

Workgroup: BESS
Internet-Draft:
draft-ietf-bess-extended-evpn-optimized-ir-06
Published: 23 December 2023
Intended Status: Standards Track
Expires: 25 June 2024
Authors: W. Lin, Ed. S. Sivaraj
 Juniper Networks, Inc. Juniper Networks, Inc.
 V. Garg J. Rabadan
 Juniper Networks, Inc. Nokia

Extended Procedures for EVPN Optimized Ingress Replication

Abstract

In the Virtualization Overlay (NVO) network with Ethernet VPN (EVPN), optimized ingress replication uses Assisted-Replication (AR) to achieve more efficient delivery of Broadcast and Multicast (BM) traffic. An AR-LEAF, which is a Network Virtualization Edge (NVE) device, forwards received BM traffic from its tenant system to an AR-REPLICATOR. The AR-REPLICATOR then replicates it to the remaining AR-LEAFs in the network. However, when replicating the packet on behalf of its multihomed AR-LEAF, an AR-REPLICATOR may face challenges in retaining the source IP address or including the expected Ethernet Segment Identifier (ESI) label that is required for EVPN split-horizon filtering. This document extends the optimized ingress replication procedures to address such limitations. The extended procedures specified in this document allow the support of EVPN multihoming on the AR-LEAFs as well as optimized ingress replication for the rest of the EVPN NVO network.

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.

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 25 June 2024.

Copyright Notice

Copyright (c) 2023 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. [Terminology](#)
2. [Introduction](#)
 - 2.1. [Background](#)
 - 2.1.1. [EVPN Multihoming and split-horizon Filtering Rule](#)
 - 2.2. [Optimized-IR and the Need to Maintain the Original Source IP address or Include the ESI Label](#)
3. [Solution](#)
 - 3.1. [AR-REPLICATOR Announcing Multihoming Assistant Capability for Optimized-IR](#)
 - 3.2. [Multihomed AR-LEAF and Extended-MH AR-REPLICATOR](#)
 - 3.3. [The Benefit of the Extended Optimized-IR Procedure](#)
 - 3.4. [Support for Mixed AR-REPLICATORS](#)
4. [Extended Optimized-IR Procedure for Supporting Extended-MH AR-REPLICATOR](#)
 - 4.1. [AR-LEAF Procedure](#)
 - 4.1.1. [Control Plane Procedure for AR-LEAF](#)
 - 4.1.2. [Forwarding Procedure for AR-LEAF](#)
 - 4.2. [AR-REPLICATOR Procedure](#)
 - 4.2.1. [Control Plane Procedure for AR-REPLICATOR](#)
 - 4.2.2. [Forwarding Procedure for AR-REPLICATOR](#)
 - 4.3. [RNVE Procedure](#)
5. [AR-LEAF's Peer multihomed NVE in the Extended Optimized-IR Procedure](#)
6. [Multicast Flags Extended Community](#)
7. [IANA Considerations](#)

- [8. Security Considerations](#)
- [9. Acknowledgements](#)
- [10. Normative References](#)
- [Authors' Addresses](#)

1. Terminology

A-D: Auto-Discovery

AR-IP Tunnel: An overlay tunnel with a destination IP address of AR-IP that an AR-REPLICATOR advertises in its REPLICATE-AR route.

BD: Broadcast Domain, as defined in [RFC7432]

EVI: EVPN Instance

RNVE: Regular Network Virtualized Edge router that performs the procedure specified in [RFC8365]

This document makes use of the terminology specified in [EVPN-AR]. It also uses the terminology specified in [RFC7432] and [RFC8365].

2. Introduction

2.1. Background

2.1.1. EVPN Multihoming and split-horizon Filtering Rule

This section gives a brief overview of the existing split-horizon filtering rules used for EVPN multihoming.

[RFC7432] defines the split-horizon filtering rule based on ESI label for EVPN multihoming with MPLS encapsulation, and this filtering rule also applies for EVPN with IP-based encapsulation for MPLS, such as MPLS over GRE or MPLS over UDP. [RFC8365] defines the split-horizon filtering rule based on "Local-Bias" for EVPN multihoming with VXLAN encapsulation.

When EVPN is used in an NVO network, a Tenant System (TS) may connect to a set of Network Virtualization Edge (NVE) devices through a multihomed Ethernet segment (ES). The split-horizon filtering rule for EVPN all-active multihoming ensures that a Broadcast, Unknown unicast or Multicast (BUM) packet received from an ES that is part of a multihomed ES is not looped back to the multihomed TS through an egress NVE connected to the same multihomed ES. When using EVPN with VXLAN encapsulation, the split-horizon filtering rule is applied by the egress NVE based on the source IP address of the BUM packet received from an overlay tunnel. The egress PE identifies the ingress NVE through the source IP address. The egress NVE does not forward the BUM packet received from an

overlay tunnel to the multihomed Ethernet segment that it has in common with the ingress NVE.

For EVPN with MPLS over IP tunnel, the split-horizon filtering rule is based on the ESI label. For ingress replication, an ESI label is downstream assigned per multihomed ES. The ingress NVE MUST include the ESI label, assigned by the egress PE, when it forwards a BUM packet to the egress NVE if the BUM traffic is from the AC that is part of the multihomed ES associated with that ESI label. The egress NVE does not forward the BUM packet it received from an overlay tunnel to the multihomed ES if the ESI label is allocated by the egress NVE for that multihomed ES.

2.2. Optimized-IR and the Need to Maintain the Original Source IP address or Include the ESI Label

[[EVPN-AR](#)] specifies an optimized ingress replication solution for the delivery of BM traffic within a bridge domain. It defines the control plane and forwarding plane procedures for AR-REPLICATOR, AR-LEAF and RNVE. To support EVPN AR-LEAF multihoming, [[EVPN-AR](#)] recommends the implementation of split-horizon filtering based on the "Local-Bias" procedures for EVPN NVO network using either 24-bit VNI or MPLS label.

To support EVPN all-active multihoming based on "Local-Bias" procedures, when an AR-REPLICATOR performs assisted replication on behalf of a multihomed AR-LEAF, the AR-REPLICATOR MUST use the source IP address of the ingress AR-LEAF for packet received on the AR-IP tunnel. This ensures that other remote NVEs, when receiving a packet from its AR-REPLICATOR, can perform the regular split-horizon filtering based on the source IP address.

To support EVPN all-active multihoming with MPLSoGRE or MPLSoUDP, sometimes it is desirable to continue using the existing split-horizon filtering rule based on [RFC7432] procedures. In this case, when performing assisted replication on behalf of a multihomed AR-LEAF, an AR-REPLICATOR MUST include the ESI label advertised by a remote NVE for that multihomed ES.

However, due to either implementation complexity or hardware limitations, an AR-REPLICATOR may be unable to retain the source IP address or include the ESI label when replicating the packet to the remote NVEs on behalf of a multihomed AR-LEAF. In such circumstances, a remote NVE, upon receiving the packet, is unable to utilize the existing split-horizon filtering rules to prevent the looping of BM traffic required for all-active multihoming.

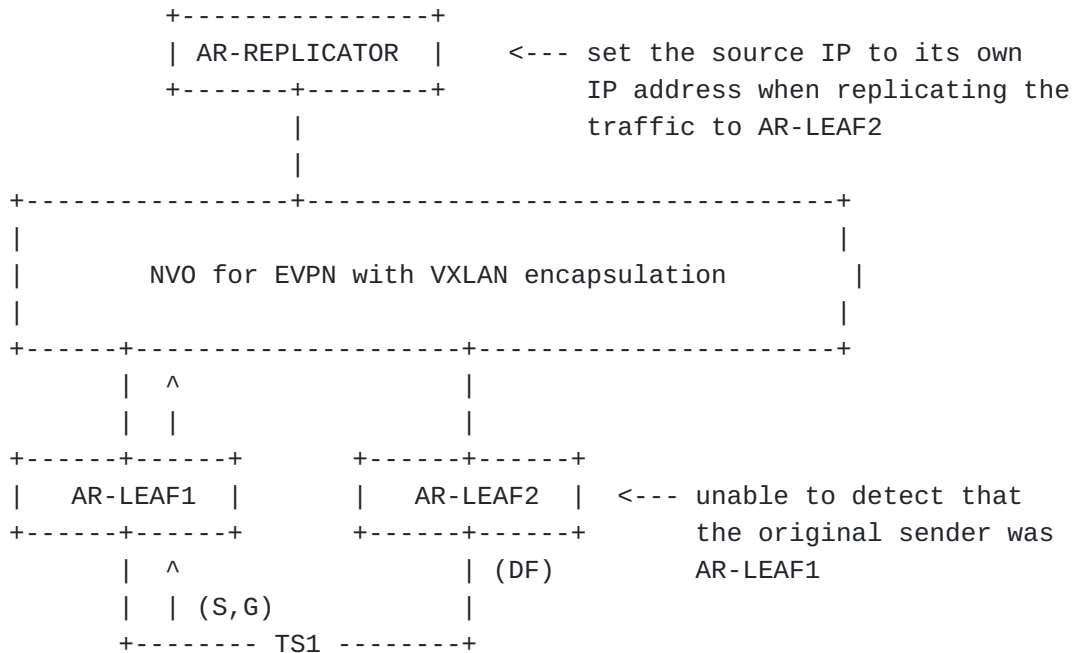


Figure 1: AR-Replicator and the VXLAN Source IP Address

For instance, let's consider a scenario with VXLAN encapsulation, as illustrated in Figure 1, where TS1 is multihomed to both AR-LEAF1 and AR-LEAF2 through a multihomed ES. When AR-LEAF1 receives an IP multicast packet from TS1, it forwards the packet to its AR-REPLICATOR, setting the source IP address to AR-LEAF1's IR-IP and the destination IP address to the AR-IP of the AR-REPLICATOR.

As the AR-REPLICATOR is unable to retain the source IP address from the packet it received over the AR-IP tunnel, it replaces it with one of its own IP addresses when replicating the packet to other NVEs. Upon receiving the packet from AR-REPLICATOR, AR-LEAF2 checks the source IP address, but it cannot identify AR-LEAF1 as the original sender.

In cases where AR-LEAF2 functions as the Designated Forwarder (DF) for the multihomed ES linked to TS1, it proceeds to forward the packet to TS1. This results in the same IP multicast packet being looped back to TS1.

The issue can also occur in EVPN with MPLS over an IP network when an AR-REPLICATOR is unable to include the ESI label to the remote NVE for the multihomed ES when the split-horizon filtering rule based on [RFC7432] is used.

3. Solution

This document extends the procedures defined in [EVPN-AR] to support EVPN multihoming on AR-LEAFs. It addresses the limitations or challenges where an NVE serves as an AR-REPLICATOR but cannot retain the source IP address or include an ESI label for its AR-LEAF due to hardware constraints or implementation complexity.

This document presents a solution for EVPN over IP-based networks, which uses an NVO tunnel with either a 24-bit VNI or an MPLS label. In order to prevent BM traffic looping while achieving optimized ingress replication, this solution relies on the use of either [RFC7432] or "Local-Bias" split-horizon filtering rules. We refer to the procedures specified in this document as the extended Optimized-IR procedures. These extended Optimized-IR procedures are also compatible with RNVE.

3.1. AR-REPLICATOR Announcing Multihoming Assistant Capability for Optimized-IR

An AR-REPLICATOR announces its AR-REPLICATOR role through the control plane. A REPLICATOR-AR route, as it is specified in the [EVPN-AR], is an Inclusive Multicast Ethernet Tag (IMET) route that an AR-REPLICATOR originates for its AR-IP and corresponding AR-replication tunnel.

If an AR-REPLICATOR cannot or chose not to retain the source IP address or include the expected ESI label for its multihomed AR-LEAFs, it MUST inform other NVEs in the control plane through the use of EVPN Multicast Flags Extended Community as follow: a) the AR-REPLICATOR MUST set the "Extended-MH-AR" flag, as it is specified in Section 6, in the multicast flags extended community, and b) it MUST attach this community to the REPLICATOR-AR route it originates. We call such an AR-REPLICATOR an Extended-MH AR-REPLICATOR.

An Extended-MH AR-REPLICATOR supports extended Optimized-IR procedures defined in this document for its multihomed AR-LEAFs. An Extended-MH AR-REPLICATOR keeps track of its AR-LEAF's multihomed peer. An Extended-MH AR-REPLICATOR can perform assisted replication for an AF-LEAF to other NVEs that are not attached to the same multihomed ES as the AR-LEAF. An Extended-MH AR-REPLICATOR does not perform assisted replication for its AR-LEAF to other NVEs that have a multihomed ES in common with the AR-LEAF. The changes in the control plane and forwarding plan procedures for an Extended-MH AR-REPLICATOR are further explained in detail in Section 5.2.

If an AR-REPLICATOR originates a REPLICATOR-AR route without a multicast flags extended community or with the Extended-MH-AR flag unset, it is considered to be multihoming assistant capable. An MH-

capable-assistant AR-REPLICATOR can perform assisted replication for its single-homed AR-LEAF as well as multihomed AR-LEAF.

3.2. Multihomed AR-LEAF and Extended-MH AR-REPLICATOR

An AR-LEAF follows the control plane and forwarding plane procedures specified in [[EVPN-AR](#)]. In addition, if a multihomed AR-LEAF detects that one of its AR-REPLICATORS is an Extended-MH AR-REPLICATOR based on the processing of its REPLICATOR-AR route, the multihomed AR-LEAF follows the extended Optimized-IR procedures specified in this document. With the extended Optimized-IR procedures, within the same BD, the multihomed AR-LEAF will use the regular ingress replication procedure to deliver a copy of a BM packet received from its local AC to each of the remote NVEs that has a multihomed ES in common with it. In this way, the egress NVE can use the regular split-horizon filtering rule defined in [RFC7432] or [RFC8365] to prevent the BM traffic to be looped through the egress NVE to the source of origin. The extended procedures required for an AR-LEAF is further specified in detail in section 5.

Please note that for an AR-LEAF, the additional forwarding procedures specified above apply to BM packets that originate from any of its ACs within the same BD. These ACs can either be a single-homed ES or be part of a multihomed ES. It may also apply to Unknown unicast traffic. This is to ease the burden of an Extended-MH AR-REPLICATOR as it may be unable to detect whether a packet received on its AR-IP tunnel was originally received from a single-homed or multihomed ES.

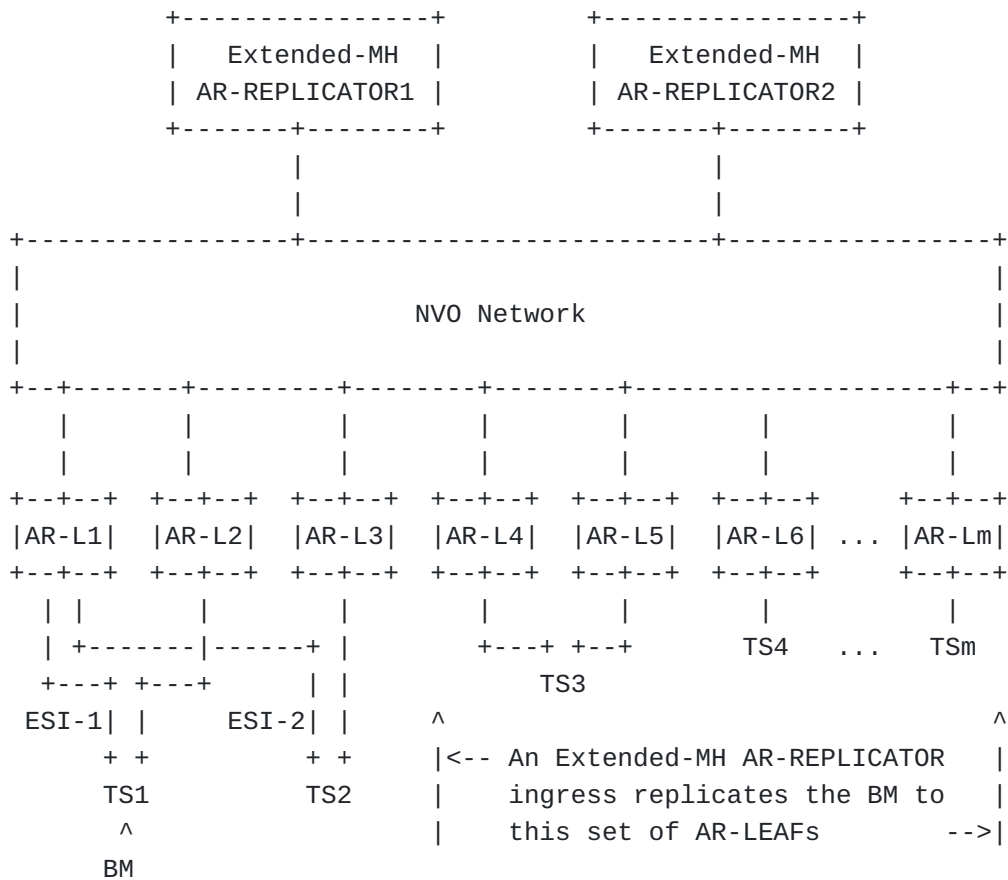


Figure 2: Extended Optimized-IR Model

Consider an EVPN NVO network in Figure 2, the tenant domain consists of a set of m AR-LEAFs in BD X. For brevity, we use "AR-L" to represent "AR-LEAF" in Figure 2. TS1 is multihomed to AR-LEAF1 and AR-LEAF2 in BD X via a multihomed ES, ES1. Similarly, TS2 is multihomed to AR-LEAF1 and AR-LEAF3 in BD X through another multihomed ES, ES2. Additionally, there are two Extended-MH AR-REPLICATORS in the same tenant domain: AR-REPLICATOR1 and AR-REPLICATOR2.

AR-LEAF1 will detect that its AR-REPLICATORS are Extended-MH AR-REPLICATORS through the Extended-MH-AR flag within the EVPN multicast flags extended community. This extended community is signaled by the AR-REPLICATORS through their REPLICATOR-AR routes. Following the normal EVPN procedure, AR-LEAF1 will also detect that both AR-LEAF2 and AR-LEAF3 have a multihomed ES in common with it. AR-LEAF1 will use regular ingress replication to send the BM traffic it receives from its access to both AR-LEAF2 and AR-LEAF3. AR-LEAF1

will rely on one of its AR-REPLICATORS to send the BM traffic to AR-LEAF4, AR-LEAF5, AR-LEAF6, ..., and AR-LEAFm.

3.3. The Benefit of the Extended Optimized-IR Procedure

The extended Optimized-IR solution specified in this document greatly reduces the implementation complexity of an AR-REPLICATOR or helps to overcome the limitation of an AR-REPLICATOR. It allows EVPN multihoming on AR-LEAFs while adhering to existing multihoming procedures and split-horizon filtering rules. Consequently, it frees AR-REPLICATORS from the requirements of multihoming assisted replication. For EVPN with MPLS over IP-based encapsulation, an NVE can continue to use the split-horizon filtering rule based on the ESI label. Furthermore, it still allows the support of efficient Optimized-IR for the rest of an EVPN NVO network.

For example, in a typical NVO network, a TS is most likely multihomed to two or a small set of NVEs for redundancy. In an NVO network comprises many NVEs, the AR-REPLICATOR is still responsible for replicating the BM packet to the most of NVEs functioning as AR-LEAFs. Therefore, it gets the advantage of optimized ingress replication for the majority of its NVO network.

3.4. Support for Mixed AR-REPLICATORS

When there are mixed MH-capable-assistant AR-REPLICATORS and Extended-MH AR-REPLICATORS in the same tenant domain, all AR capable NVEs MUST follow the extended Optimized-IR procedures as long as one of the AR-REPLICATORS is an Extended-MH AR-REPLICATOR.

In situations where there are different types of AR-REPLICATORS, all MH-capable-assistant AR-REPLICATORS SHALL be provisioned administratively to behave as Extended-MH AR-REPLICATORS. In such cases, each AR-REPLICATOR originates its REPLICATOR-AR route with the Extended-MH-AR flag set in the

The procedure for using mixed AR-REPLICATORS is beyond the scope of this document.

4. Extended Optimized-IR Procedure for Supporting Extended-MH AR-REPLICATOR

4.1. AR-LEAF Procedure

This section covers the extended Optimized-IR procedures required for an AR-LEAF when at least one of the AR-REPLICATORS is an Extended-MH AR-REPLICATOR. It is assumed that an AR-LEAF follows the procedures defined in [[EVPN-AR](#)] unless otherwise specified.

4.1.1. Control Plane Procedure for AR-LEAF

An AR-LEAF detects whether an AR-REPLICATOR is capable of performing multihoming assisted replication through the Extended-MH-AR flag in the multicast flags extended community carried in the REPLICATOR-AR route. An AR-REPLICATOR originating a REPLICATOR-AR route without a multicast flags extended community or with the Extended-MH-AR flag unset is considered to be multihoming assistant capable.

If an AR-LEAF does not have a locally attached segment that is part of a multihomed ES, it does not need to follow any additional extended Optimized-IR procedure, and we can proceed directly to Section 4.2.

If selective assistant-replication is used for the EVI, selective AR-LEAFs that share the same multihomed ES MUST select the same primary AR-REPLICATOR and the same backup AR-REPLICATOR, if there is one. This can be achieved through either manual configuration on each multihomed selective AR-LEAF or by other methods that are beyond the scope of this document. Each selective AR-LEAF follows the procedures defined in the [\[EVPN-AR\]](#) to send its corresponding leaf-AD routes to its AR-REPLICATOR.

An AR-LEAF follows the normal procedures defined in [\[RFC7432\]](#) when it originates a type-4 ES route and type-1 Ethernet A-D routes for its locally attached segment that is a part of a multihomed ES.

In addition, an AR-LEAF builds a peer-multihomed-flood-list for each BD it attaches to. As per the standard EVPN procedures defined in [\[RFC7432\]](#), an AR-LEAF discovers the ESI of each multihomed ES that every remote NVE connects to. For a given BD, an AR-LEAF constructs a peer-multihomed-flood-list that consists of its peer multihomed NVEs in that BD that have at least one multihomed ES in common with it. An AR-LEAF may consider a common multihomed ES that it shares with a remote NVE in a BD specific scope or an EVI scope. Please refer to Section 5 for details.

4.1.2. Forwarding Procedure for AR-LEAF

Suppose a multihomed AR-LEAF detects through a control plane procedure that one or more of its AR-REPLICATORS are Extended-MH AR-REPLICATORS. In that case, in addition to following the forwarding procedures defined in [\[EVPN-AR\]](#), it will use regular ingress replication to send the BM packet received from one of its ACs to each NVE in that BD's peer-multihomed-flood-list.

If there are no more AR-REPLICATORS within the tenant domain, the AR-LEAF will revert back to its regular IR behavior, as defined in [\[RFC7432\]](#). An AR-LEAF will follow the regular EVPN procedures when it receives a packet from an overlay tunnel.

4.2. AR-REPLICATOR Procedure

This section describes the additional procedures for an AR-REPLICATOR when there is at least one AR-REPLICATOR in the same tenant domain that is an Extended-MH AR-REPLICATOR.

It is also assumed that an AR-REPLICATOR follows the procedures defined in [[EVPN-AR](#)] unless specified otherwise.

4.2.1. Control Plane Procedure for AR-REPLICATOR

An NVE that performs an AR-REPLICATOR role follows the control plane procedures for AR-REPLICATOR defined in [[EVPN-AR](#)].

In addition, if an AR-REPLICATOR is an Extended-MH AR-REPLICATOR or if it is administratively provisioned to behave as an Extended-MH AR-REPLICATOR, it SHALL attach a multicast flags extended community to its REPLICATOR-AR route with the Extended-MH-AR flag set.

An AR-REPLICATOR also discovers whether another AR-REPLICATOR is an Extended-MH AR-REPLICATOR based on the multicast flags extended community. If at least one AR-REPLICATOR is an Extended-MH AR replicator, then the rest of AR-REPLICATORS SHALL fall back to support the extended procedures specified in this document.

When there are mixed AR-REPLICATORS, this document recommends that all MH-capable-assistant AR-REPLICATORS SHOULD fall back to behave as Extended-MH AR-REPLICATORS through administrative provisioning.

An Extended-MH AR-REPLICATOR builds a multihomed list for each BD that its AR-LEAF attaches to. We refer to such a multihomed list as an AR-LEAF's multihomed-list. Per normal EVPN procedures defined in [[RFC7432](#)], an AR-REPLICATOR imports the Ethernet A-D per EVI route, the alias route, originated by each remote NVE in the same tenant domain. For a given BD that an AR-LEAF belongs to, an AR-LEAF's multihomed-list consists of all the NVEs in that BD that have at least one multihomed ES in common with the said AR-LEAF. Please also refer to Section 5 for the common multihomed ES an AR-LEAF shares with its remote NVE.

Consider the EVPN NVO network described in Figure 2. Both AR-LEAF1 and AR-LEAF2 originate their Ethernet A-D per EVI routes for ES1. Both AR-LEAF1 and AR-LEAF3 originate their Ethernet A-D per EVI routes for ES2. As per normal EVPN procedures, each AR-REPLICATOR imports and processes Ethernet A-D per EVI routes. Each AR-REPLICATOR builds an AR-LEAF1's multihomed-list for BD X that consists of AR-LEAF2 and AR-LEAF3. Each AR-REPLICATOR also builds AR-LEAF's multihomed-lists for other AR-LEAFs.

4.2.2. Forwarding Procedure for AR-REPLICATOR

When an AR-REPLICATOR determines that it is an Extended-MH AR-REPLICATOR or determines that it SHALL fall back to become an Extended-MH AR-REPLICATOR, it MUST follow the forwarding procedures described in this section.

When an AR-REPLICATOR replicates a packet from an AR-IP tunnel to other overlay tunnels on behalf of an ingress AR-LEAF, it MUST skip any NVE that is in the multihomed-list of that ingress AR-LEAF built for the corresponding BD.

When replicating the traffic to other AR-REPLICATORS or other AR-LEAFs over an overlay tunnel, an AR-REPLICATOR does not set the source IP address to its ingress AR-LEAF's IR-IP. It is assumed under the scope of this document that no AR-LEAF shares any common multihoming ES with any AR-REPLICATOR.

When replicating the traffic to other RNVEs, an AR-REPLICATOR MUST set the source IP address to its own IR-IP. This is because an RNVE does not recognize the AR-IP.

4.3. RNVE Procedure

There is no change to the RNVE control and forwarding procedures. RNVE follows the regular ingress replication procedure defined in [RFC7432].

5. AR-LEAF's Peer multihomed NVE in the Extended Optimized-IR Procedure

For the extended Optimized-IR procedures specified in this document, a multihomed AR-LEAF MAY keep track of the common multihomed ES it shares with other remote NVEs on a per BD specific scope or on a per EVI scope. Correspondingly, an Extended-MH AR-REPLICATOR MUST also use the same scheme to keep track of the common multihomed ES that its AR-LEAF shares with other remote NVEs. All multihomed AR-LEAFs and all AR-REPLICATORS within the same EVI MUST use the same scheme to keep track of the common multihomed ES that an AR-LEAF shares with other remote NVEs. This consistency can be enforced through a manual configuration.

A multihomed AR-LEAF maintains a peer-multihomed-flood-list for each BD it attaches to. If the common multihomed ES is tracked on a per EVI scope, the peer-multihomed-flood-list of an AR-LEAF for a particular BD X will include all the NVEs in BD X that have at least one common multihomed ES with it. This is regardless of whether each common multihomed ES has BD X or not. If the common multihomed ES is tracked on a per BD specific scope, for a given BD X, each common multihomed ES MUST contain BD X.

RFC 7432 allows the Ethernet A-D route to be advertised at different granularities. If the Ethernet A-D per EVI route is advertised at the granularity of per ES per EVI, the common multihomed ES shared among NVEs SHALL be tracked on a per EVI scope.

6. Multicast Flags Extended Community

The EVPN multicast flags extended community is defined in [RFC9251]. This transitive extended community has a bit vector for its Flags field. An AR Replicator utilizes one bit for the Extended-MH-AR flag, which is designated E in the Flags bit vector below.

```

      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=0x09  |   Flags (2 Octets)   |E|M|I|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Reserved=0         |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The Extended-MH-AR flag is used by the AR-REPLICATOR. By setting this flag, the AR-REPLICATOR signals to other NVEs that it is an Extended-MH AR Replicator and supports the extended Optimized-IR procedures specified in this document.

7. IANA Considerations

IANA has opened the Flags registry for EVPN multicast Extended Community. IANA has allocated bit 13 in the Flags registry field for the Extended-MH-AR flag specified in this document.

Bit Value	Name	Reference
13	Extended-MH-AR	This document

8. Security Considerations

The security consideration in [RFC7432], [RFC8365] and [EVPN-AR] apply to this document.

9. Acknowledgements

The authors would like to thank Eric Rosen and Jeffrey Zhang for their valuable comments and feedbacks. The authors would also like to thank Aldrin Isaac for his useful discussion, insight on this subject. Special thanks to Nicolai Leymann and Thomas Fossati for their thorough reviews and valuable inputs that greatly enhanced the document.

10. Normative References

- [EVPN-AR] Rabadan, J., Ed., "Optimized Ingress Replication solution for Ethernet VPN (EVPN)", internet-draft ietf-bess-evpn-optimized-ir-12.txt, July 2022, <<https://www.ietf.org/archive/id/draft-ietf-bess-evpn-optimized-ir-12.txt>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [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>>.
- [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>>.

Authors' Addresses

Wen Lin (editor)
Juniper Networks, Inc.

Email: wlin@juniper.net

Selvakumar Sivaraj
Juniper Networks, Inc.

Email: ssivaraj@juniper.net

Vishal Garg
Juniper Networks, Inc.

Email: vishalg@juniper.net

Jorge Rabadan
Nokia

Email: jorge.rabadan@nokia.com