

Updates on VPN Prefix ORF for BGP-4

[draft-ietf-idr-vpn-prefix-orf-06](#)

Wei Wang (China Telecom)

Aijun Wang (China Telecom)

Haibo Wang (Huawei)

Gyan Mishra (Verizon)

Shunwan Zhuang (Huawei)

Jie Dong (Huawei)

IETF 119, March 2024

Updates compared to -v00

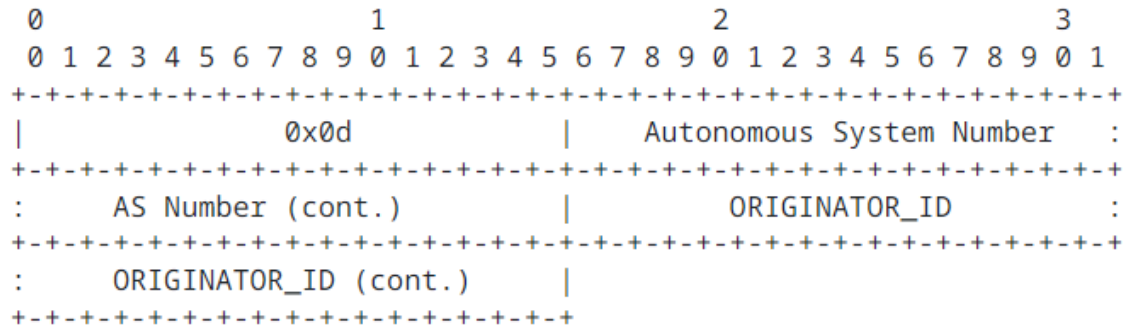
- ✓ Newly defined Source PE Extended Community to carry the source.
- ✓ Refined the processing method of the received ORF entry by the receiver.
- ✓ The length of the VRF Prefix limit was changed to 4 bytes.
- ✓ Added relevant implementation content.
- ✓ Added details of operation process of VPN Prefix ORF mechanism on sender.
- ✓ Refined the IANA Consideration
- ✓ Modified some typos.

Main update 1

We usually use NEXT_HOP to identify the source, but it may **not** useful in the following scenarios:

- a PE may have multiple addresses so that its BGP peer will receive several different NEXT_HOP from the same source.
- In Option B inter-domain scenario, the ASBR will change the NEXT_HOP.

ORIGINATOR_ID is a non-transitive attribute generated by RR to identify the source, but ORIGINATOR_ID cannot be advertised outside the local AS. To cover the above scenarios, we define a new Extended Community: **Source PE Extended Community(SPE EC)** to transmit the identifier of source. Its value can be set by **source PE/RR/ASBR**. Once set and attached with the BGP UPDATE message, **its value should not be changed along the advertisement path.**



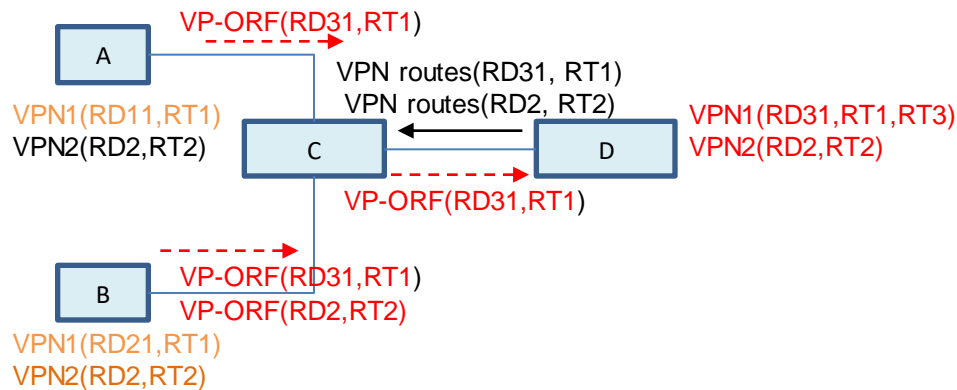
The format of SPE EC

For the RR/ASBR, it should perform as following:

- Check the existence of the SPE EC. If it exists, does not change it.
- If SPE EC does not exist, check the existence of ORIGINATOR_ID. If it exists, put it into SPE EC.
- If ORIGINATOR_ID does not exist, put the router-id of source PE into SPE EC.

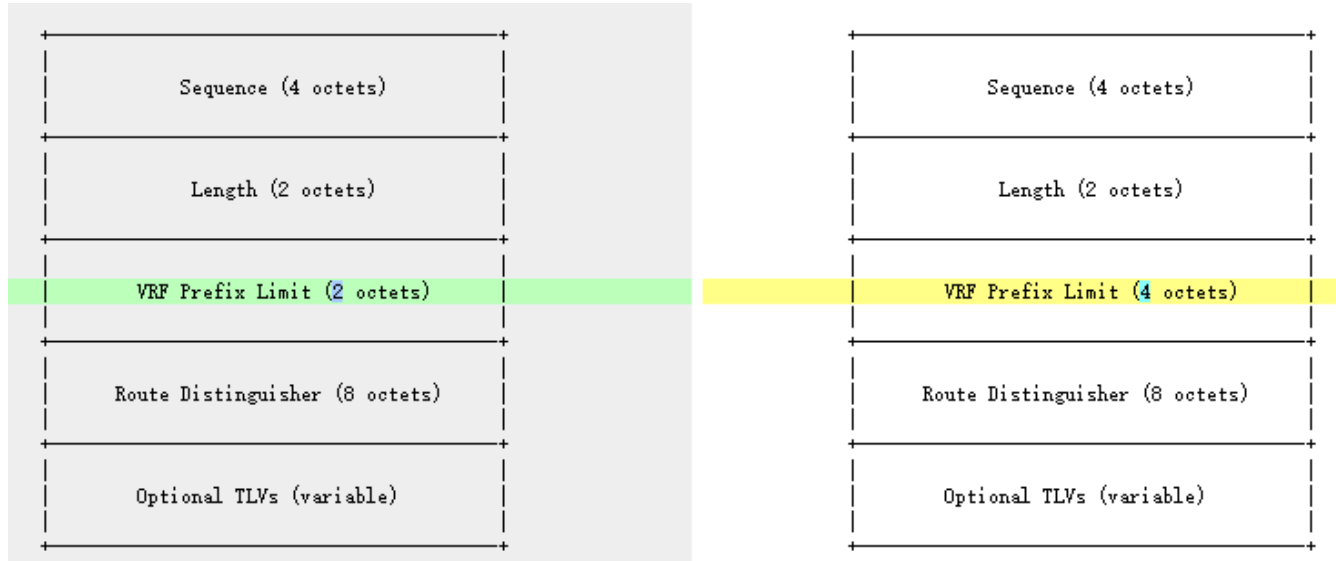
Main update 2

In Section 7, the processing method of the received ORF entry by the receiver was further refined, and some new contents was added on how the receiver determines whether to send the received ORF entry to the upstream device.



- D sent excessive VPN routes (RD31, RT1) and VPN routes(RD2, RT2), which caused **VPN1 on A** and **VPN1, VPN2 on B** overflowed.
- As the VPN Prefix ORF mechanism is triggered, A sent **VP-ORF(RD31,RT1)** to C, and B sent **VP-ORF(RD31,RT1)** and **VP-ORF(RD2,RT2)** to C.
- Due to C only has 2 downstream devices, and both of them sent **VP-ORF(RD31,RT1)** to C. After checking the received VP-ORF entries, C should generate a VP-ORF(RD31,RT) and send it to its upstream device D.
- VP-ORF(RD2,RT2) should not be sent to D, because C only received it from B.

Main update 3



The length of the VRF Prefix limit was changed to **4 bytes**. Its length was originally set to 2 bytes, which can only support a maximum of **65535** private network routes, far less than the actual number of routes supported by the device.

Main update 4

We added relevant implementation content to section 10.2.

Currently, **H3C** has implemented some VPN Prefix ORF mechanism related functions as follows:

- By configuring VRF Prefix limit and quota, achieve the use of RD and Source PE to control VPN routing.
- Generating, transmitting and processing Type 1 and Type 2 Source PE TLV.
- Using the Offending VPN routes process method to revoke all routes.

Besides, we also implemented the following functions based on the open-source BGP implementation (FRR):

- VPN Prefix ORF mechanism triggered based on VRF limit in intra-domain and inter-domain scenarios.
- RD based VPN routing filtering in intra-domain and inter-domain scenarios.

Main update 5

We updated the content of IANA Consideration based on the comments from IANA.

```
under "Border Gateway Protocol (BGP) Parameters"
Registry: "VPN Prefix ORF TLV"
Registration Procedure(s): First Come, First Served
Value range:0-255, value 0 is reserved.
+-----+-----+-----+
| Registry          | Type   | Meaning          |
+-----+-----+-----+
| IPv4 Source PE TLV | 1(suggested) | IPv4 address for source PE. |
+-----+-----+-----+
| IPv6 Source PE TLV | 2(suggested) | IPv6 address for source PE. |
+-----+-----+-----+
| Source PE Extended | 3(suggested) | Source PE Extended |
| Community TLV      |             | Community for source PE |
+-----+-----+-----+
| Route Target TLV  | 4(suggested) | Route Target of the |
|                   |             | offending VPN routes |
+-----+-----+-----+
```

```
Under "BGP Transitive Extended Community Types:"
Registry: "Source PE Extended Community" type
0x0d(suggested)          Source PE Extended Community
```

We need feedbacks about:

- 1) Suggestions for the updates.
- 2) There are already two related implementations, can we consider WG Last Call?
- 3) There are already two related implementations, can we apply to change this draft to Standard Track?

wangw36@chinatelecom.cn
wangaj3@chinatelecom.cn
rainsword.wang@huawei.com
gyan.s.mishra@verizon.com
zhuangshunwan@huawei.com
jie.dong@huawei.com
IETF119