

BESS Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 12, 2022

M. MacKenzie, Ed.
P. Brissette
Cisco
S. Matsushima
Softbank
July 11, 2021

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

Abstract

This document brings the machinery and solution providing higher network availability and load balancing benefits of EVPN Multi-Chassis Link Aggregation Group (MC-LAG) to various L3 services delivered by EVPN.

Requirements Language

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 [RFC 2119](#) [[RFC2119](#)] and [RFC 8174](#) [[RFC8174](#)].

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 January 12, 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

EVPN L3MH

July 2021

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
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	5
1.4.	Problems with supporting multiple subnets on same ES in all active mode	6
1.5.	Acronyms	6
1.6.	Requirements	8
2.	Solution	8
2.1.	Mapping of L3VRF to EVPN EVI	10
2.2.	Mapping for L3 Interface to ESI	11
2.3.	Mapping for L3 Sub-Interface to Attachment Circuit ID	11
2.4.	Route sync for ARP/ND	11
2.4.1.	Local adjacency (ARP/ND) learning	11
2.4.2.	Remote ARP/ND learning	12
2.5.	Route sync for IGMP	12
2.5.1.	Local IGMP Join/Leave learning	13
2.5.2.	Remote IGMP Join/Leave learning	13
2.6.	Customer Subnet Route sync using Route-type(5)	13
2.7.	Mapping for VLAN to ETAG	14
3.	Extensions to RT-2, RT-5, RT-7 and RT-8	14
4.	Convergence Considerations	14
5.	Overall Advantages	15
6.	Security Considerations	15
7.	IANA Considerations	15
8.	References	15
8.1.	Normative References	15
8.2.	Informative References	16
	Authors' Addresses	16

[1.](#) Introduction

Resilient L3VPN service to a CE requires multiple service PEs to run a MC-LAG mechanism, which previously required a proprietary ICL control plane link between them.

This proposed extension to [\[RFC7432\]](#) brings EVPN based MC-LAG all-active multi-homing load-balancing to various services (L2 and L3) delivered by EVPN. Although this solution is also applicable to some L2 service use cases, (example Centralized Gateway) this document will focus on the L3VPN [\[RFC4364\]](#) use case to provide examples.

EVPN MC-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 will be limited to the redundant service PEs sharing a Ethernet Segment Identifier (ESI). This will be used to synchronize ARP/ND, multicast Join/Leave, and IGP routes replacing need for ICL link.




Figure 1: EVPN MC-LAG Topology

Figure 1 shows a MC-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. 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, but 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. MC-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 will advertise layer-3 reach-ability towards the layer-3 core and will 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 its possible for only one service PE to populate its ARP/ND cache. Take for example PE1 and PE2 from Figure 1. If CE1 ARP/ND response happens to always hash over I1 towards PE1, then PE2 ARP/ND table will be empty. Since unicast traffic from remote PEs can be received by either service PE, traffic that reaches the service PE2 will not find an ARP entry matching the host IP address and traffic will drop until ARP/ND resolves the adjacency.

If the CEs hash implementation always calculates the ARP/ND response towards PE1, the resolution on PE2 will never happen and traffic load balanced to PE2 will black-hole.

The route sync solution is described in [Section 2.4](#)

[1.2.](#) Problems with multicast from core to CE

Similar to the unicast behavior above, multicast IGMP join messages from CE to LAG link may always hash to a single PE.

When PIM runs on both redundant layer-3 PEs that both service multicast for the same access segment, PIM elects only one of the PEs

as a PIM Designated Router (DR) using PIM DR election algorithm [RFC7761]. The PIM DR is responsible for tracking local multicast listeners and forwarding traffic to those listeners. The PIM DR is also responsible for sending local Join/Prune messages towards the RP or source.

For example, if in Figure 1 PE2 is designated PIM-RP, but CE IGMP join messages are hashed to I1 towards PE1, then multicast traffic will not be attracted to this service pair as PE2 will not send PIM Join on behalf of CE.

In order to ensure that the PIM DR always has all the MCAST route(s) and able to forward PIM Join/Prune message towards RP, BGP-EVPN multicast route-sync will be leveraged to synchronize MCAST route(s) learned to the DR.


When a fail-over occurs, multicast states would be pre-programmed on the newly elected DR service PE and assumes responsibility for the routing and forwarding of all the traffic.

The multicast route sync solution is described in [Section 2.5](#)

[1.3.](#) Problems with IGP adjacencies over the LAG port

A layer-3 CE device/router that connects to the redundant PEs may establish an IGP adjacency on the bundle port. In this case, the adjacency will be formed to one of the PEs and IGP customer route(s) will only be present on that PE.

This prevents the load-balancing benefits of redundant PEs from supporting this use case, as only one PE will be aware and advertising the customer routes to the core.



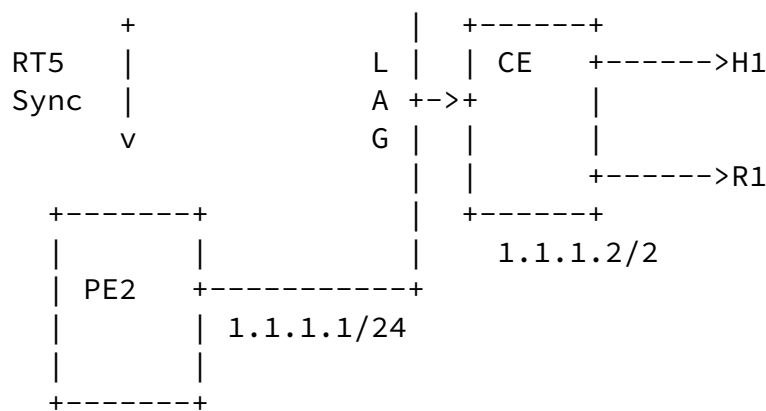


Figure 2: IGP Adjacency over LAG Port

Figure 2 provides an example of this use case, where CE 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 will only be aware of the service from PE1, and cannot load balance through PE2 as well.

Further study is required in order to support the case of BGP PE to CE protocols.

A solution to this is described in [Section 2.6](#)

[1.4.](#) Problems with supporting multiple subnets on same ES in all active mode

In the case where the L3 service is L3VPN such as [\[RFC4364\]](#), it is likely the CE device could be a layer-2 switch supporting multiple subnets through the use of VLANs. In addition, each VLAN may be associated with a different customer VRF.

When ARP/ND routes are synchronized between the PEs for ARP proxy support using RT-2, a similar problem is encountered as described by Section 1.1 of [\[I-D.sajassi-bess-evpn-ac-aware-bundling\]](#). The PE receiving RT-2 is unable to determine which sub-interface the ARP/ND entry is associated with.

When IGMP routes are synchronized between the PEs using RT-7 and RT-

8, a similar problem is encountered as described by Section 1.2 of [[I-D.sajassi-bess-evpn-ac-aware-bundling](#)]. The PE receiving RT-7 and RT-8 is unable to determine which sub-interface the IGMP join is associated with.

This document proposes to use the solution defined by Section 4 of [[I-D.sajassi-bess-evpn-ac-aware-bundling](#)] to solve both these cases. All route sync messages (RT-2, RT-5, RT-7, RT-8) will carry an Attachment Circuit Identifier Extended Community to signal which sub-interface the routes were learnt on.

[1.5.](#) Acronyms

BD: Broadcast Domain. As per [[RFC7432](#)], an EVI consists of a single or multiple BDs. In case of VLAN-bundle and VLAN-aware bundle service model, an EVI contains multiple BDs.

DF: Designated Forwarder

DR: 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'.

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

ICL: Inter Chassis Link

IGMP: Internet Group Management 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.

L3AA All-Active Redundancy Mode for Layer 3 services. When all PEs attached to an Ethernet segment are allowed to forward known unicast traffic to/from that Ethernet segment for a given VLAN, then the Ethernet segment is defined to be operating in All-Active redundancy mode.

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).

PE: Provider Edge.

PIM: Protocol Independent Multicast

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

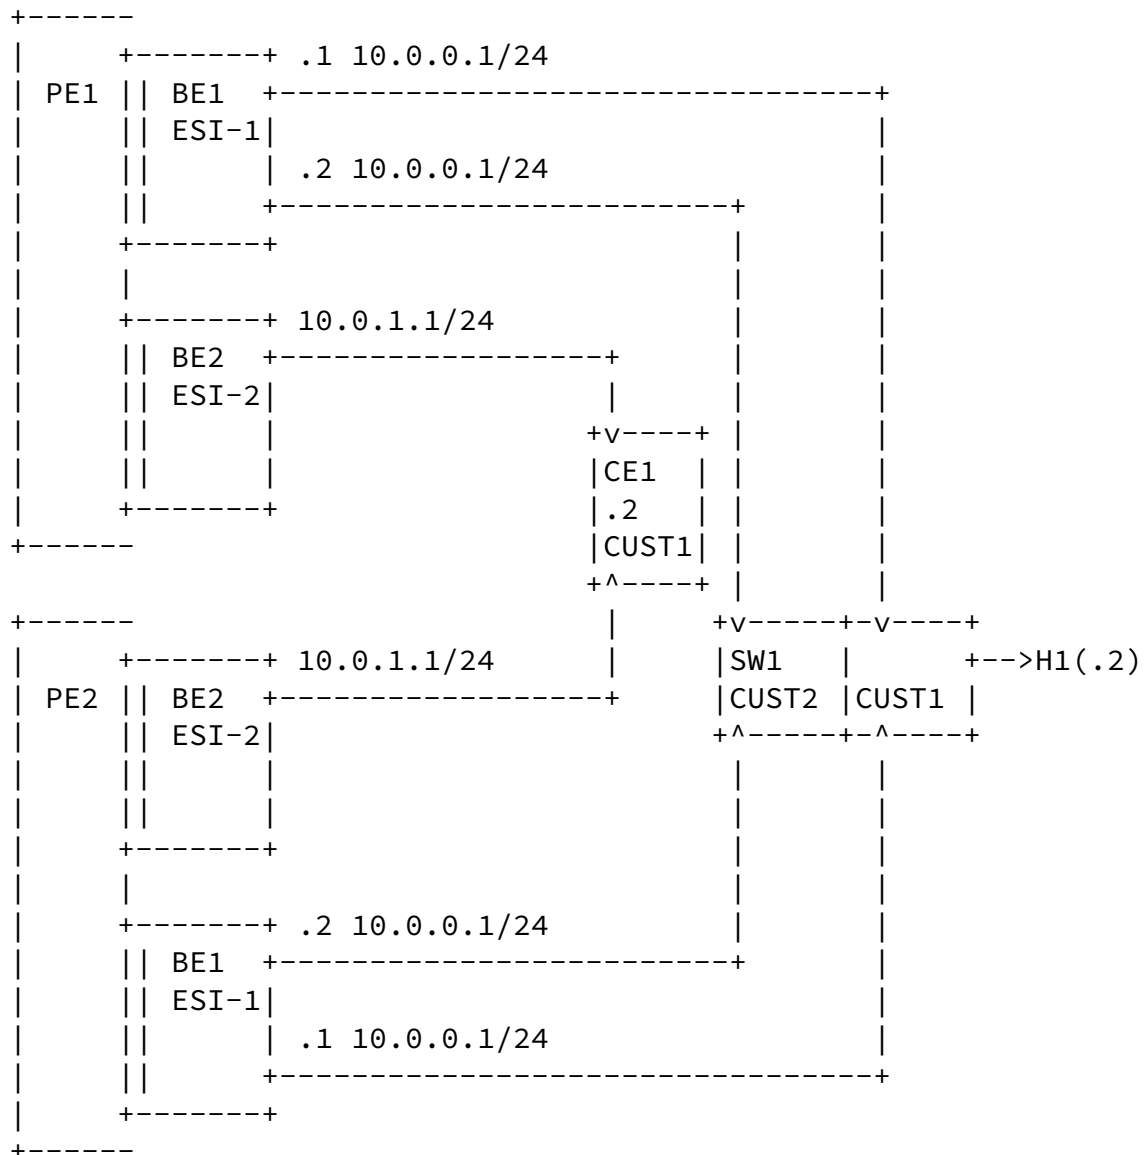
RT-5: EVPN route type 5, i.e., IP Prefix route, as defined in Section 3 of [[I-D.ietf-bess-evpn-prefix-advertisement](#)]

RT-7: EVPN route type 7, i.e., Multicast Join Synch Route, as defined in Section 9.2 of [[I-D.ietf-bess-evpn-igmp-mld-proxy](#)]

RT-8: EVPN route type 8, i.e., Multicast Leave Synch Route, as defined in Section 9.3 of [[I-D.ietf-bess-evpn-igmp-mld-proxy](#)]

1. The multi-homing solution MUST support Layer-3 access interface
2. The multi-homing solution MUST support Layer-3 access sub-interface
3. The solution MUST support unicast and multicast VPN services
4. The solution SHOULD support igp synchronization
5. The solution SHOULD support unicast and multicast GRT services
6. The solution MUST support all-active load-balancing mode
7. The solution MAY support single-active load-balancing mode
8. The solution MUST support port-active load-balancing mode

2. Solution



PE(1,2):

CUST1-VRF: EVI 1

CUST2-VRF: EVI 2

SW1:

CUST1-Subnet1: 10.0.0.2/24 (VLAN 1)

CUST2-Subnet1: 10.0.0.2/24 (VLAN 2)

CE1:

CUST1-Subnet2 10.0.1.2/24

Figure 3: ARP/ND MAC-IP route-sync over different VRF(s)

Consider the Figure 3 topology, where 2 AC aware bundling service interfaces are supported. On first bundling interface BE1, PE1 and PE2 share a LAG interface with switch 1 (SW1) and have 2 separate (but overlapping) customer 1 and customer 2 subnets. CUST1 Subnet 1 is resolving over sub-interface VLAN 1 (.1), and CUST2 Subnet 1 is resolving over sub-interface VLAN 2 (.2).

On second bundling interface BE2, both PEs share a LAG interface with Customer Edge device 1 (CE1) and only a single Customer (CUST1) subnet on native VLAN.

Main interface BE1 on PE1 and PE2 is shared by customer 1 and 2, and represented by ESI-1.

Main interface BE2 on PE1 and PE2 is only used by customer 1, and represented by ESI-2.

If we focus on CUST1 for now, there are 2 cases visible.

Case 1: For CE 1, if its ARP responses hash towards PE2, then PE1 will be unaware of its presence. For PE2 to synchronize this information to PE1, in addition to CE1 IP address (10.0.1.2) and MAC address (m1), 2 additional unique identifiers are needed. 1. IP-VRF. CUST 1 VRF is represented by EVI ID 1 2. Interface. BE2 Interface is represented by ESI-2

Case 2: For Host 1 (H1), if its ARP responses hash towards PE2, then PE1 will be unaware of its presence. For PE2 to synchronize this information to PE1, then in addition to H1 IP address (10.0.0.2) and MAC address (m2), 3 additional unique identifiers are required. 1. IP-VRF. CUST 1 VRF is represented by EVI ID 1 2. Main Interface. BE1 Interface is represented by ESI-1 3. Sub-Interface. Subnet/VLAN 1 is represented by Attachment Circuit ID 1.

[2.1.](#) Mapping of L3VRF to EVPN EVI

A separate EVPN instance will be configured to each layer-3 VRF and be marked for route-sync only. Each L3-VRF will have a unique associated EVI ID. The multi-homed peer PEs MUST have the same

configured EVI to layer-3 VRF mapping. This mapping also extends to the GRT, where a unique EVI ID can be assigned to support non VPN layer-3 services. Mis-configuration detection across peering PEs are left for further study.

When an EVPN instance is created as route-sync only, a MAC-VRF table is created to store all advertised routes. Local MAC learning may be disabled as this feature does not require MAC-only RT-2 advertisements.

This EVI is applicable to the multi-homed peer PEs only

The EVPN instance will be responsible for populating the following layer-3 VRF tables from remotely synced routes from peer PE

- o ARP/ND
- o IGMP
- o IP (for customer subnets learned from IGP adjacency)

In the example Figure 3, route-syncs from VRF CUST1 will have EVI-RT BGP Extended Community (EC) with EVI 1, and VRF CUST2 will have EVI 2.

[2.2.](#) Mapping for L3 Interface to ESI

The ESI represents the L3 LAG interface between PE and CEs. This ESI is signalled 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 others common ES.

In the example Figure 3, route-syncs from interface BE1 have ES-Import RT EC with ESI 1

[2.3.](#) Mapping for L3 Sub-Interface to Attachment Circuit ID

The Attachment Circuit ID represents the sub-interface subnet on the L3 LAG interface between PE and CEs. The AC-ID is signalled using RT-2, RT-5, RT-7 and RT-8 by attaching Attachment Circuit ID Extended community as described in Section 6.1 of [\[I-D.sajassi-bess-evpn-ac-aware-bundling\]](#).

In the example Figure 3, route-syncs from sub-interface BE1.1 (VLAN1) have Attachment-Circuit-ID EC with ID 1

[2.4.](#) Route sync for ARP/ND

This document proposes solving the issue described in [Section 1.1](#) using RT-2 IP/MAC route sync as described in [Section 10 of \[RFC7432\]](#) with a modification described below.

[2.4.1.](#) Local adjacency (ARP/ND) learning

Local ARP/ND learning will trigger a RT-2 route sync to any peer PE. There is no need for local MAC learning or sync over the L3 interface, only adjacencies. The MAC-only RT-2 route SHOULD not be advertised to peer PE.

[Section 9.1 of \[RFC7432\]](#) describes different mechanisms to learn adjacency routes locally.

- o An ARP/ND Sync route MUST carry exactly one ES-Import Route Target extended community, the one that corresponds to the ES on which the ARP or ND was received.
- o It MUST also carry exactly one EVI-RT EC, the one that corresponds to the EVI on which the ARP or ND was received. The EVI maps the layer-3 VRF See Section 9.5 of [\[I-D.ietf-bess-evpn-igmp-mld-proxy\]](#) for details on how to encode and construct the EVI-RT EC.
- o If the case where PE supports AC aware bundling, it MUST also carry one Attachment Circuit ID Extended Community. The circuit ID maps the sub-interface (or subnet) this route was received. For details on how to encode and construct this Extended Community, see section 6.1 of [\[I-D.sajassi-bess-evpn-ac-aware-bundling\]](#).

[2.4.2.](#) Remote ARP/ND learning

When consuming a remote layer-3 RT-2 sync route:

- o BGP only imports layer-3 sync route(s) when both ES-Import and EVI-RT extended communities match those locally configured

- o The layer-3 VRF is derived from the matching EVI
- o The main interface is derived from the ESI
- o The VLAN / sub-interface is derived from the AC-ID provided in the Attachment-Circuit-ID extended community
- o The combination of ES Import and EVI RT will allow BGP to import layer-3 sync route(s) to only PE(s) that have are attached to the same ESI and have the respective EVI.

[2.5.](#) Route sync for IGMP

This document proposes solving the issue described in [Section 1.2](#) using RT-7 and RT-8 route sync as described by [\[I-D.ietf-bess-evpn-igmp-mld-proxy\]](#).

Local IGMP join and leave will trigger a RT-7/8 route sync to peer PE.

[2.5.1.](#) Local IGMP Join/Leave learning

An IGP Join or Leave will trigger a RT-7/8 route sync to any peer PE.

[Section 9.1 of \[RFC7432\]](#) describes different mechanisms to learn adjacency routes locally.

- o An Multicast Join or Leave Sync route MUST carry exactly one ES-Import Route Target extended community, the one that corresponds to the ES on which the IGMP Join or Leave was received.
- o It MUST also carry exactly one EVI-RT EC, the one that corresponds to the EVI on which the IGMP Join or Leave was received. The EVI maps the layer-3 VRF See Section 9.5 of [\[I-D.ietf-bess-evpn-igmp-mld-proxy\]](#) for details on how to encode and construct the EVI-RT EC.
- o If the case where PE supports AC aware bundling, it MUST also

carry one Attachment Circuit ID Extended Community. The circuit ID maps the sub-interface (or subnet) this route was received. For details on how to encode and construct this Extended Community, see section 6.1 of [\[I-D.sajassi-bess-evpn-ac-aware-bundling\]](#).

- o The combination of ES Import and EVI RT will allow BGP to import Multicast Join and Leave sync route(s) to only PE(s) that have are attached to the same ESI and have the respective EVI.

[2.5.2.](#) Remote IGMP Join/Leave learning

When consuming a remote multicast RT-7 or RT-8 sync route:

- o BGP only imports multicast sync route(s) when both ES-Import and EVI-RT extended communities match those locally configured
- o The layer-3 VRF is derived from the matching EVI
- o The main interface is derived from the ESI
- o The VLAN / sub-interface is derived from the AC-ID provided in the Attachment-Circuit-ID extended community

[2.6.](#) Customer Subnet Route sync using Route-type(5)

Section 3 of [\[I-D.ietf-bess-evpn-prefix-advertisement\]](#) provides a mechanism to synchronize layer-3 customer subnets between the PEs in order to solve problem described in [Section 1.3](#).

Using Figure 2 as example, if PE1 forms the IGP adjacency with CE, it will be the only PE with knowledge of the customer subnet R1. BGP on PE1 will then advertise R1 to remote PEs using L3-VPN signalling.

Although PE2 has the same ES connection to the CE, and could provide load balancing to remote PEs, due to it not having formed an IGP adjacency with CE it is not aware of the customer subnet R1.

This can be solved by PE1 signaling R1 to PE2 using a RT-5 synch route. BGP on PE2 can then advertise this customer subnet R1 towards the core as if it was locally learned through IGP, and provide load-

balancing from the remote PEs.

The route-type(5) will carry the ESI as well as the gateway address GW (prefix next-hop address).

The same mapping mechanism will be used as for Route and IGMP sync, where EVI will determine the L3-VRF, ESI carried with route-type(5) will provide the main interface, and the gateway address will provide the nexthop.

[2.7.](#) Mapping for VLAN to ETAG

Another possible signalling of VLAN/sub-interface between service PE peers is to use the Ethernet Tag (ETAG) ID value in RT-2, RT-5, RT-7 and RT-8 as apposed to the Attachment Circuit Extended Community.

This will not work with vlan-aware bundling mode, but as that is a layer2 mode this should not prevent ETAGs use for L3 services.

[3.](#) Extensions to RT-2, RT-5, RT-7 and RT-8

This document proposes extending the usecase of Extended communities already defined in other drafts for the route types RT-2, RT-5, RT-7 and RT-8.

- o EVI-RT Extended Community as defined in Section 9.5 of [[I-D.ietf-bess-evpn-igmp-mld-proxy](#)].
- o Attachment Circuit ID Extended Community as defined in Section 6.1 of [[I-D.sajassi-bess-evpn-ac-aware-bundling](#)].

[4.](#) Convergence Considerations

[5.](#) Overall Advantages

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

- o Open standards based per interface all-active redundancy mechanism that eliminates the need to run ICCP and LDP.
- o Agnostic of underlay technology (MPLS, VXLAN, SRv6) and associated services (L3, L3-VPN)
- o Replaces legacy MC-LAG ICCP-based solution, and offers following additional benefits:
 - * Fast convergence with mass-withdraw is possible with EVPN, no equivalent in ICCP
- o Requires signalling already defined in existing EVPN RFCs [[RFC7432](#)] and drafts [[I-D.ietf-bess-evpn-igmp-mld-proxy](#)], [[I-D.sajassi-bess-evpn-ac-aware-bundling](#)], and [[I-D.ietf-bess-evpn-prefix-advertisement](#)]
- o Removes the burden of having the need for ICL link

[6.](#) Security Considerations

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

[7.](#) IANA Considerations

There are no IANA considerations.

[8.](#) References

[8.1.](#) Normative References

- [I-D.ietf-bess-evpn-igmp-mld-proxy]
 Sajassi, A., Thoria, S., Mishra, M., Patel, K., Drake, J., and W. Lin, "IGMP and MLD Proxy for EVPN", [draft-ietf-bess-evpn-igmp-mld-proxy-09](#) (work in progress), April 2021.
- [I-D.ietf-bess-evpn-prefix-advertisement]
 Rabadan, J., Henderickx, W., Drake, J., Lin, W., and A. Sajassi, "IP Prefix Advertisement in EVPN", [draft-ietf-bess-evpn-prefix-advertisement-11](#) (work in progress), May 2018.

[I-D.sajassi-bess-evpn-ac-aware-bundling]

Sajassi, A., Mishra, M., Thoria, S., Brissette, P., Rabadan, J., and J. Drake, "AC-Aware Bundling Service Interface in EVPN", [draft-sajassi-bess-evpn-ac-aware-bundling-03](#) (work in progress), February 2021.

[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>>.

[8.2](#). 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>>.

[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>>.

Authors' Addresses

Michael MacKenzie (editor)
Cisco Systems

Email: mimacken@cisco.com

Patrice Brissette
Cisco Systems

Email: pbrisset@cisco.com

Internet-Draft

EVPN L3MH

July 2021

Satoru Matsushima
Softbank

Email: satoru.matsushima@g.softbank.co.jp

MacKenzie, et al.

Expires January 12, 2022

[Page 17]