

BESS Workgroup
Internet-Draft
Intended status: Standards Track
Expires: 18 August 2025

J. Rabadan, Ed.
J. Kotalwar
S. Sathappan
Nokia
Z. Zhang
W. Lin
Juniper
14 February 2025

Multicast Source Redundancy in EVPN Networks
draft-ietf-bess-evpn-redundant-mcast-source-15

Abstract

In Ethernet Virtual Private Networks (EVPNs), IP multicast traffic replication and delivery play a crucial role in enabling efficient and scalable layer-2 and layer-3 services. A common deployment scenario involves redundant multicast sources that ensure high availability and resiliency. However, the presence of redundant sources can lead to duplicate IP multicast traffic in the network, causing inefficiencies and increased overhead. This document specifies extensions to the EVPN multicast procedures that allow for the suppression of duplicate IP multicast traffic from redundant sources. The proposed mechanisms enhance EVPN's capability to deliver multicast traffic efficiently while maintaining high availability. These extensions are applicable to various EVPN deployment scenarios and provide guidelines to ensure consistent and predictable behavior across diverse network topologies.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 18 August 2025.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	4
1.2. Background on IP Multicast Delivery in EVPN Networks	7
1.2.1. Intra-subnet IP Multicast Forwarding	7
1.2.2. Inter-subnet IP Multicast Forwarding	9
1.3. Multi-Homed IP Multicast Sources in EVPN	11
1.4. The Need for Redundant IP Multicast Sources in EVPN	12
2. Solution Overview	13
2.1. Warm Standby Solution	13
2.2. Hot Standby Solution	14
2.3. Solution Selection	14
2.4. System Support	15
3. BGP EVPN Extensions	15
3.1. Single Flow Group Flag in the Multicast Flags Extended Community	15
3.2. ESI Label Extended Community in S-PMSI A-D Routes	15
4. Warm Standby (WS) Solution for Redundant G-Sources	16
4.1. Specification	16
4.2. Warm Standby Example in an OISM Network	19
4.3. Warm Standby Example in a Single-BD Tenant Network	21
5. Hot Standby Solution for Redundant G-Sources	22
5.1. Specification	22
5.2. Extensions for the Advertisement of DCB Labels	26
5.3. Use of BFD in the HS Solution	27
5.4. Hot Standby Example in an OISM Network	28
5.4.1. Multi-Homed Redundant G-Sources	28
5.4.2. Single-Homed Redundant G-Sources	31
5.5. Hot Standby Example in a Single-BD Tenant Network	33
6. Security Considerations	33
7. IANA Considerations	33
8. Acknowledgments	34
9. Contributors	34

10. References	34
10.1. Normative References	34
10.2. Informative References	35
Authors' Addresses	37

1. Introduction

Ethernet Virtual Private Networks (EVPN) [RFC7432] support both intra-subnet and inter-subnet IP multicast forwarding. [RFC9251] outlines the procedures required to optimize the delivery of IP multicast flows when both sources and receivers are connected to the same EVPN Broadcast Domain. [RFC9625], on the other hand, defines the procedures for supporting inter-subnet IP multicast within a tenant network, where the IP multicast source and receivers of the same multicast flow are connected to different Broadcast Domains within the same tenant.

However, [RFC9251], [RFC9625], and conventional IP multicast techniques do not provide a solution for scenarios where:

1. A given multicast group carries multiple flows (i.e., multiple sources are active), and
2. Each receiver should receive only from one source.

Existing multicast solutions typically assume that there are no redundant sources sending identical flows to the same IP multicast group. In cases where redundant sources do exist, the receiver application is expected to handle duplicate packets.

In conventional IP multicast networks, such as those running Protocol Independent Multicast (PIM) [RFC7761] or Multicast VPNs (MVPN) [RFC6513], a workaround is to configure all redundant sources with the same IP address. This approach ensures that each receiver gets only one flow because:

- * The RP (Rendezvous Point) in the multicast network always creates (S,G) state for each source.
- * The Last Hop Router (LHR) may also create (S,G) state.
- * The (S,G) state binds the flow to a source-specific tree rooted at the source IP address. When multiple sources share the same IP address, the resulting source-specific trees ensure that each LHR or RP resides on at most one tree.

This workaround, which often uses anycast addresses, is suitable for warm standby redundancy solutions (Section 4). However, it is not effective for hot standby redundancy scenarios (Section 5) and introduces challenges when sources need to be reachable via IP unicast or when multiple sources with the same IP address are attached to the same Broadcast Domain. In scenarios where multiple multicast sources stream traffic to the same group using EVPN Optimized Inter-Subnet Multicast (OISM), there is not necessarily any (S,G) state created for the redundant sources. In such cases, the Last Hop Routers may only have (*,G) state, and there may not be a Rendezvous Point router to create (S,G) state.

This document extends [RFC9251] and [RFC9625] to address scenarios where IP multicast source redundancy exists. Specifically, it defines procedures for EVPN PEs (Provider Edge devices/routers) to ensure that receivers do not experience packet duplication when two or more sources send identical IP multicast flows into the tenant domain. These procedures are limited to the context of [RFC9251] and [RFC9625]; handling redundant sources in other multicast solutions is beyond the scope of this document.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

- * Broadcast Domain (BD): an emulated Ethernet, such that two systems on the same BD will receive each other's link-local broadcasts. In this document, BD also refers to the instantiation of a Broadcast Domain on an EVPN PE. An EVPN PE can be attached to one or multiple BDs of the same tenant.
- * BUM: Broadcast, Unknown unicast, and Multicast traffic.
- * Designated Forwarder (DF): as defined in [RFC7432], an Ethernet Segment may be multi-homed (attached to more than one PE). An Ethernet Segment may also contain multiple BDs, of one or more EVIs. For each such EVI, one of the PEs attached to the segment becomes that EVI's DF for that segment. Since a BD may belong to only one EVI, we can speak unambiguously of the BD's DF for a given segment.
- * Downstream PE: in this document a Downstream PE is referred to as the EVPN PE that is connected to the IP Multicast receivers and gets the IP Multicast flows from remote EVPN PEs.

- * EVI: EVPN Instance.
- * G-traffic: any frame with an IP payload whose IP Destination Address (IP DA) is a multicast group G.
- * G-source: any system sourcing IP multicast traffic to group G.
- * Hot Standby Redundancy: multicast source redundancy procedure defined in this document, by which the upstream PEs forward the redundant multicast flows to the downstream PEs, and the downstream PEs make sure only one flow is forwarded to the interested attached receivers.
- * IGMP: Internet Group Management Protocol [RFC3376].
- * Inclusive Multicast Tree or Inclusive Provider Multicast Service Interface (I-PMSI): defined in [RFC6513], in this document it is applicable only to EVPN and refers to the default multicast tree for a given BD. All the EVPN PEs that are attached to a specific BD belong to the I-PMSI for the BD. The I-PMSI trees are signaled by EVPN Inclusive Multicast Ethernet Tag (IMET) routes.
- * IMET route: EVPN Inclusive Multicast Ethernet Tag route, as in [RFC7432].
- * MLD: Multicast Listener Discovery [RFC3810].
- * MVPN: Multicast Virtual Private Networks, as in [RFC6513].
- * OISM: Optimized Inter-Subnet Multicast, as in [RFC9625].
- * PE: Provider Edge.
- * PIM: Protocol Independent Multicast, as in [RFC7761].
- * P-tunnel: The term "Provider tunnel" refers to the type of tree employed by an upstream EVPN PE to forward multicast traffic to downstream PEs. The P-tunnels supported in this document include Ingress Replication (IR), Assisted Replication (AR) [RFC9574], Bit Indexed Explicit Replication (BIER) [RFC8296], multicast Label Distribution Protocol (mLDP), and Point-to-Multi-Point Resource Reservation Protocol with Traffic Engineering extensions (P2MP RSVP-TE).
- * Redundant G-source: A host or router transmitting a Single Flow Group (SFG) within a tenant network, where multiple hosts or routers are also transmitting the same SFG. Redundant G-sources transmitting the same SFG should have distinct IP addresses;

however, they may share the same IP address if located in different Broadcast Domains (BDs) within the same tenant network. For the purposes of this document, redundant G-sources are assumed not to exhibit "bursty" traffic behavior.

- * S-ES and S-ESI: multicast Source Ethernet Segment and multicast Source Ethernet Segment Identifier. The Ethernet Segment and Ethernet Segment Identifier associated to a G-source.
- * Selective Multicast Tree or Selective Provider Multicast Service Interface (S-PMSI): As defined in [RFC6513], this term refers to a multicast tree to which only the PEs interested in a specific Broadcast Domain (BD) belong. In the context of this document, it is specific to EVPN. Two types of EVPN S-PMSIs are supported:
 - S-PMSIs with Auto-Discovery Routes: These S-PMSIs require the upstream PE to advertise S-PMSI Auto-Discovery (S-PMSI A-D) routes, as described in [RFC9572]. Downstream PEs interested in the multicast traffic join the S-PMSI tree following the procedures specified in [RFC9572].
 - S-PMSIs without Auto-Discovery Routes: These S-PMSIs do not require the advertisement of S-PMSI A-D routes. Instead, they rely on the forwarding information provided by Inclusive Multicast Ethernet Tag (IMET) routes. Upstream PEs forward IP multicast flows only to downstream PEs that advertise Selective Multicast Ethernet Tag (SMET) routes for the specific flow. These S-PMSIs are supported exclusively with the following P-tunnel types: Ingress Replication (IR), Assisted Replication (AR), and Bit Indexed Explicit Replication (BIER).
- * SFG (Single Flow Group): A multicast group that represents traffic containing a single flow. Multiple sources, which may have the same or different IP addresses, can transmit traffic for an SFG. An SFG can be represented in two forms:
 - (*,G): Indicates that any source transmitting multicast traffic to group G is considered a redundant G-source for the SFG.
 - (S,G): Indicates that S is a prefix of any length. In this representation, a source is deemed a redundant G-source for the SFG if its address matches the specified prefix S.
- * SMET route: Selective Multicast Ethernet Tag route, as in [RFC9251].

- * (S,G) and (*,G): used to describe multicast packets or multicast state. S stands for Source (IP address of the multicast traffic) and G stands for the Group or multicast destination IP address of the group. An (S,G) multicast packet refers to an IP packet with source IP address "S" and destination IP address "G", and it is forwarded on a multicast router if there is a corresponding state for (S,G). A (*,G) multicast packet refers to an IP packet with "any" source IP address and a destination IP address "G", and it is forwarded on a multicast router based on the existence of the corresponding (*,G) state. The document uses variations of these terms. For example, (S1,G1) represents the multicast packets or multicast state for source IP address "S1" and group IP address "G1".
- * Upstream PE: In this document, an Upstream PE refers to the EVPN PE that is either directly connected to the IP multicast source or is the PE closest to the source. The Upstream PE receives IP multicast flows through local Attachment Circuits (ACs).
- * Warm Standby Redundancy: A multicast source redundancy mechanism defined in this document, wherein the upstream PEs connected to redundant sources within the same tenant ensure that only one source of a given flow transmits multicast traffic to the interested downstream PEs at any given time.

This document also assumes familiarity with the terminology of [RFC7432], [RFC4364], [RFC6513], [RFC6514], [RFC9251], [RFC9625], [RFC9136], and [RFC9572].

1.2. Background on IP Multicast Delivery in EVPN Networks

IP multicast facilitates the delivery of a single copy of a packet from a source (S) to a group of receivers (G) along a multicast tree. In an EVPN tenant domain, the multicast tree can be constructed where the source (S) and the receivers for the multicast group (G) are either connected to the same Broadcast Domain (BD) or to different Broadcast Domains. The former scenario is referred to as "Intra-subnet IP Multicast forwarding", while the latter is referred to as "Inter-subnet IP Multicast forwarding".

1.2.1. Intra-subnet IP Multicast Forwarding

When the source S1 and the receivers interested in G1 are connected to the same Broadcast Domain (BD), the EVPN network can deliver IP multicast traffic to the receivers using two different approaches, as illustrated in Figure 1:

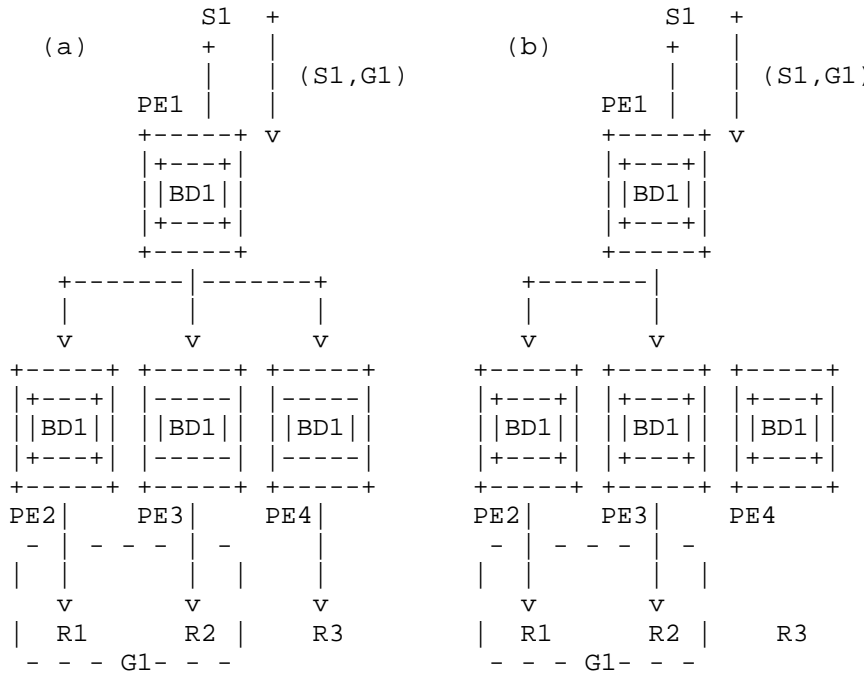


Figure 1: Intra-subnet IP Multicast

* Model (a): IP Multicast Delivery as BUM Traffic

The upstream PE sends the IP Multicast flows to all downstream PEs, even to PEs with non-interested receivers, such as, e.g., PE4 in Figure 1. To optimize this behavior, downstream PEs can snoop IGMP/MLD messages from receivers to build Layer 2 multicast state. For instance, PE4 could avoid forwarding (S1, G1) to R3, since R3 has not expressed interest in (S1, G1).

* Model (b): Optimized Delivery with S-PMSI

Model (b) in Figure 1 uses a "Selective Provider Multicast Service Interface (S-PMSI)" to optimize the delivery of the (S1, G1) flow.

- For example, if PE1 uses "Ingress Replication (IR)", it will forward (S1, G1) only to downstream PEs that have issued a "Selective Multicast Ethernet Tag (SMET)" route for (S1, G1), such as PE2 and PE3.

- If PE1 uses a P-tunnel type other than IR (e.g., Assisted Replication (AR) or Bit Indexed Explicit Replication (BIER)), PE1 will advertise an "S-PMSI Auto-Discovery (A-D)" route for (S1,G1). Downstream PEs such as PE2 and PE3 will then join the corresponding multicast tree to receive the flow.

Procedures for Model (b) are specified in [RFC9251].

1.2.2. Inter-subnet IP Multicast Forwarding

When the sources and receivers are connected to different BDs within the same tenant domain, the EVPN network can deliver IP multicast traffic using either Inclusive or Selective Multicast Trees, as illustrated in Figure 2 with models (a) and (b), respectively.

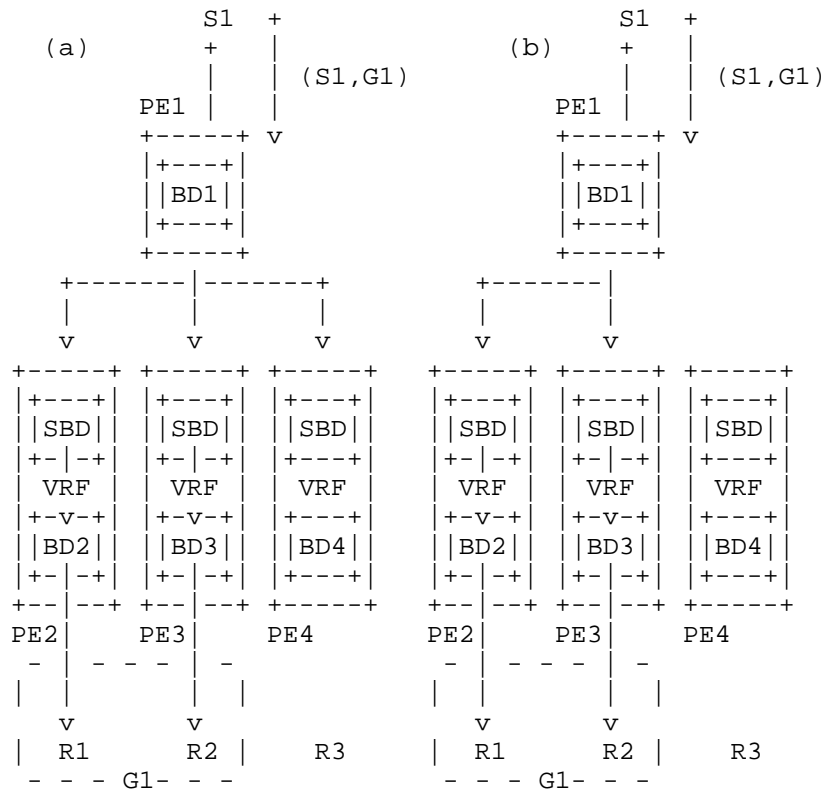


Figure 2: Inter-subnet IP Multicast

As defined in [RFC9625], inter-subnet multicast forwarding in EVPN is optimized by ensuring IP multicast flows are sent within the context of the source BD. If a downstream PE is not connected to the source BD, the IP multicast flow is delivered to the Supplementary Broadcast Domain (SBD), as shown in Figure 2.

* Inclusive and Selective Multicast Trees

Model (a): Inclusive Multicast Tree

In this model, the Inclusive Multicast Tree for BD1 on PE1 delivers (S1,G1) to all downstream PEs, such as PE2, PE3, and PE4, in the context of the SBD. Each downstream PE then locally routes the flow to its Attachment Circuits, ensuring delivery to interested receivers.

Model (b): Selective Multicast Tree

In this model, PE1 optimizes forwarding by delivering (S1,G1) only to downstream PEs that explicitly indicate interest in the flow via Selective Multicast Ethernet Tag (SMET) routes. If the P-tunnel type is "Ingress Replication (IR)", "Assisted Replication (AR)", or "Bit Indexed Explicit Replication (BIER)", PE1 does not need to advertise an S-PMSI A-D route. Downstream PEs join the multicast tree based on the SMET routes advertised for (S1,G1).

* Advantages of [RFC9625]

[RFC9625] extends the procedures defined in [RFC9251] to support both intra- and inter-subnet multicast forwarding for EVPN. It ensures that every upstream PE attached to a source is aware of all downstream PEs within the same tenant domain that have interest in specific flows. This is achieved through the advertisement of SMET routes with the SBD Route Target, which are imported by all upstream PEs.

* Elimination of Complexity

The inter-subnet multicast mechanism defined in [RFC9625] eliminates the need for: Rendezvous Points (RP), Shared trees, Upstream Multicast Hop selection, or other complex conventional multicast routing techniques.

By leveraging the EVPN framework, inter-subnet multicast forwarding achieves efficient delivery without introducing unnecessary overhead or dependencies on classic IP multicast protocols.

1.3. Multi-Homed IP Multicast Sources in EVPN

Unlike conventional multicast routing technologies, multi-homed PEs connected to the same source do not create IP multicast packet duplication when utilizing a multi-homed Ethernet Segment. Figure 3 illustrates this scenario, where two multi-homed PEs (PE1 and PE2) are attached to the same source S1. The source S1 is connected via a Layer 2 switch (SW1) to an all-active Ethernet Segment (ES-1), with a Link Aggregation Group (LAG) extending to PE1 and PE2.

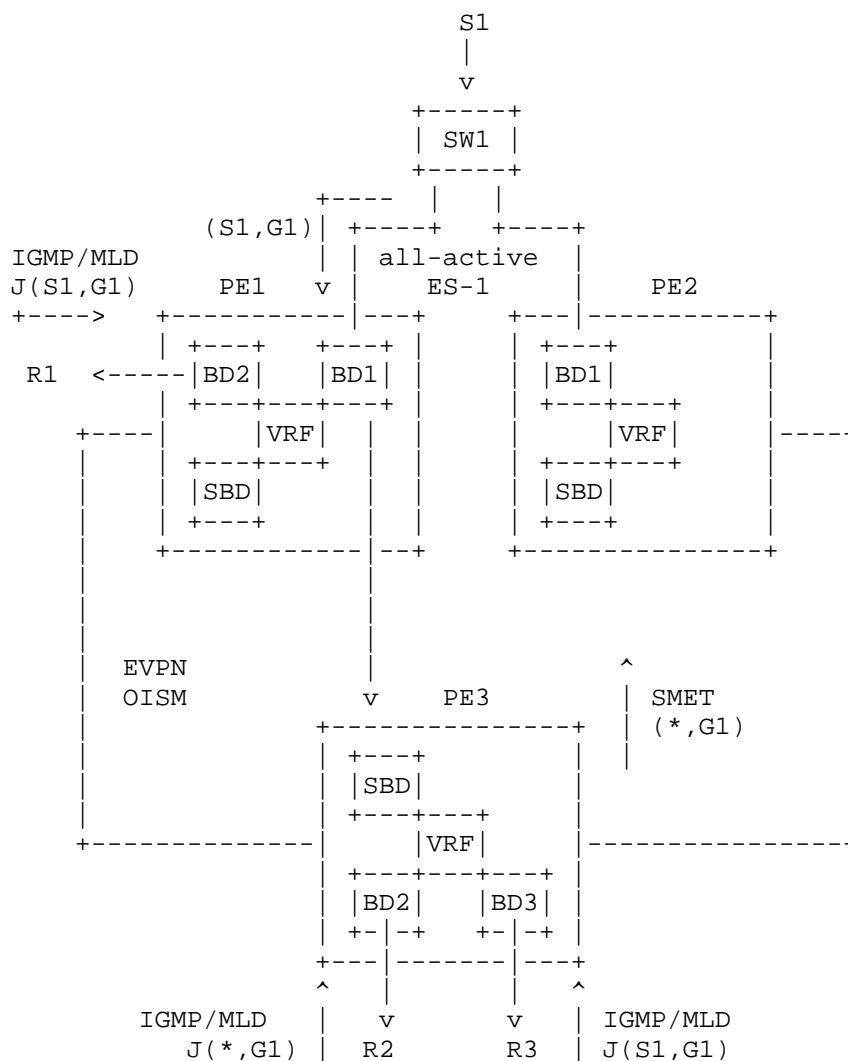


Figure 3: All-active Multi-homing and OISM

When S1 transmits the (S1,G1) flow, SW1 selects a single link within the all-active Ethernet Segment to forward the flow, as per [RFC7432]. In this example, assuming PE1 is the receiving PE for Broadcast Domain BD1, the multicast flow is forwarded once BD1 establishes multicast state for (S1,G1) or (*,G1). In Figure 3:

- * Receiver R1 receives (S1,G1) directly via the IRB (Integrated Routing and Bridging) interface, defined in [RFC9135], following the procedures in [RFC9625].
- * Receivers R2 and R3, upon issuing IGMP/MLD reports, trigger PE3 to advertise an SMET (*,G1) route. This creates multicast state in PE1's BD1, enabling PE1 to forward the multicast flow to PE3's SBD. PE3 subsequently delivers the flow to R2 and R3 as defined in [RFC9625].

Requirements for Multi-Homed IP Multicast Sources:

- * When IP multicast source multi-homing is needed, EVPN multi-homed Ethernet Segments MUST be used.
- * EVPN multi-homing ensures that only one upstream PE forwards a given multicast flow at a time, preventing packet duplication at downstream PEs.
- * The SMET route for a multicast flow ensures that all upstream PEs in the multi-homed Ethernet Segment maintain state for the flow. This allows for immediate failover, as the backup PE can seamlessly take over forwarding in case of an upstream PE failure.

This document assumes that multi-homed PEs connected to the same source always utilize multi-homed Ethernet Segments.

1.4. The Need for Redundant IP Multicast Sources in EVPN

While multi-homing PEs to the same IP multicast G-source provides a certain level of resiliency, multicast applications are often critical in operator networks, necessitating a higher level of redundancy. This document assumes the following:

- a. Redundant G-sources: redundant G-sources for an SFG may exist within the EVPN tenant network. A redundant G-source is defined as a host or router transmitting an SFG stream in a tenant network where another host or router is also sending traffic to the same SFG.

- b. G-source placement: redundant G-sources may reside in the same BD or in different BDs of the tenant network. There must be no restrictions on the locations of receiver systems within the tenant.
- c. G-source attachment to EVPN PEs: redundant G-sources may be either single-homed to a single EVPN PE or multi-homed to multiple EVPN PEs.
- d. Packet duplication avoidance: the EVPN PEs must ensure that receiver systems do not experience duplicate packets for the same SFG.

This framework ensures that EVPN networks can effectively support redundant multicast sources while maintaining high reliability and operational efficiency.

2. Solution Overview

An SFG can be represented as (*,G) if any source transmitting multicast traffic to group G is considered a redundant G-source. Alternatively, this document allows an SFG to be represented as (S,G), where the source IP address S is a prefix of variable length. In this case, a source is deemed a redundant G-source for the SFG if its address falls within the specified prefix. The use of variable-length prefixes in source advertisements via S-PMSI A-D routes is permitted in this document only for the specific application of redundant G-sources.

This document describes two solutions for handling redundant G-sources:

- * Warm Standby Solution
- * Hot Standby Solution

2.1. Warm Standby Solution

The Warm Standby solution is an upstream PE-based solution, where downstream PEs do not participate in the procedures. In this solution, all upstream PEs connected to redundant G-sources for an SFG (*,G) or (S,G) elect a "Single Forwarder (SF)" among themselves. After the Single Forwarder is elected, the upstream PEs apply Reverse Path Forwarding checks to the multicast state for the SFG:

- * Non-Single Forwarder Behavior: a non-Single Forwarder upstream PE discards all (*,G) or (S,G) packets received over its local Attachment Circuit.

- * **Single Forwarder Behavior:** the Single Forwarder accepts and forwards (*,G) or (S,G) packets received on a single local Attachment Circuit for the SFG. If packets are received on multiple local Attachment Circuits, the Single Forwarder discards packets on all but one. The selection of the Attachment Circuit for forwarding is a local implementation detail.

In the event of a failure of the Single Forwarder, a new Single Forwarder is elected among the upstream PEs. This election process requires BGP extensions on existing EVPN routes, which are detailed in Section 3 and Section 4.

2.2. Hot Standby Solution

The Hot Standby solution relies on downstream PEs to prevent duplication of SFG packets. Upstream PEs, aware of locally connected G-sources, append a unique Ethernet Segment Identifier (ESI) label to multicast packets for each SFG. Downstream PEs receive SFG packets from all upstream PEs attached to redundant G-sources and avoid duplication by performing a Reverse Path Forwarding check on the (*,G) state for the SFG:

- * **Packet Filtering:** a downstream PE discards (*,G) packets received from the "wrong G-source."
- * **Wrong G-source Identification:** the "wrong G-source" is identified using an ESI label that differs from the ESI label associated with the selected G-source.
- * **ESI Label Usage:** in this solution, the ESI label is used for "ingress filtering" at the downstream PE, rather than for "egress filtering" as described in [RFC7432]. In [RFC7432], the ESI label indicates which egress Attachment Circuits must be excluded when forwarding BUM traffic. Here, the ESI label identifies ingress traffic that should be discarded by the downstream PE.

Control plane and data plane extensions to [RFC7432] are required to support ESI labels for SFGs forwarded by upstream PEs. Upon failure of the selected G-source, the downstream PE switches to a different G-source and updates its Reverse Path Forwarding check for the (*,G) state. These extensions and procedures are described in Section 3 and Section 5.

2.3. Solution Selection

Operators should select a solution based on their specific requirements:

- * The Warm Standby solution is more bandwidth-efficient but incurs longer failover times in the event of a G-source or upstream PE failure. Additionally, only the upstream PEs connected to redundant G-sources for the same SFG need to support the new procedures in the Warm Standby solution.
- * The Hot Standby solution is recommended for scenarios requiring fast failover times, provided that the additional bandwidth consumption (due to multiple transmissions of SFG packets to downstream PEs) is acceptable.

2.4. System Support

This document does not mandate support for both solutions on a single system. If one solution is implemented, support for the other is OPTIONAL.

3. BGP EVPN Extensions

This document introduces the following BGP EVPN extensions:

3.1. Single Flow Group Flag in the Multicast Flags Extended Community

A new Single Flow Group (SFG) flag is defined within the Multicast Flags Extended Community. This flag is requested from the IANA registry for "Multicast Flags Extended Community Flag Values". The SFG flag is set in S-PMSI A-D routes that carry (*,G) or (S,G) Single Flow Group information in the NLRI (BGP EVPN Network Layer Reachability Information).

3.2. ESI Label Extended Community in S-PMSI A-D Routes

The Hot Standby solution requires the advertisement of one or more ESI Label Extended Communities [RFC7432] alongside the S-PMSI A-D routes. These extended communities encode the ESI values associated with an S-PMSI A-D (*,G) or (S,G) route that advertises the presence of a Single Flow Group.

Key considerations include:

1. When advertised with the S-PMSI A-D routes, only the ESI Label value in the extended community is relevant to the procedures defined in this document.
2. The Flags field within the extended community MUST be set to '0x00' on transmission and MUST be ignored on reception.

[RFC7432] specifies the use of the ESI Label Extended Community in conjunction with the A-D per ES route. This document extends the applicability of the ESI Label Extended Community by allowing its inclusion multiple times (with different ESI Label values) alongside the EVPN S-PMSI A-D route. These extensions enable the precise encoding and advertisement of Single Flow Group-related information, facilitating efficient multicast traffic handling in EVPN networks.

4. Warm Standby (WS) Solution for Redundant G-Sources

This section specifies the Warm Standby (WS) solution for handling redundant multicast sources (G-sources). Note that while the examples use IPv4 addresses, the solution supports both IPv4 and IPv6 sources.

4.1. Specification

The Warm Standby solution follows these general procedures:

1. Configuration of the upstream PEs

Upstream PEs, potentially connected to redundant G-sources, are configured to recognize:

- * The multicast groups that carry an SFG in the tenant domain.
- * The local Broadcast Domains that may host redundant G-sources

The SFG configuration applies to either 'any' source, i.e., (*) or to a specific 'source prefix' (e.g., "192.0.2.0/30" or "2001:db8::/120"). For instance, if the IPv4 prefix is 192.0.2.0/30, the sources 192.0.2.1 and 192.0.2.2 are considered redundant G-sources for the SFG, while 192.0.2.10 is not. In a different example for IPv6, if the prefix is 2001:db8::/120, sources 2001:db8::1 or 2001:db8::2 are considered redundant G-sources for the SFG, but 2001:db8:1::1 is not.

Example Configuration (Figure 4):

- * PE1 is configured to recognize G1 as an SFG for any source, with potential redundant G-sources attached to Broadcast Domains BD1 and BD2.
- * Alternatively, PE1 may recognize G1 as an SFG for sources within 192.0.2.0/30 (or 2001:db8::/120), with redundant G-sources in BD1 and BD2.

2. Signaling the location of a G-source for an SFG

Upon receiving the first IP multicast packet for a configured SFG on a Broadcast Domain, an upstream PE (e.g., PE1):

- * MUST advertise an S-PMSI A-D route for the SFG:
 - An (*,G) route if the SFG is configured for any source.
 - An (S,G) route (where S can have any prefix length) if the SFG is configured for a source prefix.
- * MUST include the following attributes in the S-PMSI A-D route:
 - Route Targets (RTs): the Supplementary Broadcast Domain Route Target (SBD-RT), if applicable, and the Broadcast Domain Route Target (BD-RT) of the Broadcast Domain receiving the traffic. The SBD-RT is needed so that the route is imported by all PEs attached to the tenant domain in an OISM solution.
 - Multicast Flags Extended Community: that MUST include the SFG flag to indicate that the route conveys an SFG.
 - Designated Forwarder Election Extended Community: specifies the algorithm and preferences for the Single Forwarder election, using the Designated Forwarder election defined in [RFC8584].
- * Advertises the route:
 - Without a PMSI Tunnel Attribute if using Ingress Replication (IR), Assisted Replication (AR), or Bit Indexed Explicit Replication (BIER).
 - With a PMSI Tunnel Attribute for other tunnel types.
- * MUST withdraw the S-PMSI A-D route when the SFG traffic ceases. A timer is RECOMMENDED to detect inactivity and trigger route withdrawal.

3. Single Forwarder Election on the upstream PEs

If an upstream PE receives one or more S-PMSI A-D routes for the same SFG from remote PEs, it performs Single Forwarder Election based on the Designated Forwarder Election Extended Community.

- * Two routes are considered part of the same SFG if they are advertised for the same tenant and match on the following fields:

- Multicast Source Length
- Multicast Source
- Multicast Group Length
- Multicast Group

* Election Rules:

1. A consistent Designated Forwarder Election Algorithm MUST be used across all upstream PEs for the Single Forwarder election. In OISM networks, the Default Designated Forwarder Election Algorithm MUST NOT be used if redundant G-sources are attached to Broadcast Domains with different Ethernet Tags.
2. In case of a mismatch in the Designated Forwarder Election Algorithm or capabilities, the tie-breaker is the lowest PE IP address (as advertised in the Originator Address field of the S-PMSI A-D route).

4. Reverse Path Forwarding Checks on Upstream PEs

All PEs with a local G-source for an SFG apply a Reverse Path Forwarding check to the (*,G) or (S,G) state based on the Single Forwarder election result:

1. Non-Single Forwarder PEs: MUST discard all (*,G) or (S,G) packets received on local Attachment Circuits.
2. Single Forwarder PEs: MUST forward (*,G) or (S,G) packets received on one (and only one) local Attachment Circuit.

Key Features of the Warm Standby Solution:

- * The solution ensures redundancy for SFGs without requiring upgrades to downstream PEs (where no redundant G-sources are connected).
- * Existing procedures for non-SFG G-sources remain unchanged.
- * Redundant G-sources can be either single-homed or multi-homed. Multi-homing does not alter the above procedures.

Examples of the Warm Standby solution are provided in Section 4.2 and Section 4.3.

4.2. Warm Standby Example in an OISM Network

Figure 4 illustrates an example where S1 and S2 are redundant G-sources for the Single Flow Group (*,G1).

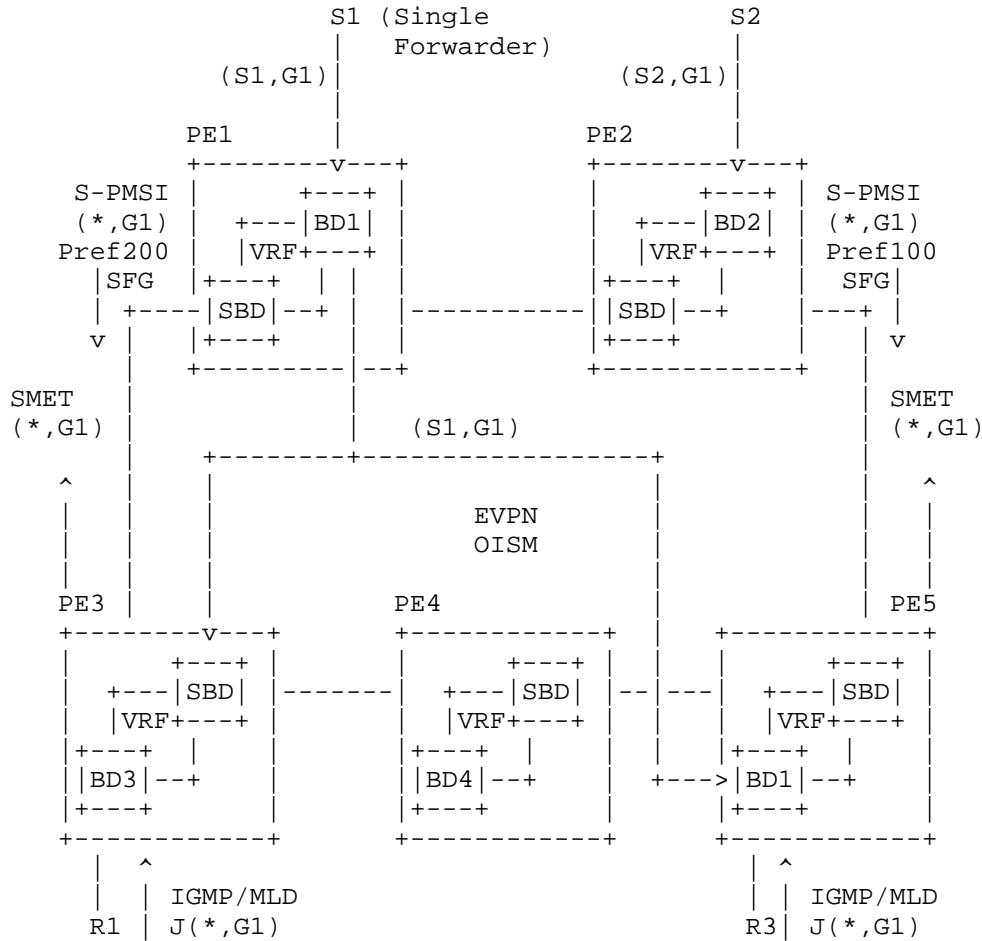


Figure 4: Warm Standby Solution for Redundant G-Sources

The Warm Standby procedure is as follows:

1. Configuration of the upstream PEs (PE1 and PE2)

- * PE1 and PE2 are configured to recognize G1 as a Single Flow Group for any source.

- * Redundant G-sources for G1 may be attached to Broadcast Domains BD1 (for PE1) and BD2 (for PE2).

2. Signaling the location of S1 and S2 for (*,G1)

Upon receiving traffic for G1 on a local Attachment Circuit:

- * PE1 and PE2 originate S-PMSI A-D routes for (*,G1), including:
 - The Supplementary Broadcast Domain Route Target (SBD-RT),
 - The Designated Forwarder Election Extended Community, and
 - The SFG flag in the Multicast Flags Extended Community.
- * These routes indicate the presence of a Single Flow Group.

3. Single Forwarder Election

- * Based on the Designated Forwarder Election Extended Community, PE1 and PE2 perform Single Forwarder election.
- * Assuming they use Preference-based Election [I-D.ietf-bess-evpn-pref-df], PE1 (with a higher preference) is elected as the Single Forwarder for (*,G1).

4. Reverse Path Forwarding check on the PEs attached to a redundant G-source

- a. Non-Single Forwarder Behavior: PE2 discards all (*,G1) packets received on its local Attachment Circuit.
- b. Single Forwarder Behavior: PE1 forwards (*,G1) packets received on one (and only one) local Attachment Circuit.

The outcome:

- * Upon receiving IGMP/MLD reports for (*,G1) or (S,G1), downstream PEs (PE3 and PE5) issue SMET routes to pull the multicast Single Flow Group traffic from PE1 only.
- * In the event of a failure of S1, the Attachment Circuit connected to S1, or PE1 itself, the S-PMSI A-D route for (*,G1) is withdrawn by PE1.
- * As a result, PE2 is promoted to Single Forwarder, ensuring continued delivery of (*,G1) traffic.

4.3. Warm Standby Example in a Single-BD Tenant Network

Figure 5 illustrates an example where S1 and S2 are redundant G-sources for the Single Flow Group (*,G1). In this case, all G-sources and receivers are connected to the same Broadcast Domain (BD1), and there is no Supplementary Broadcast Domain (SBD).

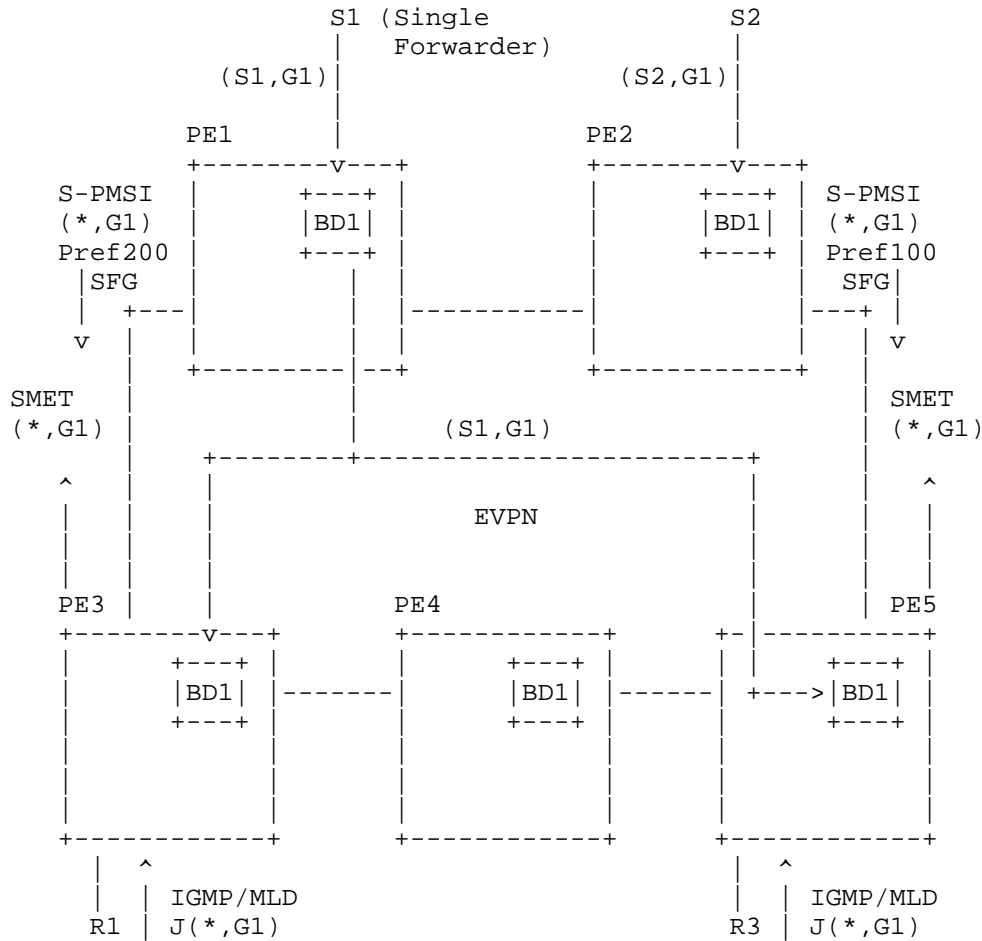


Figure 5: WS Solution for Redundant G-Sources in the same BD

The procedures for the Warm Standby solution in this example are identical to those described in Section 4.2, with the following distinction:

- * Signaling the S-PMSI A-D Routes

- Upon receiving traffic for the Single Flow Group (*,G1), PE1 and PE2 advertise S-PMSI A-D routes.
- These routes include only the BD1-RT (Broadcast Domain 1 Route Target) as there is no Supplementary Broadcast Domain (SBD) in this scenario.

This example represents a specific sub-case of the broader procedure detailed in Section 4.2, adapted to a single Broadcast Domain environment. The absence of an SBD simplifies the configuration, as all signaling remains within the context of BD1.

5. Hot Standby Solution for Redundant G-Sources

This section specifies the Hot Standby solution for handling redundant multicast sources (G-sources). The solution supports both IPv4 and IPv6 sources.

5.1. Specification

The Hot Standby solution is designed for scenarios requiring fast failover in the event of a G-source or upstream PE failure. It assumes that additional bandwidth consumption in the tenant network is acceptable. The procedure is as follows:

1. Configuration of PEs

- * Upstream PEs are configured to identify Single Flow Groups in the tenant domain. This includes groups for any source or a source prefix containing redundant G-sources.
- * Each redundant G-source MUST be associated with an Ethernet Segment on the upstream PEs. This applies to both single-homed and multi-homed G-sources. For both, single-homed and multi-homed G-sources, ESI labels are essential for Reverse Path Forwarding checks on downstream PEs. The term S-ESI is used to denote the ESI associated with a redundant G-source.
- * Unlike the Warm Standby solution, the Hot Standby solution requires downstream PEs to support the procedure.
- * Downstream PEs:
 - Do not need explicit configuration for Single Flow Groups or their ESIs (since they get that information from the upstream PEs).

- Dynamically select an ESI for each Single Flow Group based on local policy (hence different downstream PEs may select different Ethernet Segment Identifiers) and program a Reverse Path Forwarding check to discard (*,G) or (S,G) packets from other ESIs.

2. Signaling the location of a G-source for a given SFG and its association to the local Ethernet Segments

An upstream PE configured for Hot Standby procedures:

- * MUST advertise an S-PMSI A-D route for each Single Flow Group. These routes:
 - Use the Broadcast Domain Route Target (BD-RT) and, if applicable, the Supplementary Broadcast Domain Route Target (SBD-RT) so that the routes are imported in all the PEs of the tenant domain.
 - MUST include ESI Label Extended Communities to convey the S-ESI labels associated with the Single Flow Group. These ESI-labels match the labels advertised by the EVPN A-D per ES routes for each S-ES.
- * MAY include a PMSI Tunnel Attribute, depending on the tunnel type, as specified in the Warm Standby procedure.
- * MUST trigger the S-PMSI A-D route advertisement based on the SFG configuration (and not based on the reception of traffic).

3. Distribution of DCB ESI-labels and G-source ES routes

- * Upstream PEs advertise corresponding EVPN routes:
 - EVPN Ethernet Segment (ES) routes for the local S-ESIs. ES routes are used for regular Designated Forwarder Election for the S-ES. This document does not introduce any change in the procedures related to the EVPN ES routes.
 - A-D per EVI and A-D per ES routes for tenant-specific traffic. If the SBD exists, the EVPN A-D per EVI and A-D per ES routes MUST include the route target SBD-RT since they have to be imported by all the PEs in the tenant domain.
- * ESI Label Procedures:

- The EVPN A-D per ES routes convey the S-ESI labels that the downstream PEs use to implement Reverse Path Forwarding checks for SFGs.
- All packets for a given G-source MUST carry the same S-ESI label. For example, if two redundant G-sources are multi-homed to PE1 and PE2 via S-ES-1 and S-ES-2, PE1 and PE2 MUST allocate the same ESI label "Lx" for S-ES-1 and they MUST allocate the same ESI label "Ly" for S-ES-2. In addition, Lx and Ly MUST be different.
- S-ESI labels are allocated as Domain-wide Common Block (DCB) labels and follow the procedures in [RFC9573]. In addition, the PE indicates that these ESI labels are DCB labels by using the extensions described in Section 5.2.

4. Processing of EVPN A-D per ES/EVI routes and Reverse Path Forwarding check on the downstream PEs

The EVPN A-D per ES/EVI routes are received and imported in all the PEs in the tenant domain. Downstream PEs process received EVPN A-D per ES/EVI routes based on their configuration:

- * The PEs attached to the same Broadcast Domain of the route target BD-RT that is included in the EVPN A-D per ES/EVI routes process the routes as in [RFC7432] and [RFC8584]. If the receiving PE is attached to the same Ethernet Segment as indicated in the route, [RFC7432] split-horizon procedures are followed and the Designated Forwarder Election candidate list is modified as in [RFC8584] if the Ethernet Segment supports the AC-DF (Attachment Circuit influenced Designated Forwarder) capability.
- * The PEs that are not attached to the Broadcast Domain identified by the route target BD-RT but are attached to the Supplementary Broadcast Domain of the received route target SBD-RT, MUST import the EVPN A-D per ES/EVI routes and use them for redundant G-source mass withdrawal, as explained later.
- * Upon importing EVPN A-D per ES routes corresponding to different S-ESes, a PE MUST select a primary S-ES based on local policy, and add a Reverse Path Forwarding check to the (*,G) or (S,G) state in the Broadcast Domain or Supplementary Broadcast Domain. This Reverse Path Forwarding check discards all ingress packets to (*,G)/(S,G) that are not received with the ESI-label of the primary S-ES.

5. G-traffic forwarding for redundant G-sources and fault detection

* Traffic Forwarding with S-ESI Labels:

- When there is an existing (*,G) or (S,G) state for the SFG with output interface list entries associated with remote EVPN PEs, the upstream PE will add an S-ESI label to the bottom of the stack when forwarding G-traffic received on an S-ES. This label is allocated from a domain-wide common block as described in Step 3.
- If Point-to-multipoint or BIER PMSIs are used, this procedure does not introduce new data path requirements on the upstream PEs, apart from allocating the S-ESI label from the domain-wide common block as per [RFC9573]). However, when Ingress Replication or Assisted Replication are employed, this document extends the procedures defined in [RFC7432]. In these scenarios, the upstream PE pushes the S-ESI labels on packets not only destined for PEs sharing the ES but also for all PEs within the tenant domain. This ensures that downstream PEs receive all the multicast packets from the redundant G-sources with an S-ESI label, regardless of the PMSI type or local ESes. Downstream PEs will discard any packet carrying an S-ESI label different from the primary S-ESI label (associated with the selected primary S-ES), as outlined in Step 4.

* Handling Route Withdrawals and Fault Detection

- If the last EVPN A-D per EVI or the last EVPN A-D per ES route for the primary S-ES is withdrawn, the downstream PE will immediately select a new primary S-ES and update the Reverse Path Forwarding check accordingly.
- For scenarios where the same S-ES is used across multiple tenant domains by the upstream PEs, the withdrawal of all the EVPN A-D per-ES routes associated with an S-ES enables a mass withdrawal mechanism. This allows the downstream PE to simultaneously update the Reverse Path Forwarding check for all tenant domains that rely on the same S-ES.

* Removal of Reverse Path Forwarding Checks on S-PMSI Withdrawal

- The withdrawal of the last EVPN S-PMSI A-D route for a given (*,G) or (S,G) that represents an SFG SHOULD result in the downstream PE removing the S-ESI label-based Reverse Path Forwarding check for that (*,G) or (S,G).

This document supports the use of Context Label Space ID Extended Communities, as described in [RFC9573], for scenarios where S-ESI labels are allocated within context label spaces. When the context label space ID extended community is advertised along with the ESI label in an EVPN A-D per ES route, the ESI label is from a context label space identified by the Domain-wide Common Block (DCB) label in the Extended Community.

5.2. Extensions for the Advertisement of DCB Labels

Domain-wide Common Block Labels are specified in [RFC9573] and this document makes use of them as outlined in Section 5.1. [RFC9573] assumes that Domain-wide Common Block labels are applicable only to Multipoint-to-Multipoint, Point-to-Multipoint, or BIER tunnels. Additionally, it specifies that when a PMSI label is a Domain-wide Common Block label, the ESI label used for multi-homing is also a Domain-wide Common Block label.

This document extends the use of DCB-allocated ESI labels with the following provisions:

- a. DCB-allocated ESI labels MAY be used with Ingress Replication tunnels, and
- b. DCB-allocated ESI labels MAY be used by PEs that do not use DCB-allocated PMSI labels.

These control plane extensions are indicated in the EVPN A-D per ES routes for the relevant S-ESs by:

1. Adding the ESI-DCB-flag (Domain-wide Common Block flag) to the ESI Label Extended Community, or
2. Adding the Context Label Space ID extended community

The encoding of the DCB-flag within the ESI Label Extended Community is shown below:




Figure 6: ESI Label Extended Community

This document defines the DCB-flag as follows:

- * Bit 5 of the Flags octet in the ESI Label Extended Community is defined as the ESI-DCB-flag by this document.
- * When the ESI-DCB-flag is set, it indicates that the ESI label is a DCB label.

Criteria for identifying a DCB label:

An ESI label is considered a DCB label if either of the following conditions is met:

- a. The ESI label is encoded in an ESI Label Extended Community with the ESI-DCB-flag set.
- b. The ESI label is signaled by a PE that has advertised a PMSI label that is a DCB label.

As in [RFC9573] this document also permits the use of context label space ID extended community. When this extended community is advertised along with the ESI label in an EVPN A-D per ES route, it indicates that the ESI label is from a context label space identified by the DCB label in the Extended Community.

5.3. Use of BFD in the HS Solution

In addition to utilizing the state of the EVPN A-D per EVI, EVPN A-D per ES or S-PMSI A-D routes to adjust the Reverse Path Forwarding checks for (*,G) or (S,G) as discussed in Section 5.1, the Bidirectional Forwarding Detection (BFD) protocol MAY be employed to monitor the status of the multipoint tunnels used to forward the SFG packets from redundant G-sources.

BFD integration:

- * The BGP-BFD Attribute is advertised alongside the S-PMSI A-D or Inclusive Multicast Ethernet Tag routes, depending on whether Inclusive PMSI or Selective PMSI trees are being utilized.
- * The procedures outlined in [I-D.ietf-mppls-p2mp-bfd] are followed to bootstrap multipoint BFD sessions on the downstream PEs.

5.4. Hot Standby Example in an OISM Network

This section describes the Hot Standby model applied in an Optimized Inter-Subnet Multicast (OISM) network. Figure 7 and Figure 8 illustrate scenarios with multi-homed and single-homed redundant G-sources, respectively.

5.4.1. Multi-Homed Redundant G-Sources

Scenario (Figure 7):

- * S1 and S2 are redundant G-sources for the Single Flow Group (*,G1), connected to Broadcast Domain BD1.
- * S1 and S2 are all-active multi-homed to upstream PEs (PE1 and PE2).
- * Receivers are connected to downstream PEs (PE3 and PE5) in Broadcast Domains BD3 and BD1, respectively.
- * S1 and S2 are connected to the multi-homing PEs using a LAG. Multicast traffic can traverse either link.
- * In this model, downstream PEs receive duplicate G-traffic for (*,G1) and must use Reverse Path Forwarding checks to avoid delivering duplicate packets to receivers.

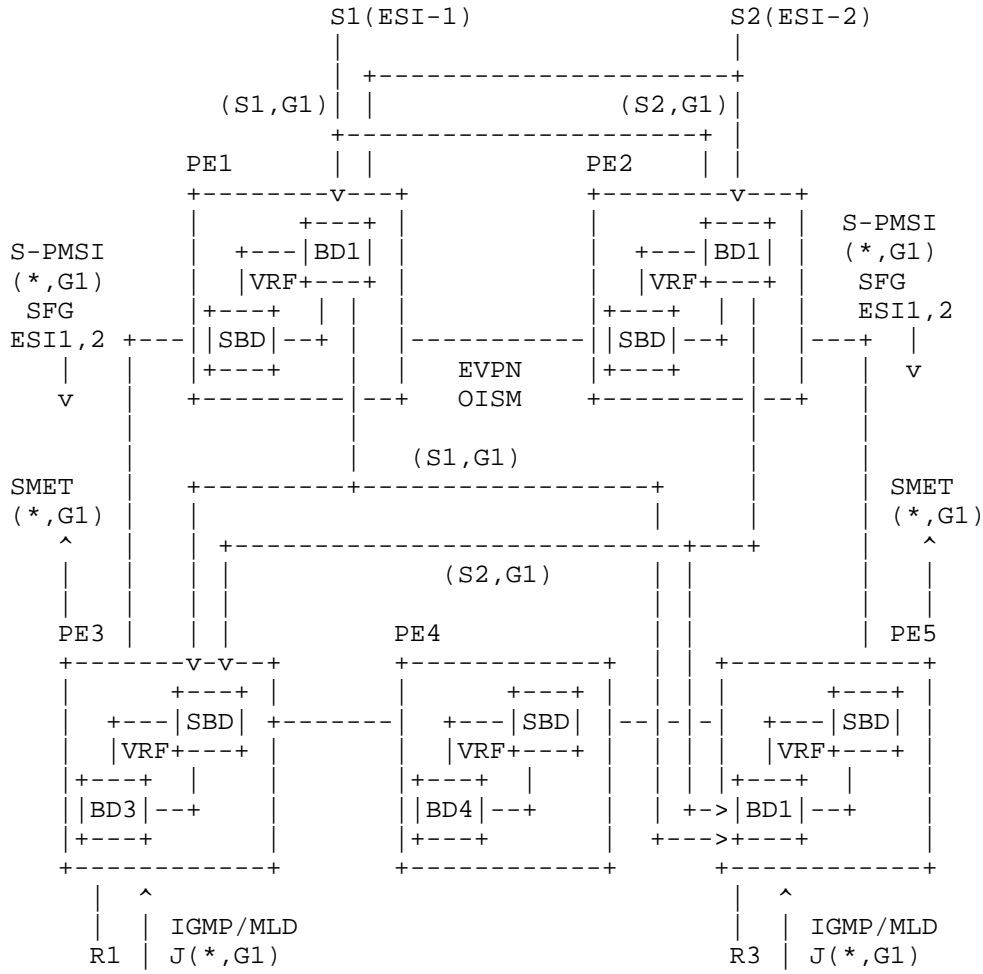


Figure 7: HS Solution for Multi-homed Redundant G-Sources in OISM

The procedure is as follows:

1. Configuration of the PEs:

- * PE1 and PE2 are configured to recognize (*,G1) as a Single Flow Group.
- * Redundant G-sources use S-ESIs: ESI-1 for S1 and ESI-2 for S2.

- * The Ethernet Segments (ES-1 and ES-2) are configured on both PEs. ESI-labels are allocated from a Domain-wide Common Block (DCB) [RFC9573] - ESI-label-1 for ESI-1 and ESI-label-2 for ESI-2.
- * The downstream PEs, PE3, PE4 and PE5 are configured to support Hot Standby mode and select the G-source with, e.g., lowest ESI value.

2. Advertisement of the EVPN routes:

- * PE1 and PE2 advertise S-PMSI A-D routes for (*,G1), including:
 - Route Targets: BD1-RT and SBD-RT.
 - ESI Label Extended Communities for ESI-label-1 and ESI-label-2.
 - The SFG flag indicating that (*,G1) represents a Single Flow Group.
- * EVPN ES and A-D per ES/EVI routes are also advertised for ESI-1 and ESI-2. These include SBD-RT for downstream PE import. The EVPN A-D per ES routes contain ESI-label-1 for ESI-1 (on both PEs) and ESI-label-2 for ESI-2 (also on both PEs).

3. Processing of EVPN A-D per ES/EVI routes and Reverse Path Forwarding check on Downstream PEs:

- * PE1 and PE2 receive each other's ES and A-D per ES/EVI routes. Designated Forwarder Election and programming of the ESI-labels for egress split-horizon filtering follow, as specified in [RFC7432] and [RFC8584].
- * PE3/PE4 import the EVPN A-D per ES/EVI routes in the SBD, PE5 imports them in BD1.
- * As downstream PEs, PE3 and PE5 use the EVPN A-D per ES/EVI routes to program Reverse Path Forwarding checks.
- * The primary S-ESI for (*,G1) is selected based on local policy (e.g., lowest ESI value), and therefore packets with ESI-label-2 are discarded if ESI-label-1 is selected as the primary label.

4. Traffic forwarding and fault detection:

- * PE1 receives (S1,G1) traffic and forwards it with ESI-label-1 in the context of BD1. This traffic passes Reverse Path Forwarding checks on downstream PEs (PE3 and PE5, since PE4 has no local interested receivers) and is delivered to receivers.
- * PE2 receives (S2,G1) traffic and forwards it with ESI-label-2. This traffic fails the Reverse Path Forwarding check on PE3 and PE5 and is discarded.
- * If the link between S1 and PE1 fails, PE1 withdraws the EVPN ES and A-D routes for ESI-1. S1 forwards the (S1,G1) traffic to PE2 instead. PE2 continues forwarding (S2,G1) traffic using ESI-label-2 and now also forwards (S1,G1) with ESI-label-1. The Reverse Path Forwarding checks do not change in PE3/PE5.
- * If all links to S1 fail, PE2 also withdraws the EVPN ES and A-D routes for ESI-1 and downstream PEs update the Reverse Path Forwarding checks to accept ESI-label-2 traffic.

5.4.2. Single-Homed Redundant G-Sources

Scenario (Figure 8):

- * S1 is single-homed to PE1 using ESI-1, and S2 is single-homed to PE2 using ESI-2.
- * The scenario is a subset of the multi-homed case. Only one PE advertises EVPN A-D per ES/EVI routes for each S-ESI.




Figure 7 and Figure 8 demonstrate the application of the Hot Standby solution, ensuring seamless traffic forwarding while avoiding duplication in the presence of redundant G-sources.

5.5. Hot Standby Example in a Single-BD Tenant Network

The Hot Standby procedures described in Section 5.4 apply equally to scenarios where the tenant network comprises a single Broadcast Domain (e.g., BD1), irrespective of whether the redundant G-sources are multi-homed or single-homed. In such cases:

- * The advertised routes do not include the Supplementary Broadcast Domain Route Target (SBD-RT).
- * All procedures are confined to the single Broadcast Domain (BD1).

The absence of the SBD simplifies the configuration and limits the scope of the Hot Standby solution to BD1, while maintaining the integrity of the procedures for managing redundant G-sources.

6. Security Considerations

The same Security Considerations described in [RFC9625] are valid for this document.

From a security perspective, out of the two methods described in this document, the Warm Standby method is considered lighter in terms of control plane and therefore its impact is low on the processing capabilities of the PEs. The Hot Standby method adds more burden on the control plane of all the PEs of the tenant with sources and receivers.

7. IANA Considerations

IANA is requested to allocate bit 4 in the Multicast Flags Extended Community registry that was introduced by [RFC9251]. This bit indicates that a given (*,G) or (S,G) in an S-PMSI A-D route is associated with an SFG. This bit is called "Single Flow Group" bit and it is defined as follows:

Bit	Name	Reference
4	Single Flow Group	This Document

Table 1

IANA is requested to allocate bit 5 in the ESI Label Extended Community Flags registry that was introduced by [I-D.ietf-bess-evpn-mh-split-horizon]. This bit is the ESI-DCB flag and indicates that the ESI label contained in the ESI Label Extended Community is a Domain-wide Common Block label. This bit is defined as follows:

Bit	Name	Reference
5	ESI-DCB Flag	This Document

Table 2

8. Acknowledgments

The authors would like to thank Mankamana Mishra, Ali Sajassi, Greg Mirsky and Sasha Vainshtein for their review and valuable comments. Special thanks to Gunter van de Velde for his excellent review, which significantly enhanced the document's readability.

9. Contributors

In addition to the authors listed on the front page, the following people have significantly contributed to this document:

Eric C. Rosen

Email: erosen52@gmail.com

10. References

10.1. Normative References

- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.

- [RFC9251] Sajassi, A., Thoria, S., Mishra, M., Patel, K., Drake, J., and W. Lin, "Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Proxies for Ethernet VPN (EVPN)", RFC 9251, DOI 10.17487/RFC9251, June 2022, <<https://www.rfc-editor.org/info/rfc9251>>.
- [RFC9625] Lin, W., Zhang, Z., Drake, J., Rosen, E., Ed., Rabadan, J., and A. Sajassi, "EVPN Optimized Inter-Subnet Multicast (OISM) Forwarding", RFC 9625, DOI 10.17487/RFC9625, August 2024, <<https://www.rfc-editor.org/info/rfc9625>>.
- [RFC8584] Rabadan, J., Ed., Mohanty, S., Ed., Sajassi, A., Drake, J., Nagaraj, K., and S. Sathappan, "Framework for Ethernet VPN Designated Forwarder Election Extensibility", RFC 8584, DOI 10.17487/RFC8584, April 2019, <<https://www.rfc-editor.org/info/rfc8584>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC9573] Zhang, Z., Rosen, E., Lin, W., Li, Z., and IJ. Wijnands, "MVPN/EVPN Tunnel Aggregation with Common Labels", RFC 9573, DOI 10.17487/RFC9573, May 2024, <<https://www.rfc-editor.org/info/rfc9573>>.
- [I-D.ietf-bess-evpn-mh-split-horizon]
Rabadan, J., Nagaraj, K., Lin, W., and A. Sajassi, "BGP EVPN Multi-Homing Extensions for Split Horizon Filtering", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-mh-split-horizon-11, 17 August 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-mh-split-horizon-11>>.

10.2. Informative References

- [RFC9136] Rabadan, J., Ed., Henderickx, W., Drake, J., Lin, W., and A. Sajassi, "IP Prefix Advertisement in Ethernet VPN (EVPN)", RFC 9136, DOI 10.17487/RFC9136, October 2021, <<https://www.rfc-editor.org/info/rfc9136>>.

- [RFC9572] Zhang, Z., Lin, W., Rabadan, J., Patel, K., and A. Sajassi, "Updates to EVPN Broadcast, Unknown Unicast, or Multicast (BUM) Procedures", RFC 9572, DOI 10.17487/RFC9572, May 2024, <<https://www.rfc-editor.org/info/rfc9572>>.
- [I-D.ietf-bess-evpn-pref-df] Rabadan, J., Sathappan, S., Lin, W., Drake, J., and A. Sajassi, "Preference-based EVPN DF Election", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-pref-df-13, 9 October 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-pref-df-13>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [I-D.ietf-mpls-p2mp-bfd] Mirsky, G., Mishra, G. S., and D. E. Eastlake, "Bidirectional Forwarding Detection (BFD) for Multipoint Networks over Point-to-Multi-Point MPLS Label Switched Path (LSP)", Work in Progress, Internet-Draft, draft-ietf-mpls-p2mp-bfd-09, 6 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-p2mp-bfd-09>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC9135] Sajassi, A., Salam, S., Thoria, S., Drake, J., and J. Rabadan, "Integrated Routing and Bridging in Ethernet VPN (EVPN)", RFC 9135, DOI 10.17487/RFC9135, October 2021, <<https://www.rfc-editor.org/info/rfc9135>>.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.

- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", RFC 8296, DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.
- [RFC9574] Rabadan, J., Ed., Sathappan, S., Lin, W., Katiyar, M., and A. Sajassi, "Optimized Ingress Replication Solution for Ethernet VPNs (EVPNs)", RFC 9574, DOI 10.17487/RFC9574, May 2024, <<https://www.rfc-editor.org/info/rfc9574>>.

Authors' Addresses

Jorge Rabadan (editor)
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America
Email: jorge.rabadan@nokia.com

Jayant Kotalwar
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085 USA
Email: jayant.kotalwar@nokia.com

Senthil Sathappan
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085 USA
Email: senthil.sathappan@nokia.com

Zhaohui Zhang
Juniper Networks
Email: zzhang@juniper.net

Wen Lin
Juniper Networks
Email: wlin@juniper.net