Multicast Source Redundancy in EVPN Networks

draft-skr-bess-evpn-redundant-mcast-source-02

Jorge Rabadan (Nokia)
Jayant Kotalwar (Nokia)
Senthil Sathappan (Nokia)
Eric Rosen (individual)
Jeffrey Zhang (Juniper)
Wen Lin (Juniper)

IETF109, Nov 2020
Online

Agenda

Short refresh

What’s new in rev 02

Conclusions and Next Steps

The Goal – a solution for Multicast Redundancy

That works in any EVPN network

In any redundancy scenario for a given multicast flow:

- Multi-homed Source
- Redundant Single-Homed Sources
- Redundant Multi-Homed Sources

And any EVPN tenant domain configuration:

- Sources and Receivers in the same BD
- Sources and Receivers in different BD of the same tenant
- A mix of the two above

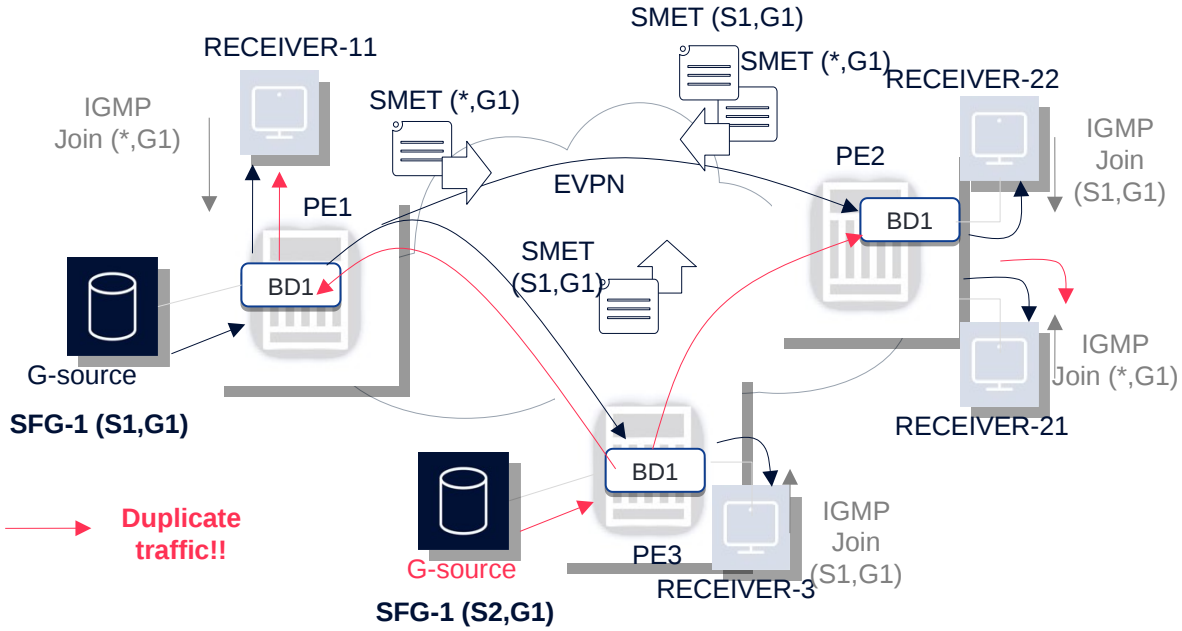

And avoids packet duplication on the receiver systems

Assuming that there may be multiple Redundant Sources sending the same Single Flow Group (SFG) to the network

NOTE: Single Flow Group (SFG)

A multicast group address G which represents traffic that contains only a single flow (e.g., G1)
Multiple sources may be transmitting an SFG (e.g., S1 and S2)

Warm Standby (WS) Solution Details

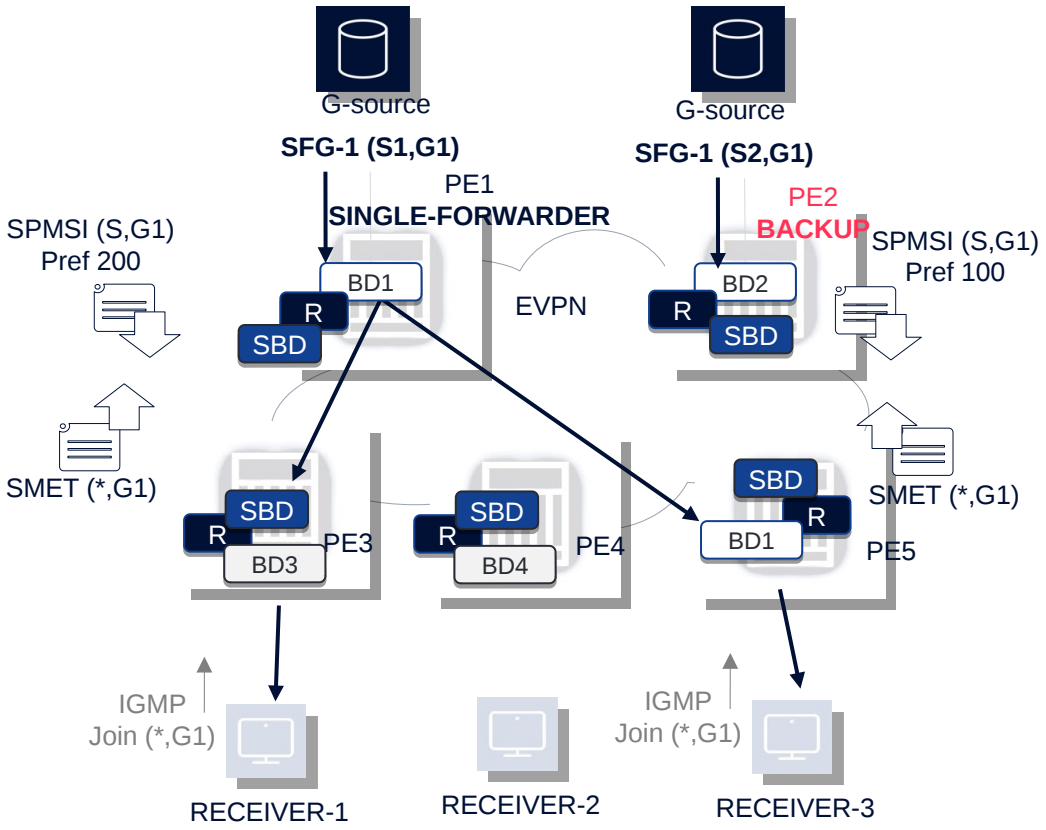

1. Config on PE1 and PE2 only

PE1 and PE2 configured to know that:

- G1 is an SFG, represented as (*,G) or (Sn,G) – Sn is a prefix
- Redundant G-sources for the SFG may exist in BD1 or BD2

2. Signaling the location of G-Sources for (Sn,G1)

Upon receiving SFG for G1, PE1/PE2 originate S-PMSI (Sn,G1) routes that are imported by all PEs. Include DF Election EC and SFG flag.

3. SF Election

PE1/PE2 elect a SF based on the DF Election EC information.

4. RPF check programmed in PE1 and PE2

Non-SF PEs discard any (Sn,G1) packets on a local AC
SF PE accepts (Sn,G1) over at most one local AC

5. Only the Single Forwarder (SF) forwards the SFG

Assuming Downstream PEs have local receivers for the SFG and send SMET(*,G1) or (S,G1) routes (with S contained in Sn)

Hot Standby (HS) Solution Details

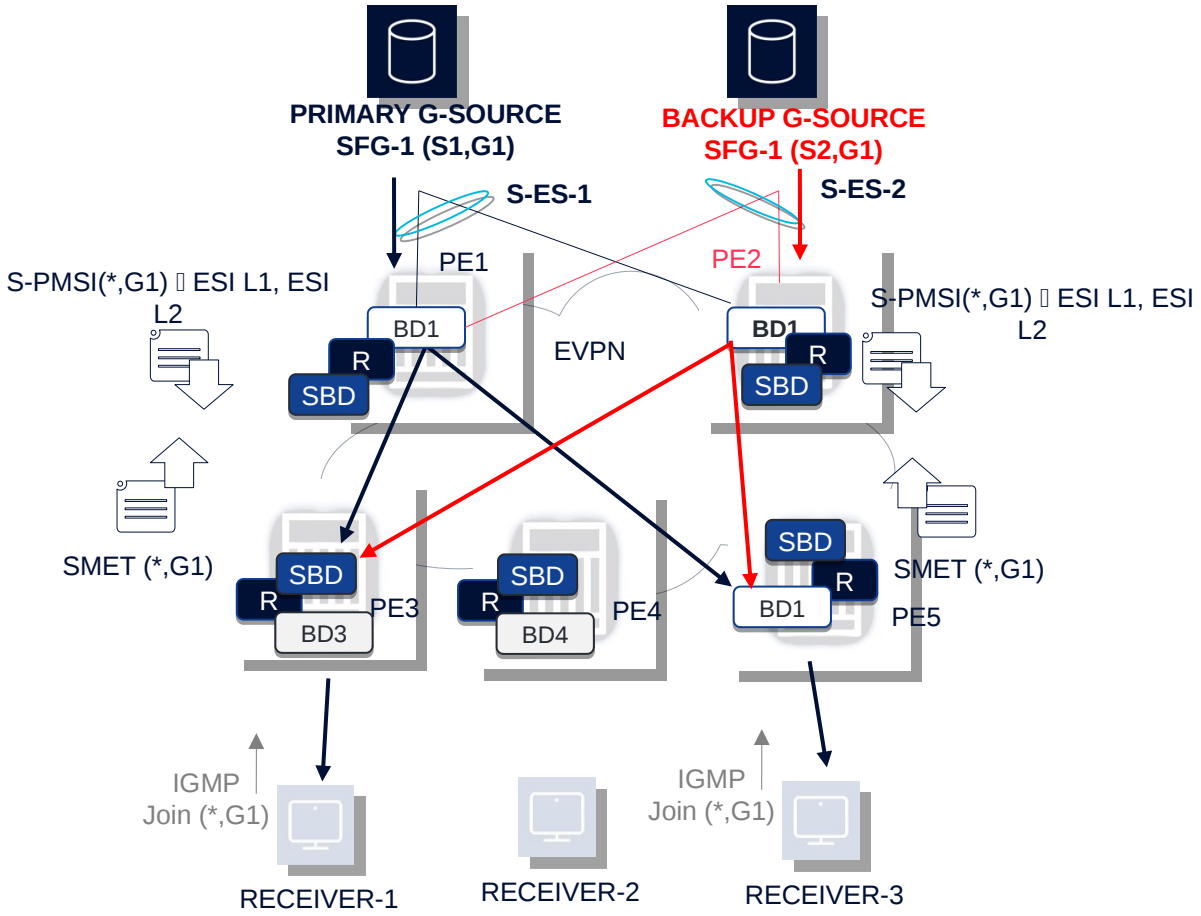

S-ES – Ethernet Segment associated to a G-Source

1. Configuration on all PEs

PE1 and PE2 configured to know that

- (*,G1) is an SFG
- S-ES-1 and S-ES-2 are attached to the G-Sources for (*,G1)

PE3/PE4/PE5 configured with HS mode

2. Signaling the location of G-Sources and S-ESI association

PE1/PE2 send S-PMSI(*,G1)(ESI L1,ESI L2) incl. SFG flag

PE1/PE2 advertise AD per-ES routes with DCB allocated ESI-labels matching the ones in S-PMSI routes, i.e., ESI L1 for S-ES-1 and ESI L2 for S-ES-2 (on both PEs, via DCB)

3. Processing AD per-ES routes and RPF check programming

PE1/PE2 follow regular multi-homing procedures.

Downstream PEs import S-PMSI and AD per-ES routes. They program RPF checks, e.g., PE3 discards traffic with ESI L2.

4. G-traffic forwarding and fault detection

PE1 and PE2 forward G-traffic with ESI L1 and ESI L2 respectively. Only one flow passes the RPF check and is delivered.

- A link failure does not change the RPF check programming
- A complete ES failure or node failure changes RPF check on downstream PEs
- Fault detection based on AD per-ES or per-EVI withdrawal. BFD possible too.

What’s new in rev 02

Clean up

Typos

Improved Terminology section

Added optional use of BFD with reference to draft-ietf-bess-evpn-bfd

BFD MAY be used to determine the status of the tunnels used to forward the SFG from the redundant G-sources

BGP-BFD attribute advertised with S-PMSI A-D or IMET routes

Security Considerations

Added

Conclusions and next steps

The authors ask for WG adoption

Thank you