

Internet Engineering Task Force (IETF)
Request for Comments: 9573
Updates: 6514, 7432, 7582
Category: Standards Track
ISSN: 2070-1721

Z. Zhang
Juniper Networks
E. Rosen
Individual
W. Lin
Juniper Networks
Z. Li
Huawei Technologies
IJ. Wijnands
Individual
May 2024

MVPN/EVPN Tunnel Aggregation with Common Labels

Abstract

The Multicast VPN (MVPN) specifications allow a single Point-to-Multipoint (P2MP) tunnel to carry traffic of multiple IP VPNs (referred to as VPNs in this document). The EVPN specifications allow a single P2MP tunnel to carry traffic of multiple Broadcast Domains (BDs). These features require the ingress router of the P2MP tunnel to allocate an upstream-assigned MPLS label for each VPN or for each BD. A packet sent on a P2MP tunnel then carries the label that is mapped to its VPN or BD (in some cases, a distinct upstream-assigned label is needed for each flow.) Since each ingress router allocates labels independently, with no coordination among the ingress routers, the egress routers may need to keep track of a large number of labels. The number of labels may need to be as large as, or larger than, the product of the number of ingress routers times the number of VPNs or BDs. However, the number of labels can be greatly reduced if the association between a label and a VPN or BD is made by provisioning, so that all ingress routers assign the same label to a particular VPN or BD. New procedures are needed in order to take advantage of such provisioned labels. These new procedures also apply to Multipoint-to-Multipoint (MP2MP) tunnels. This document updates RFCs 6514, 7432, and 7582 by specifying the necessary procedures.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9573>.

Copyright Notice

Copyright (c) 2024 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction
1.1.	Requirements Language
1.2.	Terminology
2.	Problem Description
3.	Proposed Solutions
3.1.	MP2MP Tunnels
3.2.	Segmented Tunnels
3.3.	Summary of Label Allocation Methods
4.	Specifications
4.1.	Context-Specific Label Space ID Extended Community
4.2.	Procedures
5.	Security Considerations
6.	IANA Considerations
7.	References
7.1.	Normative References
7.2.	Informative References
	Acknowledgements
	Contributors
	Authors' Addresses

1. Introduction

A Multicast VPN (MVPN) can use Point-to-Multipoint (P2MP) tunnels (set up by RSVP-TE, Multipoint LDP (mLDP), or PIM) to transport customer multicast traffic across a service provider's backbone network. Often, a given P2MP tunnel carries the traffic of only a single VPN. However, there are procedures defined that allow a single P2MP tunnel to carry traffic of multiple VPNs. In this case, the P2MP tunnel is called an "aggregate tunnel". The Provider Edge (PE) router that is the ingress node of an aggregate P2MP tunnel allocates an "upstream-assigned MPLS label" [RFC5331] for each VPN, and each packet sent on the P2MP tunnel carries the upstream-assigned MPLS label that the ingress PE has bound to the packet's VPN.

Similarly, an EVPN can use P2MP tunnels (set up by RSVP-TE, mLDP, or PIM) to transport Broadcast, Unknown Unicast, or Multicast (BUM) traffic across the provider network. Often, a P2MP tunnel carries the traffic of only a single Broadcast Domain (BD). However, there are procedures defined that allow a single P2MP tunnel to be an aggregate tunnel that carries traffic of multiple BDs. The procedures are analogous to the MVPN procedures -- the PE router that is the ingress node of an aggregate P2MP tunnel allocates an upstream-assigned MPLS label for each BD, and each packet sent on the P2MP tunnel carries the upstream-assigned MPLS label that the ingress PE has bound to the packet's BD.

An MVPN or EVPN can also use Bit Index Explicit Replication (BIER) [RFC8279] to transmit VPN multicast traffic [RFC8556] or EVPN BUM traffic [BIER-EVPN]. Although BIER does not explicitly set up P2MP tunnels, from the perspective of an MVPN/EVPN, the use of BIER transport is very similar to the use of aggregate P2MP tunnels. When BIER is used, the PE transmitting a packet (the "Bit-Forwarding Ingress Router" (BFIR) [RFC8279]) must allocate an upstream-assigned MPLS label for each VPN or BD, and the packets transmitted using BIER transport always carry the label that identifies their VPN or BD. (See [RFC8556] and [BIER-EVPN] for details.) In the remainder of this document, we will use the term "aggregate tunnels" to include both P2MP tunnels and BIER transport.

When an egress PE receives a packet from an aggregate tunnel, it must look at the upstream-assigned label carried by the packet and must interpret that label in the context of the ingress PE. Essentially, for each ingress PE, the egress PE has a context-specific label space [RFC5331] that matches the default label space from which the ingress PE assigns the upstream-assigned labels. When an egress PE looks up the upstream-assigned label carried by a given packet, it looks it up in the context-specific label space for the ingress PE of the packet. How an egress PE identifies the ingress PE of a given packet depends on the tunnel type.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

Familiarity with MVPN/EVPN protocols and procedures is assumed. Some terms are listed below for convenience.

VPN: Virtual Private Network. In this document, "VPN" specifically refers to an IP VPN [RFC4364].

BUM [RFC7432]: Broadcast, Unknown Unicast, or Multicast (traffic).

BD [RFC7432]: Broadcast Domain.

EC [RFC4360]: Extended Community.

PMSI [RFC6513]: Provider Multicast Service Interface. A pseudo-overlay interface for PEs to send certain overlay/customer multicast traffic via underlay/provider tunnels. It includes Inclusive/Selective PMSIs (I/S-PMSIs) (often referred to as x-PMSIs). A PMSI is instantiated by the underlay/provider tunnel.

Inclusive PMSI (I-PMSI): A PMSI that enables traffic to be sent to all PEs of a VPN/BD. The underlay/provider tunnel that instantiates the I-PMSI is referred to as an inclusive tunnel.

Selective PMSI (S-PMSI): A PMSI that enables traffic to be sent to a subset of PEs of a VPN/BD. The underlay/provider tunnel that instantiates the S-PMSI is referred to as a selective tunnel.

Aggregate Tunnel: A tunnel that instantiates x-PMSIs of multiple MVPNs or EVPN BDs.

IMET [RFC7432]: Inclusive Multicast Ethernet Tag. An EVPN-specific name for an I-PMSI Auto-Discovery (A-D) route.

PTA [RFC6514]: PMSI Tunnel Attribute. A BGP attribute that may be attached to a BGP-MVPN/EVPN x-PMSI A-D route.

ASBR: Autonomous System Border Router.

RBR: Regional Border Router. A border router between segmentation regions that participates in segmentation procedures.

(C-S,C-G): A Customer/overlay <S,G> multicast flow.

(C-*,C-G): Customer/overlay <*,G> multicast flows.

(C-*,C-*): All Customer/overlay multicast flows.

ES: Ethernet Segment.

ESI [RFC7432]: ES Identifier.

ESI Label [RFC7432]: A label that identifies an ES.

SRGB [RFC8402]: Segment Routing (SR) Global Block. The set of global segments in the SR domain. In SR-MPLS [RFC8660], an SRGB is a local property of a node and identifies the set of local labels reserved for global segments.

DCB: Domain-wide Common Block. A common block of labels reserved on all nodes in a domain.

Context-Specific Label Space [RFC5331]: A router may maintain additional label spaces besides its default label space. When the label at the top of the stack is not from the default label space, there must be some context in the packet that identifies which one of those additional label spaces is to be used to look up the label; hence, those label spaces are referred to as context-specific label spaces.

Upstream Assigned [RFC5331]: When the label at the top of the label stack is not assigned by the router receiving the traffic from its default label space, the label is referred to as upstream assigned. Otherwise, it is downstream assigned. An upstream-assigned label must be looked up in a context-specific label space specific for the assigner.

2. Problem Description

Note that the upstream-assigned label procedures may require a very large number of labels. Suppose that an MVPN or EVPN deployment has 1001 PEs, each hosting 1000 VPNS/BDS. Each ingress PE has to assign 1000 labels, and each egress PE has to be prepared to interpret 1000 labels from each of the ingress PEs. Since each ingress PE allocates labels from its own label space and does not coordinate label assignments with others, each egress PE must be prepared to interpret 1,000,000 upstream-assigned labels (across 1000 context-specific label spaces -- one for each ingress PE). This is an evident scaling problem.

So far, few if any MVPN/EVPN deployments use aggregate tunnels, so this problem has not surfaced. However, the use of aggregate tunnels is likely to increase due to the following two factors:

- * In an EVPN, a single customer ("tenant") may have a large number of BDS, and the use of aggregate RSVP-TE or mLDP P2MP tunnels may become important, since each tunnel creates state at the intermediate nodes.
- * The use of BIER as the transport for an MVPN/EVPN is becoming more

and more attractive and feasible.

A similar problem also exists with EVPN ESI labels used for multihoming. A PE attached to a multihomed ES advertises an ESI label in its Ethernet A-D per ES route. The PE imposes the label when it sends frames received from the ES to other PEs via a P2MP/BIER tunnel. A receiving PE that is attached to the source ES will know from the ESI label that the packet originated on the source ES and thus will not transmit the packet on its local Attachment Circuit to that ES. From the receiving PE's point of view, the ESI label is (upstream) assigned from the source PE's label space, so the receiving PE needs to maintain context-specific label tables, one for each source PE, just like the VPN/BD label case above. If there are 1001 PEs, each attached to 1000 ESs, this can require each PE to understand 1,000,000 ESI labels. Notice that the issue exists even when no P2MP tunnel aggregation (i.e., one tunnel used for multiple BDs) is used.

3. Proposed Solutions

The number of labels could be greatly reduced if a central entity in the provider network assigned a label to each VPN, BD, or ES and if all PEs used that same label to represent a given VPN, BD, or ES. Then, the number of labels needed would just be the sum of the number of VPNs, BDs, and/or ESs.

One method of achieving this is to reserve a portion of the default label space for assignment by a central entity. We refer to this reserved portion as the DCB of labels. This is analogous to the concept of using identical SRGBs on all nodes, as described in [RFC8402]. A PE that is attached (via L3VPN Virtual Routing and Forwarding (VRF) interfaces or EVPN Attachment Circuits) would know by provisioning which label from the DCB corresponds to which of its locally attached VPNs, BDs, or ESs.

For example, all PEs could reserve a DCB [1000~2000], and they would all be provisioned so that label 1000 maps to VPN 0, label 1001 maps to VPN 1, and so forth. Now, only 1000 labels instead of 1,000,000 labels are needed for 1000 VPNs.

In this document, "domain" is defined loosely; it simply includes all the routers that share the same DCB, and it only needs to include all PEs of an MVPN/EVPN.

The "domain" could also include all routers in the provider network, making it not much different from a common SRGB across all the routers. However, that is not necessary, as the labels used by PEs for the purposes defined in this document will only rise to the top of the label stack when traffic arrives at the PEs. Therefore, it is better to not include internal P routers in the "domain". That way, they do not have to set aside the same DCB used for the purposes defined in this document.

In some deployments, it may be impractical to allocate a DCB that is large enough to contain labels for all the VPNs/BDs/ESs. In this case, it may be necessary to allocate those labels from one or a few context-specific label spaces that are independent of each PE. For example, if it is too difficult to have a DCB of 10,000 labels across all PEs for all the VPNs/BDs/ESs that need to be supported, a separate context-specific label space can be dedicated to those 10,000 labels. Each separate context-specific label space is identified in the forwarding plane by a label from the DCB (which does not need to be large). Each PE is provisioned with the label-space-identifying DCB label and the common VPN/BD/ES labels allocated from that context-specific label space. When sending traffic, an ingress PE imposes all necessary service labels (for the VPN/BD/ES)

first, then imposes the label-space-identifying DCB label. From the label-space-identifying DCB label, an egress PE can determine the label space where the inner VPN/BD/ES label is looked up.

The MVPN/EVPN signaling defined in [RFC6514] and [RFC7432] assumes that certain MPLS labels are allocated from a context-specific label space for a particular ingress PE. In this document, we augment the signaling procedures so that it is possible to signal that a particular label is from the DCB, rather than from a context-specific label space for an ingress PE. We also augment the signaling so that it is possible to indicate that a particular label is from an identified context-specific label space that is not for an ingress PE.

Notice that the VPN/BD/ES-identifying labels from the DCB or from those few context-specific label spaces are very similar to Virtual eXtensible Local Area Network (VXLAN) Network Identifiers (VNIs) in VXLANs. Allocating a label from the DCB or from a context-specific label space and communicating the label to all PEs is not different from allocating VNIs and is especially feasible with controllers.

3.1. MP2MP Tunnels

Multipoint-to-Multipoint (MP2MP) tunnels present the same problem (Section 2) that can be solved the same way (Section 3), with the following additional requirement.

Per [RFC7582] ("Multicast Virtual Private Network (MVPN): Using Bidirectional P-Tunnels"), when MP2MP tunnels are used for an MVPN, the root of the MP2MP tunnel is required to allocate and advertise "PE Distinguisher Labels" (Section 4 of [RFC6513]). These labels are assigned from the label space used by the root node for its upstream-assigned labels.

It is REQUIRED by this document that the PE Distinguisher Labels allocated by a particular node come from the same label space that the node uses to allocate its VPN-identifying labels.

3.2. Segmented Tunnels

There are some additional issues to be considered when an MVPN or EVPN is using "tunnel segmentation" (see [RFC6514], [RFC7524], and Sections 5 and 6 of [RFC9572]).

3.2.1. Selective Tunnels

For selective tunnels that instantiate S-PMSIs (see Sections 2.1.1 and 3.2.1 of [RFC6513] and Section 4 of [RFC9572]), the procedures outlined above work only if tunnel segmentation is not used.

A selective tunnel carries one or more particular sets of flows to a particular subset of the PEs that attach to a given VPN or BD. Each set of flows is identified by an S-PMSI A-D route [RFC6514]. The PTA of the S-PMSI route identifies the tunnel used to carry the corresponding set of flows. Multiple S-PMSI routes can identify the same tunnel.

When tunnel segmentation is applied to an S-PMSI, certain nodes are "segmentation points". A segmentation point is a node at the boundary between two segmentation regions. Let's call these "region A" and "region B". A segmentation point is an egress node for one or more selective tunnels in region A and an ingress node for one or more selective tunnels in region B. A given segmentation point must be able to receive traffic on a selective tunnel from region A and label-switch the traffic to the proper selective tunnel in region B.

Suppose that one selective tunnel (call it "T1") in region A is carrying two flows, Flow-1 and Flow-2, identified by S-PMSI routes Route-1 and Route-2, respectively. However, it is possible that in region B, Flow-1 is not carried by the same selective tunnel that carries Flow-2. Let's suppose that in region B, Flow-1 is carried by tunnel T2 and Flow-2 by tunnel T3. Then, when the segmentation point receives traffic from T1, it must be able to label-switch Flow-1 from T1 to T2, while also label-switching Flow-2 from T1 to T3. This implies that Route-1 and Route-2 must signal different labels in the PTA. For comparison, when segmentation is not used, they can all use the common per-VPN/BD DCB label.

In this case, it is not practical to have a central entity assign domain-wide unique labels to individual S-PMSI routes. To address this problem, all PEs can be assigned their own disjoint label blocks in those few context-specific label spaces; each PE will independently allocate labels for a segmented S-PMSI from its own assigned label block. For example, PE1 allocates from label block [101~200], PE2 allocates from label block [201~300], and so on.

Allocating from disjoint label blocks can be used for VPN/BD/ES labels as well, though it does not address the original scaling issue, because there would be 1,000,000 labels allocated from those few context-specific label spaces in the original example, instead of just 1000 common labels.

3.2.2. Per-PE/Region Tunnels

Similarly, for segmented per-PE (MVPN (C-*,C-*) S-PMSI or EVPN IMET) or per-AS/region (MVPN Inter-AS I-PMSI or EVPN per-region I-PMSI) tunnels [RFC6514] [RFC7432] [RFC9572], labels need to be allocated per PMSI route. In the case of a per-PE PMSI route, the labels should be allocated from the label block allocated to the advertising PE. In the case of a per-AS/region PMSI route, different ASBRs/RBRs attached to the same source AS/region will advertise the same PMSI route. The same label could be used when the same route is advertised by different ASBRs/RBRs, though that requires coordination; a simpler way is for each ASBR/RBR to allocate a label from the label block allocated to itself (see Section 3.2.1).

In the rest of this document, we call the label allocated for a particular PMSI a "(per-)PMSI label", just like we have (per-)VPN/BD/ES labels. Notice that using a per-PMSI label in the case of a per-PE PMSI still has the original scaling issue associated with the upstream-assigned label, so per-region PMSIs are preferred. Within each AS/region, per-PE PMSIs are still used, though they do not go across borders and per-VPN/BD labels can still be used.

Note that when a segmentation point re-advertises a PMSI route to the next segment, it does not need to re-advertise a new label unless the upstream or downstream segment uses ingress replication.

3.2.3. Alternative to Per-PMSI Label Allocation

Per-PMSI label allocation in the case of segmentation, whether for S-PMSIs or per-PE/region I-PMSIs, is applied so that segmentation points are able to label-switch traffic without having to do IP or Media Access Control (MAC) lookups in VRFs (the segmentation points typically do not have those VRFs at all). Alternatively, if the label scaling becomes a concern, the segmentation points could use (C-S,C-G) lookups in VRFs for flows identified by the S-PMSIs. This allows the S-PMSIs for the same VPN/BD to share a VPN/BD-identifying label that leads to lookups in the VRFs. That label needs to be different from the label used in the per-PE/region I-PMSIs though, so that the segmentation points can label-switch other traffic (not identified by those S-PMSIs). However, this moves the scaling

problem from the number of labels to the number of (C-S/*,C-G) routes in VRFs on the segmentation points.

3.3. Summary of Label Allocation Methods

In summary, labels can be allocated and advertised in the following ways:


1. A central entity allocates per-VPN/BD/ES labels from the DCB. PEs advertise the labels with an indication that they are from the DCB.
2. A central entity allocates per-VPN/BD/ES labels from a few common context-specific label spaces and allocates labels from the DCB to identify those context-specific label spaces. PEs advertise the VPN/BD labels along with the context-identifying labels.
3. A central entity assigns disjoint label blocks from a few context-specific label spaces to each PE and allocates labels from the DCB to identify those context-specific label spaces. A PE independently allocates a label for a segmented S-PMSI from its assigned label block and advertises the label along with the context-identifying label.

Option 1 is simplest, but it requires that all the PEs set aside a common label block for the DCB that is large enough for all the VPNs/BDs/ESs combined. Option 3 is needed only for segmented selective tunnels that are set up dynamically. Multiple options could be used in any combination, depending on the deployment situation.

4. Specifications

4.1. Context-Specific Label Space ID Extended Community

The Context-Specific Label Space ID Extended Community (EC) is a new Transitive Opaque EC with the following structure:



ID-Type: A 2-octet field that specifies the type of Label Space ID. In this document, the ID-Type is 0, indicating that the ID-Value field is a label.

ID-Value: A 4-octet field that specifies the value of the Label Space ID. When it is a label (with ID-Type 0), the most significant 20-bit portion is set to the label value.

This document introduces a DCB-flag (Bit 47 as assigned by IANA) in the "Additional PMSI Tunnel Attribute Flags" BGP Extended Community [RFC7902].

In the remainder of this document, when we say that a BGP-MVPN/EVPN A-D route carries a DCB-flag or has a DCB-flag attached to it, we mean the following:

- * The route carries a PTA and its Flags field has the Extension bit set, AND
- * The route carries an "Additional PMSI Tunnel Attribute Flags" EC

and its DCB-flag is set.

4.2. Procedures

The protocol and procedures specified in this section MAY be used when BIER or P2MP/MP2MP tunnel aggregation is used for an MVPN/EVPN or when BIER/P2MP/MP2MP tunnels are used with EVPN multihoming. When these procedures are used, all PE routers and segmentation points MUST support the procedures. How to ensure this behavior is outside the scope of this document.

Via methods outside the scope of this document, each VPN/BD/ES is assigned a label from the DCB or one of those few context-specific label spaces, and every PE that is part of the VPN/BD/ES is aware of the assignment. The ES label and the BD label MUST be assigned from the same label space. If PE Distinguisher Labels are used [RFC7582], they MUST be allocated from the same label space as well.

In the case of tunnel segmentation, each PE is also assigned a disjoint label block from one of those few context-specific label spaces, and it allocates labels for its segmented PMSI routes from its assigned label block.

When a PE originates/re-advertises an x-PMSI/IMET route, the route MUST carry a DCB-flag if and only if the label in its PTA is assigned from the DCB.

If the VPN/BD/ES/PMSI label is assigned from one of those few context-specific label spaces, a Context-Specific Label Space ID EC MUST be attached to the route. The ID-Type in the EC is set to 0, and the ID-Value is set to a label allocated from the DCB and identifies the context-specific label space. When an ingress PE sends traffic, it imposes the DCB label that identifies the context-specific label space after it imposes the label (that is advertised in the Label field of the PTA in the x-PMSI/IMET route) for the VPN/BD and/or the label (that is advertised in the ESI Label EC) for the ESI, and then imposes the encapsulation for the transport tunnel.

When a PE receives an x-PMSI/IMET route with the Context-Specific Label Space ID EC, it MUST place an entry in its default MPLS forwarding table to map the label in the EC to a corresponding context-specific label table. That table is used for the next label lookup for incoming data traffic with the label signaled in the EC.

Then, the receiving PE MUST place an entry for the label that is in the PTA or ESI Label EC in either the default MPLS forwarding table (if the route carries the DCB-flag) or the context-specific label table (if the Context-Specific Label Space ID EC is present) according to the x-PMSI/IMET route.

An x-PMSI/IMET route MUST NOT carry both the DCB-flag and the Context-Specific Label Space ID EC. A received route with both the DCB-flag set and the Context-Specific Label Space ID EC attached MUST be treated as withdrawn. If neither the DCB-flag nor the Context-Specific Label Space ID EC is attached, the label in the PTA or ESI Label EC MUST be treated as the upstream-assigned label from the label space of the source PE, and procedures provided in [RFC6514] and [RFC7432] MUST be followed.

If a PE originates two x-PMSI/IMET routes with the same tunnel, it MUST ensure that one of the following scenarios applies, so that the PE receiving the routes can correctly interpret the label that follows the tunnel encapsulation of data packets arriving via the tunnel:

- * They MUST all have the DCB-flag,

- * They MUST all carry the Context-Specific Label Space ID EC,
- * None of them have the DCB-flag, or
- * None of them carry the Context-Specific Label Space ID EC.

Otherwise, a receiving PE MUST treat the routes as if they were withdrawn.

5. Security Considerations

This document allows three methods (Section 3.3) of label allocation for MVPN PEs [RFC6514] or EVPN PEs [RFC7432] and specifies corresponding signaling and procedures. The first method (Option 1) is the equivalent of using common SRGBs [RFC8402] from the regular per-platform label space. The second method (Option 2) is the equivalent of using common SRGBs from a third-party label space [RFC5331]. The third method (Option 3) is a variation of the second in that the third-party label space is divided into disjoint blocks for use by different PEs, where each PE will use labels from its respective block to send traffic. In all cases, a receiving PE is able to identify one of the few label forwarding tables to forward incoming labeled traffic.

[RFC6514], [RFC7432], [RFC8402], and [RFC5331] do not list any security concerns related to label allocation methods, and this document does not introduce any new security concerns in this regard.

6. IANA Considerations

IANA has made the following assignments:

- * Bit 47 (DCB) in the "Additional PMSI Tunnel Attribute Flags" registry:

Bit Flag	Name	Reference
47	DCB	RFC 9573

Table 1

- * Sub-type 0x08 for "Context-Specific Label Space ID Extended Community" in the "Transitive Opaque Extended Community Sub-Types" registry:

Sub-Type Value	Name	Reference
0x08	Context-Specific Label Space ID Extended Community	RFC 9573

Table 2

IANA has created the "Context-Specific Label Space ID Type" registry within the "Border Gateway Protocol (BGP) Extended Communities" group of registries. The registration procedure is First Come First Served [RFC8126]. The initial assignment is as follows:

Type Value	Name	Reference
0	MPLS Label	RFC 9573

1-254	Unassigned	
255	Reserved	

Table 3

7. References

7.1. 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>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.
- [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>>.
- [RFC7524] Rekhter, Y., Rosen, E., Aggarwal, R., Morin, T., Grosclaude, I., Leymann, N., and S. Saad, "Inter-Area Point-to-Multipoint (P2MP) Segmented Label Switched Paths (LSPs)", RFC 7524, DOI 10.17487/RFC7524, May 2015, <<https://www.rfc-editor.org/info/rfc7524>>.
- [RFC7582] Rosen, E., Wijnands, IJ., Cai, Y., and A. Boers, "Multicast Virtual Private Network (MVPN): Using Bidirectional P-Tunnels", RFC 7582, DOI 10.17487/RFC7582, July 2015, <<https://www.rfc-editor.org/info/rfc7582>>.
- [RFC7902] Rosen, E. and T. Morin, "Registry and Extensions for P-Multicast Service Interface Tunnel Attribute Flags", RFC 7902, DOI 10.17487/RFC7902, June 2016, <<https://www.rfc-editor.org/info/rfc7902>>.
- [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>>.

7.2. Informative References

- [BIER-EVPN] Zhang, Z., Przygienda, A., Sajassi, A., and J. Rabadan, "EVPN BUM Using BIER", Work in Progress, Internet-Draft, draft-ietf-bier-evpn-14, 2 January 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-evpn-14>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.

2006, <<https://www.rfc-editor.org/info/rfc4364>>.

- [RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, DOI 10.17487/RFC5331, August 2008, <<https://www.rfc-editor.org/info/rfc5331>>.
- [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>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8556] Rosen, E., Ed., Sivakumar, M., Przygienda, T., Aldrin, S., and A. Dolganow, "Multicast VPN Using Bit Index Explicit Replication (BIER)", RFC 8556, DOI 10.17487/RFC8556, April 2019, <<https://www.rfc-editor.org/info/rfc8556>>.
- [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>>.
- [RFC9572] Zhang, Z., Lin, W., Rabadan, J., Patel, K., and A. Sajassi, "Updates to EVPN Broadcast, Unknown Unicast, or Multicast (BUM) Procedures", RFC 9572, DOI 10.17487/RFC9572, May 2024, <<https://www.rfc-editor.org/info/rfc9572>>.

Acknowledgements

The authors thank Stephane Litkowski, Ali Sajassi, and Jingrong Xie for their reviews of, comments on, and suggestions for this document.

Contributors

The following individual also contributed to this document:

Selvakumar Sivaraj
Juniper Networks
Email: ssivaraj@juniper.net

Authors' Addresses

Zhaohui Zhang
Juniper Networks
Email: zzhang@juniper.net

Eric Rosen
Individual
Email: erosen52@gmail.com

Wen Lin
Juniper Networks
Email: wlin@juniper.net

Zhenbin Li
Huawei Technologies
Email: lizhenbin@huawei.com

IJsbrand Wijnands
Individual
Email: ice@braindump.be