

BESS Workgroup
Internet-Draft
Intended status: Standards Track
Expires: April 12, 2021

J. Rabadan, Ed.
K. Nagaraj
Nokia
W. Lin
Juniper
A. Sajassi
Cisco
October 9, 2020

EVPN Multi-Homing Extensions for Split Horizon Filtering
draft-ietf-bess-evpn-mh-split-horizon-00

Abstract

Ethernet Virtual Private Network (EVPN) is commonly used along with Network Virtualization Overlay (NVO) tunnels. The EVPN multi-homing procedures may be different depending on the NVO 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 others, E.g., VXLAN tunnels. The current 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 extends the EVPN Multi-Homing procedures so that an operator can decide the Split Horizon procedure for a given NVO 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 April 12, 2021.

Copyright Notice

Copyright (c) 2020 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | | |
|-----------------------------|---|--------------------|
| 1. | Introduction | 2 |
| 1.1. | Conventions and Terminology | 5 |
| 2. | BGP EVPN Extensions | 7 |
| 2.1. | The Split Horizon Type (SHT) | 7 |
| 2.2. | Use of the Split Horizon Type In A-D Per ES Routes | 8 |
| 2.3. | ESI Label Value In A-D Per ES Routes | 9 |
| 2.4. | Backwards Compatibility With [RFC8365] NVEs | 10 |
| 3. | Procedures for NVEs Supporting Multiple Encapsulations . . . | 10 |
| 4. | Security Considerations | 12 |
| 5. | IANA Considerations | 13 |
| 6. | References | 13 |
| 6.1. | Normative References | 13 |
| 6.2. | Informative References | 13 |
| Appendix A. | Acknowledgments | 14 |
| Appendix B. | Contributors | 15 |
| | Authors' Addresses | 15 |

[1.](#) Introduction

Ethernet Virtual Private Network (EVPN) is commonly used along with Network Virtualization Overlay (NVO) tunnels and specified in [\[RFC8365\]](#). The EVPN multi-homing procedures may be different depending on the NVO 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 [\[RFC7510\]](#), and its procedures described in [\[RFC7432\]](#). Local Bias is used by non-MPLS NVO tunnels, E.g., VXLAN tunnels, and it is described in [\[RFC8365\]](#).

As a refresher:

- o 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 multi-homing. The ingress NVE adds a label corresponding to the source Ethernet Segment (aka 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 multi-homing to avoid transient loops for in-flight packets when the egress NVE takes over as DF for an Ethernet Segment.

- o Local Bias for non-MPLS NVO tunnels [[RFC8365](#)]

Since non-MPLS NVO 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 multi-homing. This mechanism is called Local Bias and relies on the NVO 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 multi-homed ESs. When the egress NVE receives a BUM frame encapsulated in a non-MPLS NVO packet, 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 Ethernet Segments (regardless of the DF 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 multi-homing, since the ingress NVE brings operationally down the ACs for which it is non-DF (hence local replication to non-DF ACs cannot be done). This means transient in-flight BUM packets may be looped back to the originating site by new elected DF egress NVEs.

[RFC8365] states that Local Bias is used only for non-MPLS NVO tunnels, and ESI Label based Split Horizon for MPLS NVO tunnels. However, MPLS NVO tunnels, such as MPLSoGRE or MPLSoUDP, 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 non-MPLS NVO 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:

- o 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.
- o In an SRv6 tunnel, the source IP address also identifies the ingress NVE, however, by default and as described in [\[I-D.ietf-bess-srv6-services\]](#) 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 of the service 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.

| Tunnel Encapsulation | Default Split Horizon Type (SHT) | Supports Local Bias | Supports ESI Label |
|----------------------|----------------------------------|---------------------|--------------------|
| VXLAN | Local Bias | Yes | No |
| NVGRE | Local Bias | Yes | No |
| MPLS | ESI Label filtering | No | Yes |
| MPLSoGRE | ESI Label filtering | Yes | Yes |
| MPLSoUDP | ESI Label filtering | Yes | Yes |
| GENEVE | Local Bias (no ESI Lb) | Yes | Yes |
| | ESI Label (if ESI lb) | | |
| 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 Ethernet Segment. 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 multi-homed hosts.

This document extends the EVPN Multi-Homing 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 NVO tunnels that support both methods. Non-MPLS NVO tunnels that do not support both methods, E.g., VXLAN or NVGRE, will keep following [\[RFC8365\]](#) procedures.

1.1. 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.

- o BUM: Broadcast, Unknown unicast and Multicast traffic.
- o ES and ESI: Ethernet Segment and Ethernet Segment Identifier.
- o A-D per ES route: refers to the EVPN Ethernet Auto-Discovery per Ethernet Segment route defined in [\[RFC7432\]](#).
- o AC: Attachment Circuit.
- o NVE: Network Virtualization Edge device.
- o 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.
- o 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.
- o MPLSoUDP: Multi-Protocol Label Switching over User Datagram Protocol, [\[RFC7510\]](#)
- o MPLSoGRE: Multi-Protocol Label Switching over Generic Network Encapsulation, [\[RFC4023\]](#).
- o MPLSoX: refers to MPLS over any IP encapsulation. Examples are MPLSoUDP or MPLSoGRE.
- o GENEVE: Generic Network Virtualization Encapsulation, [\[I-D.ietf-nvo3-geneve\]](#).
- o VXLAN: Virtual eXtensible Local Area Network, [\[RFC7348\]](#).
- o NVGRE: Network Virtualization Using Generic Routing Encapsulation, [\[RFC7637\]](#).
- o 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).
- o 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.

- o Designated Forwarder (DF): as defined in [\[RFC7432\]](#), an ethernet segment may be multi-homed (attached to more than one PE). An ethernet segment 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.
- o 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.

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 Ethernet Segment. 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 Ethernet Segment advertise an A-D per ES route for the ES, including this extended community that conveys the information for the multi-homing mode (All-active or Single-Active), as well as the ESI Label to be used (if needed).

```

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

```

Figure 1: ESI Label extended community

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

- o A value of 0 means that the multi-homed Ethernet Segment is operating in All-Active mode.
- o A value of 1 means that the multi-homed Ethernet Segment is operating in Single-Active mode.

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 non-MPLS NVO tunnels, and ESI Label based Split Horizon for MPLS NVO 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 |

- o 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 NVO 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).
- o SHT = 01 indicates that the advertising NVE intends to use Local Bias procedures in the Ethernet Segment for which the AD per-ES route is advertised.
- o SHT = 10 indicates that the advertising NVE intends to use the ESI Label based Split Horizon method procedures in the Ethernet Segment for which the AD per-ES route is advertised.

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

The following must be observed:

- o 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.
- o An SHT value different than 00 expresses the intend 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.
- o 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.
- o 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 [\[RFC7606\]](#).

- o 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 MPLS NVO tunnels, E.g., MPLSoGRE or MPLSoUDP, will advertise A-D per ES route(s) for the ES along with the [\[RFC5512\]](#) BGP Encapsulation extended community [\[I-D.ietf-idr-tunnel-encaps\]](#) 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 [\[I-D.ietf-idr-tunnel-encaps\]](#) BGP tunnel encapsulation extended community. 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:

- o 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).
- o 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 Ethernet Segment 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 to 00 (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.

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 [\[I-D.ietf-idr-tunnel-encaps\]](#) with all EVPN routes. This section

clarifies the multi-homing 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 [[I-D.ietf-idr-tunnel-encaps](#)] BGP Encapsulation extended communities, 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:

- o An A-D per ES route for ES-x may be advertised with {VXLAN, NVGRE} encapsulations.
- o An A-D per ES route for ES-y may be advertised with {MPLS, MPLSoUDP, MPLSoGRE} encapsulations (or a subset).
- o 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:
- o the EVIs corresponding to (RT1..RTi) support VXLAN,
 - o the ones for (RTi+1..RTm) (with $i < m$) support MPLSoUDP with Local Bias,
 - o 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:

- o A-D per ES route group 1, including (RT1..RTi), with encapsulation {VXLAN}, SHT = 00. The ESI Label MAY be zero.
- o A-D per ES route group 2, including (RTi+1..RTm), with encapsulation {MPLSoUDP}, SHT = 01. The ESI Label MAY be zero.
- o 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 [[I-D.ietf-nvo3-geneve](#)] 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 Multi-Homing 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

IANA is requested to allocate the SHT bits (6 and 7) in the Flags Octet of the EVPN ESI Label extended community. This field is called "Split Horizon Type" bits.

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>>.
- [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>>.
- [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>>.

6.2. Informative References

- [I-D.ietf-bess-evpn-geneve] Boutros, S., Sajassi, A., Drake, J., Rabadan, J., and S. Aldrin, "EVPN control plane for Geneve", [draft-ietf-bess-evpn-geneve-01](#) (work in progress), June 2020.
- [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>>.

- [RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute", [RFC 5512](#), DOI 10.17487/RFC5512, April 2009, <<https://www.rfc-editor.org/info/rfc5512>>.
- [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>>.
- [I-D.ietf-nvo3-geneve]
Gross, J., Ganga, I., and T. Sridhar, "Geneve: Generic Network Virtualization Encapsulation", [draft-ietf-nvo3-geneve-16](#) (work in progress), March 2020.
- [I-D.ietf-idr-tunnel-encaps]
Patel, K., Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", [draft-ietf-idr-tunnel-encaps-19](#) (work in progress), September 2020.
- [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>>.
- [I-D.ietf-bess-srv6-services]
Dawra, G., Filsfils, C., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "SRv6 BGP based Overlay services", [draft-ietf-bess-srv6-services-04](#) (work in progress), July 2020.

[Appendix A](#). Acknowledgments

Appendix B. Contributors

Authors' Addresses

Jorge Rabadan (editor)
Nokia
777 Middlefield Road
Mountain View, CA 94043
USA

Email: jorge.rabadan@nokia.com

Kiran Nagaraj
Nokia
701 Middlefield Road
Mountain View, CA 94043
USA

Email: kiran.nagaraj@nokia.com

Wen Lin
Juniper Networks

Email: wlin@juniper.net

Ali Sajassi
Cisco Systems, Inc.

Email: sajassi@cisco.com

