

# No Further Fast Reroute for SRv6 Service SID

**draft-liu-bess-srv6-service-sid-nffrr-flag-02**

Presenter : Changwang Lin (New H3C Technologies)

Co-Authors : Yisong Liu (China Mobile)

Changwang Lin (New H3C Technologies)

Yao Liu (ZTE)

IETF-121, November 2024

# Introduction

In the multi-homing scenarios:

- Multiple egress PEs can implement local protection when the CE side link fails;

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 No-Further-FRR service SID, and defines new flags for them when advertising through BGP messages.

# Updates after IETF-116 and IETF-118

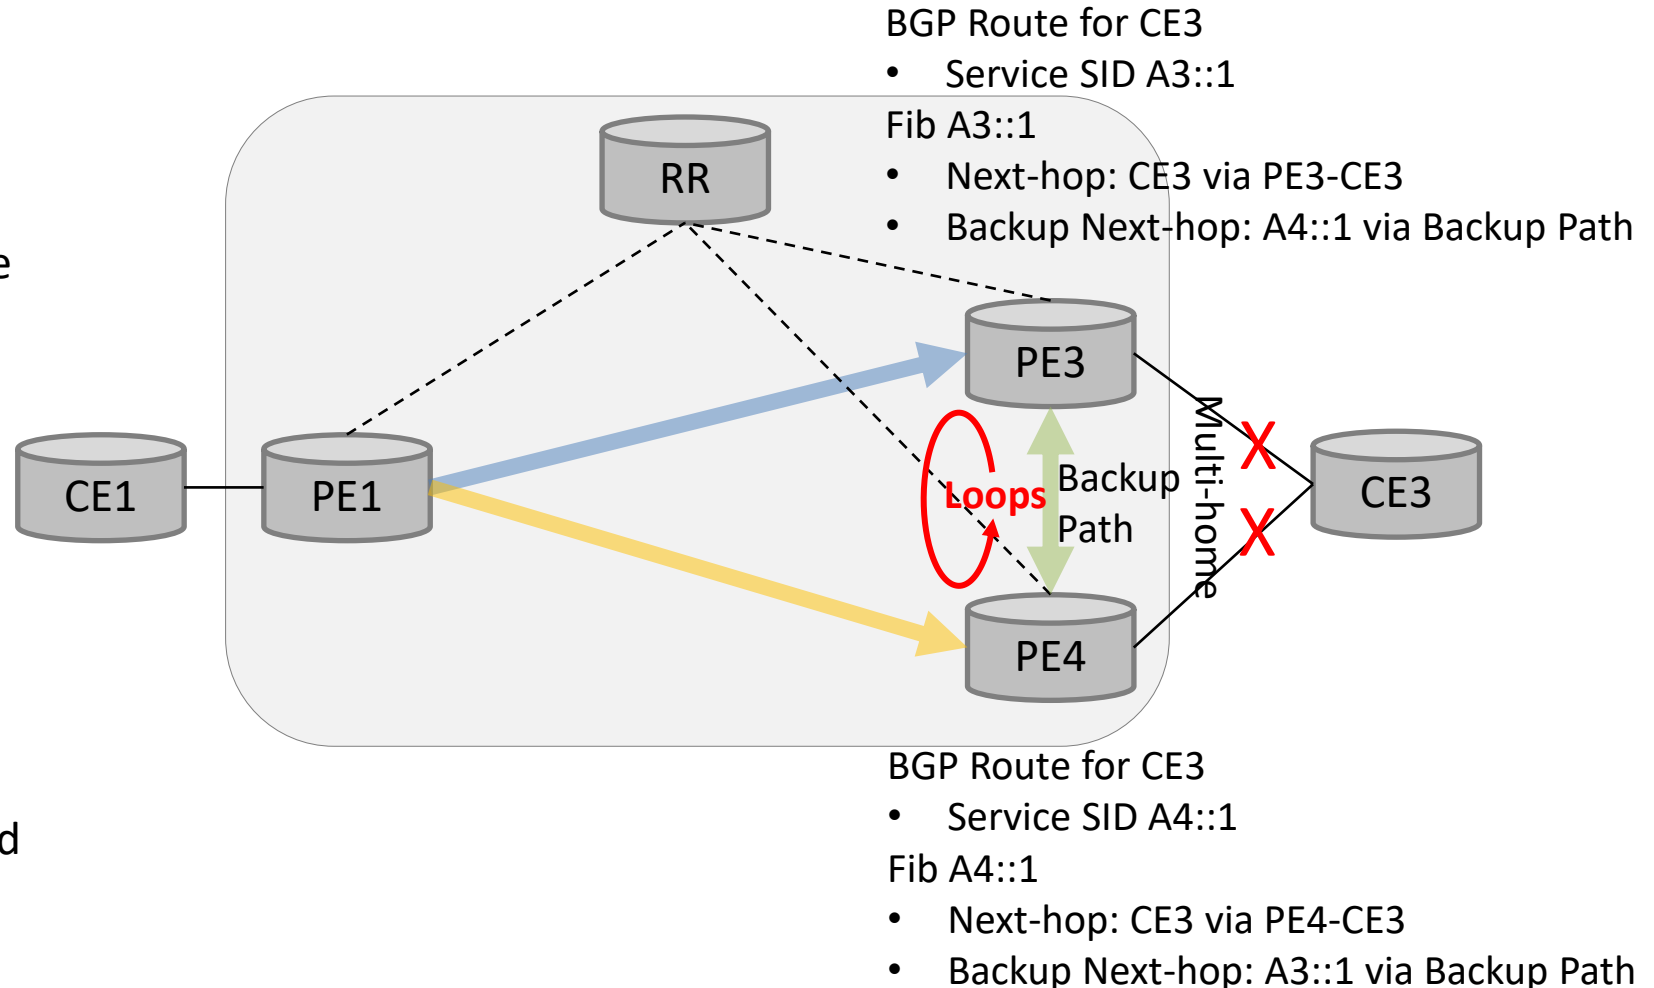
- Presented at IETF-116 and IETF-118
- Revise the draft according to received comments.
- Change "Bypass SID" to "No-Further-FRR SID".
- The draft 'draft-liu-bess-multihome-srv6-service-sid-flag' has been split into two drafts: 'draft-liu-bess-srv6-service-sid-nffrr-flag' and 'draft-liu-bess-srv6-service-sid-anycast-flag-01'.
- This draft focuses on scenarios with no further fast reroute.
- Added the L2VPN scenario; the solution is applicable to both L3VPN and L2VPN.
- Add Section 3.1 "Consideration for EVPN Single-Active Mode".

# Use Case : Egress Fast Reroute

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

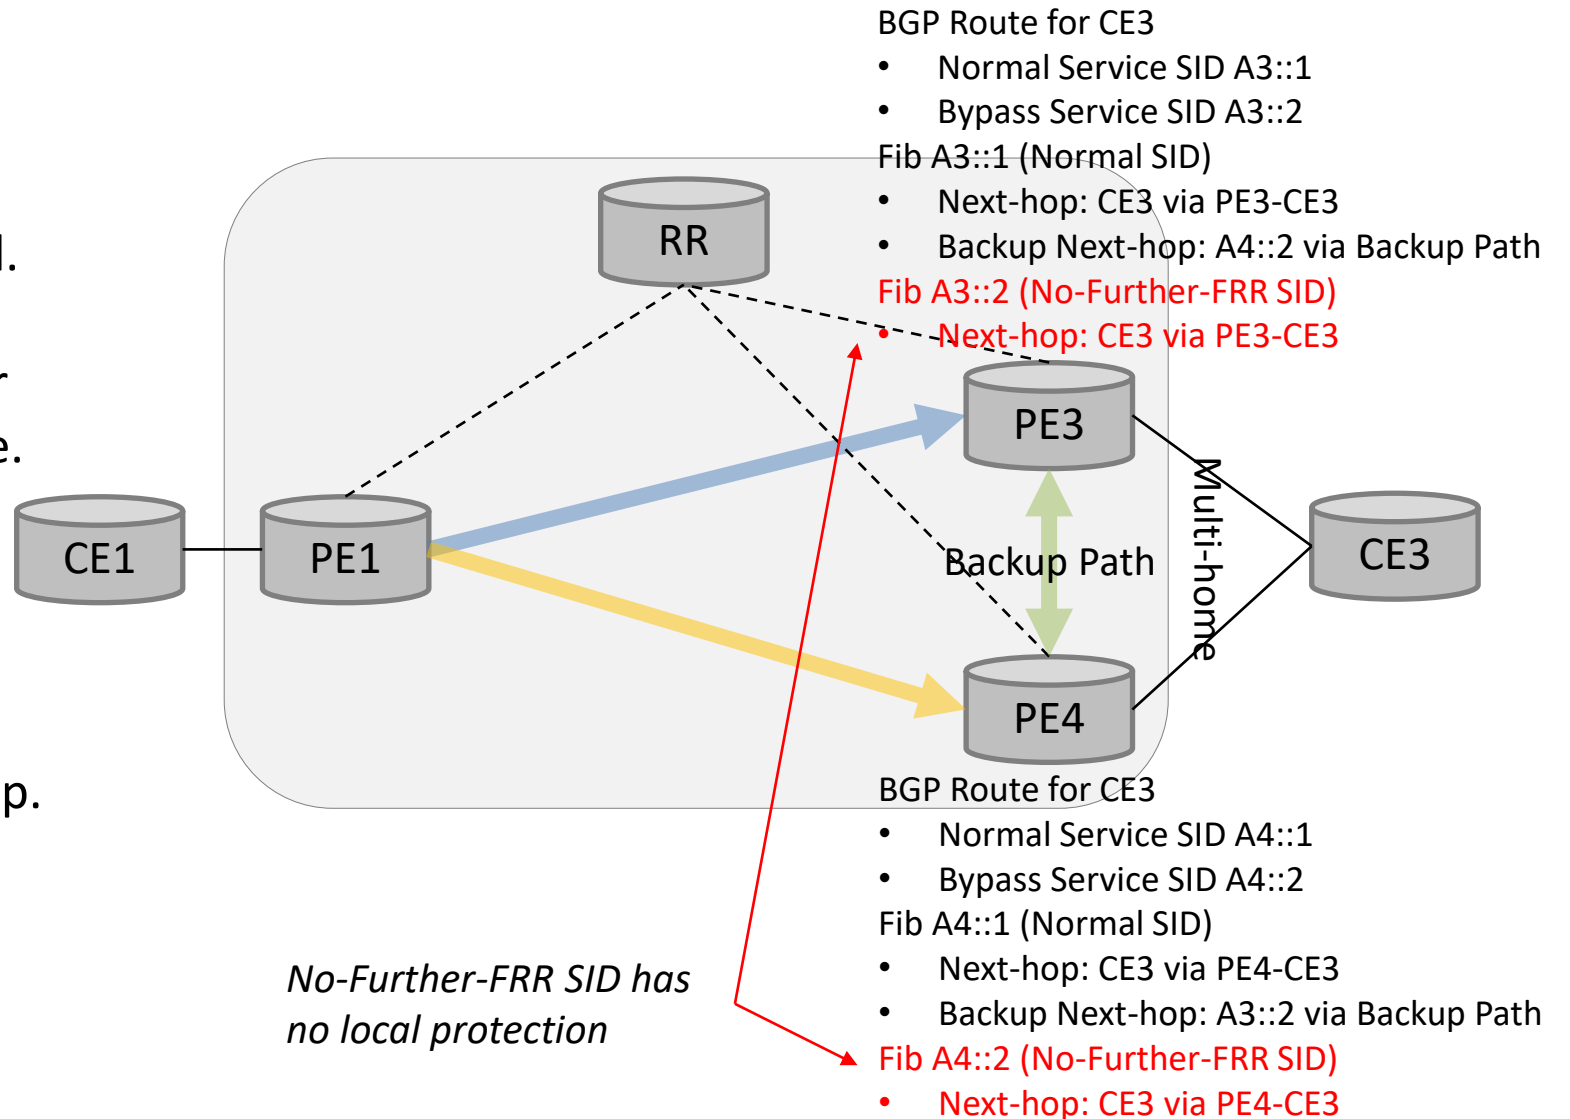
When PE3-CE3 and PE4-CE3 fail at the same time:

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.

# BGP Extensions

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

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| SRv6 Service |   SRv6 Service   |           |           |
| Sub-TLV     |   Sub-TLV     |           |           |
| Type=1     |   Length     |   RESERVED1 |           |
+-----+-----+-----+-----+-----+-----+-----+
| SRv6 SID Value (16 octets) |           //
+-----+-----+-----+-----+-----+-----+-----+
| Svc SID Flags | SRv6 Endpoint Behavior |   RESERVED2 |           |
+-----+-----+-----+-----+-----+-----+-----+
| SRv6 Service Data Sub-Sub-TLVs |           //
+-----+-----+-----+-----+-----+-----+-----+

```

This draft defines a new flag in the SRv6 Service SID Flags field:

```

0 1 2 3 4 5 6 7
+-----+-----+-----+-----+
|N|           |
+-----+-----+-----+-----+
N-flag: No-Further-FRR flag. When set, the associated SID has no
fast reroute protection.

```

# 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 Backup-flag for SRv6 SID. ([8.1 of RFC9352] [9.1 of RFC9513])  
(3.4 of Segment routing architecture RFC 8402 Eligible for Protection Flag) .
- 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.

# Compared with Section 7.2 of [draft-burdet-bess-evpn-fast-reroute-08]

## Section 7.2 of draft-burdet-bess-evpn-fast-reroute-08

- It is necessary to extend a new flavor, though flavor is actually a local behavior.
- Currently, the solution only applies to L2VPN.
- Due to compatibility issues, it's still necessary to send two Service SIDs.

## This draft:

- Using SID flags is simpler than defining new ones for each existing behavior.
- Aligns with existing Segment routing architecture (3.4 of RFC 8402 Eligible for Protection Flag) and IGP( [8.1 of RFC9352] [9.1 of RFC9513]) that uses flags to indicate whether there is a backup.
- Supports both L3VPN and L2VPN.

# Optimization of Assigning Only One SRv6 Service SID

First SRv6 SID: SRv6 SID 1 with Arg 0

Second SRv6 SID: SRv6 SID 1 with Arg 1 (not zero)

with N Flag: no-frr-flag

SRv6 SID 1 with Arg 0: has local fast repair

SRv6 SID 1 with Arg 1 (not zero) and N flag has no fast reroute protection.

We can assign one SID, then assign one arg to distinguish no fast reroute.

Apply both L3VPN and L2VPN.

# Next Steps

- More discussion with Section 7.2 of draft-burdet-bess-evpn-fast-reroute-08.
- Ask for WG adoption.
- Any questions or comments are Welcomed.

**Thanks**