

BESS Workgroup
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2026

J. Rabadan, Ed.
S. Sathappan
M. Gautam
Nokia
P. Brissette
Cisco Systems
W. Lin
Juniper
2 March 2026

Domain Path (D-PATH) for Ethernet VPN (EVPN) Interconnect Networks
draft-ietf-bess-evpn-dpath-04

Abstract

The BGP Domain PATH (D-PATH) attribute is defined for Inter-Subnet Forwarding (ISF) BGP Sub-Address Families that advertise IP prefixes. When used along with EVPN IP Prefix routes or IP-VPN routes, it identifies the domain(s) through which the routes have passed and that information can be used by the receiver BGP speakers to detect routing loops or influence the BGP best path selection. This document extends the use of D-PATH so that it can also be used along with other EVPN route types.

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 3 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	2
1.1. D-PATH to Prevent Loops for EVPN Routes	3
1.2. Add Path Visibility and Influence Best Path Selection for EVPN Routes	4
2. Conventions used in this document	5
3. Terminology	5
4. Use of Domain Path Attribute (D-PATH) with EVPN routes	7
4.1. D-PATH and Best Path Selection for MAC/IP Advertisement routes	9
4.2. D-PATH and Best Path Selection for Ethernet A-D per EVI routes	10
4.3. D-PATH and Best Path Selection for Inclusive Multicast Ethernet Tag routes	10
4.4. Loop Detection	10
4.5. Error Handling	11
5. Use-Case Examples	11
6. Operational Considerations	14
7. Security Considerations	14
8. IANA Considerations	15
9. Acknowledgments	15
10. Contributors	15
11. References	15
11.1. Normative References	15
11.2. Informative References	16
Authors' Addresses	16

1. Introduction

The BGP Domain PATH (D-PATH) attribute [I-D.ietf-bess-evpn-ipvpn-interworking] is defined for Inter-Subnet Forwarding (ISF) BGP Sub-Address Families that advertise IP prefixes. When used along with EVPN IP Prefix routes or IP-VPN routes, it identifies the domain(s) through which the routes have passed and that information can be used by the receiver BGP speakers to detect routing loops or influence the BGP best path selection. This document extends the use of D-PATH so that it can also be used along with other EVPN route types.

The D-PATH attribute can be used to prevent control plane loops for EVPN routes, or to provide full path visibility of all the EVPN Interconnect Gateways through which a route has gone and modify the best path selection based on it. Some use cases in which D-PATH can be used along with (non-IP Prefix) EVPN routes follow, but the use cases are not limited to the ones described in this section.

1.1. D-PATH to Prevent Loops for EVPN Routes

Figure 1 illustrates an EVPN Interconnect case where EVPN MAC/IP Advertisement routes can be looped indefinitely. The three Gateways (GW1, GW2 and GW3) and PE1 in the diagram are attached to the same EVPN Broadcast Domain (BD1). However, BD1 is extended throughout three different domains that are interconnected by the Gateways, which follow [RFC9014] procedures. Suppose a host with MAC address M1 is learned on GW1 and GW1 advertises an EVPN MAC/IP Advertisement route for M1 into Domain-1 and Domain-2. When the route gets imported by GW2 and GW3 and later exported into Domain-3, GW2 and GW3 may re-originate each other's route for M1 back into Domain-1 and Domain-2, respectively, creating a loop. D-PATH can be used by the Gateways when re-originating the route between Domains, to identify the Domains through which the route for M1 has gone. When GW1 receives an EVPN MAC/IP Advertisement route for M1 that contains a D-PATH with a domain-id locally assigned, GW1 identifies the route as "looped".

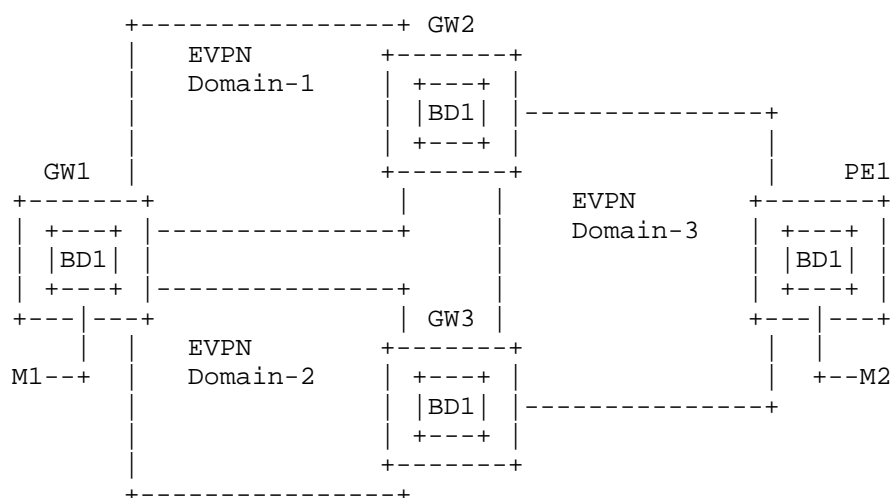


Figure 1: Loops for EVPN routes

It is important to note that the example above illustrates the use of D-PATH with EVPN MAC/IP Advertisement routes to prevent control plane loops. Data plane loops caused by BUM traffic - such as broadcast frames sent from M1 in Figure 1 - are prevented through other mechanisms. For instance, the use of Interconnect Ethernet Segments, as described in [RFC9014], ensures that BUM traffic is forwarded by only one of the Gateways connected to the same Domain.

Similar examples to those in Figure 1 can also occur with EVPN VPWS services on the Gateways and PEs, where preventing loops for re-originated AD per EVI routes is necessary. D-PATH delivers the end-to-end path visibility required to avoid such loops.

1.2. Add Path Visibility and Influence Best Path Selection for EVPN Routes

Figure 2 illustrates another [RFC9014] EVPN Interconnect case where, in addition to using D-PATH to prevent EVPN MAC/IP Advertisement route loops when re-originating routes between domains, the D-PATH attribute can also influence the best path selection for the routes. For example, if all the Gateways in the diagram are attached to the same BD1, an EVPN MAC/IP Advertisement route for MAC address M1 advertised by GW1 is advertised into Domain-1. Two routes for M1 will arrive at GW3 with different route distinguishers and BGP Next Hops. If D-PATH is used by all the Gateways, the two routes arriving at GW3 will have a different sequence of domain-ids in the D-PATH attribute. GW3 can use the length of the D-PATH as a way of influencing the selection (i.e., the shortest D-PATH route is selected). D-PATH improves the path visibility of the route since it provides information about all the Domains through which the route has passed.

As mentioned in Section 1.1, data plane loops caused by BUM traffic are handled by other mechanisms - such as Interconnect Ethernet Segments [RFC9014] - which are outside the scope of this document.

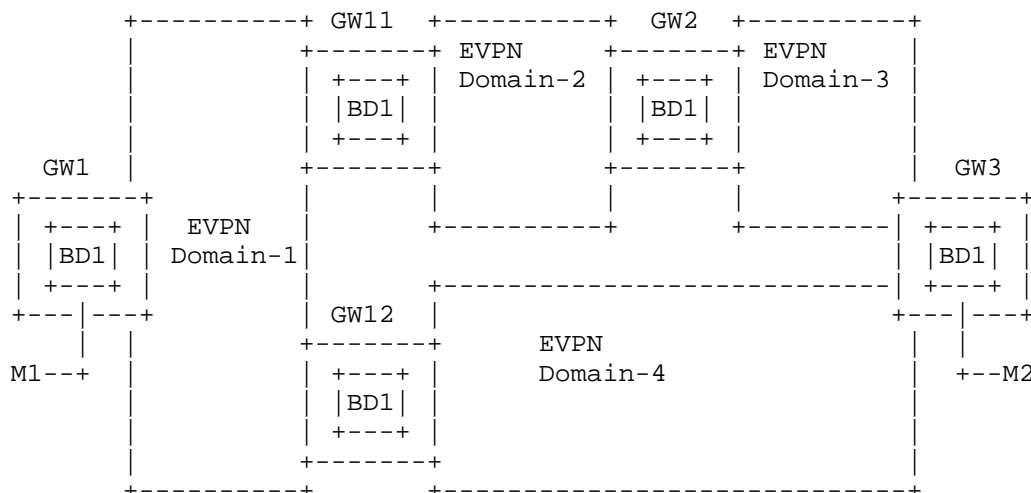


Figure 2: Influence Best Path Selection for EVPN routes

2. Conventions used in this document

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.

3. Terminology

This section summarizes the terminology that is used throughout the rest of the document.

- * AC: Attachment Circuit or logical interface associated to a given Broadcast Domain. To determine the AC on which a packet arrived, the PE examines the combination of a physical port and VLAN tags (where the VLAN tags can be individual c-tags, s-tags or ranges of both).
- * BD and BT: a Broadcast Domain and Bridge Table, as defined in [I-D.ietf-bess-rfc7432bis]. A BD is a group of PEs attached to the same EVPN layer-2 multipoint service. A BT is the instantiation of a Broadcast Domain in a PE. When there is a single Broadcast Domain in a given EVI, the MAC-VRF in each PE contains a single BT. When there are multiple BTs within the same MAC-VRF, each BT is associated to a different Ethernet Tag. The EVPN routes specific to a BT indicate which Ethernet Tag the route corresponds to. This document does not distinguish between a

"Broadcast Domain" and a "Bridge Table", and will use the terms interchangeably (or will use the acronym "BD" to refer to either). Also, the way the BDs are grouped into MAC-VRFs depending on the service model (VLAN-based versus VLAN Bundle or VLAN-aware Bundle) is not relevant to the procedures specified in this document.

- * BUM: Broadcast, Unknown unicast and Multicast traffic.
- * ES and ESI: Ethernet Segment and Ethernet Segment Identifier, as defined in [I-D.ietf-bess-rfc7432bis].
- * EVPN Domain: two PEs are in the same EVPN Domain if they are attached to the same service and the packets between them do not require a data path lookup of the inner frame (e.g., in the BD of a MAC-VRF) in any intermediate router. An EVPN Domain Gateway PE is always configured with multiple Domain identifiers (EVPN Domain-ID) in the MAC-VRF or VPWS that connects those EVPN Domains, each EVPN Domain-ID representing an EVPN Domain.

Example: Figure 3 illustrates an example where PE1 and PE2 belong to different EVPN Domains since packets between them (for flows between hosts with MAC addresses M1 and M2) require a MAC lookup in two of the gateways that are connecting the three EVPN Domains. E.g., if frames from M1 to M2 go through PE1, GW1, GW3 and PE2, a MAC lookup is performed at GW1 and GW3.

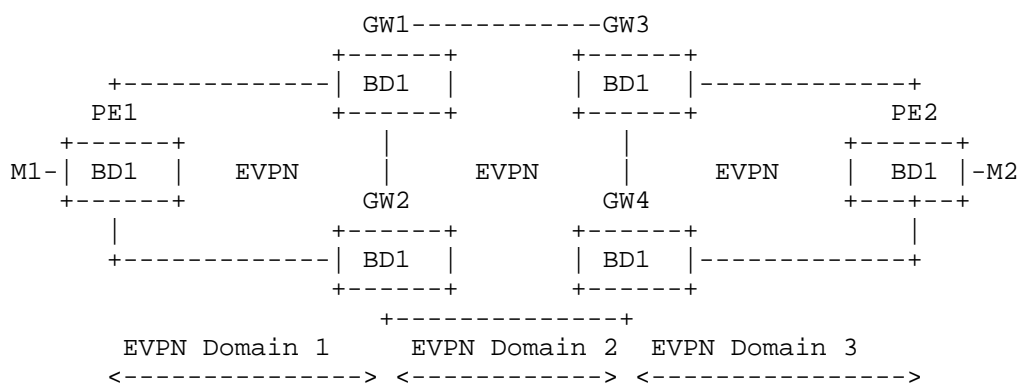


Figure 3: EVPN Domain Interconnect Example

- * EVPN Domain Gateway: a PE where a service (BD or VPWS instance) is instantiated and is attached to two or more EVPN Domains. This document uses the term "EVPN Domain Gateway" or simply "Gateway". An example of EVPN Domain Gateway is a PE following the procedures in section 4.4 or section 4.6 of [RFC9014]. In the example in Figure 3, GW1 and GW2 connect EVPN Domains 1 and 2, whereas GW3

and GW4 connect EVPN Domains 2 and 3. GW1 and GW2 import the MAC/IP Advertisement route for M1 coming from the EVPN Domain 1 into the MAC-VRF for BD1, and re-originate it into EVPN Domain 2. Likewise, GW3 and GW4 import the route into their MAC-VRF and re-advertise it into EVPN Domain 3.

- * MAC-VRF: a MAC Virtual Routing and Forwarding table, as defined in [I-D.ietf-bess-rfc7432bis]. It is also the instantiation of an EVI (EVPN Instance) in a PE.

4. Use of Domain Path Attribute (D-PATH) with EVPN routes

This document extends the use of the D-PATH attribute, as specified in [I-D.ietf-bess-evpn-ipvpn-interworking], to allow it to be advertised and processed in conjunction with the following EVPN route types:

- * EVPN MAC/IP Advertisement routes that are not used for Inter-Subnet Forwarding (ISF). If an EVPN MAC/IP Advertisement route is used for ISF, as specified in [RFC9135], the procedures for the D-PATH advertisement and processing are defined in [I-D.ietf-bess-evpn-ipvpn-interworking].
- * EVPN A-D per EVI routes that are used for EVPN VPWS [RFC8214].
- * EVPN Inclusive Multicast Ethernet Tag routes [I-D.ietf-bess-rfc7432bis].

The advertisement of D-PATH in the routes specified above is disabled by default and MUST be explicitly enabled by configuration.

The following EVPN routes SHOULD NOT include D-PATH:

- * A-D per EVI routes not used for EVPN VPWS. Examples include A-D per EVI routes used as specified in [I-D.ietf-bess-rfc7432bis], or [I-D.ietf-bess-evpn-ip-aliasing], which are not associated with EVPN VPWS.
- * ES routes, as specified in [I-D.ietf-bess-rfc7432bis].

The use of D-PATH with EVPN route types other than those specified in this document is outside the scope of this document.

The use of D-PATH with EVPN IP Prefix routes is specified in [I-D.ietf-bess-evpn-ipvpn-interworking]. When D-PATH is used with EVPN route types other than IP Prefix routes, the attribute is characterized as follows:

1. D-PATH is composed of a sequence of Domain segments following the format specified in [I-D.ietf-bess-evpn-ipvpn-interworking] where each Domain is represented as <DOMAIN-ID:ISF_SAFI_TYPE>. In this specification, DOMAIN-ID is an EVPN Domain identifier configured in an EVPN Domain Gateway and ISF_SAFI_TYPE is set to either 70 (EVPN) or 0 (local route). To simplify the explanation, this document represents the domains for EVPN routes as <Domain-ID:TYPE>.
2. D-PATH identifies the sequence of EVPN Domains the route has gone through, with the last <Domain-ID:TYPE> entry (rightmost) identifying the first PE or EVPN Domain Gateway that added the D-PATH attribute.
3. For non-Inter Subnet Forwarding EVPN MAC/IP Advertisement routes or EVPN A-D per EVI routes [I-D.ietf-bess-rfc7432bis], D-PATH SHOULD be added/modified by a EVPN Domain Gateway that re-originate the route between EVPN Domains and MAY be added by a PE or EVPN Domain Gateway that originates the route, as follows:
 - a. An EVPN Domain Gateway that interconnects two EVPN Domains "X" and "Y", and receives a route on a EVPN Domain identified by a Domain-ID "X" SHOULD append a domain <X:EVPN> to the existing (or newly added) D-PATH attribute when re-originate the route to EVPN Domain "Y". The route is re-originated if it is first imported in a MAC-VRF (or VPWS instance), the MAC (or Ethernet Tag) is active, and policy allows the re-export of the route to a BGP neighbor.
 - b. An EVPN Domain Gateway MAY also add the D-PATH attribute for locally learned MACs or MAC/IP pairs. In this case, the domain added would be <A:0>, where "A" is the Domain-ID configured on the Gateway's MAC-VRF that is specific to local routes (MAC/IP learned via local AC), and "0" is the TYPE of the EVPN Domain and indicates that the route is locally originated and not re-originated after receiving it from a BGP-EVPN neighbor. The EVPN Domain-ID for local routes MAY be shared by all the EVPN Domain Gateways of the same redundancy group for local routes, or each EVPN Domain Gateway of the redundancy group can use its own Domain-ID.
 - c. A PE that is connected to a single EVPN Domain (therefore the PE is not an EVPN Domain Gateway) MAY add D-PATH with a domain <B:0>, where "B" is the Domain-ID configured on the PE's MAC-VRF (or VPWS) for locally learned MAC/IPs (or Ethernet Tag IDs for VPWS). "0" is the TYPE that indicates the route is not re-advertised, but originated in the PE.

4. For EVPN Inclusive Multicast Ethernet Tag routes, an EVPN Domain Gateway must not re-originate routes between Domains as specified in [RFC9014]. An EVPN Domain Gateway originates an EVPN Inclusive Multicast Ethernet Tag route per Domain to which it is attached, in order to attract BUM traffic from one Domain to the others. Accordingly, only the procedure described in bullet 3.b above applies to EVPN Domain Gateways. Specifically, an EVPN Domain Gateway MAY attach a D-PATH attribute of the form <A:0> to the EVPN Inclusive Multicast Ethernet Tag routes it originates for its attached EVPN Domains, where "A" is the locally configured Domain-ID and "0" is the TYPE indicating that the route is locally originated and not re-originated across EVPN Domains. When two or more EVPN Domain Gateways belonging to the same redundancy group interconnect EVPN Domains "X" and "Y", and D-PATH is used for EVPN Inclusive Multicast Ethernet Tag routes, it is RECOMMENDED that the D-PATH attribute be added with the same local Domain-ID and applied in only one of the two Domains (either "X" or "Y", but not both). In such cases, all Gateways within the same redundancy group MUST select the same Domain in which the D-PATH attribute is applied.
5. On received EVPN routes, D-PATH is processed and used for loop detection (Section 4.4) as well as to influence the best path selection (Section 4.1, Section 4.2 and Section 4.3).
6. A Gateway PE SHOULD support the removal of the D-PATH attribute on import and on export, based on configuration.

4.1. D-PATH and Best Path Selection for MAC/IP Advertisement routes

When two (or more) MAC/IP Advertisement routes with the same route key (regardless of whether the RDs are identical or different) are received, the best path selection algorithm is used to select and install only one route. The best path selection for MAC/IP Advertisement routes is specified in [I-D.ietf-bess-rfc7432bis], in section 7.13.1, and this document modifies the algorithm by including the D-PATH comparison across EVPN MAC/IP Advertisement routes after tie-breaking rule 5 in [I-D.ietf-bess-rfc7432bis] section 7.13.1, which removes from consideration routes that are not tied for higher degree of preference.

If none of the tie-breaking rules up to (and including) rule 5 produces a single route, the router compares the D-PATH attribute in the remaining candidate routes:

1. The routes with the shortest D-PATH are preferred, hence routes not tied for the shortest D-PATH are removed. Routes without D-PATH are considered zero-length D-PATH.

2. Then routes with the numerically lowest left-most Domain-ID are preferred (only the Domain-ID is compared, and not the ISF_SAFI_TYPE). Hence, routes not tied for the numerically lowest left-most Domain-ID are removed from consideration. When comparing two Domain-IDs, the two six byte values are compared starting with the Global Admin field. The lowest value in the first differing byte between the two six byte values, is considered to belong to the "numerically lowest Domain-ID".

If the steps above do not produce a single route, then the rest of the rules in [I-D.ietf-bess-rfc7432bis] follow.

4.2. D-PATH and Best Path Selection for Ethernet A-D per EVI routes

When two (or more) EVPN A-D per EVI routes with the same route key (regardless of whether the RDs are identical or different) are received for a Virtual Private Wire Service (VPWS), the best path selection algorithm is applied to select a single route. The best path selection for EVPN A-D per EVI routes is specified in [I-D.ietf-bess-rfc7432bis], in section 7.13.2, and this document modifies the algorithm by including the D-PATH comparison across EVPN A-D per EVI routes in the same way Section 4.1 does it for EVPN MAC/IP Advertisement routes. That is, rules 1 and 2 of Section 4.1 are interleaved between rules 5 and 6 of [I-D.ietf-bess-rfc7432bis].

4.3. D-PATH and Best Path Selection for Inclusive Multicast Ethernet Tag routes

When two (or more) EVPN Inclusive Multicast Ethernet Tag routes with the same route key (regardless of whether the RDs are identical or different) are received for a MAC-VRF, the best path selection algorithm is used and a single route is programmed. The selection algorithm follows [I-D.ietf-bess-rfc7432bis] the same D-PATH comparison steps as in Section 4.1 interleaved between rules 5 and 6 of [I-D.ietf-bess-rfc7432bis].

4.4. Loop Detection

An EVPN route received by a PE with a D-PATH attribute that contains one or more of its locally associated Domain-IDs for the MAC-VRF or VPWS instance is considered to be a looped route. A looped route MUST NOT be re-originated to a different domain and SHOULD be flagged as "looped".

EVPN Inclusive Multicast Ethernet Tag looped routes MUST NOT be installed, where "install" in this document means "create forwarding state". An EVPN MAC/IP Advertisement looped route or an A-D per EVI looped route (in EVPN VPWS services) MAY be installed if selected as the best route.

For instance, in the example of Figure 3, assuming PE1 advertises M1's MAC/IP and does not add the D-PATH attribute, the EVPN Domain Gateway GW1 receives two MAC/IP Advertisement routes for M1's MAC/IP:

- * A MAC/IP Advertisement route with next hop PE1 and no D-PATH.
- * A MAC/IP Advertisement route with next hop GW2 and D-PATH={length=1,<6500:1:EVPN>}, assuming that the Domain-ID for EVPN Domain 1 is 6500:1.

In this case, EVPN Domain Gateway GW1 flags the MAC/IP Advertisement route with D-PATH as "looped", and does not install the MAC in the BD, and does not re-originate the route back to EVPN Domain 1 (since the route includes one of GW1's Domain-IDs). In case the MAC/IP Advertisement route with next hop PE1 is withdrawn, GW1 may install the route with next hop GW2 and D-PATH <6500:1:EVPN>; this may help speed up convergence in case of failures.

The procedures described in this section, based on D-PATH, can be used along with the Ethernet Segment Identifier of the received routes as a way to detect looped routes on EVPN domain gateways attached to an Interconnect Ethernet Segment as in [RFC9014]. An EVPN domain gateway MUST NOT re-originate a received EVPN MAC/IP route or EVPN A-D per EVI route with an Ethernet Segment Identifier value that matches the value of a local Ethernet Segment, irrespective of the D-PATH Domain-IDs.

4.5. Error Handling

The error handling for the D-PATH attribute is described in [I-D.ietf-bess-evpn-ipvpn-interworking] and apply to this document. Although this document extends the applicability of D-PATH to EVPN routes that are not used for Inter-Subnet Forwarding (non-ISF), the same error-handling procedures apply.

5. Use-Case Examples

This section illustrates the use of D-PATH in EVPN routes with examples.

Figure 4 and Figure 5 illustrate an integrated interconnect solution for an EVPN BD, as described in section 4.4 and section 4.6 of [RFC9014]. GW1 and GW2 are EVPN Domain Gateways connecting two EVPN Domains identified by D-PATH domain {1:1:EVPN} and {1:2:EVPN}, respectively. Received Ethernet A-D routes, ES routes, and Inclusive Multicast routes from the routers in one EVPN Domain are consumed and processed by GW1 and GW2, but not re-originated to the other EVPN Domain. However, MAC/IP Advertisement routes received by GW1 and GW2 in one EVPN Domain are processed and, if installed, re-originated into the other EVPN Domain.

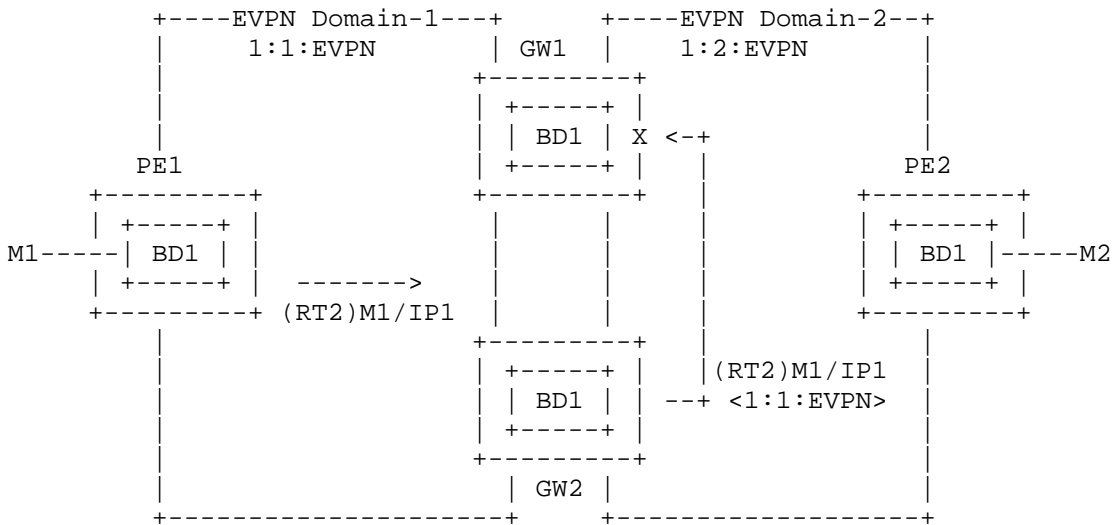


Figure 4: EVPN Interconnect

Consider the example of Figure 4, where PE1 advertises a MAC/IP Advertisement route for M1/IP1. The route is processed and installed by GW1 and GW2 in BD1, and both re-originate the routes into the EVPN Domain-2. By using D-PATH in GW1 and GW2, when the route is received on PE2, PE2 has the visibility of the EVPN Domains through which the route has gone, and can also use the D-PATH for best path selection. In addition, GW1 and GW2 can compare the D-PATH of the incoming routes with their local list of EVPN Domain-IDs, and detect looped routes if any of the local EVPN Domain-IDs matches a domain in the received D-PATH. This procedure prevents the re-origination of the route back into EVPN Domain-1. For example, when GW1 receives the MAC/IP Advertisement route for M1/IP1 with D-PATH <1:1:EVPN>, GW1 identifies the route as looped and it does not re-originate it back to Domain-1. The M1/IP1 route with Next Hop PE1 is installed. If M1/IP1 with Next Hop PE1 is withdrawn, GW1 MAY install the route M1/IP1 with Next Hop GW2, as specified in Section 4.4.

The example of Figure 5 illustrates how GW1 and GW2 can also have local ACs in BD1 and learn local MAC (or MAC/IP) addresses from devices connected to the ACs.

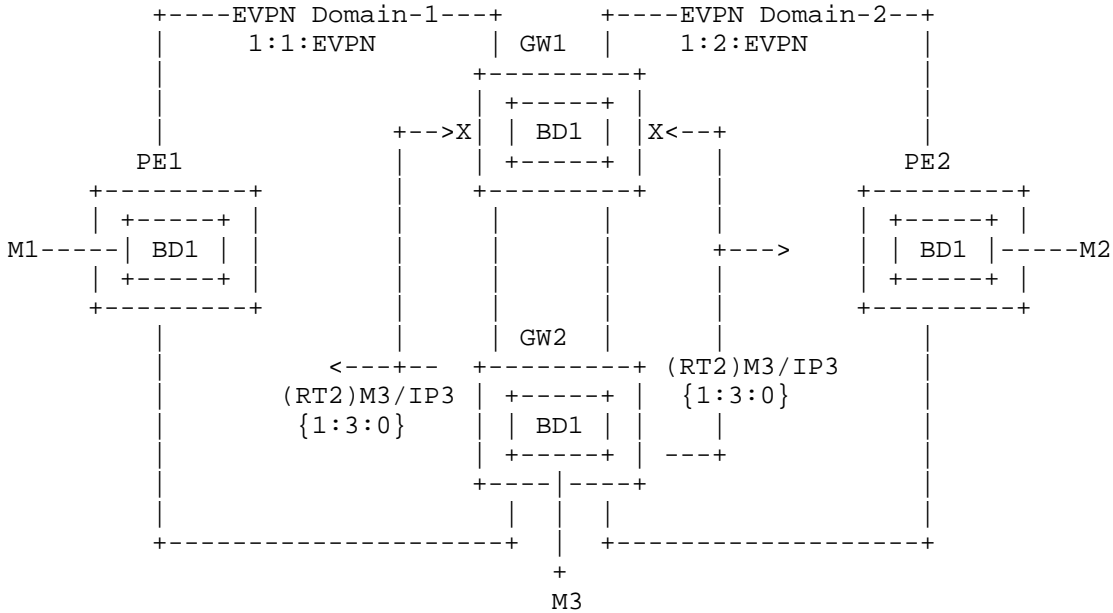


Figure 5: EVPN Interconnect local AC

Assuming GW2 learns M3/IP3 via local AC, GW2 advertises a MAC/IP Advertisement route for M3/IP3 into each of the EVPN Domains that GW2 is connected to. As described in Section 4, GW2 can advertise these two MAC/IP Advertisement routes with a configured EVPN Domain-ID for local MAC/IPs routes that can be shared with GW1. Consider this EVPN Domain-ID is 1:3 and it is configured on both, GW1 and GW2. When GW2 advertises the route into each EVPN Domain, it adds the D-PATH attribute with a domain {1:3:0}. These routes are flagged by GW1 as "looped" since 1:3 is configured as a local EVPN Domain-ID in GW1. In addition, PE1 and PE2 receive the routes with the D-PATH and they have the visibility of the origin of the routes, in this case local EVPN Domain Gateway routes. This information can be used to influence the best path selection in case of multiple routes for M3/IP3 are received on PE1 or PE2 for BD1.

As an alternative solution to configuring the same EVPN Domain-ID for local routes on both EVPN Domain Gateways, GW2 can be configured with EVPN Domain-ID 1:3 for local routes, and GW1 can use a different EVPN Domain-ID, e.g., 1:4. In this case, GW2 advertises the route for M3/IP3 into each EVPN Domain as before, but now GW1 does not flag the

route as "looped" since 1:3 is not on the list of GW1's local EVPN Domain-IDs. GW1 receives the routes from both EVPN Domains, and GW1 selects the route from e.g., EVPN Domain-1. GW1 then installs the route in its BD and re-originates the route into EVPN Domain-2 with D-PATH {1:1:EVPN, 1:3:0}. When PE2 receives two routes for M3/IP3, one from GW2 with D-PATH {1:3:0} and another from GW1 with D-PATH {1:1:EVPN, 1:3:0}, PE2 uses best path selection and choose to send its traffic to GW2. Also GW2 receives the route for M3/IP3 from GW1 and mark it as "looped" since that route conveys its own EVPN Domain-IDs 1:1 and 1:3.

In a nutshell, the use of D-PATH in MAC/IP Advertisement routes helps prevent loops and influences the best path selection so that PEs choose the shortest paths to the destination PEs.

6. Operational Considerations

This document modifies the best-path selection algorithm for EVPN routes that include D-PATH. Consequently, an upgraded PE and a non-upgraded PE may select different best paths when processing multiple EVPN routes carrying D-PATH. Although selecting different MAC/IP Advertisement, A-D per EVI, or Inclusive Multicast Ethernet Tag routes for a given EVPN service does not introduce forwarding loops, it may lead to inconsistent behavior among PEs depending on their upgrade status. Therefore, enabling D-PATH is RECOMMENDED only when all PEs participating in the EVPN service (Broadcast Domain or VPWS) have been upgraded.

In deployments where EVPN Domain Gateways are redundantly interconnecting the same two domains, enabling D-PATH is RECOMMENDED only when all such Gateways for a given EVPN service have been upgraded. If some Gateways remain non-upgraded, D-PATH may not prevent forwarding loops or packet duplication. In such cases, alternative loop-prevention mechanisms (without D-PATH) are assumed to be in use. Those mechanisms are outside the scope of this document.

7. Security Considerations

Most of the security considerations included in [I-D.ietf-bess-evpn-ipvpn-interworking] related to D-PATH apply to this document.

In particular, a correct use of the D-PATH will prevent control plane and data plane loops in the network, however an incorrect configuration of the DOMAIN-IDs or an inconsistent support of D-PATH on the Gateway PEs may lead to the detection of false route loops, dropping packets or may result in inconsistent and sub-optimal

forwarding. As an additional security mechanism, a PE following this specification that receives an EVPN route from a non-upgraded PE should discard the route via policy if the route contains the D-PATH attribute.

8. IANA Considerations

None.

9. Acknowledgments

The authors would like to thank Jeff Haas and Alex Nichol for their detailed review and valuable feedback on this document.

10. Contributors

In addition to the authors included on the front page, the following people contributed significantly:

Vinod Prabhu, Nokia

Bin Wen, Comcast

11. References

11.1. 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>>.
- [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>>.
- [I-D.ietf-bess-evpn-ipvpn-interworking] Rabadan, J., Sajassi, A., Rosen, E. C., Drake, J., Lin, W., Uttaro, J., and A. Simpson, "Interconnecting EVPN and IPVPN Domains", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-ipvpn-interworking-16, 28 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-ipvpn-interworking-16>>.

[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>>.

[RFC9014] Rabadan, J., Ed., Sathappan, S., Henderickx, W., Sajassi,
A., and J. Drake, "Interconnect Solution for Ethernet VPN
(EVPN) Overlay Networks", RFC 9014, DOI 10.17487/RFC9014,
May 2021, <<https://www.rfc-editor.org/info/rfc9014>>.

[RFC8214] Boutros, S., Sajassi, A., Salam, S., Drake, J., and J.
Rabadan, "Virtual Private Wire Service Support in Ethernet
VPN", RFC 8214, DOI 10.17487/RFC8214, August 2017,
<<https://www.rfc-editor.org/info/rfc8214>>.

11.2. Informative References

[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>>.

[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>>.

Authors' Addresses

J. Rabadan (editor)
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America
Email: jorge.rabadan@nokia.com

S. Sathappan
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America
Email: senthil.sathappan@nokia.com

M. Gautam
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America
Email: mallika.gautam@nokia.com

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

W. Lin
Juniper
United States of America
Email: wlin@juniper.net