

Updates on VPN Prefix ORF for BGP-4

[draft-ietf-idr-vpn-prefix-orf-08](#)

Wei Wang (China Telecom)

Aijun Wang (China Telecom)

Haibo Wang (Huawei)

Gyan Mishra (Verizon)

Jie Dong (Huawei)

IETF 121, March 2024

Updates compared to -v06


- ✓ Modified some descriptions according to Keyur's suggestions, especially the content in Section 4 and 7.
- ✓ Modified the format of SPE EC.
- ✓ Newly defined Source AS TLV.
- ✓ Deleted the Section 10.4 “Results of Experiments”.
- ✓ Added Section 13 “Contributors”.
- ✓ Changed “Source PE Extended Community TLV” to “Source PE Identifier TLV”.

Main update 1

We refined the contents of whole draft based on the comments from Keyur, especially in Section 4 and 7.

On PE, each VRF has a prefix limit. When the PE receives VPN routes from its BGP peer, due to the received VPN routes may belong to different VPN and carry the corresponding RDs, the PE should extract the VPN route information from BGP UPDATE message which contains VPN routes related to BGP optimal routing. PE can determine the target VRFs of the received VPN route based on the RT of the VPN route and the RT-import of VRFs. Then, the PE should sequentially determine whether each target VRF will exceed the limit after importing the received VPN routes. If a target VRF exceeds the limit which is caused by the VPN routes carrying a certain RD and the other target VRFs have not overflow, PE should not trigger the VPN Prefix ORF mechanism, and only performs VPN route filtering for the target VRF and stop importing VPN routes carrying the specific RD. If a target VRF exceeds the limit and there is no other VRFs need these VPN routes, the PE should trigger the VPN Prefix ORF mechanism and send a BGP ROUTE-REFRESH message contains the corresponding VPN Prefix ORF entry to its peer, which will generate a VPN routes filtering strategy for the VRF. And if the "Offending VPN routes process method" bit is 1, the receiver of VPN Prefix ORF entry should withdraw the extra VPN routes according to the value of VRF Prefix Limit, RD, RT and information in optional TLVs in the entry, and stop sending the corresponding VPN routes to the sender. If the target VRF no longer exceeds the limit, the relevant VPN routing filtering strategy needs to be deleted.

When importing VPN routes to a VRF, it is necessary to determine whether there is a VPN routes filtering strategy on the PE for that VRF. If a VPN routes filtering strategy for a certain VRF which is overflow already exists on the PE, VPN routes that comply with this strategy should not be imported, and should be discarded.



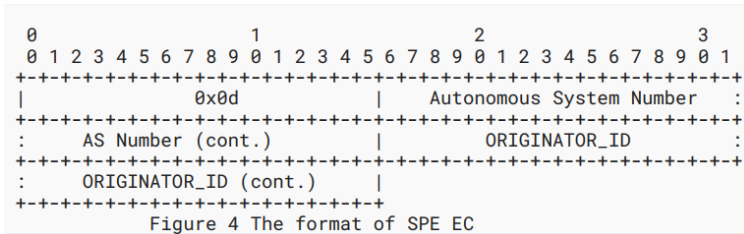
```
S01. PE extracts the VPN route information from BGP UPDATE message, and
determines the target VRFs for the received VPN routes based on the
RT of the VPN route and the RT-import of VRFs.
S02. If (a target VRF exceeds the limit which caused by VPN routes carry a
certain RD) {
S03.     If (not all the other target VRFs have overflowed) {
S04.         PE should not trigger the VPN Prefix ORF mechanism, and only
performs VPN route filtering for the target VRF, stopping
the import of VPN routes carrying the specific RD.
S05.     } else {
S06.         PE should trigger the VPN Prefix ORF mechanism and send a BGP
ROUTE-REFRESH message contains the corresponding VPN Prefix
ORF entry to its peer.
S07.     }
S08. }
```

The detail procedures for “No quota value is set on PE”

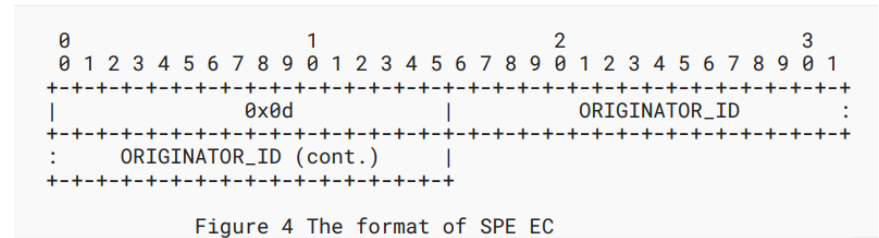
We have changed the description of the mechanism to the form shown in the picture on the right, making it more concise and readable.

Main update 2

In Section 5, the AS number field in Source PE Extended Community (SPE EC) is deleted. SPE EC only carries the identifier of source PE.



Before



After

In Section 6.2, we defined a **Source AS TLV** to carry the AS number of source PE:

- Type = 4
- Length = 4 octets
- Value = the value of Source AS in Source AS Extended Community as defined in RFC6514.

Other updates

- ❑ Section 10.4 “Results of Experiments” was deleted. The results of experiments and implementation will be posted on Protocol Implementation Reports .
- ❑ Due to the limit on the number of authors, Shunwan Zhuang was moved to the Section 13 “Contributors”.
- ❑ To avoid misunderstanding, "Source PE Extended Community TLV" was changed to "Source PE Identifier TLV".

Source PE Identifier TLV	3(suggested)	ORIGINATOR_ID in Source PE Extended Community for source PE
--------------------------	--------------	---

Source PE Identifier TLV is a TLV of VPN Prefix ORF, to carry the ORIGINATOR_ID in Source PE Extended Community for source PE.

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

Source PE Extended Community is a newly defined transitive extended community, to transmit the identifier of source between different AS.

Updates on Protocol Implementation Reports

We added relevant implementation content to Protocol Implementation Reports (<https://wiki.ietf.org/group/idr/implementations/draft-ietf-idr-vpn-prefix-orf>).

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.

We need feedbacks about:

Any suggestions or comments are welcome!

wangw36@chinatelecom.cn
wangaj3@chinatelecom.cn
rainsword.wang@huawei.com
gyan.s.mishra@verizon.com
zhuangshunwan@huawei.com
jie.dong@huawei.com
IETF121