

BGP Enabled Services
Internet-Draft
Intended status: Standards Track
Expires: 5 December 2026

P. Narasimhaprasad
M. Rumuly
A. Shashidhar
Arista Networks
3 June 2026

EVPN VLAN Tag Extended Community
draft-pavann-bess-evpn-vlan-tag-00

Abstract

Ethernet Virtual Private Network (EVPN), as defined in RFC 7432, provides a control plane for distributing MAC and IP address bindings using BGP MAC/IP Advertisement routes. In several deployment scenarios, policy decisions depend not only on MAC/IP bindings but also on VLAN encapsulation information, particularly in environments using - IEEE 802.1Q VLAN tagging and IEEE 802.1ad provider bridging (QinQ).

However, RFC 7432 does not define a mechanism to propagate VLAN information associated with MAC/IP bindings. This document defines a new EVPN Extended Community that carries VLAN identifiers to enable consistent policy enforcement across EVPN Provider Edge devices. This document does not modify EVPN route selection or forwarding behavior defined in RFC 7432.

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 5 December 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	2
2. Terminology	3
2.1. Requirements Language	3
2.2. Terms and Abbreviations	3
3. The EVPN VLAN-Tag Extended Community	3
3.1. Use of the EVPN VLAN Tag Extended Community	5
3.1.1. Transmission of the EVPN VLAN Tag Extended Community	5
3.1.2. Reception of the EVPN VLAN Tag Extended Community	6
4. IANA Considerations	6
5. Security Considerations	6
6. References	6
6.1. Normative References	6
Authors' Addresses	7

1. Introduction

[RFC7432] defines MAC/IP Advertisement Route, which can optionally carry IPv4 or IPv6 addresses associated with a MAC address. The MPLS Label1 field is encoded as 3 octets, where the high-order 20 bits contain the label value. The MPLS Label1 must be downstream assigned, and it is associated with the MAC address being advertised by the advertising PE. The advertising PE uses this label for applying policies on the received route and performs forwarding based on the destination MAC address toward the CE.

In case any routing policies need to be made on an additional label such as a VLAN identifier, then the information conveyed in the EVPN MAC/IP Advertisement Route may not be enough for the remote PE to apply the right policies and eventually make the correct forwarding decision towards the CE.

A new Extended Community that is advertised along with an EVPN MAC/IP Advertisement Route and carries information relevant to the VLAN identifiers so that an EVPN PE can apply the right policies based on this information is defined in this document.

2. Terminology

2.1. 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.

2.2. Terms and Abbreviations

EVPN

Ethernet Virtual Private Network. A BGP-based protocol for building VPNs and network overlays. Defined in [RFC7432] for use with MPLS encapsulation, and extended in [RFC8365] for use with alternative encapsulations, including VXLAN

PE

Provider Edge device

VLAN

Virtual LAN. A single bridging domain local to a device.

CE

Customer Edge device

BD

Broadcast Domain

ARP

Address Resolution Protocol

ND

Neighbor Discovery protocol, specified in [RFC4861]

3. The EVPN VLAN-Tag Extended Community

This document defines a transitive EVPN Extended Community (Type field value of 0x06) with a Sub-Type of TBD, as allocated by IANA. It is advertised along with EVPN MAC/IP Advertisement routes that carry an IPv4 or IPv6 address.

```

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=TBD  |  Flags  |  Reserved  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  VLANID3    |          VLANID2          |          VLANID1          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Flags field

```

0 1 2 3 4 5
+-----+
|Stacke.|VL.|
+-----+

```

The Flags field is a composite value of Stacked VLAN Type and VLAN count. Stacked VLAN Type can be one of the following values:

0x01 - Host was learnt via Single-tagged IEEE 802.1Q frames. The following represents a Single-tagged 802.1Q packet format

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|DestMAC | SrcMAC | 802.1Q  | EtherType | Payload | CRC/FCS |
|          |          | TPID=0x8100 |          |          |          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

0x02 - Host was learnt via IEEE 802.1Q double-tagged frames. The following represents a 802.1Q double-tagged packet format

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|DestMAC | SrcMAC | 802.1Q  | 802.1Q  | EtherType | Payload | CRC/FCS |
|          |          | TPID=0x8100 | TPID=0x8100 |          |          |          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

0x03 - Host was learnt via IEEE 802.1ad frames. The following represents a IEEE 802.1ad double-tagged packet format

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|DestMAC | SrcMAC | 802.1Q  | 802.1Q  | EtherType | Payload | CRC/FCS |
|          |          | TPID=0x88A8 | TPID=0x8100 |          |          |          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

VLAN Count indicates the number of VLAN IDs present in the extended community. The valid range of this field must be 1-3. All other values are invalid. The receiver PE MUST treat an invalid value as an invalid Extended Community and MUST ignore it.

VLANID1, VLANID2, VLANID3 must be in the range of 1 through 4094. Any other value MUST be considered as invalid.

Reserved field MUST be set to zero on transmission and MUST be ignored on receipt.

If more than one instance of this Extended Community is present in the EVPN MAC/IP route Advertisement, then the first EVPN VLAN Extended community is considered and the rest of the EVPN VLAN Extended communities MUST be ignored.

3.1. Use of the EVPN VLAN Tag Extended Community

This section describes the relevant procedures when advertising and processing the EVPN VLAN Tag Extended Community. In this section, the term "PE" refers to a PE that supports the procedures defined in this document.

3.1.1. Transmission of the EVPN VLAN Tag Extended Community

A PE MAY learn IP-to-MAC bindings through mechanisms other than EVPN, including:

- * Static configuration via the management plane
- * ARP or Neighbor Discovery (ND) learning (including proxy ARP/ND)
- * Snooping of ARP or Neighbor Advertisement (NA) messages received from a CE

When advertising the EVPN VLAN Tag community using EVPN MAC/IP Advertisement routes:

- * If the binding is learned from a single-tagged IEEE 802.1Q frames, the PE MAY include the EVPN VLAN Tag Extended Community with Stacked VLANs = 0x01.
- * If the binding is learned from double-tagged IEEE 802.1Q frames, the PE MAY include the EVPN VLAN Tag Extended Community with: Stacked VLANs = 0x02.
- * If the binding is learned from IEEE 802.1ad frames, the PE MAY include the EVPN VLAN Tag Extended Community with: Stacked VLANs = 0x03.
- * Based on the number of VLAN IDs that are present as part of the VLANIDs field, the count field is set appropriately whose value will have a range of 1-3.

- * The Reserved bits of the VLANIDs are always set to 0.

This Extended Community does not modify procedures defined in RFC 7432, including the use of the MAC Mobility Extended Community.

3.1.2. Reception of the EVPN VLAN Tag Extended Community

Upon receiving an EVPN MAC/IP Advertisement route, a PE MUST process the EVPN VLAN Tag Extended Community as follows:

- * A route MUST NOT have more than one such Extended Community. If multiple instances are received then the first community is processed and all other instances MUST be ignored.
- * Based on the VLAN Count received, VLANIDs starting from VLANID1 are appropriately read.

4. IANA Considerations

A new transitive extended community Type of 0x06 and Sub-Type of TBD for EVPN VLAN Tag Extended Community needs to be allocated by IANA.

5. Security Considerations

This document relies on the EVPN control plane defined in [RFC9742]. Therefore, all security considerations described therein apply.

An attacker MAY exploit ARP or ND mechanisms to create excessive state in PEs within the same Broadcast Domain, potentially leading to resource exhaustion. Implementations SHOULD consider mitigation techniques, including:

- * Limiting the number of proxy ARP/ND entries per BD or per interface.
- * Monitoring and rate-limiting the creation of dynamic entries.

This document does not introduce new security considerations beyond those already present in EVPN.

6. References

6.1. 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>>.

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC9742] Clarke, J., Ed., Jethanandani, M., Ed., Wildes, C., Ed., and K. Koushik, Ed., "A YANG Data Model for Syslog Management", RFC 9742, DOI 10.17487/RFC9742, April 2025, <<https://www.rfc-editor.org/info/rfc9742>>.

Authors' Addresses

Pavan Narasimhaprasad
Arista Networks
Email: pavann@arista.com

Mason Rumuly
Arista Networks
Email: mrumuly@arista.com

Akhil Shashidhar
Arista Networks
Email: akhil@arista.com