

draft-skr-bess-evpn-redundant-mcast-source-00

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

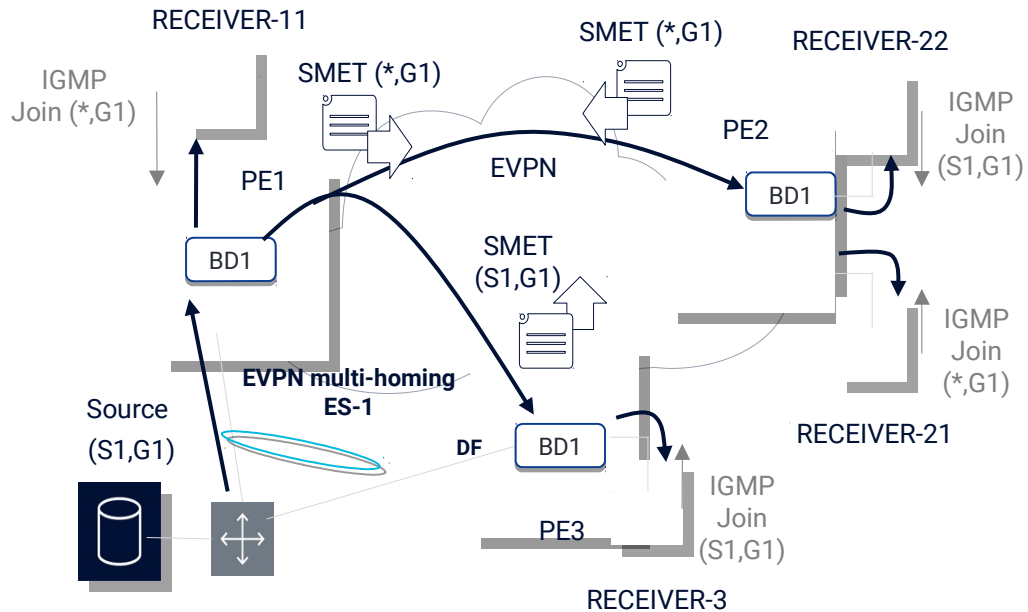
IETF103, Nov 2018
Bangkok

Multi-homed Sources use EVPN Multi-homing

And need no new specifications

Source and Receivers in the Same BD

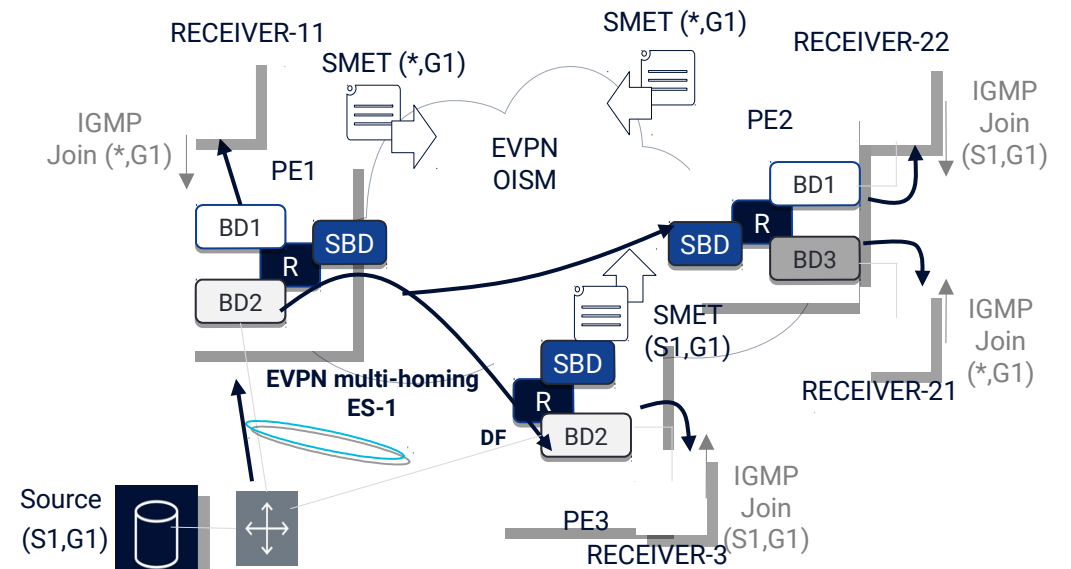
All-active MH does not create duplication
And provides redundancy with fast-failover



draft-ietf-bess-evpn-igmp-ml-d-proxy

Source and Receivers in Different BD

All-active MH does not create duplication
And provides redundancy with fast-failover



draft-ietf-bess-evpn-irb-mcast

The Goal – a solution for Multicast Redundancy

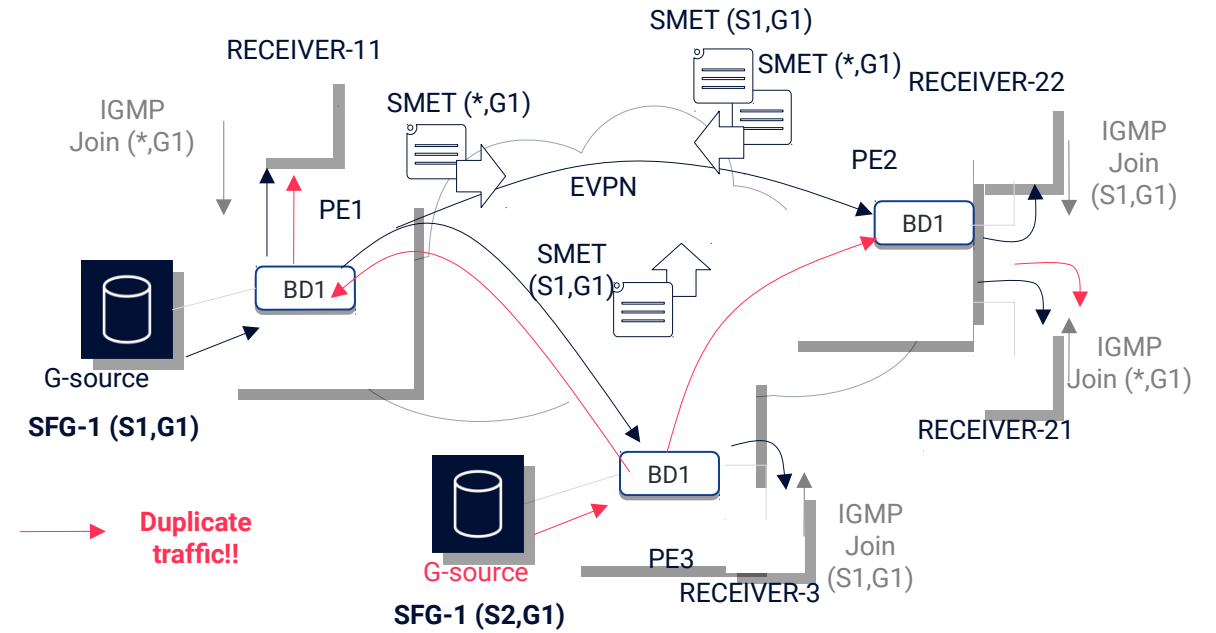
That works in any EVPN network

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:

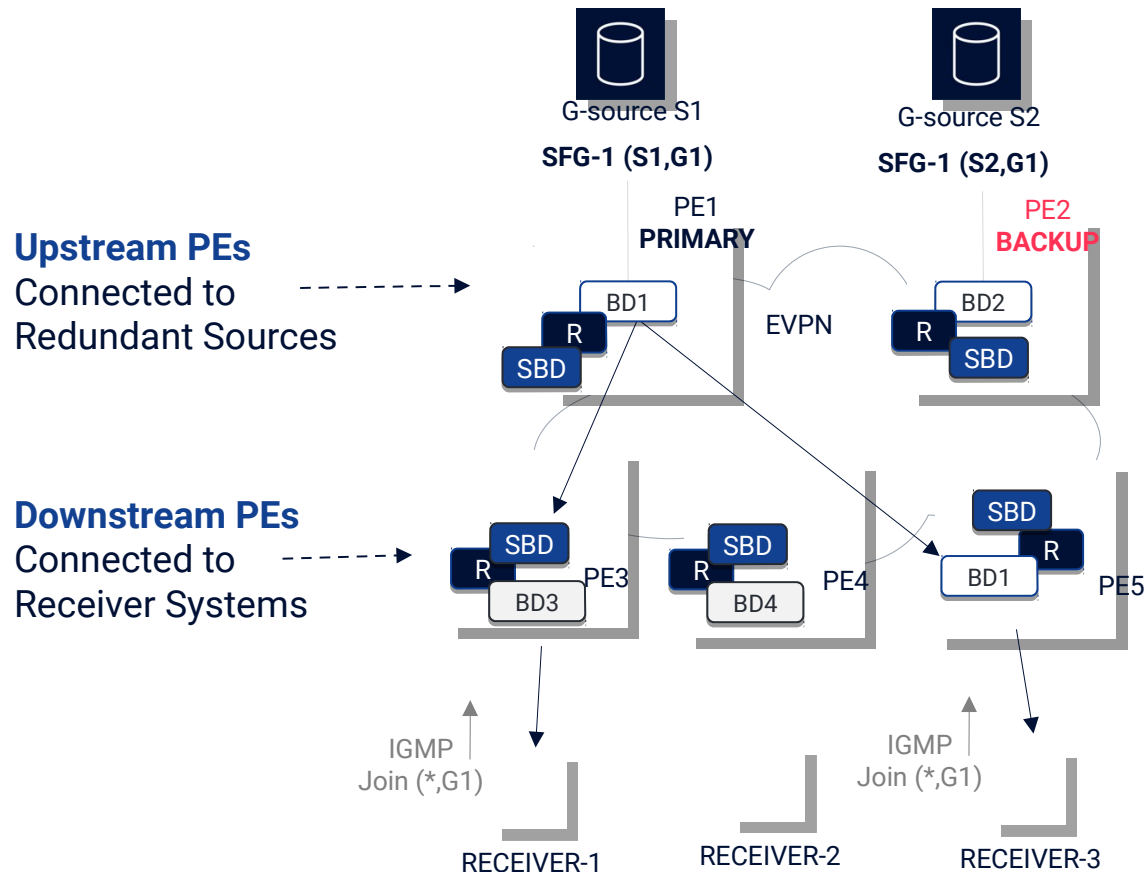
- Sources and Receivers in the same BD
- Sources and Receivers in different BD of the same tenant
- A mix of the two above



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

Terminology



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 $S2$)

NOTE

Familiarity with the following specs is assumed:

RFC7432 – EVPN

draft-ietf-bess-evpn-igmp-mld-proxy

draft-ietf-bess-evpn-irb-mcast

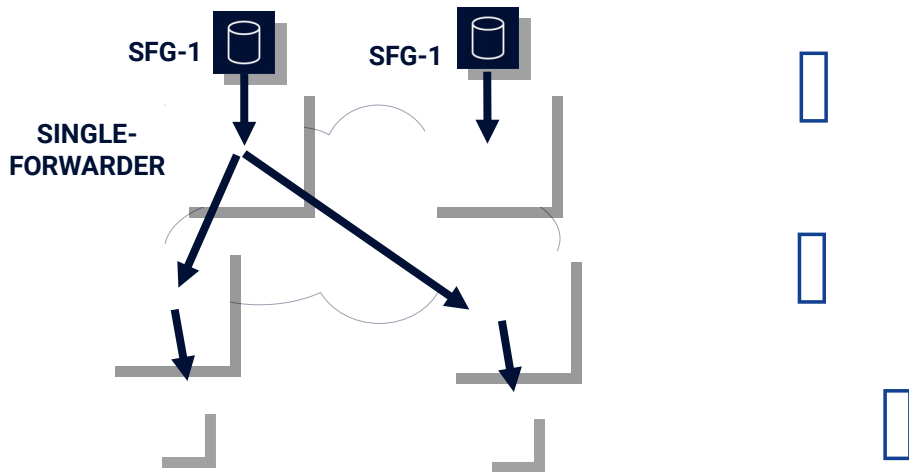
Two Redundant G-Source Solutions

WARM STANDBY SOLUTION (WS)

Avoids duplication of SFG flows in the tenant network while providing G-source redundancy for a given SFG.

Based on a **Single Forwarder (SF)** Election, only one of the Upstream PEs connected to a G-Source will forward the SFG

The Upstream PEs add an RPF check to the (*,G) state for the SFG to avoid duplication.




HOT STANDBY SOLUTION (HS)

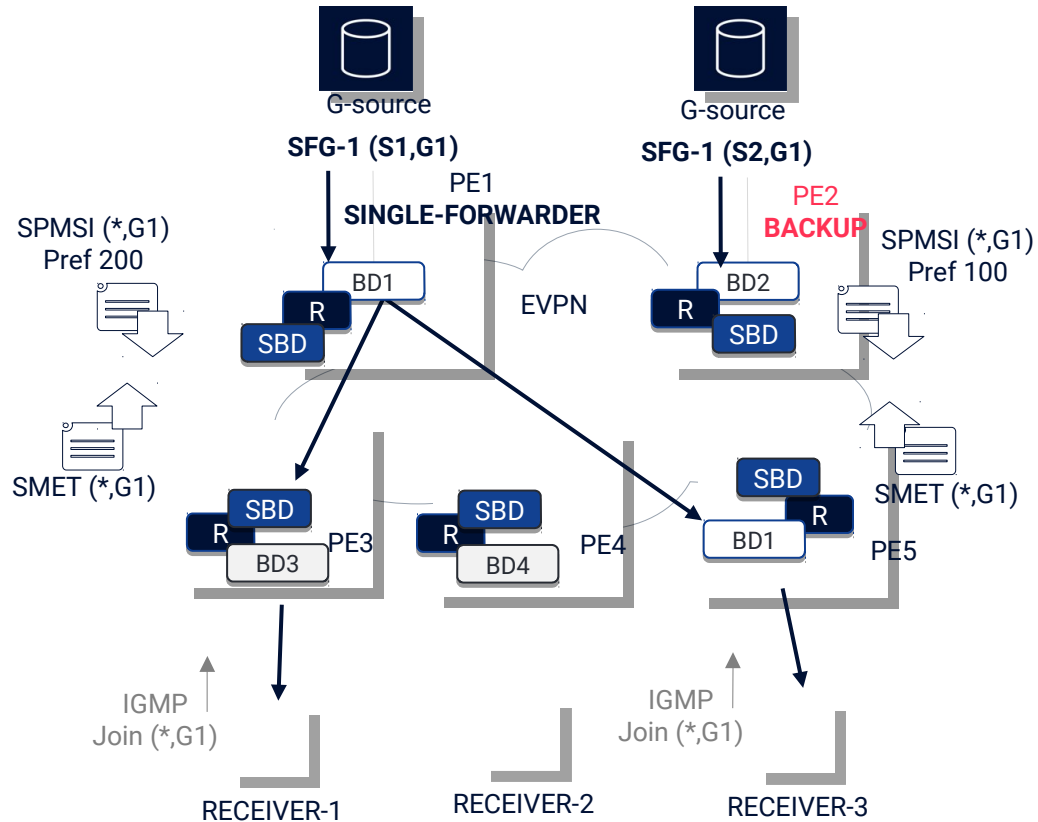
Avoids duplication of SFG flows on the receiver systems while providing G-source redundancy for a given SFG.

The Upstream PEs add a data path identification so that the Downstream PEs can add an RPF check to discard SFG traffic from the "wrong G-Source".

In case of failure in the primary G-Source, the Downstream PEs locally select a different G-Source for the SFG.



Warm Standby (WS) Solution Details



1. Config on PE1 and PE2 only

PE1 and PE2 configured to know that:

- G1 is an SFG
- Redundant G-sources for G1 may exist in BD1 or BD2

2. Signaling the location of G-Sources for G1

Upon receiving SFG for G1, PE1/PE2 originate S-PMSI (*,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.

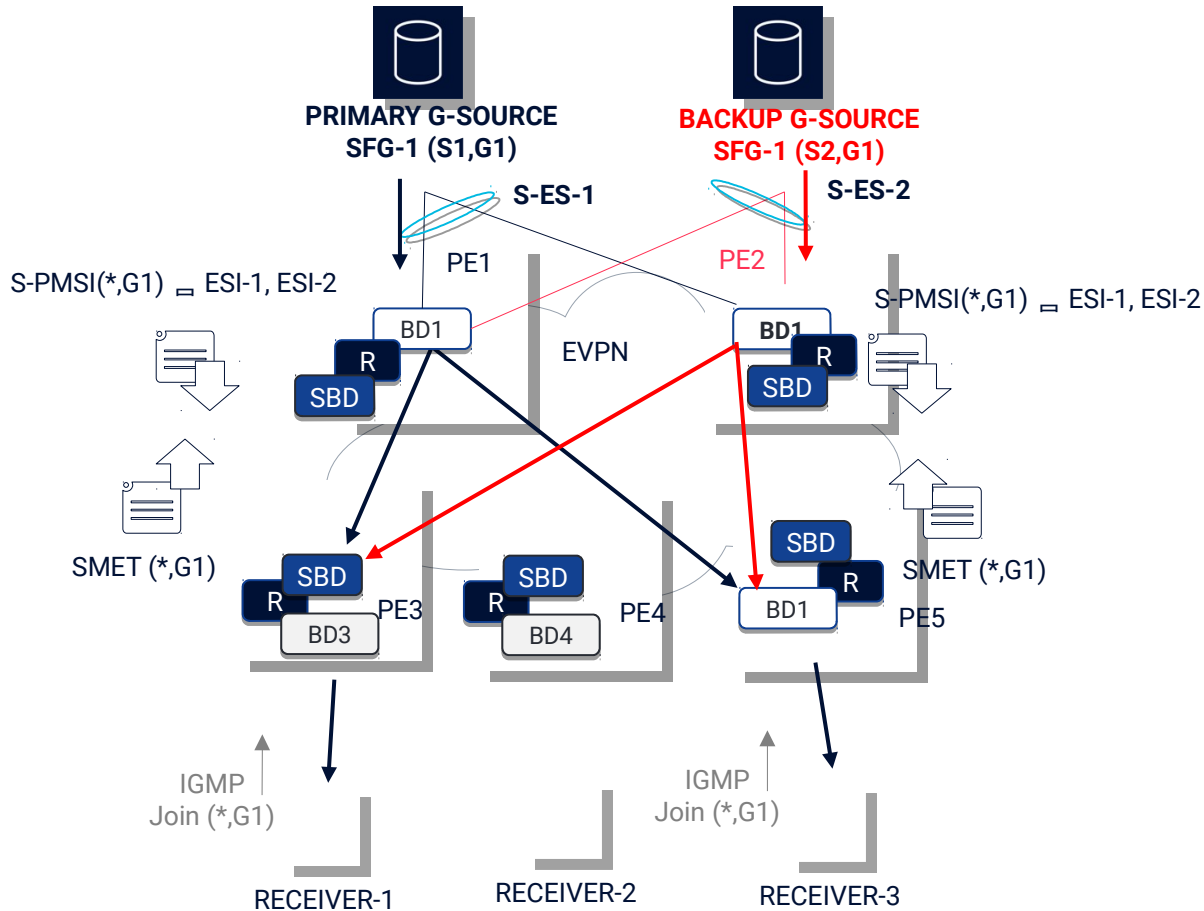
4. RPF check programmed in PE1 and PE2

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

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

Assuming Downstream PEs have local receivers for (*,G1) and send SMET(*,G1) routes.

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-1,ESI-2) incl. SFG flag
 PE1/PE2 advertise AD per-ES routes with DCB allocated ESI-labels, i.e., ESI-label-1 for S-ES-1 and ESI-label-2 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-label-2.

4. G-traffic forwarding and fault detection

PE1 and PE2 forward G-traffic with ESI-label-1 and ESI-label-2 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.

Next steps

The Authors would like to request Feedback

In particular about the attribute to encode the G-Source Identifiers (ESIs)

Thank you