

Workgroup: BESS Workgroup

Internet-Draft:

draft-ietf-bess-evpn-mh-split-horizon-07

Updates: [8365](#), [7432](#) (if approved)

Published: 4 December 2023

Intended Status: Standards Track

Expires: 6 June 2024

Authors: J. Rabadan, Ed. K. Nagaraj W. Lin A. Sajassi

 Nokia Nokia Juniper Cisco

EVPN Multi-Homing Extensions for Split Horizon Filtering

Abstract

Ethernet Virtual Private Network (EVPN) is commonly used along with Network Virtualization Overlay (NVO) tunnels, as well as MPLS and Segment Routing tunnels. The EVPN multi-homing procedures may be different depending on the tunnel type used in the EVPN Broadcast Domain. In particular, there are two multi-homing Split Horizon procedures to avoid looped frames on the multi-homed CE: ESI Label based and Local Bias. ESI Label based Split Horizon is used for MPLSoX tunnels, E.g., MPLSoUDP, whereas Local Bias is used for other tunnels, E.g., VXLAN tunnels. The existing specifications do not allow the operator to decide which Split Horizon procedure to use for tunnel encapsulations that could support both. Examples of tunnels that may support both procedures are MPLSoGRE, MPLSoUDP, GENEVE or SRv6. This document updates the EVPN Multihoming procedures in RFC8365 so that an operator can decide the Split Horizon procedure for a given tunnel depending on their own requirements.

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 6 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. Introduction](#)
 - [1.1. Split Horizon Filtering and Tunnel Encapsulations](#)
 - [1.2. Conventions and Terminology](#)
- [2. BGP EVPN Extensions](#)
 - [2.1. The Split Horizon Type \(SHT\)](#)
 - [2.2. Use of the Split Horizon Type In A-D Per ES Routes](#)
 - [2.3. ESI Label Value In A-D Per ES Routes](#)
 - [2.4. Backwards Compatibility With RFC8365 NVEs](#)
- [3. Procedures for NVEs Supporting Multiple Encapsulations](#)
- [4. Security Considerations](#)
- [5. IANA Considerations](#)
- [6. Acknowledgments](#)
- [7. Contributors](#)
- [8. References](#)
 - [8.1. Normative References](#)
 - [8.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Ethernet Virtual Private Network (EVPN) is commonly used with the following tunnel encapsulations:

*Network Virtualization Overlay (NVO) tunnels as specified in [[RFC8365](#)]

*MPLS and Segment Routing with MPLS data plane (SR-MPLS), as specified in [[RFC7432](#)]

*Segment Routing with IPv6 data plane (SRv6), as specified in [[RFC9252](#)].

The EVPN multihoming procedures may be different depending on the tunnel type used in the EVPN Broadcast Domain. In particular, there are two multihoming Split Horizon procedures to avoid looped frames on the multihomed CE: ESI Label based and Local Bias. ESI Label based Split Horizon is used for MPLS or MPLSoX tunnels, E.g., MPLSoUDP [[RFC7510](#)], and its procedures described in [[RFC7432](#)]. Local Bias is used by IP tunnels, E.g., VXLAN tunnels, and it is described in [[RFC8365](#)].

1.1. Split Horizon Filtering and Tunnel Encapsulations

EVPN supports two Split Horizon Filtering mechanisms:

*ESI Label based Split Horizon filtering [[RFC7432](#)]

When EVPN is used for MPLS transport tunnels, an MPLS label enables the Split Horizon filtering capability to support All-Active multihoming. The ingress Network Virtualization Edge (NVE) device adds a label corresponding to the source ES (an ESI label) when encapsulating the packet. The egress NVE checks the ESI label when attempting to forward a multi-destination frame out of a local ES interface, and if the label corresponds to the same site identifier (ESI) associated with that ES interface, the packet is not forwarded. This prevents the occurrence of forwarding loops for BUM traffic.

The ESI Label Split Horizon filtering SHOULD also be used with Single-Active multihoming to avoid transient loops for in-flight packets when the egress NVE takes over as Designated Forwarder for an ES.

*Local Bias [[RFC8365](#)]

Since IP tunnels (such as VXLAN or NVGRE) do not support the ESI label (or any MPLS label at all), a different Split Horizon filtering procedure must be used for All-Active multihoming. This mechanism is called Local Bias and relies on the tunnel source IP address to decide whether to forward BUM traffic to a local ES interface at the egress NVE.

In a nutshell, every NVE tracks the IP address(es) associated with the other NVE(s) with which it has shared multihomed ESs. When the egress NVE receives a BUM frame encapsulated in a IP tunnel, it examines the source IP address in the tunnel header (which identifies the ingress NVE) and filters out the frame on all local interfaces connected to ESes that are shared with the ingress NVE.

Due to this behavior at the egress NVE, the ingress NVE's behavior is also changed to perform replication locally to all

directly attached ESeS (regardless of the Designated Forwarder election state) for all BUM ingress from the access ACs. Because of this "local" replication at the ingress NVE, this approach is referred to as Local Bias.

Local Bias cannot be used for Single-Active multihoming, since the ingress NVE brings operationally down the Attachment Circuits (ACs) for which it is non-Designated Forwarder (hence local replication to non-Designated Forwarder ACs cannot be done). This means transient in-flight BUM packets may be looped back to the originating site by new elected Designated Forwarder egress NVEs.

[[RFC8365](#)] states that Local Bias is used only for IP tunnels, and ESI Label based Split Horizon for IP-based MPLS tunnels. However, IP-based MPLS tunnels, such as MPLSoGRE or MPLSoUDP, are also IP tunnels and can potentially support both procedures, since they can carry ESI Labels and they also use a tunnel IP header where the source IP address identifies the ingress NVE.

Similarly, some IP tunnels that carry an identifier of the source ES in the tunnel header, may potentially follow either procedure too. Some examples are GENEVE or SRv6:

*In a GENEVE tunnel, the source IP address identifies the ingress NVE therefore local bias is possible. Also, [[I-D.ietf-bess-evpn-geneve](#)] defines an Ethernet option TLV (Type Length Value) to encode an ESI label value.

*In an SRv6 tunnel, the source IP address also identifies the ingress NVE, however, by default, and as described in [[RFC9252](#)] the ingress PE will add information in the SRv6 packet so that the egress PE can identify the source ES of the BUM packet. That information is the ESI filtering argument (Arg.FE2) of the service Segment Identifier (SID) received on an A-D per ES route from the egress PE.

[Table 1](#) shows different tunnel encapsulations and their supported and default Split Horizon method. In the case of GENEVE, the default Split Horizon Type (SHT) depends on whether the Ethernet Option with Source ID TLV is negotiated. In the case of SRv6, the default SHT is listed as ESI label filtering in the Table, since the behavior is equivalent to that of ESI Label filtering. In this document, ESI Label filtering refers to the Split Horizon filtering based on the existence of a source ES identifier in the tunnel header.

This document classifies the tunnel encapsulations used by EVPN into:

1. IP-based MPLS tunnels

2. (SR-)MPLS tunnels
3. IP tunnels
4. SRv6 tunnels

Any other tunnel encapsulation (different from the encapsulations in [Table 1](#)) that can be classified into any of the four encapsulation groups above, supports Split Horizon based on the following rules:

*IP-based MPLS tunnels and SRv6 tunnels can support both Split Horizon filtering methods

*(SR-)MPLS tunnels only support ESI Label based Split Horizon filtering

*IP tunnels support Local Bias Split Horizon and may support ESI Label based Split Horizon, if they include a method to identify the source ESI in the header.

Tunnel Encapsulation	Default Split Horizon Type (SHT)	Supports Local Bias	Supports ESI Label
MPLSoGRE (IP-based MPLS)	ESI Label filtering	Yes	Yes
MPLSoUDP (IP-based MPLS)	ESI Label filtering	Yes	Yes
(SR-)MPLS	ESI Label filtering	No	Yes
VXLAN (IP tunnels)	Local Bias	Yes	No
NVGRE (IP tunnels)	Local Bias	Yes	No
VXLAN-GPE (IP tunnels)	Local Bias	Yes	No
GENEVE (IP tunnels)	Local Bias (no ESI Lb) ESI Label (if ESI lb)	Yes	Yes
SRv6	ESI Label filtering	Yes	Yes

Table 1: Tunnel Encapsulations and Split Horizon Types

The ESI Label method works for All-Active and Single-Active, while Local Bias only works for All-Active. In addition, the ESI Label method works across different network domains, whereas Local Bias is limited to networks with no next hop change between the NVEs attached to the same ES. However, some operators prefer the Local Bias method, since it simplifies the encapsulation, consumes less resources on the NVEs and the ingress NVE always forwards locally to other interfaces, reducing the delay to reach multihomed hosts.

This document extends the EVPN multihoming procedures so that an operator can decide the Split Horizon procedure for a given NVO tunnel depending on their own specific requirements. The choice of Local Bias or ESI Label Split Horizon is now allowed for tunnel encapsulations that support both methods, and it is advertised along with the EVPN A-D per ES route. IP tunnels that do not support both methods, E.g., VXLAN or NVGRE, will keep following [\[RFC8365\]](#) procedures.

1.2. Conventions and Terminology

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.

*AC: Attachment Circuit.

*A-D per ES route: refers to the EVPN Ethernet Auto-Discovery per ES route defined in [\[RFC7432\]](#).

*Arg.FE2: refers to the ESI filtering argument used for Split Horizon as specified in [\[RFC9252\]](#).

*Broadcast Domain (BD): an emulated ethernet, such that two systems on the same BD will receive each other's link-local broadcasts. In this document, BD also refers to the instantiation of a Broadcast Domain on an EVPN PE. An EVPN PE can be attached to one or multiple BDs of the same tenant.

*BUM: Broadcast, Unknown unicast and Multicast traffic.

*ES and ESI: Ethernet Segment and Ethernet Segment Identifier.

*Designated Forwarder (DF): as defined in [\[RFC7432\]](#), an ES may be multihomed (attached to more than one PE). An ES may also contain multiple BDs, of one or more EVIs. For each such EVI, one of the PEs attached to the segment becomes that EVI's DF for that segment. Since a BD may belong to only one EVI, we can speak unambiguously of the BD's DF for a given segment.

*EVI and EVI-RT: EVPN Instance and EVI Route Target. A group of NVEs attached to the same EVI will share the same EVI-RT.

*GENEVE: Generic Network Virtualization Encapsulation, [\[RFC8926\]](#).

*MPLS and non-MPLS NVO tunnels: refer to Multi-Protocol Label Switching (or the absence of it) Network Virtualization Overlay tunnels. Network Virtualization Overlay tunnels use an IP

encapsulation for overlay frames, where the source IP address identifies the ingress NVE and the destination IP address the egress NVE.

*MPLSoUDP: Multi-Protocol Label Switching over User Datagram Protocol, [[RFC7510](#)]

*MPLSoGRE: Multi-Protocol Label Switching over Generic Network Encapsulation, [[RFC4023](#)].

*MPLSoX: refers to MPLS over any IP encapsulation. Examples are MPLSoUDP or MPLSoGRE.

*NVE: Network Virtualization Edge device.

*NVGRE: Network Virtualization Using Generic Routing Encapsulation, [[RFC7637](#)].

*VXLAN: Virtual eXtensible Local Area Network, [[RFC7348](#)].

*VXLAN-GPE: VXLAN Generic Protocol Extension, [[I-D.ietf-nvo3-vxlan-gpe](#)].

*VNI: Virtual Network Identifier. A 24-bit identifier used by Network Virtualization Overlay (NVO) over IP encapsulations. Examples are VXLAN (Virtual Extended Local Area Network) or GENEVE (Generic Network Virtualization Encapsulation).

*SHT: Split Horizon Type, it refers to the Split Horizon method that a PE intends to use and advertises in an A-D per ES route.

*SR-MPLS: Segment Routing with an MPLS data plane, [[RFC8660](#)].

*SRv6: Segment routing with an IPv6 data plane, [[RFC8986](#)].

This document also assumes familiarity with the terminology of [[RFC7432](#)] and [[RFC8365](#)].

2. BGP EVPN Extensions

EVPN extensions are needed so that NVEs can advertise their preference for the Split Horizon method to be used in the ES. [Figure 1](#) shows the ESI Label extended community that is always advertised along with the EVPN A-D per ES route. All the NVEs attached to an ES advertise an A-D per ES route for the ES, including this extended community that conveys the information for the multihoming mode (All-active or Single-Active), as well as the ESI Label to be used (if needed).




Figure 1: ESI Label extended community

[RFC7432] defines the low-order bit of the Flags octet (bit 0) as the "Single-Active" bit:

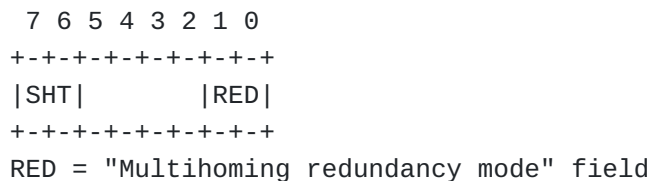
*A value of 0 means that the multihomed ES is operating in All-Active mode.

*A value of 1 means that the multihomed ES is operating in Single-Active mode.

Section 5 creates a registry for the Flags octet, where the "Single-Active" bit is the low-order bit of the RED (multihoming redundancy mode) field.

2.1. The Split Horizon Type (SHT)

[RFC8365] does not add any explicit indication about the Split Horizon method in the A-D per ES route. In this document, the [RFC8365] Split Horizon procedure is the default behavior and assumes that Local Bias is used only for IP tunnels, and ESI Label based Split Horizon for IP-based MPLS tunnels. This document defines the two high-order bits in the Flags octet (bits 6 and 7) as the "Split Horizon Type" (SHT) field, where:



SHT bit 7 6

- 0 0 --> Default SHT. Backwards compatible with [RFC8365]
- 0 1 --> Local Bias
- 1 0 --> ESI Label based filtering
- 1 1 --> reserved for future use

*SHT = 00 is backwards compatible with [RFC8365] and indicates that the advertising NVE intends to use the default or native SHT. The default SHT is shown in Table 1 for each encapsulation. An egress NVE that follows the [RFC8365] behavior and does not

support this specification will ignore the SHT bits (which is equivalent to process them as value of 00).

*SHT = 01 indicates that the advertising NVE intends to use Local Bias procedures in the ES for which the AD per-ES route is advertised.

*SHT = 10 indicates that the advertising NVE intends to use the ESI Label based Split Horizon method procedures in the ES for which the AD per-ES route is advertised.

*SHT = 11 is a reserved value, for future use.

2.2. Use of the Split Horizon Type In A-D Per ES Routes

The following behavior is observed:

*An SHT value of 01 or 10 MUST NOT be used with encapsulations that support only one SHT in [Table 1](#), and MAY be used by encapsulations that support the two SHTs in [Table 1](#).

*An SHT value different than 00 expresses the intent to use a specific Split Horizon method, but does not reflect the actual operational SHT used by the advertising NVE, unless all the NVEs attached to the ES advertise the same SHT.

*In case of inconsistency in the SHT value advertised by the NVEs attached to the same ES for a given EVI, all the NVEs MUST revert to the [\[RFC8365\]](#) behavior, and use the default SHT in [Table 1](#), irrespective of the advertised SHT.

*An SHT different from 00 MUST NOT be set if the Single-Active bit is set. A received A-D per ES route where Single-Active and SHT bits are different from zero MUST be treat-as-withdraw [\[RFC7606\]](#).

*The SHT MUST have the same value in each Ethernet A-D per ES route that an NVE advertises for a given ES and a given encapsulation (see [Section 3](#) for NVEs supporting multiple encapsulations).

As an example, egress NVEs that support IP-based MPLS tunnels, E.g., MPLSoGRE or MPLSoUDP, will advertise A-D per ES route(s) for the ES along with the BGP Encapsulation extended community [\[RFC9012\]](#) indicating the encapsulation (MPLSoGRE or MPLSoUDP) and MAY use the SHT = 01 or 10 to indicate the intent to use Local Bias or ESI Label, respectively.

An egress NVE MUST NOT use an SHT value different from 00 when advertising an A-D per ES route with encapsulation VXLAN, NVGRE, MPLS or no BGP tunnel encapsulation extended community [\[RFC9012\]](#). We

assume that, in all these cases, there is no Split Horizon method choice, and therefore the SHT value MUST be 00. A received route with one of the above encapsulation options and SHT value different from 00 SHOULD be treat-as-withdraw.

An egress NVE advertising A-D per ES route(s) for an ES with encapsulation GENEVE MAY use an SHT value of 01 or 10. A value of 01 indicates the intent to use Local Bias, irrespective of the presence of an Ethernet option TLV with a non-zero Source-ID [[I-D.ietf-bess-evpn-geneve](#)]. A value of 10 indicates the intent to use ESI Label based Split Horizon. A value of 00 indicates the default behavior in [Table 1](#), that is, use Local Bias if no ESI-Label exists in the Ethernet option TLV or no Ethernet option TLV whatsoever. Otherwise the ESI Label Split Horizon method is used.

The above procedures assume a single encapsulation supported in the egress NVE. [Section 3](#) describes additional procedures for NVEs supporting multiple encapsulations.

2.3. ESI Label Value In A-D Per ES Routes

This document also updates [[RFC8365](#)] in the value that is advertised in the ESI Label field of the ESI Label extended community, as follows:

*The A-D per ES route(s) for an ES MAY have an ESI Label value of zero if the SHT value is 01. [Section 2.2](#) specifies the cases where the SHT can be 01. An ESI Label value of zero avoids the allocation of Labels in the cases where they are not used (Local Bias).

*The A-D per ES route(s) for an ES MAY have an ESI Label value of zero for VXLAN or NVGRE encapsulations.

2.4. Backwards Compatibility With RFC8365 NVEs

As discussed in [Section 2.2](#) this specification is backwards compatible with the Split Horizon filtering behavior in [[RFC8365](#)] and a non-upgraded NVE can be attached to the same ES as other NVEs supporting this specification.

An NVE has an administrative SHT value for an ES (the one that is advertised along with the A-D per ES route) and an operational SHT value (the one that is actually used irrespective of what the NVE advertised). The administrative SHT matches the operational SHT if all the NVEs attached to the ES have the same administrative SHT.

This document assumes that an [[RFC7432](#)] or [[RFC8365](#)] implementation that does not support this document, ignores the value of all the Flags in the ESI Label extended community except for the Single-

Active bit. Based on this assumption, a non-upgraded NVE will ignore an SHT different from 00. As soon as an upgraded NVE receives at least one A-D per ES route for the ES with SHT value of 00, it MUST revert its operational SHT to the default Split Horizon method, as in [Table 1](#), and irrespective of its administrative SHT.

As an example, consider an NVE attached to ES N that receives two A-D per ES routes for N from different NVEs, NVE1 and NVE2. If the route from NVE1 has SHT = 00 and the one from NVE2 an SHT = 01, the NVE MUST use the default Split Horizon method in [Table 1](#) as operational SHT, irrespective of its administrative SHT.

All the NVEs attached to an ES with operational SHT value of 10 MUST advertise a valid non-zero ESI Label. If the operational SHT value is 01, the ESI Label MAY be zero. If the operational SHT value is 00, the ESI Label MAY be zero only if the default encapsulation supports Local Bias only and the NVEs do not check the presence of a valid non-zero ESI Label.

If an NVE changes its operational SHT value from 01 (Local Bias) to 00 (Default SHT) as a result of a new non-upgraded NVE present in the ES, and it previously advertised a zero ESI Label, it MUST send an update with a non-zero valid ESI Label, unless all the non-upgraded NVEs in the ES support Local Bias only. As an example, suppose NVE1 and NVE2 use MPLSoUDP as encapsulation, they are attached to the same Ethernet Segment ES1 and advertise an SHT value of 01 (Local Bias) and a zero ESI label value. Suppose NVE3 does not support this specification and joins ES1, therefore advertises an SHT of 00 (default). Upon receiving NVE3's A-D per ES route, NVE1 and NVE2 MUST send an update of their A-D per ES route for ES1 with a non-zero valid ESI label value. The assumption is that NVE3 supports only the default ESI label based Split Horizon filtering.

3. Procedures for NVEs Supporting Multiple Encapsulations

As specified by [\[RFC8365\]](#), an egress NVE that supports multiple data plane encapsulations (I.e., VXLAN, NVGRE, MPLS, MPLSoUDP, GENEVE) needs to indicate all the supported encapsulations using BGP Encapsulation extended communities defined in [\[RFC9012\]](#) with all EVPN routes. This section clarifies the multihoming Split Horizon behavior for NVEs advertising and receiving multiple BGP Encapsulation extended communities along with the A-D per ES routes. This section uses a notation of {x,y} to indicate the encapsulations advertised in BGP Encapsulation extended communities [\[RFC9012\]](#), with x and y being different encapsulation values.

It is important to remember that an NVE MAY advertise multiple A-D per ES routes for the same ES (and not only one), each route conveying a number of Route Targets (RT). We refer to the total

number of Route Targets in a given ES as RT-set for that ES. Any of the EVIs represented in the RT-set will have its RT included in one (and only one) A-D per ES route for the ES. When multiple A-D per ES routes are advertised for the same ES, each route MUST have a different Route Distinguisher.

As per [[RFC8365](#)], an NVE that advertises multiple encapsulations in the A-D per ES route(s) for an ES MUST advertise encapsulations that use the same Split Horizon filtering method in the same route. For example:

*An A-D per ES route for ES-x may be advertised with {VXLAN, NVGRE} encapsulations.

*An A-D per ES route for ES-y may be advertised with {MPLS, MPLSoUDP, MPLSoGRE} encapsulations (or a subset).

*But an A-D per ES route for ES-z MUST NOT be advertised with {MPLS, VXLAN} encapsulations.

This document extends this behavior as follows:

- a. An A-D per ES route for ES-x may be advertised with multiple encapsulations where some support a single Split Horizon method. In this case, the SHT value MUST be 00. As an example, {VXLAN, NVGRE}, {VXLAN, GENEVE} or {MPLS, MPLSoGRE, MPLSoUDP} can be advertised in an A-D per ES route. In all those cases SHT MUST be 00.
- b. An A-D per ES route for ES-y may be advertised with multiple encapsulations where all of them support both Split Horizon methods. In this case the SHT value MAY be 01 if the desired method is Local Bias, or 10 if ESI Label based. For example, {MPLSoGRE, MPLSoUDP, GENEVE} (or a subset) may be advertised in an A-D per ES route with SHT value of 01. The ESI Label value in this case MAY be zero.
- c. If ES-z with RT-set composed of (RT1, RT2, RT3.. RTn) supports multiple encapsulations that require a different Split Horizon method, a different A-D per ES route (or group of routes) per Split Horizon method MUST be advertised. For example, consider n RTs in ES-z and:
 - *the EVIs corresponding to (RT1..RTi) support VXLAN,
 - *the ones for (RTi+1..RTm) (with i<m) support MPLSoUDP with Local Bias,
 - *and the ones for (RTm+1..RTn) (with m<n) support GENEVE with ESI Label based Split Horizon.

In this case, three groups of A-D per ES routes MUST be advertised for ES-z:

*A-D per ES route group 1, including (RT1..RTi), with encapsulation {VXLAN}, SHT = 00. The ESI Label MAY be zero.

*A-D per ES route group 2, including (RTi+1..RTm), with encapsulation {MPLSoUDP}, SHT = 01. The ESI Label MAY be zero.

*A-D per ES route group 3, including (RTm+1..RTn), with encapsulation {GENEVE}, SHT = 10. The ESI Label MUST have a valid value, different from zero, and the Ethernet option [[RFC8926](#)] MUST be advertised.

As per [[RFC8365](#)], it is the responsibility of the operator of a given EVI to ensure that all of the NVEs in that EVI support a common encapsulation. If this condition is violated, it could result in service disruption or failure.

4. Security Considerations

The same security considerations described in [[RFC7432](#)] relevant to multihoming apply to this document.

In addition, this document modifies the [[RFC8365](#)] procedures for Split Horizon filtering, providing the operator with a choice between Local Bias and ESI Label based filtering for the tunnels that support both methods. A misconfiguration of the desired SHT to be used may result in a forwarding behavior that is different from the intended one. Other than that, this document describes procedures so that all the PEs or NVEs attached to the same ES agree on a common SHT method, therefore an attacker changing the configuration of the SHT should not cause traffic disruption, only a change in the forwarding behavior.

5. IANA Considerations

This document creates a registry called "EVPN ESI Label Extended Community Flags" for the 1-octet Flags field in the ESI Label Extended Community. New registrations will be made through the "RFC Required" procedure defined in [[RFC8126](#)]. Initial registrations are made for the "Multihoming redundancy mode" field in bits 0 and 1, as follows:

RED	Multihoming redundancy mode
00	All-Active mode
01	Single-Active mode

Table 2

In addition, this document requests the registration of the "Split Horizon Type" field in bits 6 and 7 of the Flags Octet of the EVPN ESI Label extended community. This field is called "Split Horizon Type" bits and it is defined as follows:

SHT	Split Horizon Type value
00	Default SHT
01	Local Bias
10	ESI Label based filtering
11	Reserved

Table 3

6. Acknowledgments

The authors would like to thank Anoop Ghanwani, Gyan Mishra and Jeffrey Zhang for their review and useful comments.

7. Contributors

8. References

8.1. 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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay

Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.

8.2. Informative References

- [I-D.ietf-bess-evpn-geneve] Boutros, S., Sajassi, A., Drake, J., Rabadan, J., and S. Aldrin, "EVPN control plane for Geneve", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-geneve-06, 26 May 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-geneve-06>>.
- [RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <<https://www.rfc-editor.org/info/rfc7348>>.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, DOI 10.17487/RFC4023, March 2005, <<https://www.rfc-editor.org/info/rfc4023>>.
- [RFC7637] Garg, P., Ed. and Y. Wang, Ed., "NVGRE: Network Virtualization Using Generic Routing Encapsulation", RFC 7637, DOI 10.17487/RFC7637, September 2015, <<https://www.rfc-editor.org/info/rfc7637>>.
- [RFC7510] Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", RFC 7510, DOI 10.17487/RFC7510, April 2015, <<https://www.rfc-editor.org/info/rfc7510>>.
- [RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <<https://www.rfc-editor.org/info/rfc8926>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.

[RFC8660]

Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

[RFC8986]

Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

[I-D.ietf-nvo3-vxlan-gpe] Maino, F., Kreeger, L., and U. Elzur, "Generic Protocol Extension for VXLAN (VXLAN-GPE)", Work in Progress, Internet-Draft, draft-ietf-nvo3-vxlan-gpe-13, 4 November 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-nvo3-vxlan-gpe-13>>.

Authors' Addresses

Jorge Rabadan (editor)
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America

Email: jorge.rabadan@nokia.com

Kiran Nagaraj
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America

Email: kiran.nagaraj@nokia.com

Wen Lin
Juniper Networks

Email: wlin@juniper.net

Ali Sajassi
Cisco Systems, Inc.

Email: sajassi@cisco.com