

SRv6 Service SID Flag Extension for Multihoming SRv6 BGP Services

draft-liu-bess-multihome-srv6-service-sid-flag-02

Yisong Liu (China Mobile)
Changwang Lin (New H3C Technologies)
Mengxiao Chen (New H3C Technologies)
Yao Liu (ZTE)

IETF-118, November 2023

Introduction

In the multi-homing scenarios:

- Multiple egress PEs can implement local protection when the CE side link fails;
- Ingress PE can send traffic to egress PE in load balancing mode or active/standby mode.

To meet different requirements, the egress PE may need to allocate and advertise multiple service SIDs for the same service.

This draft describes the use cases for two types of service SID, No-Further-FRR service SID and anycast service SID, and defines new flags for them when advertising through BGP messages.

Updates after IETF-116

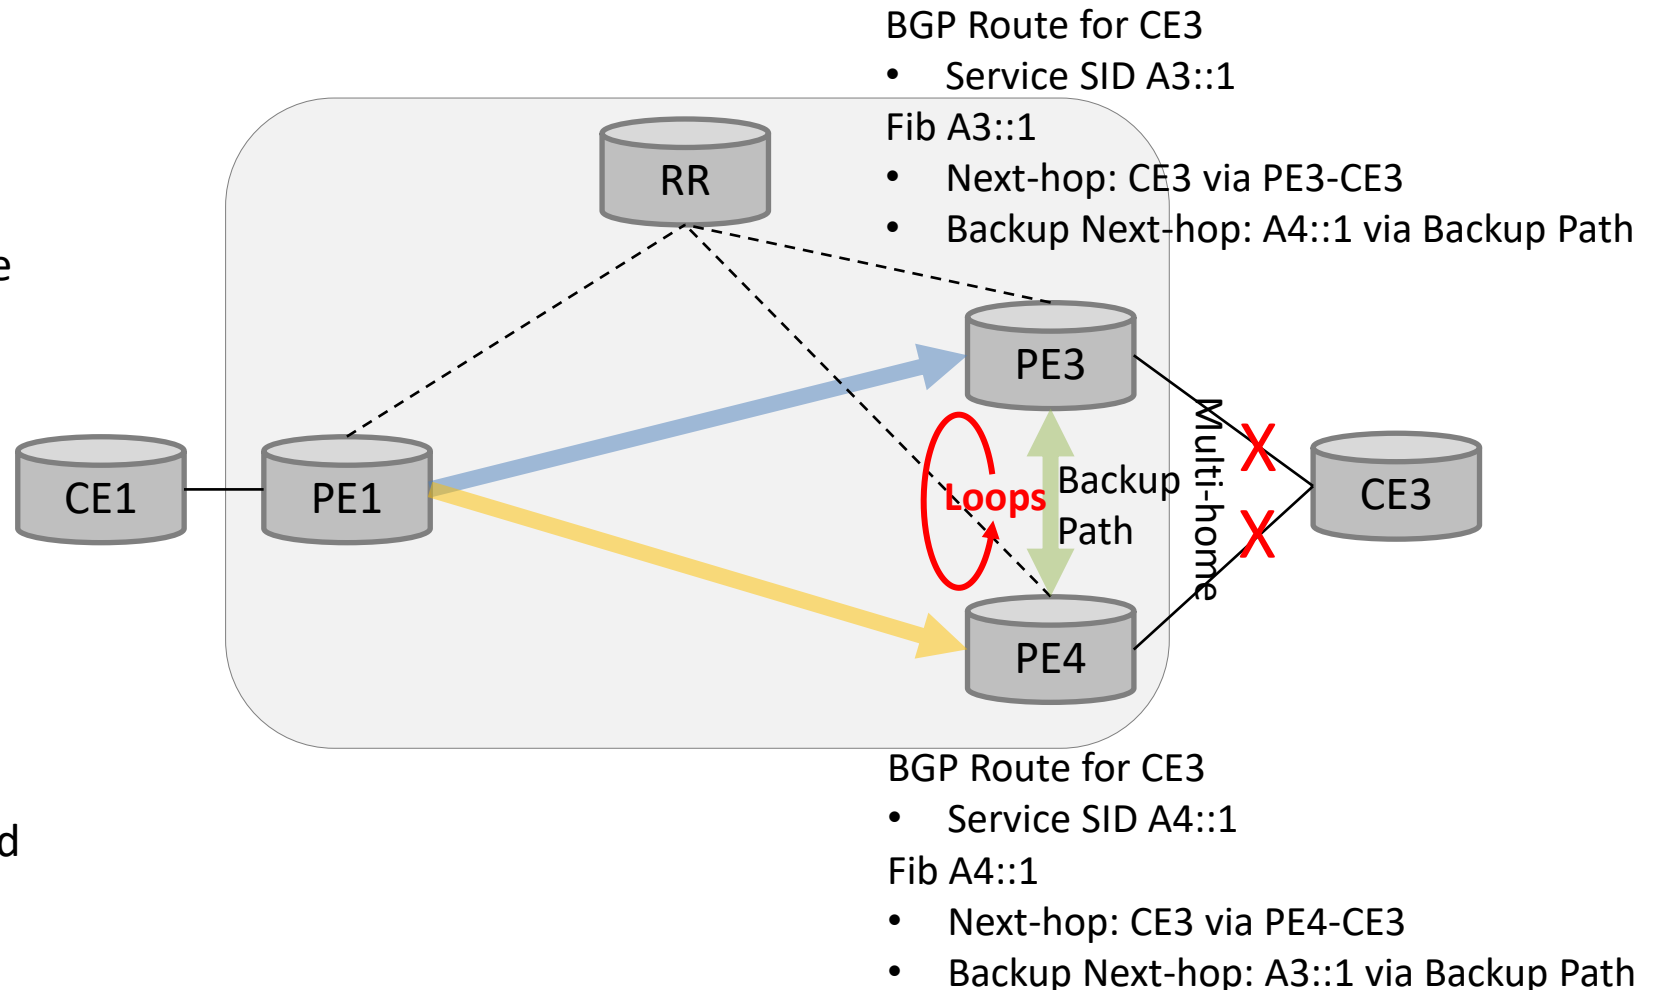
- Revise the draft according to received comments.
- Change "Bypass SID" to "No-Further-FRR SID", aligning with existing works on MPLS data plane (draft-kompella-mpls-nffrr).
- Add Section 2.1 "Consideration for EVPN Single-Active Mode".

Use Case 1: Egress Fast Reroute

PE3 and PE4 establish a backup path between them and use it as the protection of PE-CE link failure.

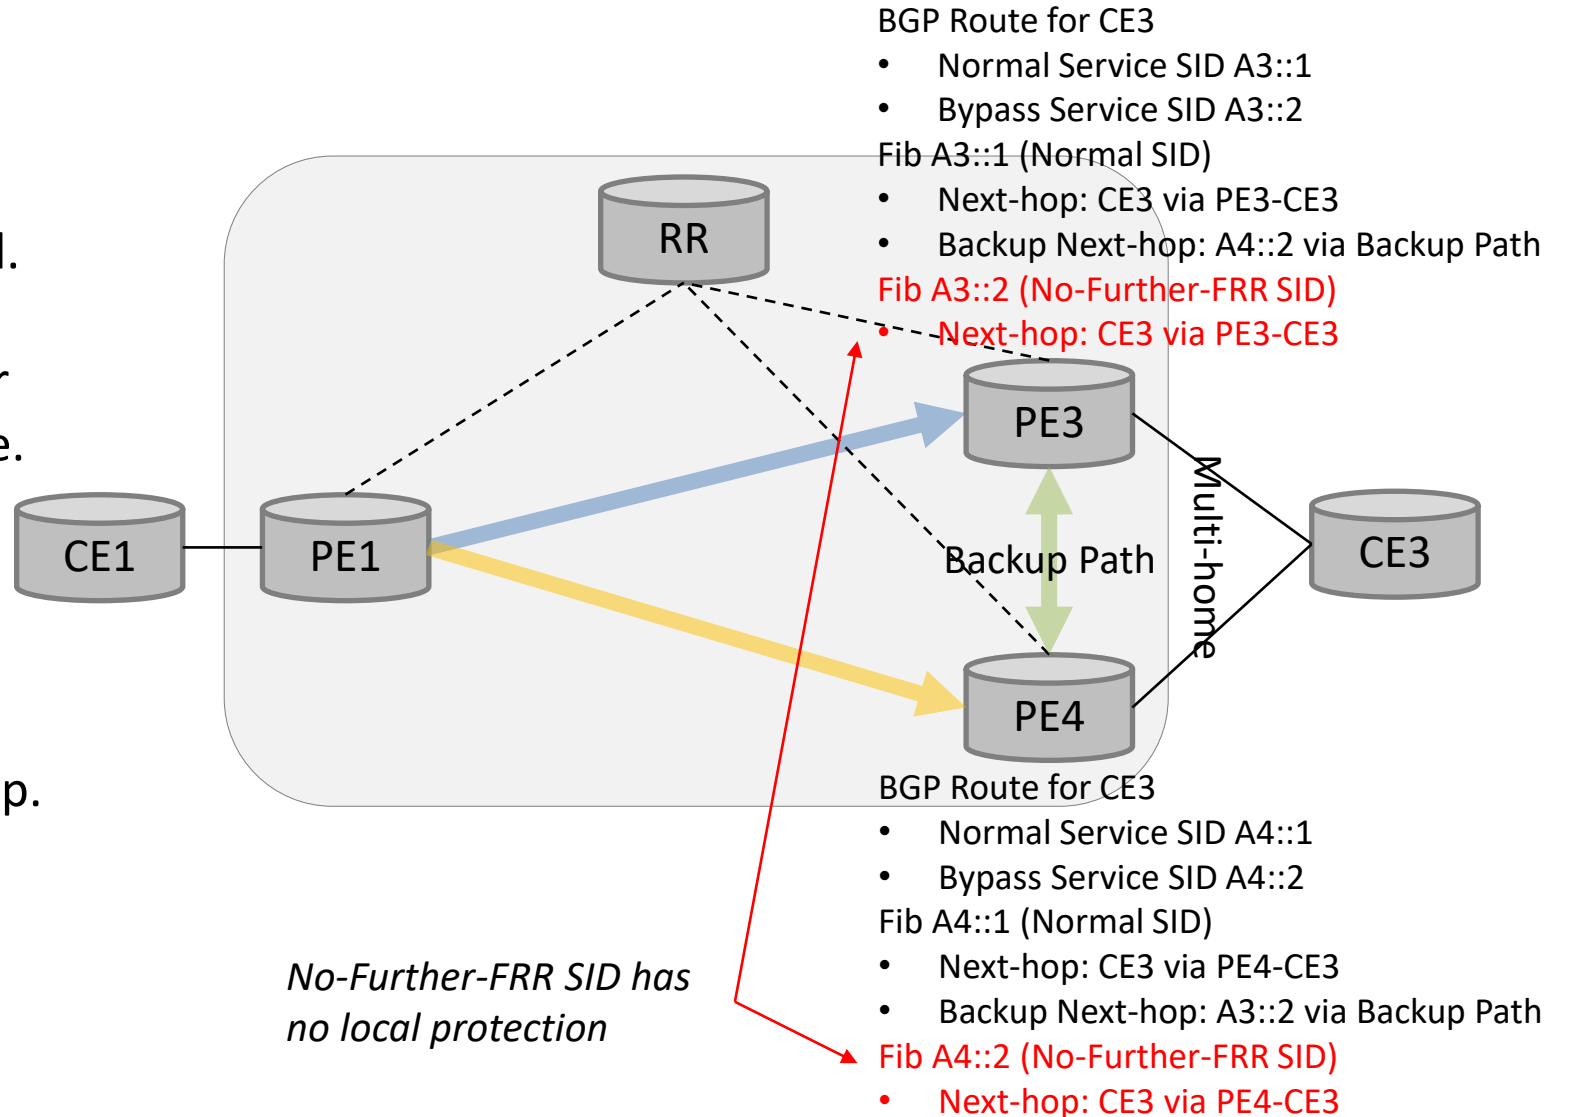
When PE3-CE3 and PE4-CE3 fail at the same time:

1. PE1 forward VPN traffic to PE3 using A3::1;
2. PE3 decapsulate A3::1, re-encapsulate A4::1, and forwards traffic to PE4;
3. PE4 decapsulate A4::1, re-encapsulate A3::1, and forwards traffic to PE3;
4. Traffic is Looping between PE3 and PE4 until routing convergence.



Solution: No-Further-FRR Service SID

- No-Further-FRR Service SID has no local protection. When PE-CE link fails, packets will be dropped. It can avoid routing loops between PE2 and PE3 when their CE side links fail at the same time.
- Egress PE advertises both the normal Service SID and No-Further-FRR Service SID to RR.
- Egress PE uses each other's No-Further-FRR Service SID as backup.

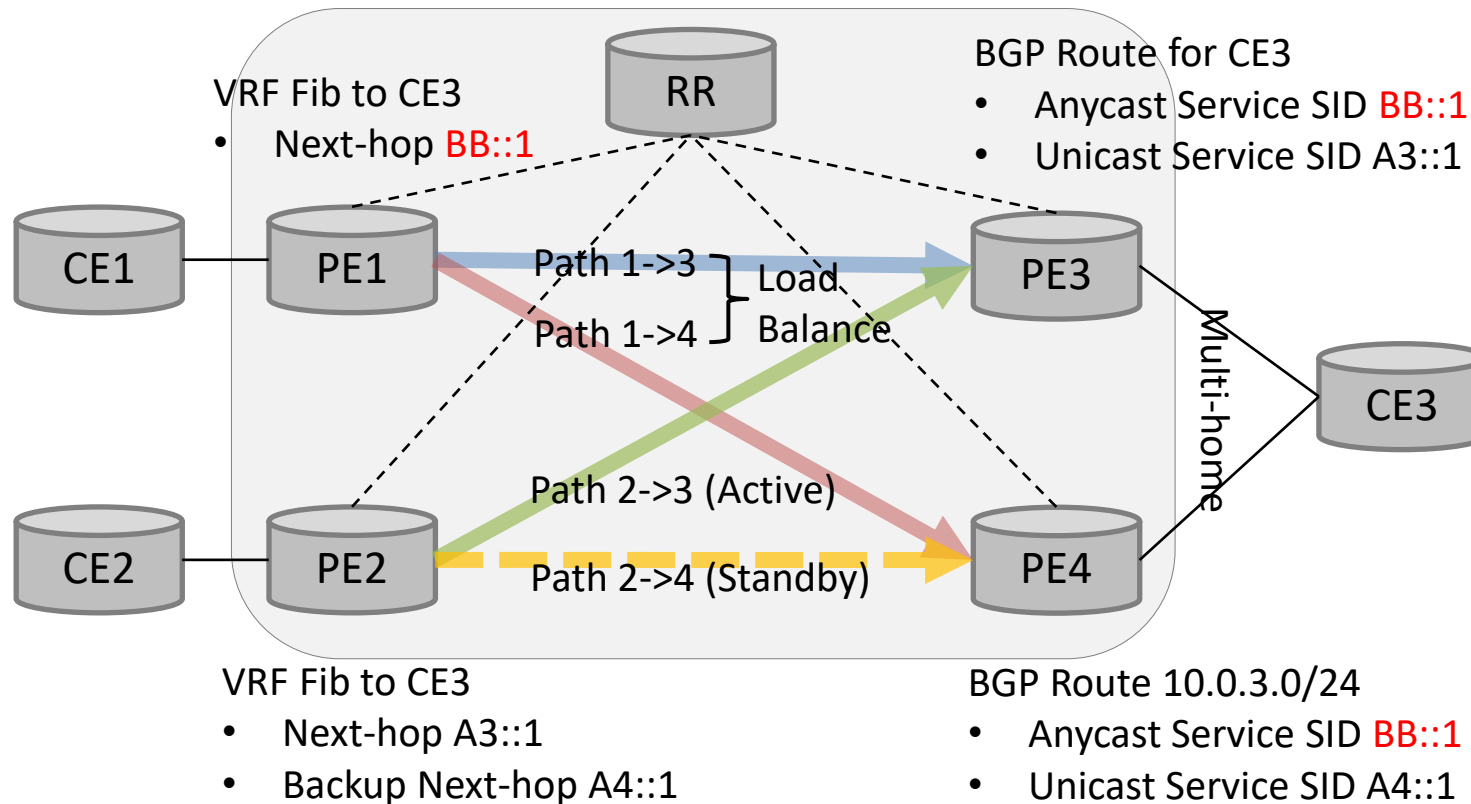


Consideration for EVPN Single-Active Mode

- The EVPN services include Designated Forwarder (DF) election procedure. In Single-Active mode, only DF is allowed to forward unicast traffic.
- The processing of the No-Further-FRR SID should apply an override to EVPN DF-Election and bypass the local blocking state on the AC, until EVPN control plane reconverges.

Use Case 2: Anycast Load Balancing

** IGP has Anycast-flag for SRv6 locator. However, the IGP Anycast-flag may be lost due to summarization.*



Problem:

- Ingress PE1 and PE2 have different traffic steering policies.

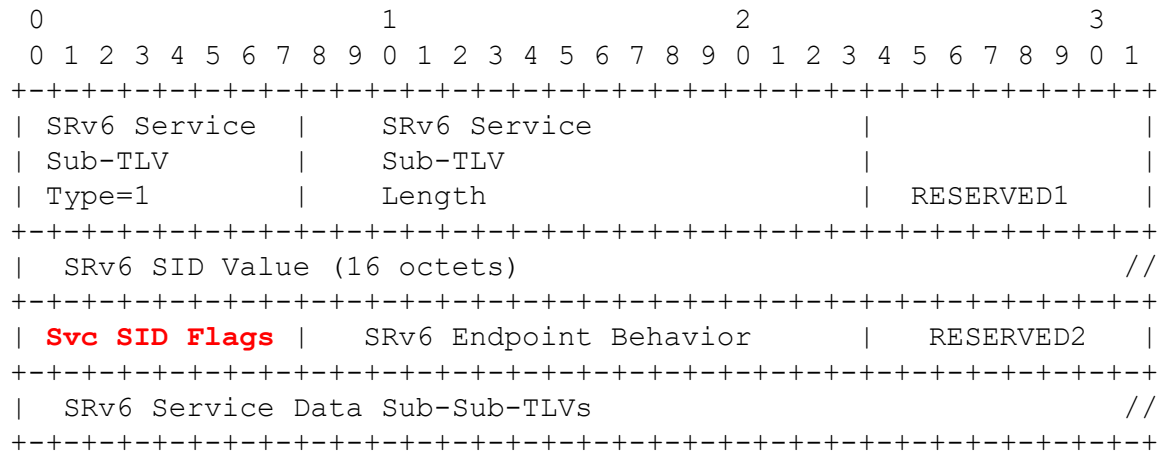
Solution:

- Egress PE3 and PE4 advertises both the anycast service SID and unicast service SID through BGP.
- Ingress PE1 and PE2 use different service SIDs to forward traffic.


- PE1 uses the anycast service SID, and the traffic can be forwarded to PE3 and PE4 in a load-balanced manner
- PE2 can deploy VPN FRR, to uses the service SID of PE3 as the primary next hop, and PE4 as the backup next hop

BGP Extensions

[RFC9252] defines the SRv6 SID Information Sub-TLV to carry SRv6 Service SID in BGP messages.



This draft defines two new flags in the SRv6 Service SID Flags field:



Backward Compatibility

About "Multiple SRv6 SIDs" in [RFC9252]:

- When multiple SRv6 SID Information Sub-TLVs are present, the ingress PE SHOULD use the SRv6 SID from the first instance of the Sub-TLV. An implementation MAY provide a local policy to override this selection.

About "SRv6 Service SID Flags" in [RFC9252]:

- Any unknown flags in the SRv6 Service SID Flags field MUST be ignored by the receiver.

When the egress PE advertises multiple service SIDs, the unicast service SID needs to be carried in the first instance of Sub-TLV. When there are PE routers not supporting the new-defined flags, the egress PE may expect those routers to use the first SID and ignore the new-defined flags.

Considerations for Using Service SID Flags

Why do we prefer to use flags rather than defining new behaviors?

- Whether to provide FRR for Service SID is local configuration on the egress node.
- IGP also has Anycast-flag and Backup-flag for SRv6 SID.
- Service SID has various behaviors, such as End.DX4, End.DT4, End.DX6, End.DT6, End.DT46. End.DX2, End.DX2V, End.DT2U, etc. Using SID flags is more simple than defining new ones for each existing behavior.

Next Steps

- Ask for WG adoption.
- Any questions or comments are Welcomed.

Thanks