

BESS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 18 September 2026

A. Sajassi
LA. Burdet, Ed.
M. Mishra, Ed.
Cisco Systems
J. Rabadan
Nokia
J. Drake
Individual
17 March 2026

AC-Aware Bundling Service Interface in EVPN
draft-ietf-bess-evpn-ac-aware-bundling-06

Abstract

An EVPN (Ethernet VPNs) provides an extensible and flexible multihoming VPN solution over an MPLS/IP network for intra-subnet connectivity among Tenant Systems and End Devices that can be physical or virtual.

EVPN multihoming with Integrated Routing and Bridging (IRB) is one of the common deployment scenarios. Some deployments requires the capability to have multiple subnets designated with multiple VLAN IDs in the single broadcast domain.

EVPN technology defines three different types of service interface which serve different requirements but none of them address the requirement of supporting multiple subnets within a single broadcast domain. In this document, we define a new service interface type to support multiple subnets in the single broadcast domain. Service interface proposed in this document will be applicable to multihoming cases only.

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 18 September 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. Problem With Unicast MAC Route	6
1.2. Problem With Multicast Route Synchronization	6
1.3. Potential Security Concern Caused By Misconfiguration	6
2. Terminology	6
3. Requirements	7
4. Solution Description	8
4.1. Control Plane Operation	8
4.1.1. MAC/IP Address Advertisement	8
4.1.1.1. Local Unicast MAC Learning	8
4.1.1.2. Remote Unicast MAC Learning	8
4.1.2. Multicast Route Advertisement	8
4.1.2.1. Local Multicast State	9
4.1.2.2. Remote Multicast State	9
4.2. Data Plane Operation	9
4.2.1. Unicast Forwarding	9
4.2.2. Multicast Forwarding	10
5. Mis-configuration Across Multihoming PEs	10
6. BGP Encoding	10
6.1. Attachment Circuit Extended Community	10

6.2. Multiple Instances of Attachment Circuit Extended Community	11
6.2.1. Updating Ethernet Tag Field	12
6.2.2. Layer 2 Attributes	12
6.2.3. AC-Influenced Designated Forwarder Election	12
6.2.4. EVPN Fast Reroute	13
6.2.5. Forward-Looking Statement	13
7. Security Considerations	14
8. IANA Considerations	14
9. References	14
9.1. Normative References	14
9.2. Informative References	14
Acknowledgements	16
Contributors	16
Authors' Addresses	16

1. Introduction

EVPN based All-Active multihoming is becoming the basic building block for providing redundancy in next generation data center deployments as well as service provider access/aggregation network. For EVPN IRB mode, there are deployments which expect to be able to support multiple subnets within a single broadcast domain. Each subnet would be differentiated by VLAN. Thus, a single IRB interface can still serve multiple subnets.

Motivation behind such deployments are

1. Manageability: The support to have multiple subnets using a single broadcast domain requires only one broadcast domain and one IRB for "N" subnets compare to the "N" broadcast domain and "N" IRB interface to manage.
2. Simplicity: It avoids extra configuration by configuring VLAN Range with single BD and IRB as compared to individual VLAN, BD, and IRB interface per subnet.

[RFC7432] defines three types of service interface. None of them provide flexibility to achieve multiple subnets within a single broadcast domain. The different types of service interface from [RFC7432] are:

1. VLAN-Based Service Interface: With this service interface, an EVPN instance consists of only a single broadcast domain (e.g., a single VLAN). Therefore, there is a one-to-one mapping between a VID on this interface and a MAC-VRF.

2. **VLAN Bundle Service Interface:** With this service interface, an EVPN instance corresponds to multiple broadcast domains (e.g., multiple VLANs); however, only a single bridge table is maintained per MAC-VRF, which means multiple VLANs share the same bridge table. The MPLS-encapsulated frames **MUST** remain tagged with the originating VID. Tag translation is **NOT** permitted. The Ethernet Tag ID in all EVPN routes **MUST** be set to 0.
3. **VLAN-Aware Bundle Service Interface:** With this service interface, an EVPN instance consists of multiple broadcast domains (e.g., multiple VLANs) with each VLAN having its own bridge table -- i.e., multiple bridge tables (one per VLAN) are maintained by a single MAC-VRF corresponding to the EVPN instance.

From the definition, it seems like VLAN Bundle Service Interface does provide flexibility to support multiple subnets within a single broadcast domain. However, the requirement is to have multiple subnets from same ES on multihoming All-Active mode; that would not work. For example, let's take the case from Figure 1 where PE1 learns MAC of H1 on VLAN 1 (subnet S1). PE1 originates EVPN MAC route, as per [RFC7432], where the Ethernet Tag would be set to 0. Incoming packets from the IRB interface, at PE2, are untagged packets. PE2 does not have any associated AC information from EVPN MAC routes advertised by PE1. PE2 can not forward traffic that is destined to H1.

This document specifies an extension to existing service interface types defined in [RFC7432] and defines AC-aware Bundling service interface. AC-aware Bundling service interface would provide a mechanism to have multiple subnets in the single broadcast domain. This extension is applicable only for multihomed EVPN PEs.

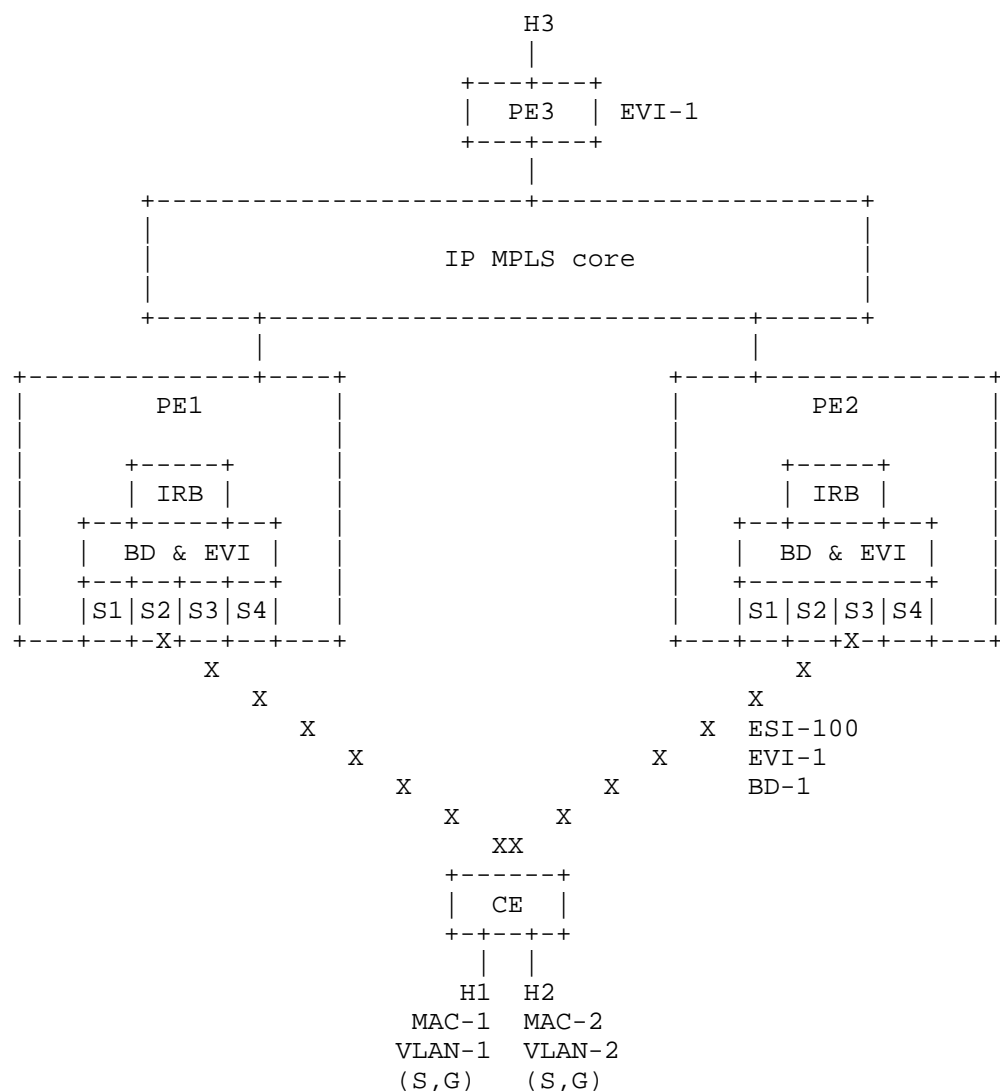


Figure 1 shows sample EVPN topology where PE1 and PE2 are multihomed PES. PE3 is remote PE participating in the same EVPN instance (EVI-1). It illustrates four subnets S1, S2, S3, and S4 where numerical value provides associated VLAN information.

1.1. Problem With Unicast MAC Route

In Figure 1, BD-1 has multiple subnets where each subnet is distinguished by VLAN 1, 2,3 and 4. PE1 learns MAC address MAC-1 from AC associated with subnet S1. PE1 uses MAC route to advertise MAC-1 presence to PE PEs. As per [RFC7432] MAC route advertisement from PE1 does not carry any context providing information about MAC address association with AC. When PE2 receives the MAC route with MAC-2, it can not determine the AC associated with this MAC.

Since PE2 could not bind MAC-1 with the correct AC when it receives data traffic destined for MAC-1, it does not know the destination AC since multiple bridge ports have the same ESI assignment.

1.2. Problem With Multicast Route Synchronization

[RFC9251] defines mechanism to synchronize multicast routes between multihome PEs. In the above case, if the receiver behind S1 sends IGMP membership request, CE could hash it to either of the PEs. When a multicast route is originated, it does not contain any AC information. Once it reaches peer PE, it does not have any information about which subnet this IGMP membership request belongs. Similarly to the unicast traffic problem, the incoming multicast traffic from IRB cannot be forwarded to the proper AC.

1.3. Potential Security Concern Caused By Misconfiguration

In the case of a single subnet per broadcast domain, there is a potential case of security issue. For example, PE1 has BD1 configured with VLAN-1 and multihome PE PE2 has BD1 configured with VLAN-2. Each of the IGMP membership requests on PE1 would be synchronized to PE2 and PE2 would process multicast routes and start forwarding multicast traffic on VLAN-2, which was not intended. Again, a similar issue can potentially be seen with unicast traffic.

2. Terminology

- * AC: Attachment circuit.
- * BD: broadcast domain. As per [RFC7432], an EVI consists of a single or multiple BDs. In case of VLAN-bundle and VLAN-based service models (see [RFC7432]), a BD is equivalent to an EVI. In the case of the VLAN-aware bundle service model, an EVI contains multiple BDs. Also, in this document, BD and subnet are equivalent terms.

- * BD Route Target: refers to the broadcast domain assigned to Route Target [RFC4364]. In the case of the VLAN-aware bundle service model, all the BD instances in the MAC-VRF share the same Route Target.
- * BT: Bridge Table. The instantiation of a BD in a MAC-VRF, as per [RFC7432].
- * Ethernet A-D route: Ethernet Auto-Discovery (A-D) route, as per [RFC7432].
- * EVI: EVPN Instance spanning the NVE/PE devices that are participating on that EVPN, as per [RFC7432].
- * EVPN: Ethernet Virtual Private Networks, as per [RFC7432].
- * IRB: Integrated Routing and Bridging interface. It connects an IP-VRF to a BD (or subnet).
- * MAC-VRF: A Virtual Routing and Forwarding Table for Media Access Control (MAC) addresses on an NVE/PE, as per [RFC7432]. A MAC-VRF is also an instantiation of an EVI in an NVE/PE.
- * PE: Provider edge device hosting EVPN instance
- * 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].
- * VLAN: The usage of VLAN refers to the 802.1Q or 802.1AD tag.
- * (S, G): Multicast membership request for source S and group G.
- * This document also assumes familiarity with the terminology of [RFC7432],[RFC8365], [RFC7365].

3. Requirements

1. A new service interface represents an attachment-circuit where multiple VLANs are configured. Each of these VLANs is represented by a different AC ID (Identifier) under a single broadcast domain.
2. Service interface MUST be applicable to multihomed PEs only.

3. Service interface MUST have an Ethernet-Segment identifier assignment.
4. New service interface handling procedures MUST be backward compatible with implementation procedures defined in [RFC7432].
5. New service interface MUST support EVPN multicast routes defined in [RFC9251] too.

4. Solution Description

4.1. Control Plane Operation

4.1.1. MAC/IP Address Advertisement

4.1.1.1. Local Unicast MAC Learning

Section 9.1 of [RFC7432] describes different mechanism to learn Unicast MAC address locally. At those PEs where AC aware bundling is supported, the MAC address is learned along with VLAN associated with AC.

MAC/IP advertisement route construction follows the mechanism defined in section 9.2.1 of [RFC7432]. An Attachment Circuit Extended Community (Section 6.1) MUST be attached to EVPN Route Type 2.

4.1.1.2. Remote Unicast MAC Learning

Presence of Attachment Circuit Extended Community (Section 6.1) MUST be ignored by non multihoming PEs. Remote PE (non-multihomed PE) MUST process MAC route as defined in [RFC7432].

Multihoming PE MUST process Attachment Circuit Extended Community (Section 6.1) to associate the remote MAC address with the appropriate AC.

From Figure 1, PE2 receives MAC route for MAC-1. It MUST get an AC-ID from the Attachment Circuit Extended Community (Section 6.1) in Route Type 2 and associate the MAC address with the specific subnet.

4.1.2. Multicast Route Advertisement

4.1.2.1. Local Multicast State

When a local multihomed PE in a given broadcast domain receives IGMP membership request on local AC, it MUST synchronize multicast state by originating multicast route defined in [RFC9251]. When Service interface is AC aware it MUST attach Attachment Circuit Extended Community (Section 6.1) along with the multicast route. For example in Figure 1 when H2 sends IGMP membership requests for (S, G), and CE hashed it to one of the PEs. Lets say PE1 received an IGMP membership request. PE1 MUST originate multicast route to synchronize multicast state with PE2. Multicast route MUST contain Attachment Circuit Extended Community (Section 6.1) along with multicast route.

PE1 MUST originate multicast route updates for any subsequent IGMP membership requests under same or different subnet attaching adequate Attachment Circuit ID Extended Community (Section 6.1).

4.1.2.2. Remote Multicast State

If multihomed PE receives a remote multicast route on the broadcast domain for a given ES, route MUST be programmed to the correct subnet. Subnet information MUST be extracted from Attachment Circuit Extended Community. That value maps to the VLAN of a local AC where the multicast route is associated with.

4.2. Data Plane Operation

4.2.1. Unicast Forwarding

Packet received from CE MUST follow the same procedure as defined in Section 13.1 of [RFC7432].

Unknown Unicast packets from a Remote PE MUST follow the procedure as per Section 13.2.1 of [RFC7432].

Known unicast Received on a remote PE MUST follow the procedure as per [RFC7432] section 13.2.2. In Figure 1, if PE3 receives a known unicast packet for destination MAC MAC-1, it MUST follow the procedure defined in Section 13.2.2 of [RFC7432].

If destination MAC lookup is performed on a known unicast packet, destination MAC lookup MUST provide VLAN and local AC information. For example, if PE2 receives a unicast packet that is destined to MAC-1 (packet might be coming from IRB or remote PE with EVPN tunnel), destination MAC lookup on PE2 MUST provide an outgoing port along with associated VLAN value.

4.2.2. Multicast Forwarding

Multicast traffic from CE and remote PE MUST follow the procedure defined in [RFC7432].

Multicast traffic received from IRB interface or EVPN tunnel, route lookup would be performed based on IGMP snooping state and traffic would be forwarded to the appropriate AC.

5. Mis-configuration Across Multihoming PEs

If there is misconfiguration of VLAN or VLAN range across multihoming PEs, the same MAC address or IGMP membership requests would be learned with different VLANs per broadcast domain. This is detected by a received Route having an ESI which is local, but an Attachment Circuit ID which does not match any Normalized VID in the local bridge domain (or vice-versa).

In this case, an error message MUST be thrown for the operator to make configuration changes. Furthermore, the errored MAC route MUST be ignored.

6. BGP Encoding

This document defines a new BGP Extended Community for EVPN and updates several existing Extended Communities.

6.1. Attachment Circuit Extended Community

A new BGP Extended Community called Attachment Circuit ID Extended Community is introduced. This new extended community is a transitive extended community with the Type field of 0x06 (EVPN) and the Sub-Type of 0x0E. It is advertised along with EVPN MAC/IP Advertisement Route (Route Type 2) per [RFC7432] for AC-Aware Bundling Service Interface. It may also be advertised along with EVPN Multicast Route (Route Type 7 and 8) as per [RFC9251]. Generically speaking, the new extended community MUST be attached to any routes that require specific VLAN identification.

The Attachment Circuit Extended Community is encoded as an 8-octet value as follows:

0		1		2		3																																																									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																																
+-----+																+-----+																+-----+																+-----+															
Type=0x06								Sub-Type=0x0E								Instance (2 octets)																																															
+-----+																+-----+																+-----+																+-----+															
								Attachment Circuit ID (4 octets)																																																							
+-----+																+-----+																+-----+																+-----+															

Figure 2: Attachment Circuit Extended Community

The Attachment Circuit ID field represents a Normalized VID, the value of which is derived using the same method as in [RFC9744].

The Instance field is used to uniquely identify an Attachment Circuit Extended Community when multiple instances are attached to a same route. When only a single instance is present on a route, this field is set to 0.

6.2. Multiple Instances of Attachment Circuit Extended Community

The procedures described in this document are backwards compatible with [RFC7432] VLAN-aware bundling mode since the Ethernet Tag ID field remains intact. This, however, may present a drawback: AC-Aware Bundle use-cases may result in multiple ACs being represented by a single EVPN route.

For instance with multicast, the same (S,G) may be used over different subnets like S1-S4 in interface ESI-100 of Figure 1. In that case, the same Route-Type 7 MUST carry multiple Attachment Circuit Extended Communities, one instance per attachment circuit / VLAN.

Similarly for unicast MAC addresses MAC-1, MAC-2 in subnets S1-S4 of same interface ESI-100 in Figure 1, separate Route Type 2 will be advertised with Attachment Circuit Extended Community according to this document.

In both cases, a single Ethernet A-D per EVI Route Type 1 is advertised.

It may also happen that the number of VLAN is fairly large, and multiple routes with different BGP Route Distinguishers may be necessary to carry the required amount of Extended Communities.

These additional Route Distinguishers or functionality implemented in concurrent specifications updating Ethernet-AD per EVI behaviour not being defined for AC-Aware Bundling, add complexity to the overall solution and implementation. The following subsections address this for known documents and provide future-looking guidance for AC-Aware Bundling support.

6.2.1. Updating Ethernet Tag Field

To remedy the possible large number of VLAN or uniqueness issues correlating multiple Attachment Circuit Extended Community instances, the Attachment Circuit ID MAY be set to 0xFFFF_FFFF, which does not correspond to a valid Normalized VID.

That value tells peer PE that the Attachment Circuit ID is carried as part of the Ethernet Tag ID field of the route. Since the key of such an EVPN route is now unique, multiple Attachment Circuit Extended Communities per route is no longer required. This however poses backward-compatibility and interoperability issues with remote PE(s) expecting a zero Ethernet Tag ID and/or with VLAN-Aware Bundle Service Interface Section 6.3 of [RFC7432].

6.2.2. Layer 2 Attributes

The use of the EVPN Layer 2 Attributes Extended Community ("L2-Attr") for bridge domains is introduced in [I-D.ietf-bess-rfc7432bis], specifically with the intent of signaling Primary and Backup states in the Control Flags field. The extended community is added to Ethernet A-D per EVI routes.

When multiple ACs / VLAN are present, DF-Election may result in different Primary (P) and Backup (B) states for each subinterface. When this is the case, supporting for AC-Aware Bundling is achieved by adding multiple L2-Attr Extended Communities on the Ethernet A-D per EVI route, each with a unique Instance. To associate this L2-Attr with a specific AC, Attachment Circuit Extended Communities are also added each with a corresponding value in their Instance field(s)

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Type=0x06      | Sub-Type=0x04 |   Control Flags (2 octets)   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|   L2 MTU (2 octets)   |   Instance (2 octets)   |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 3: Updated EVPN Layer 2 Attributes Extended Community

6.2.3. AC-Influenced Designated Forwarder Election

[RFC8584] defines procedures for peering PEs to signal AC-Influenced DF election and the desire to use AC-DF with the rest of the PEs in the ES. The procedure defined further relies on withdrawing (or not advertising) the Ethernet-AD per EVI corresponding to an AC in failed or misconfigured state.

When multiple attachment circuit / VLAN are present, each individual

AC may be misconfigured (missing) or in failed state. The Attachment Circuit Extended Community in Ethernet A-D per EVI routes is of generalized assistance, allowing to compare lists of local AC / VLAN Normalized VIDs to those received in remote routes with a same ESI (peer). A PE which receives an Ethernet A-D per EVI route without the Attachment Circuit ID corresponding to its local Normalized VID may assume the peer has misconfigured this subnet / VLAN or the AC has failed and perform AC-Influenced DF election for that AC as if the Ethernet A-D had been withdrawn.

6.2.4. EVPN Fast Reroute

[I-D.ietf-bess-evpn-fast-reroute] defines procedures for peering PEs to signal reroute labels with special disposition attributes in order to support a fast reroute mechanism for traffic loss avoidance on failure(s).

These reroute labels may be allocated at the bridge or at the AC granularity. Allocation at the bridge granularity poses no problem for AC-Aware Bundling and the Instance field may remain 0 (same value as previously Reserved=0).

However, to support AC allocation granularity and AC-Aware Bundling, the ESI Label Extended Community is updated to include an Instance field for correlation against a specific Attachment Circuit Extended Community. The Attachment Circuit ID from the extended community with matching Instance identifies the specific AC for which the reroute label is intended to be installed.

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Type=0x06      | Sub-Type=0x01 | Flags(1 octet)|      Instance  ~
+-----+-----+-----+-----+-----+-----+-----+-----+
~ (2 octets)    |      ESI Label (3 octets)      |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 4: Updated ESI Label Extended Community

6.2.5. Forward-Looking Statement

Generically speaking, any new use-case requiring a new or updated Extended Community where information pertaining to multiple ACs is to be included in a single EVPN route, MUST contain or be updated to include a 2 octet Instance field. An Attachment Circuit Extended Community MUST be included for each AC, with the same unique value in its Instance field for demultiplexing information onto the correct AC.

7. Security Considerations

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

8. IANA Considerations

IANA has made the following assignment in the "EVPN Extended Community Sub-Types" registry set up by [RFC7153].

Sub-Type Value	Name	Reference
0x0E	Attachment Circuit Extended Community	This document

Table 1

9. References

9.1. Normative References

- [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>>.
- [RFC7153] Rosen, E. and Y. Rekhter, "IANA Registries for BGP Extended Communities", RFC 7153, DOI 10.17487/RFC7153, March 2014, <<https://www.rfc-editor.org/info/rfc7153>>.
- [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>>.

9.2. Informative References

- [I-D.ietf-bess-evpn-fast-reroute] Burdet, L. A., Brissette, P., Miyasaka, T., Rabadan, J., Liu, Y., and C. Lin, "EVPN Fast Reroute", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-fast-reroute-00, 9 June 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-fast-reroute-00>>.

[I-D.ietf-bess-rfc7432bis]

Sajassi, A., Burdet, L. A., Drake, J., and J. Rabadan,
"BGP MPLS-Based Ethernet VPN", Work in Progress, Internet-
Draft, draft-ietf-bess-rfc7432bis-13, 24 June 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-bess-rfc7432bis-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>>.

[RFC7365] Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y.
Rekhter, "Framework for Data Center (DC) Network
Virtualization", RFC 7365, DOI 10.17487/RFC7365, October
2014, <<https://www.rfc-editor.org/info/rfc7365>>.

[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>>.

[RFC8365] Sajassi, A., Ed., Drake, J., Ed., Bitar, N., Shekhar, R.,
Uttaro, J., and W. Henderickx, "A Network Virtualization
Overlay Solution Using Ethernet VPN (EVPN)", RFC 8365,
DOI 10.17487/RFC8365, March 2018,
<<https://www.rfc-editor.org/info/rfc8365>>.

[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>>.

[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>>.

Acknowledgements

The authors thank Tapraj Singh and Mei Zhang and Gunter Van de Velde for their careful reviews.

Contributors

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

Patrice Brissette
Cisco Systems
Canada
Email: pbrisset@cisco.com

Samir Thoria
Cisco Systems
United States of America
Email: sthoria@cisco.com

Authors' Addresses

Ali Sajassi
Cisco Systems
Email: sajassi@cisco.com

Luc André Burdet (editor)
Cisco Systems
Canada
Email: lburdet@cisco.com

Mankamana Mishra (editor)
Cisco Systems
Email: mankamis@cisco.com

Jorge Rabadan
Nokia

Email: jorge.rabadan@nokia.com

John Drake
Individual
Email: je_drake@yahoo.com