

BESS WorkGroup
Internet-Draft
Intended status: Standards Track
Expires: 19 September 2025

N. Malhotra, Ed.
K. Ananthamurthy
A. Sajassi
L. Krattiger
Cisco Systems
J. Rabadan
Nokia
18 March 2025

Centralized Anycast Gateway in Ethernet VPN(EVPN)
draft-malhotra-bess-evpn-centralized-anycast-gw-03

Abstract

EVPN Integrated Routing and Bridging (IRB) fabrics provide a flexible and extensible method for Layer-2 and Layer-3 overlay network connectivity. [EVPN-IRB] defines operation for symmetric and asymmetric EVPN IRB using distributed anycast gateway architecture (DAG). In a DAG architecture, both bridging and first-hop routing functions for overlay subnets are located on leaf PEs, with first-hop routing provided by a distributed anycast gateway provisioned across the leaf PEs. This document describes an architecture and operation for EVPN Centralized Anycast Gateway (CAG), which allows the first-hop routing function for overlay subnets to be centralized on designated IRB GWs while the bridging function is still located on the leaf PEs. The documents also covers trade-offs of deploying a CAG as compared with DAG. It further describes operation for inter-op between CAG and DAG based EVPN-IRB network overlays.

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 19 September 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
2. Requirements Language and Terminology	5
3. Control Plane Operation	5
3.1. Requirements for L2 PE	5
3.2. Requirements for CAG PE	6
4. Data Plane Operation	7
4.1. Intra-Subnet Bridging	7
4.2. Inter-Subnet Routing	7
5. MAC/IP Mobility	8
6. CAG deployment models	8
6.1. CAG Placement as an Interconnect	8
6.1.1. Control Plane	10
6.1.2. Tunnel Adjacencies	10
6.1.3. Split Horizon domains	11
6.1.4. Data Plane	11
6.1.4.1. Inter-subnet - L2 fabric to Symmetric IRB fabric	11
6.1.4.2. Inter-subnet - Symmetric IRB fabric to L2 fabric	11
6.1.4.3. Intra-subnet - Symmetric IRB fabric to L2 fabric unicast	12
6.1.4.4. Intra-subnet - L2 fabric to Symmetric IRB fabric unicast	12
6.1.4.5. Intra-subnet - Symmetric IRB fabric to L2 fabric BUM	12
6.1.4.6. Intra-subnet - L2 fabric to Symmetric IRB fabric BUM	12
6.2. CAG Placement on DAG L2/L3PEs	12
6.2.1. Dual role of DAGs	14
6.2.2. Operation on CAG as L2/L3 GW	14
6.2.3. FHG Load-sharing and resiliency	14
6.2.4. FHG Localization	14

6.2.5. Comparison between the approaches	14
6.2.5.1. Operational simplicity	14
6.2.5.2. Scalability	15
7. Operational Considerations	15
8. Security Considerations	15
9. IANA Considerations	15
10. Acknowledgements	15
11. Contributors	15
12. Normative References	15
Authors' Addresses	16

1. Introduction

In a CAG routing architecture, overlay endpoints connect to Layer-2 EVPN gateways that provide VPN bridging function for overlay subnets. This VPN bridging function enables intra-subnet flows across the overlay while routing function between end-points in different subnets and to/from end-points external to the fabric is located on designated L2+L3 (IRB) gateways.

First-hop routing for each overlay subnet is deployed using a subnet Anycast GW that is hosted on one or more designated IRB GW nodes. A key attribute that defines this architecture is that the first-hop routing function for an overlay subnet is decoupled from the EVPN leaf that only provides intra-subnet bridging service for that subnet. This decoupling in-turn results in first-hop routing for overlay endpoints being "centralized" on designated IRB nodes. Note that the Anycast GW for each subnet is still distributed across these "centralized" IRB GW nodes.

It is common to deploy first-hop Anycast routing for all overlay subnets in a fabric on the same set of IRB nodes. While not necessarily required, as is covered later in the document, this is often done for operational simplicity and optimal routing. It is also common for this first-hop routing function to be hosted on border nodes that also act as interconnect GWs to external L2 or L2/L3 domains. Optionally, the CAG IRB nodes may also have directly connected end-points.

CAG architecture essentially uses a Layer-2 EVPN overlay and first-hop routing operation on CAG is identical to asymmetric routing as defined in [EVPN-IRB]. Figure below shows a reference L2 fabric with CAG topology that is also used to illustrate the procedures specified in later sections.

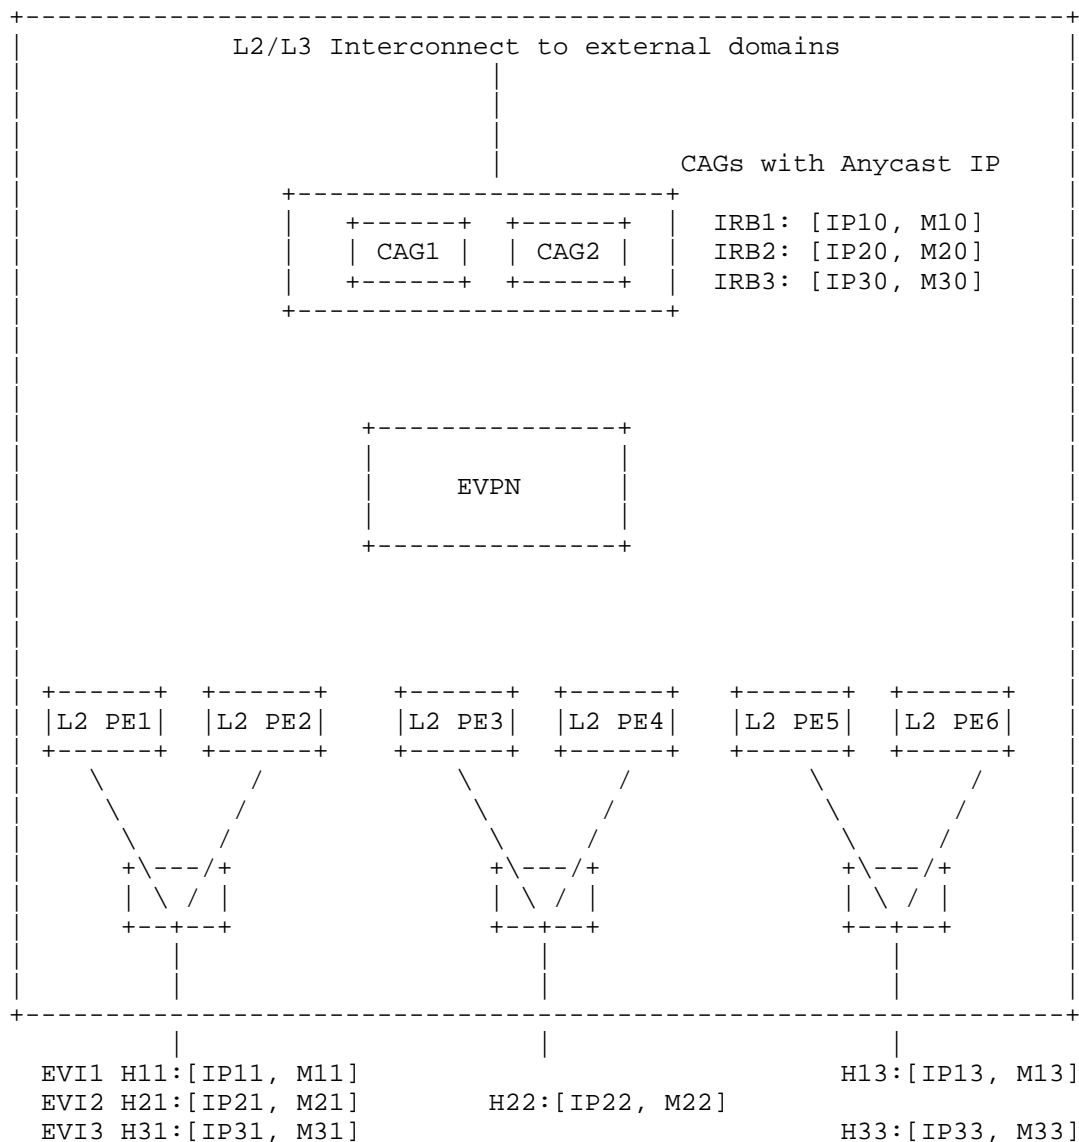


Figure 1

- * L2 PE1..L2 PE6 are L2-Only GWs.
- * IRB1/IRB2/IRB3 are IRB interfaces on CAGs for EVI1, EVI2 and EVI3.

- * Hosts H11/H13 are in EVI1, H21/H22 are in EVI2 and H31/H33 are in EVI3.
- * CAG1 and CAG2 form a redundant anycast CAG pair for EVI1, EVI2, and EVI3.

2. Requirements Language and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

- * CAG: Centralized Anycast Gateway
- * DAG: Distributed Anycast Gateway
- * EVI: An EVPN instance spanning the Provider Edge (PE) devices participating in that EVPN
- * FHR: First Hop Router
- * IRB: Integrated Routing and Bridging
- * L2-GW: Layer-2 Only Gateway
- * L2-Only GW: Layer-2 Only Gateway

3. Control Plane Operation

This section defines control plane procedures required on L2 PE and CAG PE for a CAG architecture based deployment.

3.1. Requirements for L2 PE

- * ARP/ND snooping MUST be enabled on L2-PEs. Locally connected host IP and MAC is learnt by L2 PE via ARP/ND snooping and advertised using EVPN RT-2 with a single label that represents the EVI. This enables a significant simplification in CAG operation to avoid the need for data plane learning and syncing of ARP/ND tables across redundant CAGs.
- * L2-PEs MUST handle ARP refreshes when MAC ages out or ARP ages out in order to relearn the host MAC and IP to avoid traffic loss. If a host does not respond to an ARP refresh, then the previously advertised EVPN RT-2 by the L2-PE MUST be withdrawn. If the host responds to host MAC/IP, then ARP MUST be refreshed. EVPN RT-2 SHOULD NOT be re-advertised

- * L2-PEs on receiving GW MAC/IP RT-2 from CAG, in addition to installing the MAC in the MAC-VRF, MUST also install a GW MAC/IP entry in the ARP/ND suppression cache. This enables L2-PEs to act as a proxy for CAG by using the entry in the suppression cache to respond to ARP requests from hosts for GW MAC/IP. If enabled, this avoids flooding of ARP/ND requests across the fabric and avoids duplicate ARP responses from redundant CAGs.

3.2. Requirements for CAG PE

- * A set of redundant CAG PEs that act as the FHR for the same subnet MUST be provisioned with the same anycast GW IP and MAC.
- * For each subnet for which CAG is an FHR, CAG MUST advertise Anycast GW MAC/IP using EVPN RT-2 with Default Gateway Extended Community as defined in [RFC7432]. In Figure 1, CAG1/2 advertise IRB1/2/3 MAC and IP using EVPN RT-2 with Default Gateway Extended Community with nexthop as anycast IP.
- * In case of VXLAN encapsulation, set of redundant CAG PEs provisioned as FHR for a common set of subnets MAY advertise the anycast GW MAC/IP RT-2 with an anycast VTEP IP as the next-hop. This enables the GW MAC to be installed with a single BGP path on L2 PEs and hence enables interworking with low end L2 devices that may not be capable of MAC multi-pathing. It also results in a much simplified solution that avoids a need for virtual ESI provisioning on CAG PE as well as a need for MAC multi-pathing and fast convergence procedures on CAG and L2 PEs. Note that this anycast VTEP IP is solely for the purpose of GW MAC/IP RT-2 advertisement and all directly connected end-points (single-homed or multi-homed) will continue to be advertised with a non-anycast VTEP IP as the next-hop. ESI multi-homing procedures specified in [RFC7432] apply as-is to directly connected end-points multi-homed via ESI LAG.
- * Upon receiving RT-2 advertisements from Egress EVPN L2-PE, CAG MUST import MACs into MAC-VRFs, thus establishing Host MAC reachability via Layer-2 EVPN encapsulation and tunnel to Egress L2 PE. In Figure 1, L2-GW1/2 advertise RT-2 for hosts H11, H21 and H31. CAG MUST import M11, M21 and M31 in corresponding MAC-VRFs. CAG will perform the Asymmetric IRB procedures defined in [RFC9135] with IP-to-MAC binding resolved via respective adjacency/next-hop table entry.
- * It should be noted that reachability to a remote host layer-3 adjacency is still resolved by host MAC reachability via a Layer-2 VPN tunnel to the Egress L2-PE.

4. Data Plane Operation

4.1. Intra-Subnet Bridging

Intra-subnet host to host flow is identical to that in symmetric or asymmetric distributed anycast GW based IRB deployments as defined in [EVPN-IRB]. When the ingress L2 PE receives a packet from its locally attached host, it does a destination MAC lookup in its VLAN which yields an L2 VPN and tunnel encapsulation towards the egress L2 PE. The ingress L2 PE then encapsulates the original packet and sends to the egress L2 PE. Egress L2 PE decapsulates the packet and does a destination MAC lookup in the local VLAN (MAC VRF) table identified by the received L2 VPN label. This yields a local VLAN Port. The egress L2GW then sends the decapsulated packet to the port.

4.2. Inter-Subnet Routing

Inter-subnet host to host flow destined to default GW MAC is bridged to CAG PE with the L2 VPN encapsulation learnt via default GW RT-2 from the CAG PE. CAG decapsulates the packet and does a destination MAC lookup in the local MAC VRF identified by the received L2 VPN encapsulation that results in my MAC. This yields the local IRB interface. CAG then does a destination IP lookup in IP VRF attached to the IRB interface. If the destination subnet is local/attached to the CAG node, IP lookup yields a local host adjacency on the destination subnet IRB interface. This results in host MAC rewrite, followed by host MAC lookup that results in the packet being bridged to the egress L2 PE with L2 VPN and tunnel encapsulation learnt via RT-2 from egress L2 PE. On the egress L2 PE, packet is decapsulated, local MAC VRF is resolved by the received L2 VPN encapsulation. The packet is then bridged to the local host via local MAC VRF lookup.

If the destination subnet is not local to the CAG node, IP lookup yields the destination subnet prefix that resolves via L3 VPN encapsulation and tunnel to the next-hop CAG PE that is the FHR for the destination subnet. Packet is hence routed to another CAG, where the packet is decapsulated, received L3 VPN encapsulation resolves the local IP VRF, and IP lookup now yields a local host adjacency on the destination subnet IRB interface. Packet is then bridged to the destination L2 PE with the L2 VPN encapsulation as described above.

5. MAC/IP Mobility

GW MAC/IP route advertised by CAG MUST follow default gateway best path selection procedures specified in [RFC7432bis] to ensure that GW MAC is treated as static and is always preferred over any duplicate host MAC across the L2 overlay. Beyond BGP best path selection, forwarding implementations MUST ensure that a locally learnt duplicate MAC will never overwrite the forwarding entry for the GW MAC route. Hosts attached to L2-PEs follow existing mobility procedures as defined in [RFC7432].


6. CAG deployment models

Fabrics with symmetric IRB are widely deployed. This section discusses how a CAG architecture discussed in the earlier sections may be deployed together with a symmetric IRB fabric with L2 and L3 overlays that extend across symmetric IRB and CAG based fabrics. Two deployment models are discussed:

- * CAG Placement as an Interconnect
- * CAG Placement on DAG L2/L3PEs

6.1. CAG Placement as an Interconnect

In this model, CAG is placed as an interconnect or hub between the Layer-2 fabric consisting of L2-PEs and the symmetric IRB fabric consisting of L2/L3 PEs. CAG device essentially performs the CAG function for hosts connected to L2 only fabric and in addition, also performs an L2/L3 overlay interconnect function between the L2 only fabric and the symmetric IRB fabric.



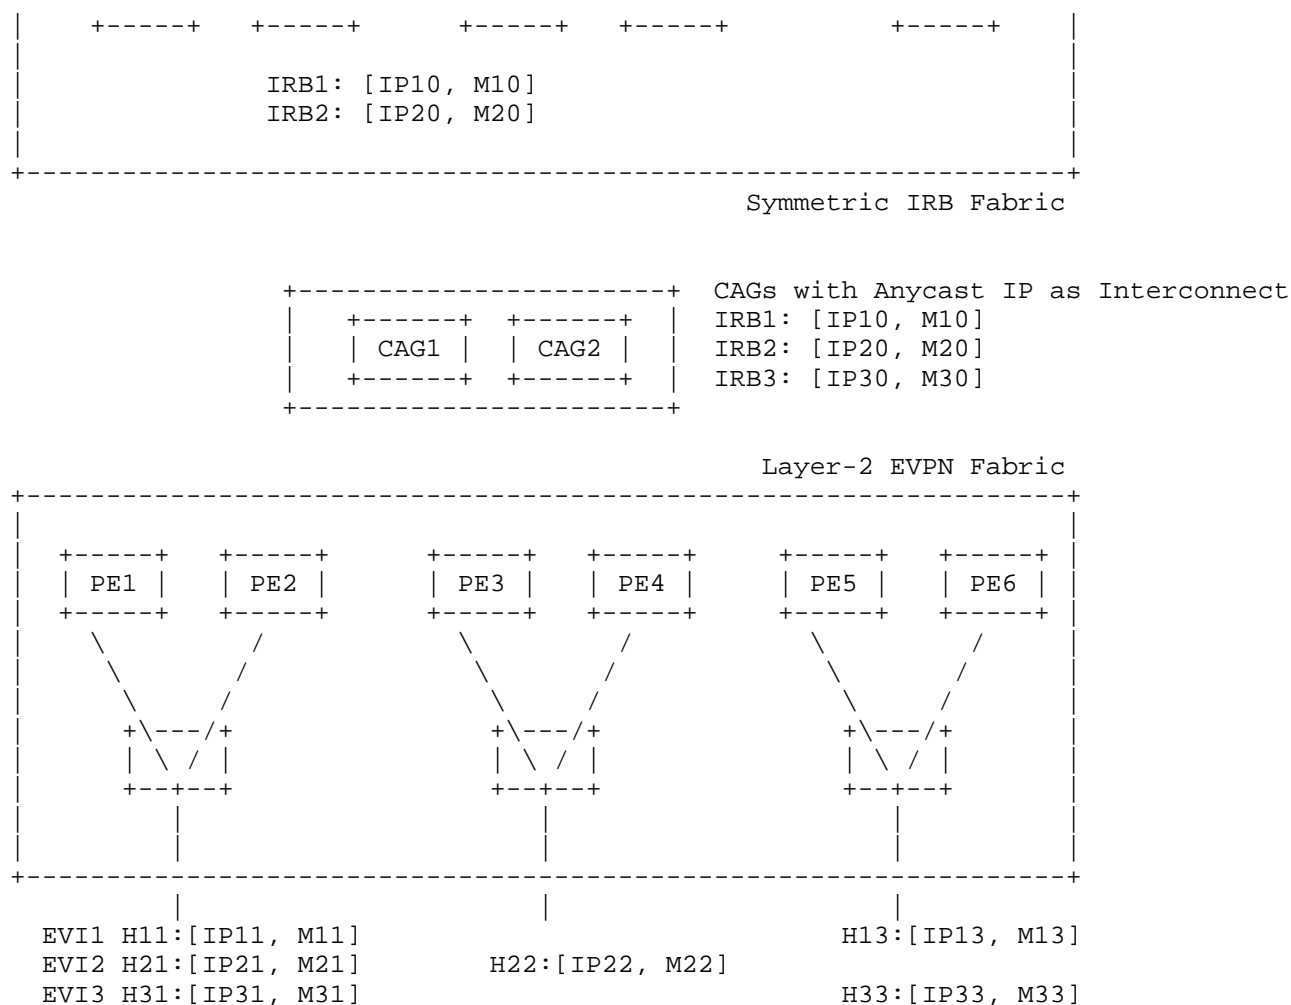


Figure 2

- * IRB1/IRB2 are IRB interfaces on L2/L3 PEs PE7- PE10.
- * PE11 is L3-only PE
- * CAG pair consisting of CAG1/CAG2 is placed an Interconnect between Layer-2 only and Symmetric IRB fabric.

6.1.1. Control Plane

In addition to the control plane operation described in Section 3, the following operations must be done at CAG to enable seamless interworking between the fabrics.

- * MAC/IP RT-2s learnt on CAGs from L2 PEs in the Layer-2 only EVPN Fabric MUST be re-originated by all CAGs towards the L2/L3 PEs and L3-only PE in the Symmetric IRB Fabric with CAG as the tunnel next-hop.
- * MAC/IP RT-2s learnt on CAGs from L2 PE in the Layer-2 only EVPN Fabric MUST be re-originated by all CAGs towards the L2/L3 PE and L3-only PE in the Symmetric IRB Fabric in symmetric IRB format with both L2 and L3 VPN local labels. Essentially, these MAC-IP routes are learnt from L2 PE with only L2 label and re-originated with locally provisioned or allocated L2 and L3 labels.
- * MAC/IP RT-2s learnt from L2/L3 PE in the Symmetric IRB fabric MUST be re-originated towards the Layer-2 only fabric with CAG as the tunnel next-hop.
- * MAC/IP RT-2s learnt from L2/L3 PE in the Symmetric IRB fabric and re-originated towards the L2 PE MAY have L3 label and L3 RTs removed resulting in re-origination with only locally assigned L2 VPN label and L2 RTs. L2 PE in the Layer-2 only fabric ignore the L3 label and RTs if they are present.

6.1.2. Tunnel Adjacencies

This results in the following overlay tunnel adjacencies at CAG:

- * L2 and L3 VPN tunnel adjacencies between CAG and L2/L3 PE in symmetric IRB fabric. In above topology, Layer-2 and Layer-3 tunnel adjacencies are formed between L2/L3 PE PE7-PE10 and CAG.
- * L3 VPN tunnel adjacencies between CAG and L3 only PE in symmetric IRB fabric. In above topology, Layer-3 tunnel adjacencies are formed between L3-only PE11 and CAG.
- * L2 VPN tunnel adjacencies between CAG and L2 PE in L2 fabric. In above topology, Layer-3 tunnel adjacencies are formed between L2 PE PE1-PE6 and CAG.

6.1.3. Split Horizon domains

CAG MUST consider the two fabrics as separate split horizon domains and hence apply split-horizon procedures specified in [RFC9014]. As a result, The BUM traffic from one domain is distributed to the other domain. For e.g. ARP requests from symmetric IRB domain are distributed to the Layer-2 domain and vice versa.

6.1.4. Data Plane

In addition to the data plane operation described in Section 4, the following operations must be done at CAG to enable seamless interworking between the fabrics.

6.1.4.1. Inter-subnet - L2 fabric to Symmetric IRB fabric

Intersubnet traffic is sent from host in the Layer-2 only fabric destined to host IP attached to the symmetric IRB fabric and GW MAC on CAG. The L2 PE performs a Gateway MAC lookup in MAC-VRF and tunnels traffic to CAG with L2 VPN label and tunnel encapsulation to CAG. CAG performs a GW MAC lookup in the MAC-VRF followed by an IP lookup in IP-VRF. This results a route to the L2/L3 PE next-hop. CAG tunnels the packet to L2/L3 PE with L3 VPN label and tunnel encapsulation to L2/L3 PE. L2/L3 PE performs a lookup in the IP-VRF which results in an adjacency to the destination host, using which traffic is routed to the attached destination host.

In the topology above, traffic sourced from H11 destined to H24 is sent to M10 and IP24. PE1/2 perform a lookup of M10 and tunnel traffic to CAG1/2. CAG1/2 perform a lookup of M10 in the MAC-VRF, followed by lookup of IP24 in IP-VRF, which results in a route to PE7/8. CAG1/2 tunnel Layer-3 traffic to PE7/8. PE7/8 perform a lookup of IP24 in IP-VRF, which results in adjacency to H24. PE7/8 bridge traffic to H24 destined to M24.

6.1.4.2. Inter-subnet - Symmetric IRB fabric to L2 fabric

Intersubnet traffic is sent from host in symmetric IRB fabric destined to host IP attached to Layer-2 only fabric and DAG MAC on L2/L3 PE. L2/L3 PE performs GW MAC lookup in MAC-VRF followed by an IP lookup in the IP-VRF. This results in a route to the CAG next-hop. L2/L3 PE tunnels the packet to the CAG with L3 VPN label and tunnel encapsulation to CAG. CAG performs destination IP lookup in IP-VRF resolved by L3 VPN label, which results in an adjacency to the destination host. CAG performs a MAC lookup of destination host MAC and tunnels traffic to L2 PE next-hop with L2 label and tunnel encapsulation to L2 PE. L2-only PE performs a MAC lookup in the MAC-VRF and bridges the packet to the destination host.

In the topology above, traffic sourced from H24 destined to H11 is sent to M20 and IP11. PE7/8 perform a lookup of M20 followed by IP11 lookup in IP-VRF and tunnels Layer-3 traffic to CAG. CAG performs a lookup of IP11 in IP-VRF, which results in adjacency to H11. CAG perform lookup M11 and tunnels traffic to PE1/2. PE1/2 perform lookup of M11 and bridge traffic to H11.

6.1.4.3. Intra-subnet - Symmetric IRB fabric to L2 fabric unicast


6.1.4.4. Intra-subnet - L2 fabric to Symmetric IRB fabric unicast

6.1.4.5. Intra-subnet - Symmetric IRB fabric to L2 fabric BUM

6.1.4.6. Intra-subnet - L2 fabric to Symmetric IRB fabric BUM

6.2. CAG Placement on DAG L2/L3PEs

In this placement approach, all L2/L3 PE in the Symmetric IRB fabric that locally host an EVI assume an additional role of a CAG for hosts in that EVI that are attached to L2 only fabric.



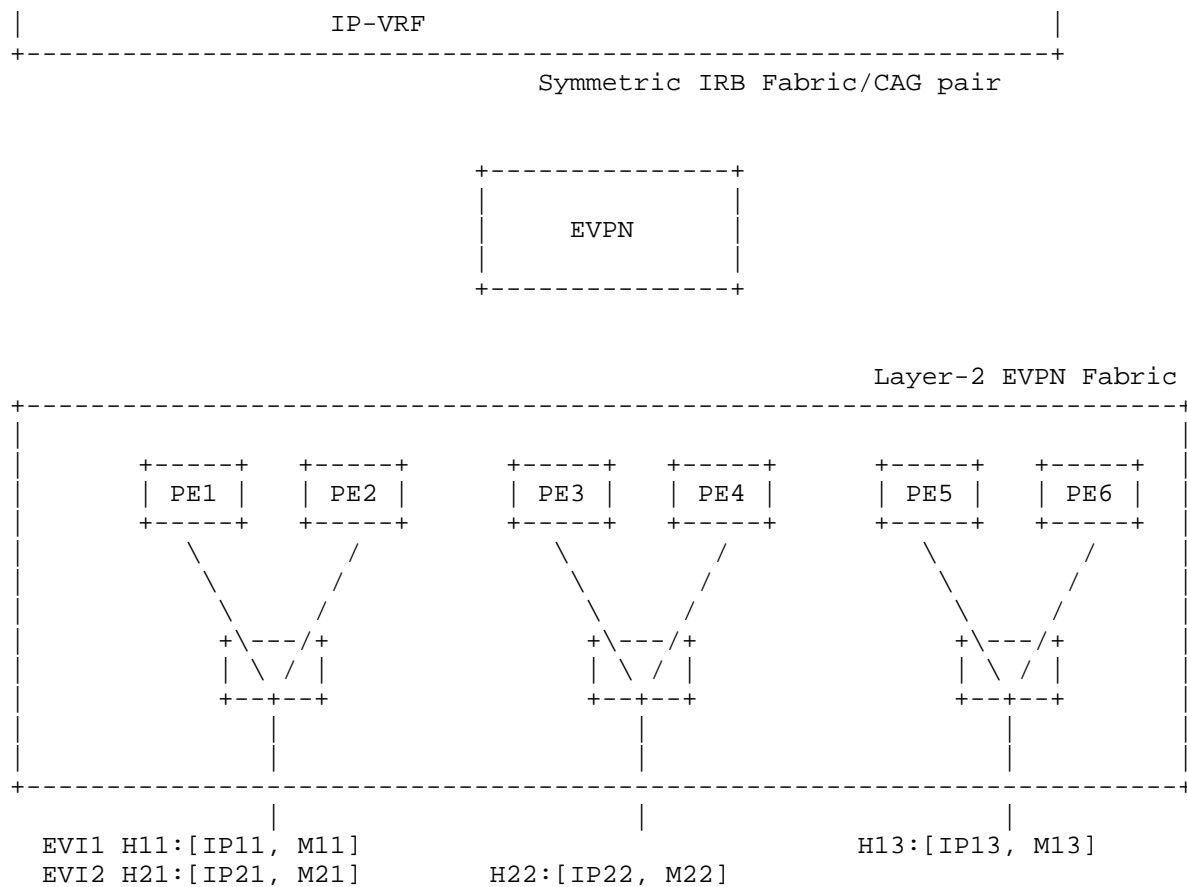


Figure 3

- * Symmetric IRB Fabric is EVPN fabric consisting of L2/L3 PEs and L3-only PEs operating in Symmetric IRB mode.
- * PE7- PE10 are L2/L3 PEs and participate in EVI/EVI2 and host IP-VRF which has subnets for EVI1/EVI2.
- * PE7 - PE10 are L2/L3 PEs and also CAG to Layer-2 only fabric.
- * IRB1/IRB2 are IRB interfaces on PE7/CAG - PE10/CAG.

6.2.1. Dual role of DAGs

In this placement method, each DAG in the symmetric IRB fabric plays a dual role of DAG and CAG. Thus, the CAG for an EVI comprises of each L2/L3 PE DAG node in the symmetric IRB fabric that hosts the EVI. Such a placement results in forming a full mesh of tunnels between the L2-only PEs and every DAG/CAG.

6.2.2. Operation on CAG as L2/L3 GW

The operation on the CAG remains the same as described in section FIX and MUST be performed by each member of the CAG pair that hosts the EVI.

6.2.3. FHG Load-sharing and resiliency

As the CAG pair is comprised of all L2/L3 PEs that host the EVI, the FHG function can be load-shared across a larger number of CAGs. As all members of the pair advertise the GW MAC/IP to all the L2-only PEs, this opens up an opportunity on the L2-only PEs for creating a scheme optimal load-sharing, load-balancing, failure resiliency and achieving faster convergence during failure events. As an example, the L2-only PEs may choose to pick one CAG per EVI as the FHG, resulting in load-sharing traffic across the pair or the L2-only PEs may pick a subset of CAGs for per EVI as FHGs which provides failure resiliency, faster convergence in failure events while providing load-sharing across the members of the CAG pair.

6.2.4. FHG Localization

If the use case requires FHG function to be localized on a subset of members of the CAG pair for a given EVI, only this subset MUST be configured to advertise the GW MAC/IP to the L2-only PEs. Thus forming two sub-clusters within the CAG cluster. The subset which does not advertise the GW MAC/IP form a CAG non-FHG sub-cluster whereas, the other cluster forms a CAG FHG sub-cluster. As the non-FHG cluster does not advertise GW MAC/IP, it does not attract inter-subnet traffic from L2-only PEs. Both sub-clusters MUST perform the operations as described in section FIX.

6.2.5. Comparison between the approaches

6.2.5.1. Operational simplicity

In the full mesh placement approach, the operation on CAGs is largely simplified as re-origination of routes as described in section FIX is not required and thus, does not require and additional functionality on the CAG for seamless interworking between the fabrics.

6.2.5.2. Scalability

The interconnect placement approach requires re-origination of routes by CAG between the two fabrics. This results in logically separated domains. As a result, segregated tunnel domains are be created, thus making this approach more scalable. In the full mesh placement, as a full mesh of tunnels is formed between the L2-only PEs and CAG, thus making this approach less scalable.

7. Operational Considerations

None

8. Security Considerations

Security considerations discussed in [RFC7432] and [RFC9135] apply to this document.

9. IANA Considerations

10. Acknowledgements

Authors would like to thank Aparna Pattekar for spending considerable time on and contributing multiple sections to rev0 of this draft.

11. Contributors

Aparna Pattekar
Cisco Systems
Email: apjoshi@cisco.com

12. Normative References

[EVPN-IRB] Sajassi, A., Salam, S., Samil, S., Drake, J., Rabadan, J., and , "Integrated Routing and Bridging in EVPN", Work in Progress, Internet-Draft, Integrated Routing and Bridging in EVPN, 26 July 2021, <<https://www.rfc-editor.org/info/rfc9135>>.

[EXTENDED-MOBILITY-EVPN-IRB] Malhotra, N., Sajassi, A., Pattekar, A., Rabadan, J., Lingala, A., Drake, J., and , "Extended Mobility Procedures for EVPN-IRB", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-irb-extended-mobility-17, 16 October 2023, <<https://www.ietf.org/archive/id/draft-ietf-bess-evpn-irb-extended-mobility-17.txt>>.

- [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>>.
- [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>>.
- [RFC7432bis]
Sajassi, A., Brudet, L.A., Rabadan, J., Drake, J., and , "BGP MPLS-Based Ethernet VPN", Work in Progress, Internet-Draft, draft-ietf-bess-rfc7432bis-08, 13 February 2024, <<https://www.ietf.org/archive/id/draft-ietf-bess-rfc7432bis-08.txt>>.
- [RFC9014] Rabadan, J., Sathappan, S., Henderickx, W., Sajassi, A., Drake, J., and , "Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks", 26 May 2021, <<https://www.rfc-editor.org/rfc/rfc9014.txt>>.

Authors' Addresses

Neeraj Malhotra (editor)
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
United States of America
Email: nmalhotr@cisco.com

Krishnaswamy Ananthamurthy
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
United States of America
Email: kriswamy@cisco.com

Ali Sajassi
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
United States of America
Email: sajassi@cisco.com

Lukas Krattiger
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
United States of America
Email: lkrattig@cisco.com

Jorge Rabadan
Nokia
777 E. Middlefield Road
Mountain View, CA 94043
United States of America
Email: jorge.rabadan@nokia.com