IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 29, 2021

W. Wang
A. Wang
China Telecom
S. Zhuang
J. Dong
H. Wang
Huawei Technologies

July 28, 2020

Route Distinguisher Outbound Route Filter (RD-ORF) for BGP-4
draft-wang-idr-rd-orf-01

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 infomation directly (e.g. routers in single-domain connect via Route Reflector, or routers in Option B/ 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 January 29, 2021.

Table of Contents

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.

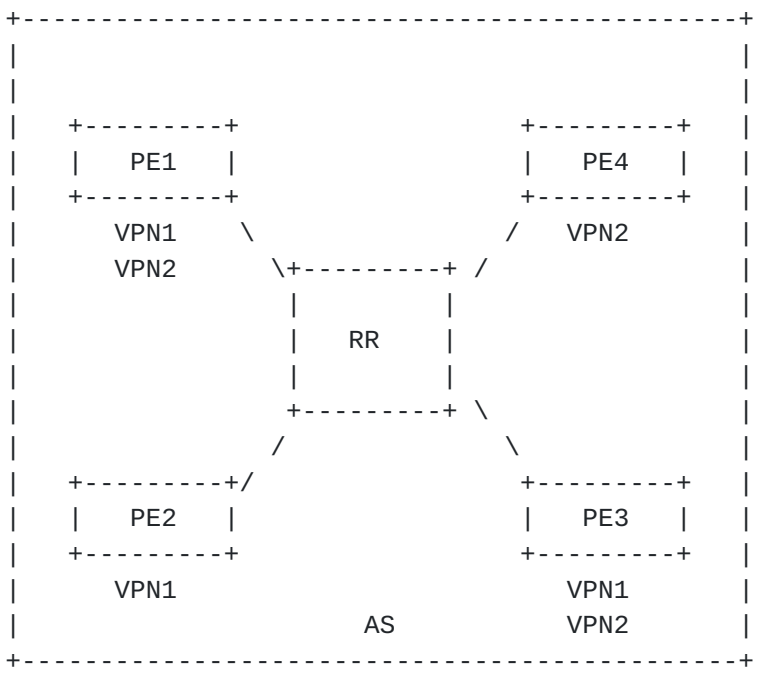

Figure 1: Single-domain scenario

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

Now, there are two 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]

However, RTC can only specify the VPN routes it want, it cannot control the route limit with a specific VRF. Using Address Prefix ORF to filter VPN routes need to pre-configuration, but it is impossible to know which device may overflow in advance.

This draft defines a new ORF-type, called the Route Distinguisher ORF (RD-ORF). Based on RD-ORF, VPN routes of a VPN can be controlled based on source RD and originator. 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 address and RD of VPN routes in this VRF, and send a RD-ORF to its BGP peer that carrys the RD and the source address. 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 infomation directly (e.g. routers in single-domain connect via Route Reflector, or routers in Option B/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:

- RD: Route Distinguisher, defined in [RFC4364]
- ORF: Outbound Route Filter, defined in [RFC5291]
- AFI: Address Family Identifier, defined in [RFC4760]
- SAFI: Subsequent Address Family Identifier, defined in [RFC4760]
- EVPN: BGP/MPLS Ethernet VPN, defined in [RFC7432]
- RR: Router Reflector, provides a simple solution to the problem of IBGP full mesh connection in large-scale IBGP implementation.
- 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:

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

- 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:

- Action (2 bits): the value is ADD, REMOVE or REMOVE-ALL

- Match (1 bit): the value is PERMIT or DENY

- Reserved (5 bits)

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

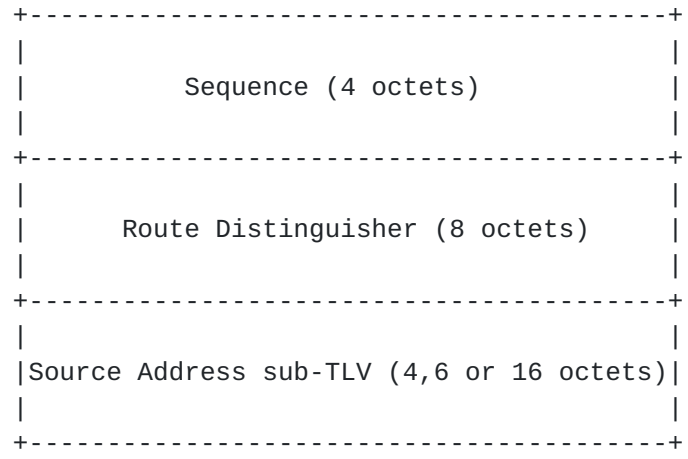

Figure 2: RD-ORF type-specific encoding

- Sequence: identifying the order in which RD-ORF is generated

- Route Distinguisher: distinguish the different user routes. The RD-ORF filters the VPN routes it tends to send based on Route Distinguisher.

- Source Address sub-TLV: the source address is TLV format, which contains the following sub-TLVs:

  * For L3 EVPN case, Gateway IP Address in EVPN RT-5 (IP Prefix Advertisement Route) can be used as source address.

    Type = 1, Length = 4 octets, value = IPv4 Gateway IP Address.

    Type = 2, Length = 16 octets, value = IPv6 Gateway IP Address.

* For L2 EVPN case, the source address can be extracted from EVPN Router's MAC Extended Community (as defined in Section 8.1 of [I-D.ietf-bess-evpn-inter-subnet-forwarding]).

  Type = 3. Length = 6 octets, value = the value field of EVPN Router's MAC Extended Community.

* For MPLS VPN case, the source address can be extracted from the Route Origin Community in BGP (as defined in Section 5 of [RFC4360]).

  Type = 4, Length = 6 octets, value = the value field of Route Origin Community.

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:

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

5. Application in single-domain scenarios

In scenario shown in Figure 1, When the VRF of VPN1 in PE1 overflows, it will find out the main source address of VPN routs in this VRF, assuming it is PE3. Then, PE1 will extract PE3's host address from BGP UPDATE message and generate a BGP ROUTE-REFRESH message contains a RD-ORF entry, and send it to RR. The entry consists of the following fields:

- AFI is set to IPv4 , IPv6 or L2 VPN
- SAFI is set to "MPLS-labeled VPN address" or "BGP EVPN"
- When-to-refresh is set to IMMEDIATE
- ORF Type is set to RD-ORF

- Length of ORF entries depends on the type of Source Address sub-TLV (21, 23 or 33 octets)
- Action is set to ADD
- Match is set to DENY
- Sequence is set to 1
- Route Distinguisher is set to RD1
- Source Address sub-TLV is set to PE3's host address

It noted that for a RD, the sequence of the first RD-ORF is equal to 1. When a PE needs to send a second RD-ORF entry associated with the same RD, the RD-ORF sequence SHOULD increment.

When RR receives the ROUTE-REFRESH message, it checks the <AFI/SAFI, ORF-Type> to determine whether it is willing to receive the entry. It also checks the Sequence, Route Distinguisher and Source Address sub-TLV, to find whether it received the latest entry or not. If the above conditions are not all met, RR will discard the entry; otherwise, RR will add the RD-ORF entry into its Adj-RIB-out, and regenerate a BGP ROUTE-REFRESH message to send this RD-ORF entry to PE3.

After receiving this ROUTE-REFRESH message that carries a RD-ORF entry, PE3 will repeat the above process to check if it is willing to receive this message. If not, PE3 will discard it; Otherwise, PE3 will add the RD-ORF entry into its Adj-RIB-out.

Before sending a VPN route (the RD is equal to RD1) toward PE1, 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.

When the VRF of VPN1 in PE1 no longer overflows, it will send RR a BGP ROUTE-REFRESH message encoded as following:

- AFI is set to IPv4 , IPv6 or L2 VPN
- SAFI is set to "MPLS-labeled VPN address" or "BGP EVPN"
- When-to-refresh is set to IMMEDIATE
- ORF Type is set to RD-ORF

- Length of ORF entries depends on the type of Source Address sub-TLV (21, 23 or 33 octets)
- Action is set to REMOVE
- Match is set to DENY
- Sequence is set to 2
- Route Distinguisher is set to RD1
- Source Address sub-TLV is set to PE3's host address

After receiving the BGP ROUTE-REFRESH message, RR will check whether it is willing to receive this entry. In this process, RR finds that it received the latest entry. Then, RR will remove the associated RD-ORF entry from its Adj-RIB-out.

6. Applications in cross-domain scenarios

6.1. Application in Option B cross-domain scenario

The Option B cross-domain scenario is shown in Figure 3:

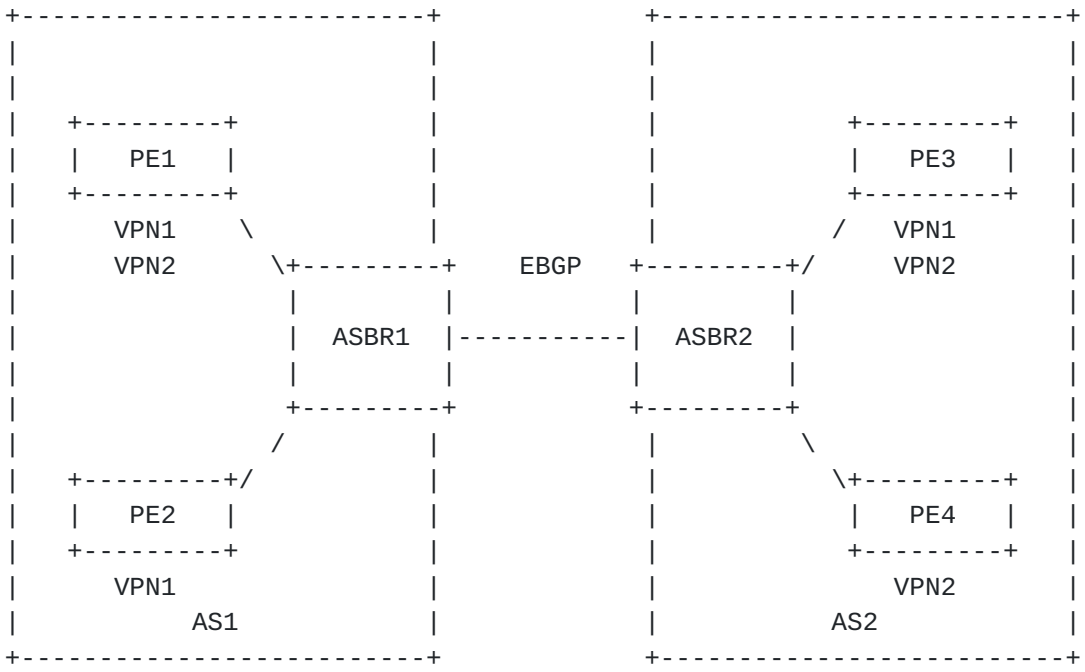

Figure 3: The Option B cross-domain scenario

In this scenario, PE1 - PE4 are responsible for maintaining VPN routing of devices in AS1 and AS2. There is a direct link between

ASBR1 and ASBR2 via EBGP. In AS1, PE1, PE2 and ASBR1 establish IBGP sessions 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 2. When the VRF of VPN1 in PE1 overflows, it will check and find out the main source address of VPN routs in this VRF is PE3. Then, PE1 will extract the next hop address associated with PE3 from BGP UPDATE message and generate a BGP ROUTE-REFRESH message contains a RD-ORF entry, and send it to ASBR1. The entry consists of the following fields:

- AFI is set to IPv4 , IPv6 or L2 VPN
- SAFI is set to "MPLS-labeled VPN address" or "BGP EVPN"
- When-to-refresh is set to IMMEDIATE
- ORF Type is set to RD-ORF
- Length of ORF entries depends on the type of Source Address sub-TLV (21, 23 or 33 octets)
- Action is set to ADD
- Match is set to DENY
- Sequence is set to 1
- Route Distinguisher is set to RD1
- Source Address sub-TLV is set to PE3's host address

When ASBR1 receives the ROUTE-REFRESH message, it checks the <AFI/ SAFI, ORF-Type> to determine whether it is willing to receive the entry. Then ASBR1 will check the Sequence. Route Distinguisher and Source Address sub-TLV to check whether it receives the latest RD-ORF entry. If the above conditions are not all met, ASBR1 will discard the entry; otherwise, ASBR1 will add the RD-ORF entry into its Adj-RIB-out and regenerate a ROUTE-REFRESH message carries the RD-ORF entry to send it to ASBR2.

After receiving the RD-ORF entry, ASBR2 will repeat the above process. If necessary, the RD-ORF entry will be transmitted toward PE3. PE3 will receive it and add the associated entries into its Adj-RIB-out.

Before sending a VPN route (the RD is equal to RD1) toward PE1, PE3 will check its Adj-RIB-out and find the RD-ORF entry prevent it from sending VPN route which carries RD1 to ASBR2. Then, it will stop sending that VPN route.

If PE1 can re-receive the route entries, it will send a ROUTE-REFRESH message to ASBR1, carrying a RD-ORF entry consists of the following fields:

- AFI is set to IPv4 , IPv6 or L2 VPN
- SAFI is set to "MPLS-labeled VPN address" or "BGP EVPN"
- When-to-refresh is set to IMMEDIATE
- ORF Type is set to RD-ORF
- Length of ORF entries depends on the type of Source Address sub-TLV (21, 23 or 33 octets)
- Action is set to REMOVE
- Match is set to DENY
- Sequence is set to 2
- Route Distinguisher is set to RD1
- Source Address sub-TLV is set to PE3's host address

When ASBR1 receives the ROUTE-REFRESH message, it checks the <AFI/ SAFI, ORF-Type> to determine whether it is willing to receive the entry. Then ASBR1 will check the Sequence. Route Distinguisher and Source Address sub-TLV to check whether it receives the latest RD-ORF entry. If the above conditions are not all met, ASBR1 will discard the entry; otherwise, ASBR1 will remove the RD-ORF entry from its Adj-RIB-out and regenerate a ROUTE-REFRESH message carries the RD-ORF entry to send it to ASBR2.

After receiving the RD-ORF entry, ASBR2 will repeat the above process. If necessary, the RD-ORF entry will be transmitted toward PE3. PE3 will receive it and remove the associated entries from its Adj-RIB-outs.

Before sending a VPN route (the RD is equal to RD1) toward PE1, PE3 will check its Adj-RIB-out and find there is no filter associated with RD1. Then, it will send that VPN route.

6.2. Application in Option C cross-domain scenario

The Option C cross-domain scenario is shown in Figure 4:

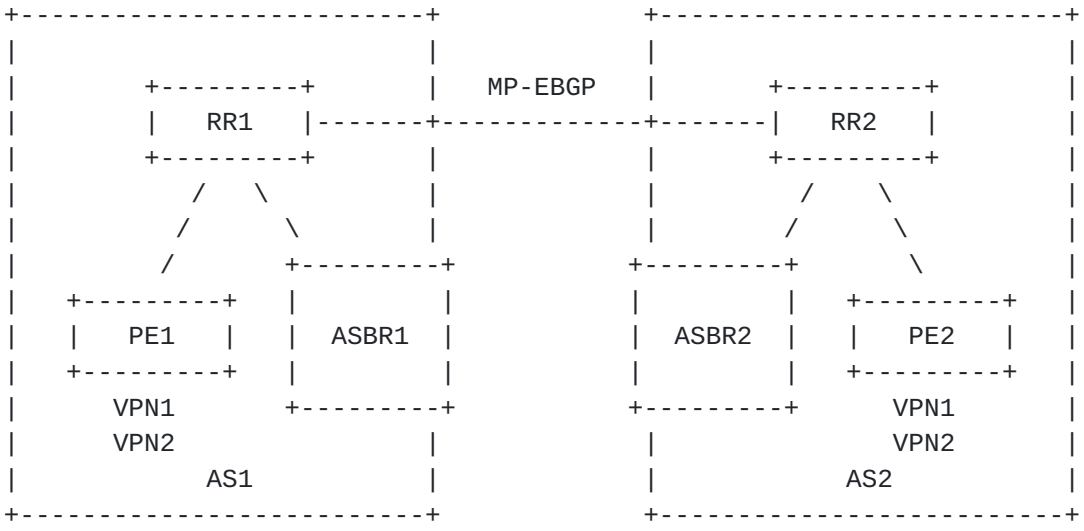

Figure 4: The Option C cross-domain scenario

In this scenario, PE1 and PE2 are responsible for maintaining VPN routing of devices 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.

IANA is requested to allocate one code point for Source Address sub-TLV for RD-ORF.

This document defines the following new RD-ORF sub-TLV types, which should be reflected in the Source Address sub-TLV for RD-ORF Code Point registry:

Type	Description
1	IPv4 L3EVPN Source Address TLV
2	IPv6 L3EVPN Source Address TLV
3	L2EVPN Source Address TLV
4	MPLS VPN Source Address TLV

9. 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-09 (work in progress), June 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

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

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

Email: rainsword.wang@huawei.com