
Stream: Internet Engineering Task Force (IETF)

RFC: [9793](#)

Category: Standards Track

Published: May 2025

ISSN: 2070-1721

Authors:

X. Xu

China Mobile

M. Chen

Huawei

K. Patel

Arrcus, Inc.

IJ. Wijnands

Individual

T. Przygienda

Juniper

Z. Zhang, Ed.

Juniper

RFC 9793

BGP Extensions for Bit Index Explicit Replication (BIER)

Abstract

Bit Index Explicit Replication (BIER) is a multicast forwarding architecture that doesn't require an explicit tree-building protocol and doesn't require intermediate routers to maintain per-tree multicast states. Some BIER-specific information and states, which are only in proportion to the number of BIER routers but not per-tree, do need to be advertised, calculated, and maintained. This document describes BGP extensions for advertising the BIER information and methods for calculating BIER states based on the advertisements.

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/rfc9793>.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions

with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
2.1. Requirements Language	3
3. BIER Path Attribute	3
3.1. BIER MPLS Encapsulation Sub-TLV	5
3.2. BIER Non-MPLS Encapsulation Sub-TLV	6
3.3. BIER Nexthop Sub-TLV	7
4. Originating/Propagating/Updating the BIER Attribute	7
5. BIFT Calculation with BGP Signaling	8
6. Example of BIER Nexthop Usage and Handling	9
7. Operational Considerations	9
8. IANA Considerations	10
9. Security Considerations	11
10. References	11
10.1. Normative References	11
10.2. Informative References	11
Acknowledgements	12
Contributors	12
Authors' Addresses	12

1. Introduction

Bit Index Explicit Replication (BIER) [RFC8279] is a multicast forwarding architecture that doesn't require an explicit tree-building protocol and doesn't require intermediate routers to maintain per-tree multicast states. It supports both direct and tunneled BIER forwarding. This document describes BGP extensions for advertising the BIER information and methods for calculating BIER states based on the advertisements. More specifically, in this document, we

define a new optional transitive BGP attribute, referred to as the "BIER attribute", to convey the BIER-specific information such as Bit-Forwarding Router Identifier (BFR-ID), BitStringLength (BSL), and so on. The signaling is to be used in a single Administrative Domain (AD), and [Section 7](#) specifies procedures to prevent the BIER attribute from "leaking out" of the domain.

2. Terminology

This document makes use of the terminology defined in [\[RFC4271\]](#) and [\[RFC8279\]](#). Some terms are listed below for convenience.

BIER: Bit Indexed Explicit Replication

BFR: Bit-Forwarding Router

BFR-ID: BFR Identifier

BSL: BitStringLength

BIFT: Bit Index Forwarding Table

BIFT-id: Bit Index Forwarding Table Identifier

BFER: Bit-Forwarding Egress Router

BFR-prefix: Each BFR is assigned a single "BFR-prefix" for each sub-domain to which it belongs. It is recommended that the BFR-prefix be a loopback address of the BFR.

NLRI: Network Layer Reachability Information [\[RFC4271\]](#)

AFI: Address Family Identifier [\[RFC4760\]](#)

SAFI: Subsequent Address Family Identifier [\[RFC4760\]](#)

2.1. Requirements Language

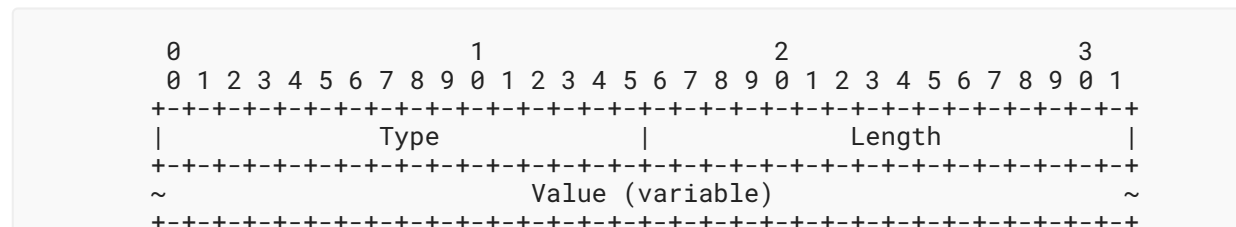
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [\[RFC2119\]](#) [\[RFC8174\]](#) when, and only when, they appear in all capitals, as shown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

3. BIER Path Attribute

This specification defines an optional, transitive BGP path attribute, referred to as the "BIER attribute". This attribute can be attached to a BGP UPDATE message by the originator for NLRIs of AFI 1 or 2 and SAFI 1, 2, or 4 to indicate the BIER-specific information of a particular BFR identified by the /32 (for IPv4) or /128 (for IPv6) host address prefix contained in the NLRI. The attachment of the BIER attribute to non-host address prefixes is not defined by this document. It may be specified in the future, for example, by [\[BIER-Prefix-Redistribute\]](#).

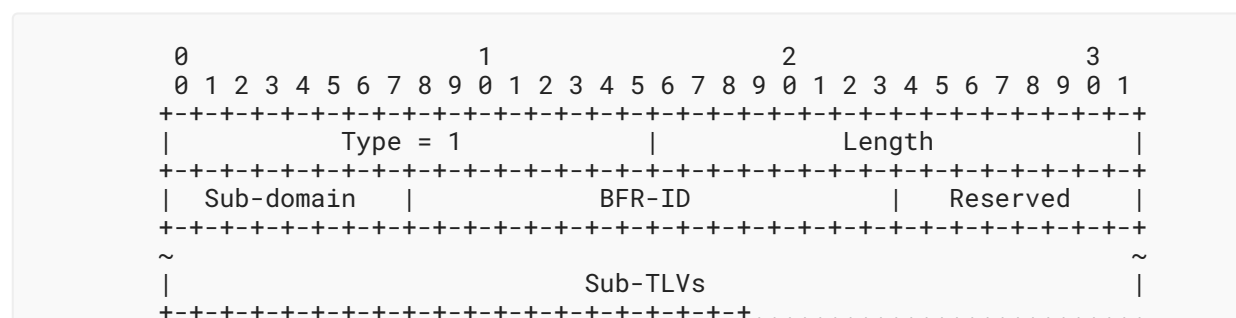
If the BIER path attribute is present, the NLRI is referred to as a "BFR-prefix". Use of the attribute with other AFIs/SAFIs is outside the scope of this document.

The BIER path attribute is an optional, transitive BGP path attribute with type code 41 and of variable length. The attribute value portion carries BIER TLVs, which are encoded as follows:



The Length field defines the length of the value portion in octets (thus, a TLV with no value portion would have a length of zero). The TLV is not padded to 4-octet alignment. Unknown and unsupported types **MUST** be preserved and propagated within the BIER attribute. The presence of unknown or unexpected TLVs **MUST NOT** result in the NLRI or the BIER attribute being considered malformed.

When creating a BIER attribute, a BFR **MUST** include one BIER TLV for every sub-domain that the prefix belongs to. The attribute type code for the BIER attribute is 41. The value field of the BIER attribute contains one or more BIER TLVs as shown below:



Type: 1

Length: 2 octets encoding the length in octets of the Value part.

Sub-domain: A 1-octet field encoding the sub-domain ID corresponding to the BFR-ID (see [RFC8279]).

BFR-ID: A 2-octet field encoding the BFR-ID (see [RFC8279]).

Reserved: **SHOULD** be set to 0 on transmission and **MUST** be ignored on reception.

Sub-TLVs: Contains one or more sub-TLVs.


The BIER TLV **MAY** appear multiple times in the BIER path attribute, one for each sub-domain. There **MUST** be no more than one BIER TLV with the same Sub-domain value; if there is, the entire BIER path attribute **MUST** be ignored.

A BIER TLV may have sub-TLVs, which may have their own sub-TLVs. All those are referred to as sub-TLVs and share the same Type space, regardless of the level.

3.1. BIER MPLS Encapsulation Sub-TLV

The BIER MPLS Encapsulation sub-TLV has the following format. It **MAY** appear multiple times in the BIER TLV.

The BIER MPLS Encapsulation sub-TLV has the following format:



Type: 2

Length: 2 octets encoding the length in octets of the Value part. The value is 4 or other (depending on sub-TLVs).

Max SI: A 1-octet field encoding the Maximum Set Identifier (see [Section 1](#) of [\[RFC8279\]](#)) used in the encapsulation for this BIER sub-domain for this BitString length.

BS Len: BitString Length. A 4-bit field encoding the supported BitString length associated with this BFR-prefix. The values allowed in this field are specified in [Section 2](#) of [\[RFC8296\]](#).

Label: A 20-bit value representing the first label in the label range.

The "label range" is the set of labels beginning with the Label and ending with (Label + (Max SI)). A unique label range is allocated for each BitString length and sub-domain-id. These labels are used for BIER forwarding, as described in [\[RFC8279\]](#) and [\[RFC8296\]](#).

The size of the label range is determined by the number of SIs ([Section 1](#) of [\[RFC8279\]](#)) that are used in the network. Each SI maps to a single label in the label range: the first label is for SI=0, the second label is for SI=1, etc.


If the label associated with the Maximum SI exceeds the 20-bit range, the BIER MPLS Encapsulation sub-TLV containing the error **MUST** be ignored.

If the same BitString length is repeated in multiple BIER MPLS Encapsulation sub-TLVs inside the same BIER TLV, all BIER MPLS Encapsulation sub-TLVs in the BIER TLV **MUST** be ignored.

Label ranges within all BIER MPLS Encapsulation sub-TLVs advertised by the same BFR **MUST NOT** overlap. If an overlap is detected, all BIER MPLS Encapsulation sub-TLVs advertised by the BFR **MUST** be ignored.

3.2. BIER Non-MPLS Encapsulation Sub-TLV

The BIER non-MPLS Encapsulation sub-TLV is used for non-MPLS encapsulation and has the following format. It **MAY** appear multiple times within a single BIER TLV. If the same BitString length is repeated in multiple BIER non-MPLS Encapsulation sub-TLVs inside the same BIER TLV, the BIER TLV **MUST** be ignored.



Type: 3

Length: 2 octets encoding the length in octets of the Value part. The value is 4 or other (depending on sub-TLVs).

Max SI: A 1-octet field encoding the Maximum Set Identifier (Section 1 of [RFC8279]) used in the encapsulation for this BIER sub-domain for this BitString length. The first BIFT-id is for SI=0, the second BIFT-id is for SI=1, etc. If the BIFT-id associated with the Maximum SI exceeds the 20-bit range, the sub-TLV **MUST** be ignored.

BS Len: BitString Length. A 4-bit field encoding the BitString length (as per [RFC8296]) supported for the encapsulation.

BIFT-id: A 20-bit field representing the first BIFT-id in the BIFT-id range.

The "BIFT-id range" is the set of 20-bit values beginning with the BIFT-id and ending with (BIFT-id + (Max SI)). These BIFT-ids are used for BIER forwarding, as described in [\[RFC8279\]](#) and [\[RFC8296\]](#).

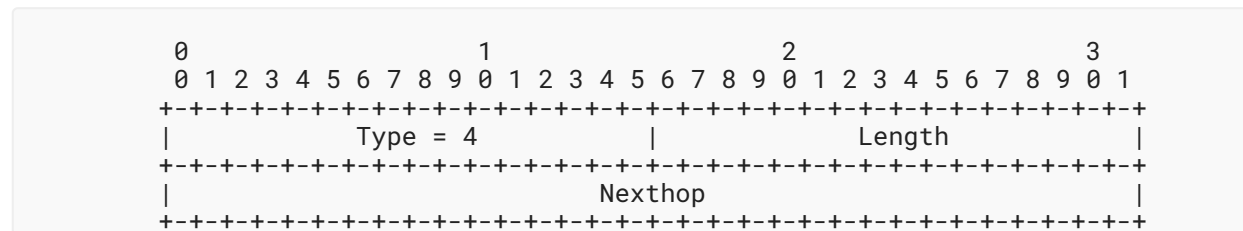
The size of the BIFT-id range is determined by the number of SIs (Section 1 of [RFC8279]) that are used in the network. Each SI maps to a single BIFT-id in the BIFT-id range: the first BIFT-id is for SI=0, the second BIFT-id is for SI=1, etc.

If the BIFT-id associated with the Maximum SI exceeds the 20-bit range, the BIER non-MPLS Encapsulation sub-TLV containing the error **MUST** be ignored.

BIFT-id ranges within all the BIER non-MPLS Encapsulation sub-TLVs advertised by the same BFR **MUST NOT** overlap. If an overlap is detected, all the BIER non-MPLS Encapsulation sub-TLVs advertised by the BFR **MUST** be ignored. However, the BIFT-id ranges may overlap across different encapsulation types and that is allowed. As an example, the BIFT-id value in the non-MPLS Encapsulation sub-TLV may overlap with the Label value in the Label range in the BIER MPLS Encapsulation sub-TLV.

3.3. BIER Nexthop Sub-TLV

The BIER Nexthop sub-TLV **MAY** be included, and it **MUST NOT** be included more than once in each of the MPLS or non-MPLS Encapsulation sub-TLVs or in the top-level BIER TLV. It is used when calculating BIFT entries, as described in [Section 5](#) and illustrated in [Section 6](#).



Type: 4

Length: 2 octets. The value is 4 if the Nexthop is an IPv4 address and 16 if the Nexthop is an IPv6 address.

Nexthop: 4 or 16 octets of an IPv4/IPv6 address.

4. Originating/Propagating/Updating the BIER Attribute

A Bit-Forwarding Egress Router (BFER) **MUST** attach a BIER attribute to its own /32 (for IPv4) or /128 (for IPv6) host BFR-prefix NLRI. The BIER attribute **MUST** include one BIER TLV for each BIER sub-domain that it supports. Each BIER TLV **MUST** include an MPLS and/or non-MPLS Encapsulation sub-TLV and **MAY** include a BIER Nexthop sub-TLV with the Nexthop set to the BIER prefix. If the BIER Nexthop sub-TLV is not included, the BIER prefix will be used by receiving BFRs as the BIER next hop when calculating BIFT.

When a BFR receives an update with the BIER path attribute, the attribute is parsed with the following validations:

- Syntactic checking based on the Length field of TLVs and sub-TLVs:
 - The total length of BIER TLVs (including the Type and Length fields) **MUST** be equal to the BIER path attribute length.
 - The total length of sub-TLVs (including the Type and Length fields) of a TLV **MUST** be equal to the length of the TLV.
- Semantic checking as per [Section 3](#).

If the syntactic checking fails, the attribute is considered malformed and the "attribute discard" action [[RFC7606](#)] for the BIER attribute **MUST** be taken. If the semantic checking passes, BIFT entries are calculated as described in [Section 5](#). Otherwise (i.e., if semantic checking fails), some or all BIER TLVs are ignored, per the rules given in [Section 3](#), and if the remaining data permits, BIFT entries are calculated per [Section 5](#).

When a BFR re-advertises a BGP NLRI with a BIER attribute, for the sub-domains that this BFR supports, in the corresponding BIER TLV, it **SHOULD** set/update the BIER Nexthop sub-TLV to use its own BIER prefix; in which case, it **MUST** replace the MPLS or non-MPLS Encapsulation sub-TLV with its own, i.e., as if the BFR is attaching the encapsulation sub-TLV for its own BIER prefix. If it does not update the BIER Nexthop sub-TLVs, it **MUST NOT** update the MPLS or non-MPLS Encapsulation sub-TLV. If it does not support a sub-domain, it **MUST NOT** update the corresponding BIER TLV.

It's possible that the BFR supports some but not all BitStringLengths (BSLs) in the received MPLS or non-MPLS Encapsulation sub-TLVs. After setting/updating the BIER Nexthop sub-TLV in the top BIER TLV to itself, for the BSLs that it does support, the BFR **MUST** remove the BIER Nexthop sub-TLV (if present) in the corresponding MPLS Encapsulation sub-TLVs. For the BSLs that it does not support:

- If a BIER Nexthop sub-TLV is included in the MPLS Encapsulation sub-TLV, it **MUST NOT** be updated.
- Otherwise, if a BIER Nexthop sub-TLV is included in the received BIER TLV, its original value (before changed for supported BSLs by this BFR) **MUST** be copied into the MPLS Encapsulation sub-TLV.
- Otherwise, a BIER Nexthop sub-TLV **MUST** be added to the MPLS Encapsulation sub-TLV with its value set to the BFR-prefix.

All impacted Length fields (e.g., the MPLS Encapsulation sub-TLV Length and the top-level BIER TLV Length) **MUST** be updated accordingly.

Since the BIER attribute is an optional, transitive BGP path attribute, a non-BFR BGP speaker could still re-advertise the received route with a BIER attribute.

Two different BFR-prefixes **MUST NOT** have the same non-zero BFR-ID in the same sub-domain. If a duplication is detected, the receiving BFR **MUST NOT** use the BFR-prefixes with the same BFR-ID for BIFT calculation for the sub-domain and an error **SHOULD** be logged.

5. BIFT Calculation with BGP Signaling

As pointed out in [RFC8279], BIFTs are derived from the unicast FIB by adding BIER-specific information.

For each sub-domain, a BFR calculates the corresponding BIFTs by going through the BIER prefixes whose BIER attribute includes a BIER TLV for the sub-domain. For a non-zero BFR-id in the BIER TLV, a BIFT entry is created or updated. The entry's BFR Neighbor (BFR-NBR) [RFC8279] is the Nexthop in the BIER Nexthop sub-TLV in the corresponding MPLS Encapsulation sub-TLV or in the top-level BIER TLV if the MPLS Encapsulation sub-TLV does not have a BIER Nexthop sub-TLV. If there is no BIER Nexthop sub-TLV at all, the entry's BFR-NBR is the BIER prefix itself. The BIER label or BIFT-id for the entry is derived from the label range in the MPLS Encapsulation sub-TLV or from the BIFT-id range in the non-MPLS Encapsulation sub-TLV.

BIER traffic is sent to the BFR-NBR either directly (BIER header directly follows a Layer 2 header) if the BFR-NBR is directly connected or via a tunnel. Notice that, if a non-BFR BGP speaker re-advertises a BIER prefix (in this case, it cannot update the BIER attribute since it is not capable), or if a BFR BGP speaker re-advertises a BIER prefix without updating the BIER Nexthop sub-TLV, the BFR receiving the prefix will tunnel BIER traffic -- the BGP speaker re-advertising the BIER prefix will not see the BIER traffic for the BIER prefix.

How the tunnel is set up and chosen is outside the scope of this document. It can be any kind of tunnel, e.g., MPLS Label Switched Path or IP/GRE, as long as the tunnel header can indicate that the payload is BIER.

6. Example of BIER Nexthop Usage and Handling

Consider a simple topology as follows:



The BFER1/2/3 each advertises a route for its loopback address with a BIER path attribute, listing one BIER TLV for each sub-domain that it is in, with a non-zero BFR-ID and an MPLS Encapsulation sub-TLV. A BIER Nexthop sub-TLV is not included in the one from BFER1 but is included in the ones from BFER2/3. The BIER Nexthop sub-TLV encodes the BFR-prefix of BFER2 and BFER3, respectively.

When BFR2 receives the route, it calculates its BIFT entries. Because the route from BFER1 does not include a BIER Nexthop, BFR2 uses BFR1's BFR-prefix as the next hop.

When BFR2 re-advertises the routes to the non-BFR, it adds a BIER Nexthop sub-TLV to the BFER1 route and updates the BIER Nexthop sub-TLV in the BFER2/3 routes, all encoding BFR2's own address. It also updates the MPLS Encapsulation sub-TLV to encode its own labels.

When the non-BFR receives the routes, since it does not support BIER, no BIER-specific action is taken and the routes are re-advertised to BFR1 with the BIER path attribute unchanged.

When BFR1 receives the routes, it calculates the BIFT entries, using BFR2's address encoded in the BIER Nexthop sub-TLV as the next hop. Because BFR2 is not directly connected, a tunnel must be used.

7. Operational Considerations

In this document, it is assumed that the BIER domain [[RFC8279](#)] is aligned with an Administrative Domain (AD), which may be composed of multiple Autonomous Systems. Use of the BIER attribute in other scenarios is outside the scope of this document.

BFR-prefixes are typically loopback addresses on the BFRs. They are distributed throughout the AD, but they do not need to be distributed outside the AD for the BIER's purposes. This is analogous to the Provider Edge router's loopback addresses that are distributed inside the AD, but they do not need to be distributed outside the AD.

If prefixes are distributed outside of the AD with the BIER attribute attached and the neighboring AD also deploying BIER, then the two BIER domains, which should be independent of each other, may be incorrectly joined together and most likely have conflicting configurations, causing security risks and operational troubles.

To prevent that, a boundary router of the AD that supports the BIER attribute **MUST** support a policy based on an External BGP (EBGP) session/group that indicates whether the attribute is allowed; by default, it is NOT allowed. If it is not allowed, the BIER attribute **MUST NOT** be sent to any EBGP peer of the session/group. If a BIER attribute is received from the peer, it **MUST** be treated exactly as if it were an unrecognized non-transitive attribute. That is, it **MUST** be quietly ignored and not passed along to other BGP peers.

8. IANA Considerations

IANA has assigned codepoint 41 to the BIER attribute in the "BGP Path Attributes" registry <<https://www.iana.org/assignments/bgp-parameters>> as follows:

Value	Code	Reference
41	BIER	RFC 9793

Table 1

IANA has created the "BGP BIER TLV and Sub-TLV Types" registry within the "Border Gateway Protocol (BGP) Parameters" registry group. The type field for the registry consists of 2 octets, with possible values from 0 to 65535 (the value 0 is reserved). The allocation policy for this field is First Come First Served [RFC8126].

The five initial values have been allocated as follows:

Value	Name	Reference
0	Reserved	RFC 9793
1	BIER TLV	RFC 9793
2	MPLS Encapsulation sub-TLV	RFC 9793
3	non-MPLS Encapsulation sub-TLV	RFC 9793
4	BIER Nexthop sub-TLV	RFC 9793

Value	Name	Reference
5-65535	Unassigned	

Table 2

9. Security Considerations

This document introduces no new security considerations beyond those already discussed in [RFC4271], [RFC8279], and the operational considerations (Section 7) of this document.

10. References

10.1. 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>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [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>>.
- [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>>.
- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", RFC 8296, DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.

10.2. Informative References

- [BIER-Prefix-Redistribute] Zhang, Z., Wu, B., Zhang, Z. J., Wijnands, I., Liu, Y., and H. Bidgoli, "BIER Prefix Redistribute", Work in Progress, Internet-Draft, draft-ietf-bier-prefix-redistribute-08, 23 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-prefix-redistribute-08>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.

[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>>.

[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>>.

Acknowledgements

Thanks to Eric Rosen and Peter Psenak for their valuable comments on this document.

Contributors

This document has the following contributor:

Zheng (Sandy) Zhang

ZTE

Email: zhang.zheng@zte.com.cn

Authors' Addresses

Xiaohu Xu

China Mobile

Email: xuxiaohu@cmss.chinamobile.com

Mach(Guoyi) Chen

Huawei

Email: mach.chen@huawei.com

Keyur Patel

Arrcus, Inc.

Email: keyur@arrcus.com

IJsbrand Wijnands

Individual

Email: ice@braindump.be

Tony Przygienda

Juniper

Email: prz@juniper.net

Zhaohui Zhang (EDITOR)

Juniper

Email: zzhang@juniper.net