake 3rd
Huawei
October 22, 2018

Expires: April 21, 2019

Controlling VPN Routes More Precisely by ORF Extension
draft-sun-bess-vpn-orf-extension-00

Abstract

[RFC 4684](#) defines Multi-Protocol BGP (MP-BGP) procedures that allow BGP speakers to exchange Route Target reachability information which can be used to build a route distribution graph in order to limit the propagation of Virtual Private Network (VPN) Network Layer Reachability Information (NLRI) between different autonomous systems or distinct clusters of the same autonomous system.

However, according to [RFC 4684](#), in some scenarios, more routes will be sent than need to be sent.

This document extends [RFC 4684](#). This extension allows a BGP speaker to advertise VPN routes more precisely.

Status of this Memo

This Internet-Draft is submitted to IETF 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 <http://www.ietf.org/lid-abstracts.html>. The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Internet-Draft

VPN ORF Extension

Table of Contents

1. Introduction.....	3
1.1 Requirements Language.....	3
1.2 Terminology.....	3
2. Solution.....	5
3. BGP Encoding.....	6
4. Security Considerations.....	7
5. IANA Considerations.....	7
Normative References.....	8
Informative References.....	8
Acknowledgments.....	8
Authors' Addresses.....	9

1. Introduction

[RFC4684] defines Multi-Protocol BGP (MP-BGP) procedures that allow BGP speakers to exchange Route Target reachability information which can be used to build a route distribution graph in order to limit the propagation of Virtual Private Network (VPN) Network Layer Reachability Information (NLRI) between different autonomous systems or distinct clusters of the same autonomous system.

However, according to the extension of BGP protocol by [RFC4684], in some scenarios, for example, when the same route targets exist in different BGP address families, more routes will be sent than need to be sent, which violates the original intention of the ORF implementation.

This document extends [RFC4684]. This extension allows a BGP speaker to advertise VPN routes more precisely when BGP speaker has the same route target in different address families.

1.1 Requirements Language

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.

[1.2](#) Terminology

AFI: Address Family Identifier(a BGP address type)

SAFI: Subsequence Address Family Identifier (a BGP address sub-type)

BGP: Border Gateway Protocol

VPN: Virtual Private Network

PE: Provider Edge device

CE: Customer Edge(router)

EVPN: Ethernet Virtual Private Network

L3VPN: Layer 3 Virtual Private Network

iBGP: Internal BGP (i.e., a BGP peering session that connects two routers within an autonomous system)

MP-BGP: MultiProtocol-Border Gateway Protocol

MPLS: MultiProtocol Label Switching

NLRI: Network Layer Reachability Information

ORF: Outbound Route Filtering

RT: Route Target

2. Solution

In the [Section 4 of \[RFC4684\]](#), the packet structure of the ORF route is defined. The ORF route prefix carries only the original AS and route-target information, and does not carry the address family information corresponding to the route-target.

Let us give an example of the problem of route advertisement in [\[RFC4684\]](#). For example, RTA and RTB are neighbors in EVPN and neighbors in L3VPN. The route-target of the EVPN instance on RTA is 100:1, the route-target of the L3VPN instance on RTA is 200:1, and the route-target of the EVPN instance on RTB is 100:1, the route-target of the L3VPN instance on RTB is 100:1. The ORF capability is enabled on both RTA and RTB. After the neighbor relationship between RTA and RTB is established, RTA sends ORF routes to inform RTB routes

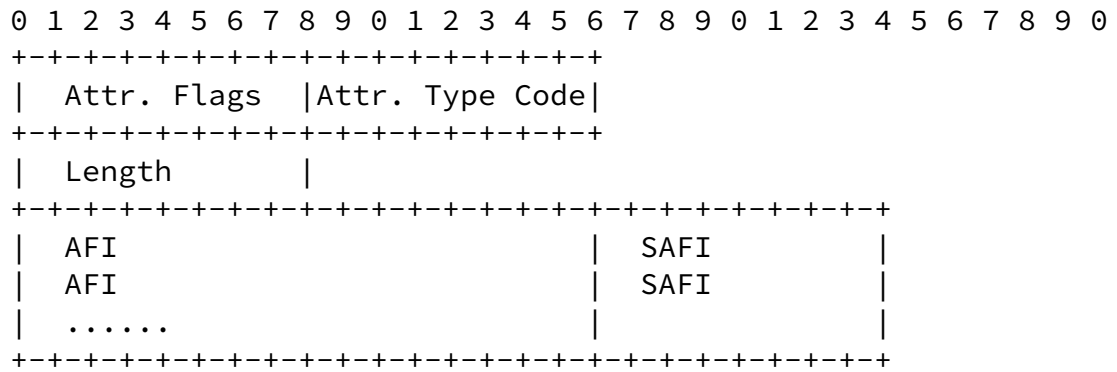
with a route- target of 100:1 and 200:1 are required. After receiving the ORF routes of RTA, RTB sends the routes with the route-target of 100:1 to RTA, including the EVPN routes with the route-target of 100:1 and L3VPN routes with route-target of 100:1. In fact, RTA is not required for L3VPN routes with route-target of 100:1.

In order to solve the problem above, [[RFC4684](#)] is extended, the RT-ORF-DOMAIN attribute is added to the ORF routes, and the address families corresponding to the route target is carried in the attribute, for example: EVPN, VPNv4, etc.;

In the above example, the ORF routes sent by RTA to RTB carry the information that RTA wants to receive the routes of EVPN address family with the route target of 100:1 and the routes of VPNv4 address family with the route target of 200:1. After receiving the ORF routes from RTA, RTB will only send the routes of the EVPN address family with the route target of 100:1 to RTA.

[3.](#) BGP Encoding

[RFC4684] defines the packet structure of the ORF route, including the route prefix and attributes. This document extends [[RFC4684](#)] and adds the RT-ORF-DOMAIN attribute to the ORF route. This attribute is composed as follows:



RT-ORF-DOMAIN is an optional transitive attribute, and the attribute type is to be assigned. The role of this attribute has been described in [Section 2](#).

[4.](#) Security Considerations

TBD

[5.](#) IANA Considerations

TBD

Internet-Draft

VPN ORF Extension

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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4684] R. Bonica, " Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS)Internet Protocol (IP) Virtual Private Networks (VPNs) ", [RFC 4684](#), November 2006.
- [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

TBD

Acknowledgments

The authors would like to thank the following for their valuable contributions of this document:

TBD

Internet-Draft

VPN ORF Extension

Authors' Addresses

Chunxia Sun
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: sunchunxia@huawei.com

Yaokun Zhang
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: zhangyaokun@huawei.com

Donald Eastlake 3rd
Huawei Technologies
1424 Pro Shop Court
Davenport, FL 33896
USA

Phone: +1-508-333-2270
Email: Donald.Eastlake@huawei.com

Internet-Draft

VPN ORF Extension

Copyright, Disclaimer, and Additional IPR Provisions

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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.

