

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 27, 2021

W. Wang
A. Wang
China Telecom
H. Wang
Huawei Technologies
G. Mishra
Verizon Inc.
S. Zhuang
J. Dong
Huawei Technologies
November 23, 2020

**Route Distinguisher Outbound Route Filter (RD-ORF) for BGP-4
draft-wang-idr-rd-orf-05**

Abstract

This draft defines a new Outbound Route Filter (ORF) type, called the Route Distinguisher ORF (RD-ORF). RD-ORF is applicable when the routers do not exchange VPN routing information directly (e.g. routers in single-domain connect via Route Reflector, or routers in Option B/Option AB/Option C cross-domain scenario).

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 May 27, 2021.

Copyright Notice

Copyright (c) 2020 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 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.

Table of Contents

1.	Introduction	2
2.	Conventions used in this document	4
3.	Terminology	4
4.	RD-ORF Encoding	5
5.	Application in single-domain scenario	6
5.1.	Addition of RD-ORF entries	6
5.1.1.	Operation process of RD-ORF mechanism on source PE .	7
5.1.2.	Operation process of RD-ORF mechanism on RR	8
5.1.3.	Operation process of RD-ORF mechanism on target PE .	9
5.2.	Withdraw of RD-ORF entries	9
6.	Applications in cross-domain scenarios	9
6.1.	Application in Option B/Option AB cross-domain scenario .	9
6.2.	Application in Option C cross-domain scenario	11
7.	Security Considerations	12
8.	IANA Considerations	12
9.	Acknowledgement	12
10.	Normative References	12
	Authors' Addresses	13

[1.](#) Introduction

With the rapid growth of network scale, Route Reflector is introduced in order to reduce the network complexity. Routers in the same Autonomous System only need to establish iBGP session with RR to transmit routes.

In VPN scenario shown in Figure 1, PE1 - PE4 establish iBGP sessions with RR to ensure the routes can be transmitted within AS100, where PE1 and PE3 maintain VRFs of VPN1 and VPN2, PE2 maintains VPN1's VRF and PE4 maintains VPN2's VRF. RR don not maintain any VRFs.




Figure 1: Single-domain scenario

When the VRF of VPN1 in PE1 overflows, due to PE1 and other PEs are not iBGP neighbors, BGP Maximum Prefix Features cannot work, so the problem on PE2 cannot be known.

Now, there are several solutions can be used to alleviate this problem:

- o Route Target Constraint (RTC) as defined in [[RFC4684](#)]
- o Address Prefix ORF as defined in [[RFC5292](#)]
- o PE-CE edge peer Maximum Prefix
- o Configure the Maximum Prefix for each VRF on edge nodes

However, there are limitations to existing solutions:

1) Route Target Constraint

RTC can only filter the VPN routes from the uninterested VRFs, if the "trashing routes" come from the interested VRF, filter on RTs will erase all prefixes from this VRF.

2) Address Prefix ORF

Using Address Prefix ORF to filter VPN routes need to pre-configuration, but it is impossible to know which prefix may cause overflow in advance.

3) PE-CE edge peer Maximum Prefix

This mechanism can only protect the edge between PE-CE, it can't be deployed within PE that peered via RR. Depending solely on the edge protection is dangerous, because if only one of the edge points being comprised/error-configured/attacked, then all of PEs within domain are under risk.

4) Configure the Maximum Prefix for each VRF on edge nodes

When a VRF overflows, it stops the import of routes and log the extra VPN routes into its RIB. However, PEs should parse the BGP updates. These processes will cost CPU cycles and further burden the overflowing PE.

This draft defines a new ORF-type, called the Route Distinguisher ORF (RD-ORF). Using RD-ORF mechanism, VPN routes of a VPN can be controlled based on source RD. This mechanism is event-driven and does not need to be pre-configured. When a VRF of a router overflows, the router will find out the main source RD of VPN routes in this VRF, and send a RD-ORF to its BGP peer that carries the RD. If a BGP speaker receives a RD-ORF from its BGP peer, it will filter the VPN routes it tends to send according to the RD-ORF entry.

RD-ORF is applicable when the routers do not exchange VPN routing information directly (e.g. routers in single-domain connect via Route Reflector, or routers in Option B/Option AB/Option C cross-domain scenario).

2. Conventions used in this document

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. Terminology

The following terms are defined in this draft:

- o RD: Route Distinguisher, defined in [[RFC4364](#)]
- o ORF: Outbound Route Filter, defined in [[RFC5291](#)]
- o AFI: Address Family Identifier, defined in [[RFC4760](#)]

- o SAFI: Subsequent Address Family Identifier, defined in [[RFC4760](#)]
- o EVPN: BGP/MPLS Ethernet VPN, defined in [[RFC7432](#)]
- o RR: Router Reflector, provides a simple solution to the problem of IBGP full mesh connection in large-scale IBGP implementation.
- o VRF: Virtual Routing Forwarding, a virtual routing table based on VPN instance.

4. RD-ORF Encoding

In this draft, we defined a new ORF type called Route Distinguisher Outbound Route Filter (RD-ORF). The ORF entries are carried in the BGP ROUTE-REFRESH message as defined in [[RFC5291](#)]. A BGP ROUTE-REFRESH message can carry one or more ORF entries, and MUST be regenerated when it is tended to be sent to other BGP peers. The ROUTE-REFRESH message which carries ORF entries contains the following fields:

- o AFI (2 octets)
- o SAFI (1 octet)
- o When-to-refresh (1 octet): the value is IMMEDIATE or DEFER
- o ORF Type (1 octet)
- o Length of ORF entries (2 octets)

A RD-ORF entry contains a common part and type-specific part. The common part is encoded as follows:

- o Action (2 bits): the value is ADD, REMOVE or REMOVE-ALL
- o Match (1 bit): the value is PERMIT or DENY
- o Reserved (5 bits)

RD-ORF also contains type-specific part. The encoding of the type-specific part is shown in Figure 2.




Figure 2: RD-ORF type-specific encoding

- o Sequence: identifying the order in which RD-ORF is generated
- o Route Distinguisher: distinguish the different user routes. The RD-ORF filters the VPN routes it tends to send based on Route Distinguisher.

Note that if the Action component of an ORF entry specifies REMOVE-ALL, the ORF entry does not include the type-specific part.

When the BGP ROUTE-REFRESH message carries RD-ORF entries, it must be set as follows:

- o The ORF-Type MUST be set to RD-ORF.
- o The AFI MUST be set to IPv4, IPv6, or Layer 2 VPN (L2VPN).
- o If the AFI is set to IPv4 or IPv6, the SAFI MUST be set to MPLS-labeled VPN address.
- o If the AFI is set to L2VPN, the SAFI MUST be set to BGP EVPN.
- o The Match field MUST be equal to DENY.

5. Application in single-domain scenario

5.1. Addition of RD-ORF entries

The operation of RD-ORF mechanism on each device is independent, each of them makes a local judgement to determine whether it needs to send RD-ORF to its peers.

In general, every VRF on PE is configured a Maximum Prefix, the trigger of RD-ORF mechanism can be set as the number of VPN routes in VRF reach 80% of the Maximum Prefix. For RR, it doesn't have VRF and the mechanism can be triggered by other conditions, such as the RR's memory/CPU utilization reaches 80%.

When the RD-ORF mechanism is triggered, the device must send an alarm information to network operators.

5.1.1. Operation process of RD-ORF mechanism on source PE

In scenario shown in Figure 1, when the VRF of VPN1 in PE1 overflows, PE1 will do analysis and calculation locally to find out the main source of VPN routes in this VRF, assuming it is PE3. Then, PE1 will resolve the corresponding RD of VPN routes from BGP UPDATE message, and generate a BGP ROUTE-REFRESH message contains a RD-ORF entry, and send it to RR. The message contains the following fields:

- o AFI is set to IPv4 , IPv6 or L2 VPN
- o SAFI is set to "MPLS-labeled VPN address" or "BGP EVPN"
- o When-to-refresh is set to IMMEDIATE
- o ORF Type is set to RD-ORF
- o Length of ORF entries depends on the type of Source Address sub-TLV (21, 23 or 33 octets)
- o Action is set to ADD
- o Match is set to DENY
- o Sequence is set to 1
- o Route Distinguisher is set to RD1

It noted that the Sequence can uniquely identifies an RD-ORF entry. All VRFs share the sequence field, and the corresponding sequence of RD-ORF sent by each VRF will be recorded on the device.

Sometimes, several VRFs in a PE may import VPN routes carries the same RT, as shown in Figure 3.

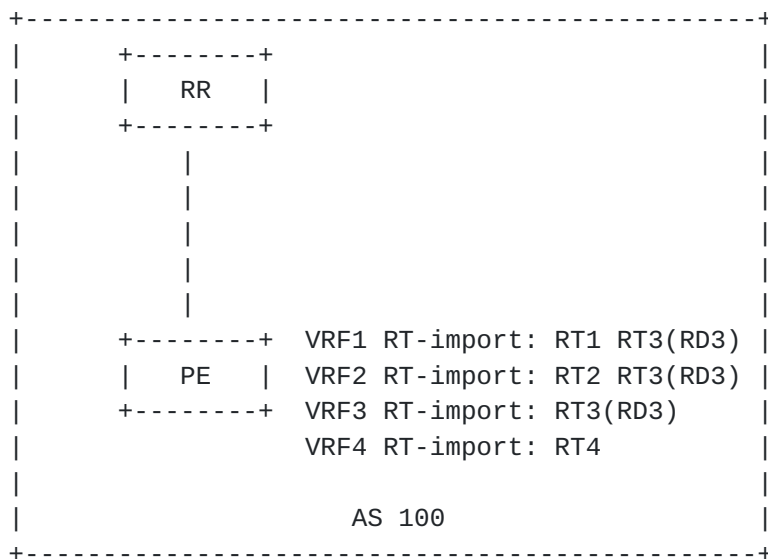


Figure 3: The scenario of several VRFs in a PE import VPN routes carries the same RT

In this scenario, VRF1, VRF2 and VRF3 import VPN routes carries RT3, which contains RD3. VRF1, VRF2 and VRF3 have different maximum prefix. When the VPN routes carrying RT3 cause the overflow of VRF3, PE will send a BGP ROUTE-REFRESH message containing a RD-ORF entry to RR, which Route Distinguisher field is equal to RD3. RR will stop sending associated VPN routes to PE. However, this will cause VRF1 to fail to receive VPN routes containing RD3.

The local determination of the PE can be used to inhibit the PE from sending RD-ORF entries. When the resources of the device are not exhausted, only prevent the overflowed VRF from importing related VPN routes without sending RD-ORF, unless all the VRFs that import the RD overflow.

5.1.2. Operation process of RD-ORF mechanism on RR

When RR receives the ROUTE-REFRESH message, it checks <AFI/SAFI, ORF-Type, Sequence, Route Distinguisher> to find whether it received the latest entry or not. If not, RR will discard the entry; otherwise, RR will add the RD-ORF entry into its Adj-RIB-out.

Before sending a VPN route toward PE1, RR will check its Adj-RIB-out and find there is a filter associated with RD1. Then, RR will stop sending that VPN route to PE1.

If the processing capacity of RR reaches the limit (e.g. RR's memory/CPU utilization reaches 80%), RR will find out the peer that sends the most routing entries to it, assuming it is PE3. Then, RR

will generate a BGP ROUTE-REFRESH message contains a RD-ORF entry based on the result of calculation, and send it to PE3.

5.1.3. Operation process of RD-ORF mechanism on target PE

After receiving the ROUTE-REFRESH message that carries a RD-ORF entry, PE3 will check if it receives the latest entry. If not, PE3 will discard it; otherwise, PE3 will add the RD-ORF entry into its Adj-RIB-out.

Before sending a VPN route toward RR, PE3 will check its Adj-RIB-out and find the RD-ORF entry prevent it from sending VPN route which carries RD1 to RR. Then, PE3 will stop sending that VPN route.

The BGP Maximum Prefix Features can be configured to protect PE-CE peering at the edge. Therefore, in general, CEs will not cause the overflow of PEs. If the boundary protection measures fail and cause the overflow, the PE can calculate and find the CEs in corresponding VRF, and break down the associated BGP sessions.

5.2. Withdraw of RD-ORF entries

When the RD-ORF mechanism is triggered, the alarm information will be generated and sent to the network operators. Operators should manually configure the network to resume normal operation. Due to devices can record the RD-ORF entries sent by each VRF, operators can find the entries needs to be withdrawn, and trigger the withdraw process as described in [[RFC5291](#)] manually. After returning to normal, the device sends withdraw ORF entries to its peers who have previously received ORF entries.

6. Applications in cross-domain scenarios

6.1. Application in Option B/Option AB cross-domain scenario

The Option B/Option AB cross-domain scenario is shown in Figure 4:




Figure 4: The Option B/Option AB cross-domain scenario

In Option B cross-domain scenario, PE1 - PE4 are responsible for maintaining VPN routing information in AS1 and AS2. There is a direct link between ASBR1 and ASBR2 via EBGP. In AS1, PE1 and PE2 establish IBGP sessions with ASBR1 to ensure the routes can be transmitted in AS1. In AS2, PE3 and PE4 establish IBGP session with ASBR2.

Due to the maintenance of VPN routes is only done by PEs. ASBRs cannot know whether the PEs' ability to handle VPN routes has reached the upper limit or not, so it needs the RD-ORF to control the number of routes.

Assume that PE1 - PE4 can transmit VPN routes through the network architecture shown in Figure 4. When the VRF of VPN1 in PE1 overflows, the RD-ORF mechanism will be implemented as follows:

- 1) PE1 will check and find out the main source of VPN routes in this VRF is PE3. Then, PE1 will resolve the corresponding RD from BGP UPDATE message, and generate a BGP ROUTE-REFRESH message contains an RD-ORF entry, and send it to ASBR1.

- 2) When ASBR1 receives the ROUTE-REFRESH message, it checks whether it receives the latest RD-ORF entry. If not, ASBR1 will discard the entry; Otherwise, ASBR1 will add the RD-ORF entry into its Adj-RIB-out.

Before sending a VPN route toward PE1, RR will check its Adj-RIB-out and find there is a filter associated with RD1. Then, ASBR1 will stop sending that VPN route.

Besides, ASBR1 will locally determine if it needs to send an RD-ORF entry to ASBR2. The judgment criteria refers to [Section 5.1.2](#).

3) If ASBR2/PE3 receives the RD-ORF entry, it will repeat the above process.

When the RD-ORF mechanism is triggered, network operators need to manually configure the network to return to resume normal operation. The withdraw of RD-ORF entries refers to [Section 5.2](#).

In Option AB cross-domain scenario, ASBRs maintain a VRF for a VPN. However, due to VPN routes in all VRFs use the same BGP session, ASBRs cannot prevent the overflow of a certain VRF by breaking down a BGP session. The operation process of RD-ORF is similar to that in Option B scenario.

6.2. Application in Option C cross-domain scenario

The Option C cross-domain scenario is shown in Figure 5:




Figure 5: The Option C cross-domain scenario

In this scenario, PE1 and PE2 are responsible for maintaining VPN routing information in AS1 and AS2. In order to reduce the complexity that full-mesh brings to the network, RR1 and RR2

establish MP-EBGP session to transmit labeled routes. In AS1, PE1 and ASBR1 establish IBGP session with RR1 to ensure the routes can be transmitted in AS1. In AS2, PE2 and ASBR2 establish IBGP session with RR2.

Due to the maintenance of VPN routes is only done by PEs. RRs cannot know whether the PEs' ability to handle VPN routes has reached the upper limit or not, so it needs the RD-ORF to control the number of routes.

The operating mechanism of RD-ORF is similar to the description in [Section 6.1](#).

7. Security Considerations

A BGP speaker will maintain the RD-ORF entries in Adj-RIB-out, this behavior consumes its memory and compute resources. To avoid the excessive consumption of resources, [\[RFC5291\]](#) specifies that a BGP speaker can only accept ORF entries transmitted by its interested peers.

8. IANA Considerations

This document defines a new Outbound Route Filter type - Route Distinguisher Outbound Route Filter (RD-ORF). The code point is from the "BGP Outbound Route Filtering (ORF) Types". It is recommended to set the code point of RD-ORF to 66.

9. Acknowledgement

Thanks Robert Raszuk, Jim Uttaro, Jakob Heitz, Jeff Tantsura, Rajiv Asati, John E Drake and Gert Doering for their valuable comments on this draft.

10. Normative References

- [I-D.ietf-bess-evpn-inter-subnet-forwarding]
Sajassi, A., Salam, S., Thoria, S., Drake, J., and J. Rabadan, "Integrated Routing and Bridging in EVPN", [draft-ietf-bess-evpn-inter-subnet-forwarding-11](#) (work in progress), October 2020.
- [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>>.

- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), DOI 10.17487/RFC4360, February 2006, <<https://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, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", [RFC 4684](#), DOI 10.17487/RFC4684, November 2006, <<https://www.rfc-editor.org/info/rfc4684>>.
- [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>>.
- [RFC5291] Chen, E. and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", [RFC 5291](#), DOI 10.17487/RFC5291, August 2008, <<https://www.rfc-editor.org/info/rfc5291>>.
- [RFC5292] Chen, E. and S. Sangli, "Address-Prefix-Based Outbound Route Filter for BGP-4", [RFC 5292](#), DOI 10.17487/RFC5292, August 2008, <<https://www.rfc-editor.org/info/rfc5292>>.
- [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, <<https://www.rfc-editor.org/info/rfc7432>>.

Authors' Addresses

Wei Wang
China Telecom
Beiqijia Town, Changping District
Beijing, Beijing 102209
China

Email: wangw36@chinatelecom.cn

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing, Beijing 102209
China

Email: wangaj3@chinatelecom.cn

Haibo Wang
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing, Beijing 100095
China

Email: rainsword.wang@huawei.com

Gyan S. Mishra
Verizon Inc.
13101 Columbia Pike
Silver Spring MD 20904
United States of America

Phone: 301 502-1347
Email: gyan.s.mishra@verizon.com

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

Email: zhuangshunwan@huawei.com

Jie Dong
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing, Beijing 100095
China

Email: jie.dong@huawei.com

