

BESS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 27 August 2026

P. Brissette, Ed.
M. MacKenzie, Ed.
Cisco
S. Matsushima
Softbank
W. Lin
Juniper
J. Rabadan
Nokia
23 February 2026

EVPN multi-homing support for L3 services
draft-ietf-bess-evpn-l3mh-proto-00

Abstract

This document describes the use of EVPN Ethernet Segment Link Aggregation Group (ES-LAG) technology to provide multi-homing redundancy for Layer 3 services. The solution synchronizes ARP/ND, multicast state, and IGP routes between redundant PEs without requiring Layer 2 constructs or proprietary Inter-Chassis Communication protocols.

Requirements Language

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.

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 27 August 2026.

Copyright Notice

Copyright (c) 2026 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. Problems with unicast load-balancing from core to CE	4
1.2. Problems with multicast from core to CE	4
1.3. Problems with IGP adjacencies over the LAG port	4
1.4. Problems with supporting multiple subnets on same ES in all active mode	5
1.5. Acronyms	6
1.6. Requirements	7
2. Solution Overview	8
3. Solution Details	8
3.1. Example Topology	8
3.2. Route Target Usage	10
3.3. EVPN Instance Usage	11
3.4. Mapping for L3 Interface to ESI	11
3.5. Mapping for L3 Sub-Interface to Ethernet Tag-id	11
3.6. ARP/ND Synchronization	12
3.6.1. Advertising ARP/ND Routes	12
3.6.2. Processing Received ARP/ND Routes	13
3.7. IGMP/MLD Synchronization	14
3.7.1. Advertising IGMP/MLD Routes	14
3.7.2. Processing Received IGMP/MLD Routes	14
3.8. IGP Route Synchronization	14
3.8.1. ESI-Based Approach	15
3.8.2. IP Gateway-Based Approach	15
4. Convergence Considerations	16
5. Overall Advantages	16
6. Security Considerations	17
7. IANA Considerations	17
8. Acknowledgments	17
9. References	17
9.1. Normative References	17
9.2. Informative References	18

Appendix A. Contributors	18
Authors' Addresses	19

1. Introduction

Resilient L3VPN service to a CE requires multiple service PEs to run a Multi-Chassis Link Aggregation Group mechanism, which previously required a proprietary ICL control plane link between them.

This document uses [RFC7432], [RFC9135] and [RFC9136] procedures to bring EVPN based ES-LAG all-active multi-homing load-balancing to L3 services focusing on the L3VPN [RFC4364] use case to provide examples.

EVPN ES-LAG is completely transparent to a CE device, and provides link and node level redundancy with load-balancing using the existing BGP control plane required by the L3 services.

For example, the L3VPN service can be MPLS, VxLAN or SRv6 based, and does not require EVPN signaling to remote neighbors. The EVPN signaling is limited to the redundant service PEs sharing a Ethernet Segment Identifier (ESI). This is used to synchronize ARP/ND, multicast Join/Leave, and IGP routes replacing need for ICL link.




Figure 1: EVPN ES-LAG Topology

Figure 1 shows an ES-LAG multi-homing topology where PE1 and PE2 are part of the same redundancy group providing multi-homing to CE1 via interfaces I1 and I2. PE1, PE2 and PE3 are attached to the same L3VPN thru the core (running [RFC4364] and/or [RFC9136] procedures).

Interfaces I1 and I2 are Bundle-Ethernet interfaces running LACP protocol. The CE device can be a layer-2 or layer-3 device connecting to the redundant PEs over a single LACP LAG port.

In the case of a layer-3 CE device, this document looks to solve the case of an IGP adjacency between PEs and CE. Further study is needed to support BGP PE to CE protocols. The core, shown as IP or MPLS enabled, provides wide range of L3 services. ES-LAG multi-homing functionality is decoupled from those services in the core and it focuses on providing multi-homing to CE.

To deliver resilient layer-3 services and provide traffic load-balancing towards the access, the two service PEs advertise layer-3 reachability towards the layer-3 core and both be eligible to receive traffic and forward towards the Access.

1.1. Problems with unicast load-balancing from core to CE

The layer-2 hashing performed by CE over its LAG port means that only one service PE may populate its ARP/ND cache. In Figure 1, if CE1 ARP/ND responses always hash to PE1, then PE2's ARP/ND table remains empty. Traffic from remote PEs can be received by either service PE. Traffic that reaches PE2 does not find an ARP entry and is dropped.

Solution: Synchronize ARP/ND entries using EVPN RT-2 routes as described in Section 3.6.

1.2. Problems with multicast from core to CE

Multicast IGMP/MLD join messages from CE may always hash to a single PE due to LAG hashing behavior. When PIM runs on both redundant PEs, PIM hello messages from each PE are not visible to the other PE because the CE cannot switch traffic between LAG members. Both PEs become PIM Designated Router (DR). However, IGMP joins for a given multicast group may hash to only one PE, so only that PE programs the multicast route and sends PIM joins.

Solution: Synchronize IGMP/MLD state using EVPN RT-7/RT-8 routes as described in Section 3.7.

1.3. Problems with IGP adjacencies over the LAG port

A layer-3 CE device connecting to redundant PEs may establish an IGP adjacency on the bundle port. The adjacency forms to only one PE, so IGP customer routes are only present on that PE. This prevents load-balancing benefits as only one PE advertises customer routes to the core.

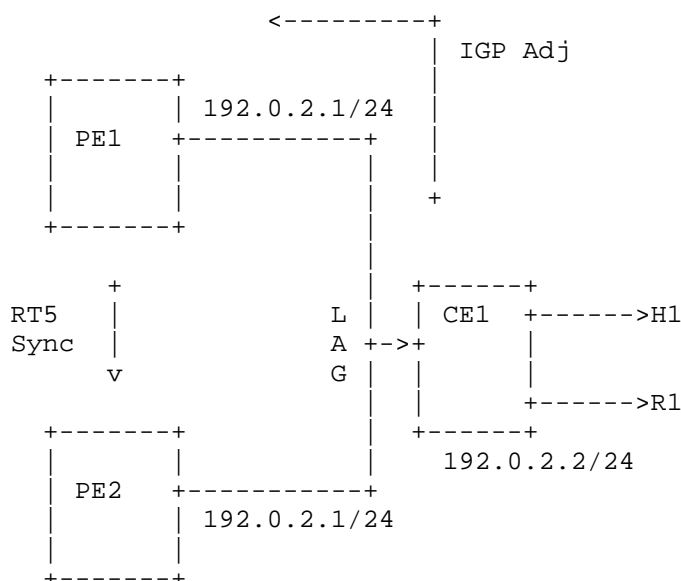


Figure 2: IGP Adjacency over LAG Port

Figure 2 provides an example of this use case, where CE1 forms an IGP adjacency with PE1 (example: ISIS or OSPF), and advertises its H1 and R1 routes into the IP-VRF of PE1. PE1 may then redistribute this IGP route into the core as an L3 service. Any remote PEs are only aware of the service from PE1, and cannot load balance through PE2 as well.

Solution: Synchronize IGP learned routes using EVPN RT-5 routes as described in Section 3.8.

Note: BGP PE-CE protocols require further study.

1.4. Problems with supporting multiple subnets on same ES in all active mode

When a CE supports multiple subnets using VLANs over a single LAG interface, each VLAN maps to a separate L3 sub-interface on the PE. When the PE synchronizes host reachability using EVPN RT-2 routes, standard RT-2 advertisements do not indicate which sub-interface (VLAN) the host belongs to. The peering PE cannot determine the correct destination sub-interface when multiple sub-interfaces share the same ESI.

The same problem occurs with IGMP/MLD route synchronization using RT-7 and RT-8.

Solution: Use the Ethernet Tag-ID field to carry the VLAN ID in all route sync messages (RT-2, RT-5, RT-7, RT-8) to identify the specific sub-interface.

Note: This document focuses on L3 sub-interfaces. Mixed L2/L3 sub-interfaces require further study.

1.5. Acronyms

BD: Broadcast Domain

BE: Bundle Ethernet Interface aka. L3 LAG interface

DF: Designated Forwarder

DR: Multicast Designated Router

EC: BGP Extended Community

ES: Ethernet Segment. When a customer site (device or network) is connected to one or more PEs via a set of Ethernet links, then that set of links is referred to as an 'Ethernet Segment'.

ESI: Ethernet Segment Identifier. A unique non-zero identifier that identifies an Ethernet Segment is called an 'Ethernet Segment Identifier'.

ES-LAG: This refers to multi-homing scenario where peering PEs, connected to same CE, are two, three or more.

ETAG: Ethernet Tag. An Ethernet tag identifies a particular broadcast domain, e.g., a VLAN. An EVPN instance consists of one or more broadcast domains.

EVI: An EVPN instance spanning the Provider Edge (PE) devices participating in that EVPN. It is used to assist a L3 VRF for route synchronization.

GRT: Global Routing Table

ICL: Inter Chassis Link

IGMP: Internet Group Management Protocol

IGP: Interior Gateway Protocol

IP-VRF: A VPN Routing and Forwarding table for IP routes on an PE.

The IP routes could be populated by EVPN and IP-VPN address families. An IP-VRF is also an instantiation of a layer 3 VPN in an PE.

MAC-VRF: A Virtual Routing and Forwarding table for Media Access Control (MAC) addresses on a PE. A MAC-VRF is also an instantiation of an EVI in a PE

MC-LAG: Multi-Chassis Link Aggregation Group (MC-LAG).

MLD: Multicast Listener Discovery.

PE: Provider Edge.

PIM: Protocol Independent Multicast.

RD: Route Distinguisher used in BGP.

RP: Multicast Rendezvous Point.

RT: Route-Targets used in BGP

RT-2: EVPN route type 2, i.e., MAC/IP advertisement route, as defined in [RFC7432].

RT-5: EVPN route type 5, i.e., IP Prefix route, as defined in Section 3 of [RFC9136].

RT-7: EVPN route type 7, i.e., Multicast Join Synch Route, as defined in Section 9.2 of [RFC9251].

RT-8: EVPN route type 8, i.e., Multicast Leave Synch Route, as defined in Section 9.3 of [RFC9251].

1.6. Requirements

The multi-homing solution described in this document satisfies the following requirements:

1. MUST support Layer-3 access interfaces
2. MUST support Layer-3 access sub-interfaces
3. MUST support unicast and multicast VPN services
4. SHOULD support IGP route synchronization
5. SHOULD support global routing table (GRT) services

6. MUST support all-active load-balancing mode
7. MAY support single-active load-balancing mode
8. MUST support port-active load-balancing mode
9. SHOULD avoid Layer 2 constructs (EVI, MAC-VRF, BD, IRB) for L3 state synchronization

2. Solution Overview

This document defines EVPN-based route synchronization mechanisms to enable all-active multi-homing for Layer 3 services. The solution uses existing EVPN route types to synchronize state between PEs sharing an Ethernet Segment:

- * RT-2 (MAC/IP routes): Synchronize ARP/ND adjacencies
- * RT-5 (IP Prefix routes): Synchronize IGP learned customer routes
- * RT-7/RT-8 (Multicast Join/Leave): Synchronize IGMP/MLD state

Key design principles:

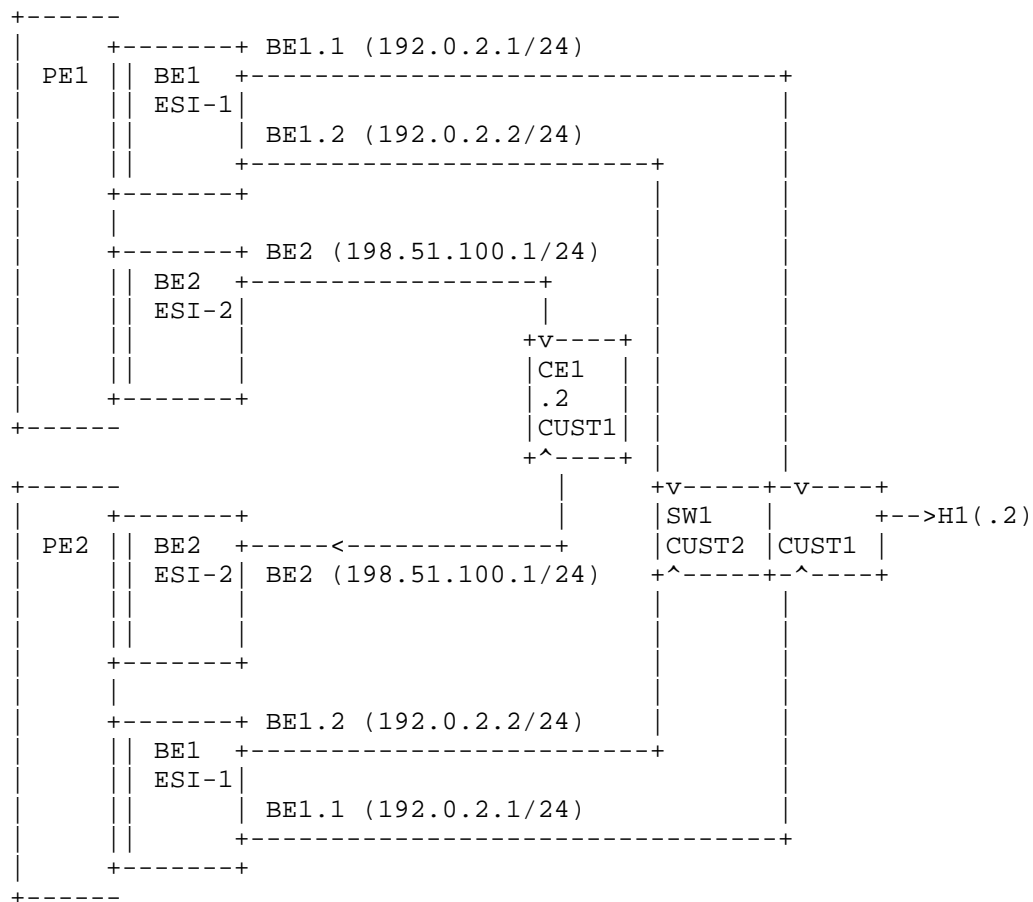
- * ESI identifies the shared LAG interface between redundant PEs
- * Ethernet Tag-ID identifies the sub-interface (VLAN) for AC-aware scenarios
- * IP-VRF Route Targets identify the VRF for route import/export
- * ES-Import RT (optional) restricts distribution to ESI-attached PEs

The following sections describe detailed procedures for each synchronization type.

3. Solution Details

3.1. Example Topology

Consider the Figure 3 topology, where two AC-aware bundling interfaces are configured. Interface BE1 on PE1 and PE2 shares a LAG with switch SW1 and supports two customer VRFs with overlapping subnets on VLAN 1 and VLAN 2. Interface BE2 supports a single customer VRF on native VLAN.



```
PE(1,2):
CUST1-VRF (IP-VRF1)
CUST2-VRF (IP-VRF2)
```

```
SW1:
CUST1-Subnet1: (192.0.2.1/24) (VLAN 1)
CUST2-Subnet1: (192.0.2.1/24) (VLAN 2)
```

```
CE1:
CUST1-Subnet: (198.51.100.1/24)
```

Figure 3: ARP/ND synchronization over different VRF(s)

In this topology:

- * BE1 (ESI-1): Shared by CUST1-VRF and CUST2-VRF with sub-interfaces VLAN 1 and 2
- * BE2 (ESI-2): Used only by CUST1-VRF on native VLAN

To synchronize state for CUST1-VRF, the solution uses:

Case 1 - Native interface (BE2 to CE1):

- * IP-VRF RT(s): Identifies CUST1-VRF
- * ESI-2: Identifies BE2 interface
- * Ethernet Tag-ID 0: Indicates native VLAN

Case 2 - Sub-interface (BE1.1 to SW1):

- * IP-VRF RT(s): Identifies CUST1-VRF
- * ESI-1: Identifies BE1 interface
- * Ethernet Tag-ID 1: Identifies VLAN 1 sub-interface

3.2. Route Target Usage

Route synchronization between peering PEs uses EVPN route types as defined in [RFC7432] and [RFC9136].

Routes SHOULD be advertised with the ES-Import Route Target to identify the Layer-3 interface for which the information must be synchronized, and with the EVI-RT Extended Community [RFC9251] to identify the routing context (IP-VRF or GRT) in which synchronization occurs. This limits route distribution to PEs attached to the same ESI and ensures that the routes are applied to the correct IP-VRF at the receiving PE.

Alternatively, synchronization routes MAY be advertised using the IP-VRF route Targets. However, this approach may cause the routes to be distributed to all remote PEs in the IP-VRF, including those that do not require the synchronization information.

In Figure 3, CUST1 routes carry IP-VRF1 RT(s) and CUST2 routes carry IP-VRF2 RT(s). When using ES-Import RT optimization, routes also carry EVI-RT Extended Community with the corresponding IP-VRF RT.

Note: When VRF Route Targets are used, routes are distributed to all PEs importing that VRF RT, not just ESI-attached PEs. However, only PEs with EVPN SAFI enabled will process these routes, effectively limiting distribution to EVPN-capable PEs.

3.3. EVPN Instance Usage

Unlike [RFC7432] MAC-VRF deployments, EVI is not required for L3 multi-homing scenarios. The Route Distinguisher (RD) MAY be auto-generated locally, and Route Targets are taken from the IP-VRF configuration.

For Global Routing Table (GRT) services, an EVPN instance MAY be assigned to provide Route Targets as required by [RFC7432]. Alternatively, users MAY explicitly configure Route Targets for GRT synchronization.

The solution synchronizes the following state types:

- * ARP/ND adjacencies (RT-2)
- * IGMP/MLD join/leave (RT-7/RT-8)
- * IGP learned routes (RT-5)

3.4. Mapping for L3 Interface to ESI

The ESI represents the L3 LAG interface between PE and CEs. This ESI is signaled using RT-4 with the ES-Import Route Target as described in Section 8.1.1 of [RFC7432] so that the service PE peers can discover each other's common ES.

In the example Figure 3, route-syncs from interface BE1 have IP-VRF RT(s) or ES-Import RT and EVI-RT EC with ESI 1 as an optimization.

3.5. Mapping for L3 Sub-Interface to Ethernet Tag-id

The Ethernet Tag-id represents the sub-interface subnet on the L3 LAG interface between PE and CEs. This apply to all route-sync types used for L3 multi-homing i.e., RT-2, RT-5, RT-7 and RT-8.

The Ethernet Tag ID encoded in synchronization routes is automatically derived from the encapsulation VLAN tags of the Layer-3 interface, following the encoding rules for single and double normalized VLAN identifiers defined in [RFC9744], as described below:

- * Untagged Layer-3 LAG interfaces use an Ethernet Tag ID value of zero.

- * Singly tagged Layer-3 LAG interfaces encode a single normalized VLAN identifier (VID) in the lower 12 bits of the Ethernet Tag ID field.
- * Doubly tagged Layer-3 LAG interfaces encode the outer normalized VID in the upper 12 bits and the inner normalized VID in the lower 12 bits of the Ethernet Tag ID field.

For synchronization to operate correctly, PEs attached to the same multi-homed CE MUST use consistent VLAN identifiers for the same multihomed CE.

In the example Figure 3, route-syncs from sub-interface BE1.1 (VLAN1) is represented by Ethernet Tag Identifier with ID 1.

3.6. ARP/ND Synchronization

This section describes procedures for synchronizing ARP/ND adjacencies between PEs using EVPN RT-2 routes, as defined in Section 10 of [RFC7432], with modifications for Layer 3 interfaces.

3.6.1. Advertising ARP/ND Routes

When a PE learns an ARP or ND adjacency on a Layer 3 interface or sub-interface, it MUST advertise an EVPN RT-2 route with non-zero ESI to synchronize the adjacency with peer PEs. Unlike Layer 2 EVPN services, MAC-only RT-2 routes MUST NOT be advertised, and Layer 2 forwarding state MUST NOT be programmed.

The RT-2 advertisement MUST include:

- * Non-zero ESI: Identifies the shared Ethernet Segment
- * IP address and MAC address of the learned adjacency
- * Ethernet Tag-ID: Set to VLAN ID for sub-interfaces, 0 for native interfaces

The RT-2 advertisement SHOULD include a label-1 value of zero and SHOULD NOT include a label-2. In addition, the RT-2 advertisement SHOULD NOT include any BGP Encapsulation Extended Communities [RFC9012].

The route MUST carry at least one of the following route target options:

- * ES-Import Route Target (instead of IP-VRF RT) to restrict distribution to ESI-attached PEs

- * EVI-RT Extended Community carrying the IP-VRF Route Target (required when using ES-Import RT) as defined in Section 9.5 of [RFC9251]

or

- * IP-VRF Route Target(s) of the associated VRF

Note: If the same ARP/ND entry exists on different LAG interfaces but uses the same subinterface normalized VLAN identifier (VID), the entry cannot be synchronized across PEs.

3.6.2. Processing Received ARP/ND Routes

A PE receiving an RT-2 synchronization route MUST:

- * Import the route only if both ES-Import RT and EVI-RT match the local configuration. Alternatively, the route MAY be imported if the IP-VRF RT matches the local IP-VRF import RT.
- * Derive the local interface from the ESI
- * Derive the sub-interface from the Ethernet Tag-ID (0 for native interface)
- * Install the adjacency in the appropriate IP-VRF and interface
- * Ignore the label value and BGP Encapsulation Extended Community value if present in the route.

A Route Reflector used to disseminate synchronization routes MUST ignore the label value carried in those routes.

The treat-as-withdraw behavior defined in [RFC 7606] is applied to EVPN MAC/IP Advertisement routes received with any of the following:

- * an ES-Import Extended Community that identifies a non-local Ethernet Segment;
- * a non-local EVI-RT; or
- * a reserved ESI, such as ESI-0 or ESI-MAC (all-FFs value).

In addition, a received RT-2 synchronization route MUST NOT trigger the programming of an ARP/ND entry if the same entry has already been learned locally on the PE.

3.7. IGMP/MLD Synchronization

This section describes procedures for synchronizing IGMP/MLD join and leave messages between PEs using EVPN RT-7 and RT-8 routes as defined in [RFC9251].

3.7.1. Advertising IGMP/MLD Routes

When a PE receives an IGMP Join or MLD Report on a Layer 3 interface or sub-interface, it MUST advertise an EVPN RT-7 route. When it receives an IGMP Leave or MLD Done, it MUST advertise an EVPN RT-8 route.

The RT-7/RT-8 advertisement MUST include:

- * Non-zero ESI: Identifies the shared Ethernet Segment
- * Multicast group and source information
- * Ethernet Tag-ID: Set to VLAN ID for sub-interfaces, 0 for native interfaces

As per Section 3.6.1, the route SHOULD carry ES-Import Route Target and EVI-RT Extended Community. Alternatively, the route MAY carry IP-VRF Route Target(s) of the associated VRF.

3.7.2. Processing Received IGMP/MLD Routes

A PE receiving an RT-7 or RT-8 synchronization route MUST:

- * Import the route only if IP-VRF Route Target matches a local VRF, OR both ES-Import RT and EVI-RT match local configuration
- * Derive the local VRF from the matching Route Target or EVI-RT
- * Derive the local interface from the ESI
- * Derive the sub-interface from the Ethernet Tag-ID
- * Install the multicast state in the appropriate VRF and interface

3.8. IGP Route Synchronization

This section describes procedures for synchronizing IGP learned customer routes between PEs using EVPN RT-5 routes as defined in Section 3 of [RFC9136].

When a CE forms an IGP adjacency on the LAG bundle, the adjacency may form to only one PE. That PE learns customer routes via IGP and must synchronize them to peer PEs so that all PEs can advertise the routes to the core and provide load-balancing.

Two approaches are defined:

1. ESI-based approach
2. IP Gateway-based approach

3.8.1. ESI-Based Approach

With the ESI-based approach, the PE learning routes via IGP advertises an RT-5 route with the ESI of the Ethernet Segment.

For [RFC4364] IP-VPN cores:

- * PE1 advertises RT-5 with non-zero ESI and IP-VPN route for R1
- * PE2 imports both routes, prefers the RT-5 due to non-zero ESI
- * PE2 treats RT-5 as a local route and advertises new IP-VPN route
- * Remote PEs receive IP-VPN routes from both PE1 and PE2 for load-balancing

For [RFC9136] EVPN IP-VRF-to-IP-VRF cores:

- * PE1 advertises RT-5 with non-zero ESI
- * PE2 synchronizes per Section 4.2 of [I-D.ietf-bess-evpn-ip-aliasing]
- * Both PEs advertise IP A-D routes for the ESI
- * Remote PEs load-balance per Section 4 of [I-D.ietf-bess-evpn-ip-aliasing]

3.8.2. IP Gateway-Based Approach

With the IP Gateway-based approach, the PE learning routes via IGP advertises an RT-5 route with the IP Gateway field set to the route's next-hop address.

For [RFC4364] IP-VPN cores:

- * PE1 advertises RT-5 with IP Gateway = nexthop and IP-VPN route for R1
- * PE2 imports both routes, prefers RT-5
- * PE2 resolves R1 via IP Gateway using synchronized ARP/ND from RT-2
- * PE2 advertises new IP-VPN route for load-balancing

For [RFC9136] EVPN IP-VRF-to-IP-VRF cores:

- * PE1 advertises RT-5 with IP Gateway = nexthop (no IP-VPN route needed)
- * PE2 imports and resolves RT-5 via synchronized ARP/ND
- * PE2 advertises RT-5 for R1
- * Remote PEs load-balance to both PE1 and PE2

4. Convergence Considerations

Entries synchronized via EVPN routes MAY be configured with a retention timer, allowing them to be retained during failure scenarios, thereby improving convergence and minimizing network churn.

The retention timer specifies how long a synchronized EVPN entry is retained after the corresponding EVPN route is withdrawn by the Layer-3 LAG ES peer. This mechanism is OPTIONAL, and the behavior for a synchronized entry is as follows:

- * When the EVPN route that originated the synchronized entry is withdrawn, the retention timer is started and the entry is retained until the timer expires.
- * For ARP/ND entries, while the retention timer is running, the PE attempts to refresh the entry by sending ARP Requests or Neighbor Solicitation messages to the IP owner.
- * For IGMP/MLD entries, while the retention timer is running, the PE attempts to refresh the entry by sending group-specific Queries for the corresponding multicast groups.

5. Overall Advantages

The use of EVPN ES-LAG all active multi-homing brings the following benefits to L3 BGP services:

- * Open standards based per interface all-active redundancy mechanism that eliminates the need to run ICCP and LDP.
- * Agnostic of underlay technology (MPLS, VXLAN, SRv6) and associated services (L3, L3-VPN).
- * Replaces legacy MC-LAG ICCP-based solution, and offers following additional benefits:
 - Fast convergence with mass-withdraw is possible with EVPN.
 - Avoid the need of a dedicated ICCP channel between peering PEs.
- * Removes the burden of having the need for ICL link and any proprietary protocols.

6. Security Considerations

The same Security Considerations described in [RFC7432] are valid for this document.

7. IANA Considerations

There are no IANA considerations.

8. Acknowledgments

The authors thank Ali Sajassi and Jeffrey Zhang for the discussions on the use case and solution options.

9. References

9.1. Normative References

- [I-D.ietf-bess-evpn-ip-aliasing]
Sajassi, A., Rabadan, J., Pasupula, S., Krattiger, L., and J. Drake, "EVPN Support for L3 Fast Convergence and Aliasing/Backup Path", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-ip-aliasing-03, 7 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-ip-aliasing-03>>.
- [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>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC9744] Sajassi, A., Ed., Brissette, P., Uttaro, J., Drake, J., Boutros, S., and J. Rabadan, "EVPN Virtual Private Wire Service (VPWS) Flexible Cross-Connect (FXC) Service", RFC 9744, DOI 10.17487/RFC9744, March 2025, <<https://www.rfc-editor.org/info/rfc9744>>.

9.2. Informative References

- [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>>.
- [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>>.

Appendix A. Contributors

The following people has contributed substantially to this document:

Jiri Chaloupka
Cisco
Email: jichalou@cisco.com

Jayashree Subramanian
Cisco
Email: jays@cisco.com

Authors' Addresses

Patrice Brissette (editor)
Cisco Systems
Email: pbrisset@cisco.com

Michael MacKenzie (editor)
Cisco Systems
Email: mimacken@cisco.com

Satoru Matsushima
Softbank
Email: satoru.matsushima@g.softbank.co.jp

Wen Lin
Juniper
Email: wlin@juniper.com

Jorge Rabadan
Nokia
Email: jorge.rabadan@nokia.com