

BIER
Internet-Draft
Intended status: Standards Track
Expires: January 13, 2021

Z. Zhang
ZTE Corporation
B. Wu
Individual
Z. Zhang
Juniper Networks
IJ. Wijnands
Cisco Systems, Inc.
Y. Liu
China Mobile
July 12, 2020

BIER Prefix Redistribute
draft-zwzw-bier-prefix-redistribute-07

Abstract

This document defines a BIER proxy function to interconnect different underlay routing protocol areas in a network. And a new BIER proxy range sub-TLV is also defined to convey BIER BFR-id information across the routing areas.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119](#).

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 January 13, 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.](#) Problem statement [2](#)
- [2.](#) Proposal [5](#)
- [3.](#) Advertisement [6](#)
 - [3.1.](#) BIER proxy range sub-TLV [7](#)
 - [3.2.](#) Proxy range sub-TLV usage [9](#)
- [4.](#) IANA Considerations [10](#)
- [5.](#) Security Considerations [10](#)
- [6.](#) Acknowledgements [10](#)
- [7.](#) References [10](#)
 - [7.1.](#) Normative References [10](#)
 - [7.2.](#) Informative References [11](#)
- Authors' Addresses [12](#)

[1.](#) Problem statement

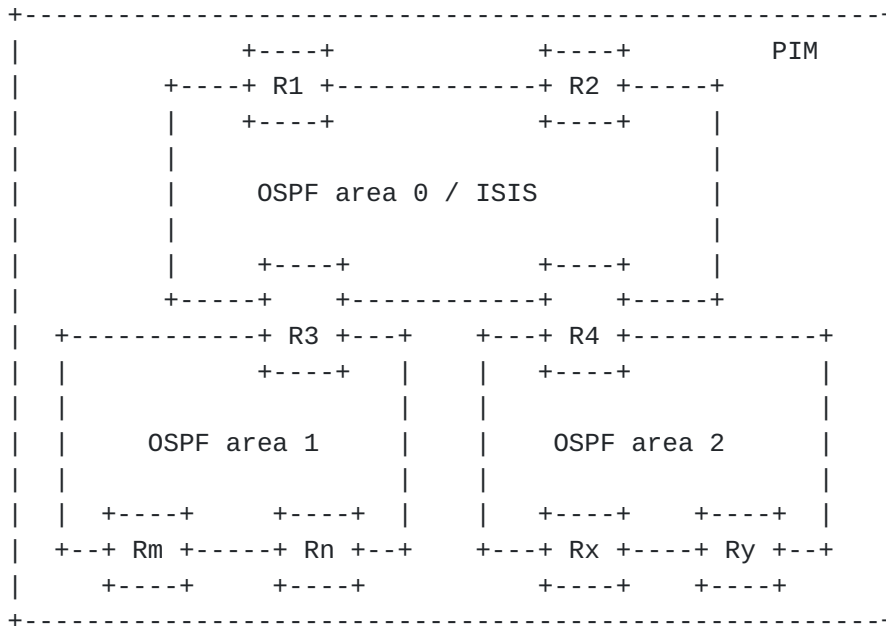


Figure 1

Figure 1 shows that there are three areas in a traditional network. In some deployment situations, different routing protocols may also be used in the network. There are just small amount of routers in each area. Currently, multicast services are provided in this kind of network by using protocol independent feature of PIM [\[RFC7761\]](#).

BIER could be a candidate multicast protocol to replace PIM to reduce multicast states in the network. BIER [\[RFC8279\]](#) is a new architecture for the forwarding of multicast data packets. It does not require a protocol for explicitly building multicast distribution trees, nor does it require intermediate nodes to maintain any per-flow state. In order to build BIER forwarding plane, BIER key parameters must be flooded in one BIER domain such as BFR-prefix, BFR-id, subdomain-id, and so on. The routing protocols which are used to flood these BIER parameters are called BIER routing underlay. The associated routing protocol extensions are defined in documents such as [\[RFC8401\]](#), [\[RFC8444\]](#), [\[I-D.ietf-bier-idr-extensions\]](#), [\[I-D.ietf-bier-ospfv3-extensions\]](#), and so on.

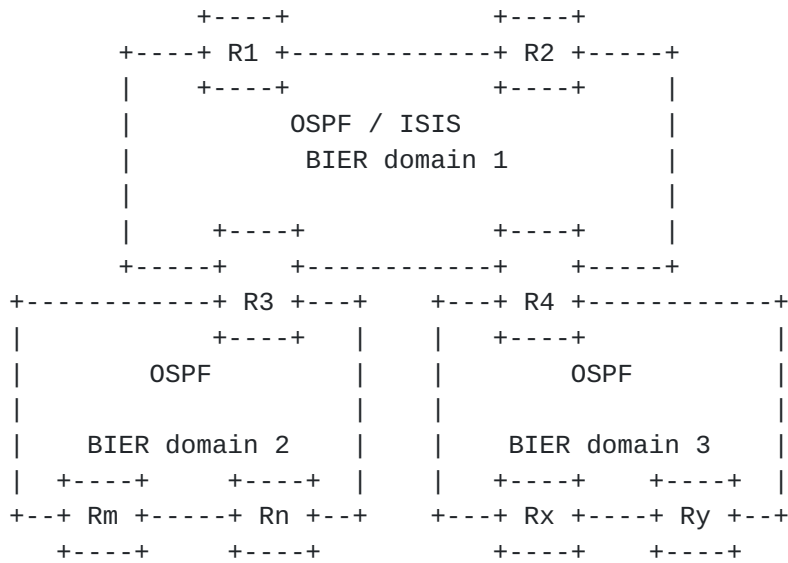


Figure 2

Based on the BIER design, a BIER domain is limited by the underlay routing protocols flooding scope. In case we want deploy BIER instead of PIM, there are several BIER domains because of different underlay routing areas limitation. Multiple encapsulating/decapsulation executions are needed to across multiple BIER domains. These executions slow down the forwarding efficiency. The border routers also need to maintain overlay state, which is undesired.

For example in figure 2, suppose that R1 and R2 connect with multicast source. Rm, Rn, Rx and Ry connect with multicast receivers. According the existed advertisement method defined in [RFC8401], [RFC8444], and so on, in BIER domain 1, R1, R2, R3 and R4 are BIER edge routers. In BIER domain 2, R3 and Rm, Rn are BIER edge routers. In BIER domain 3, R4 and Rx, Ry are BIER edge routers.

R1/R2 encapsulates BIER packet when multicast flows come into this BIER domain, R3/R4 decapsulates the BIER packet, and encapsulates them again, sends them into BIER domain 2/3. Rm, Rn, Rx and Ry decapsulates the packets and forwards them to receivers.

Due to the decapsulation and encapsulation execution in R3 and R4, the forwarding efficiency is slowed down, especially when there are not large amount of routers in each BIER domain.

[Section 2.3 in \[RFC8444\]](#) defines the duplication function across OSPF areas. Except the homogenization network, there is the hybrid network showed in figure 2 that several areas formed by different IGP protocols, and they are need to be merged into one BIER domain. In the hybrid network, necessary advertisement transform need to be

used. And further, necessary optimization method can be used to reduce the amount of the advertisement items.

2. Proposal

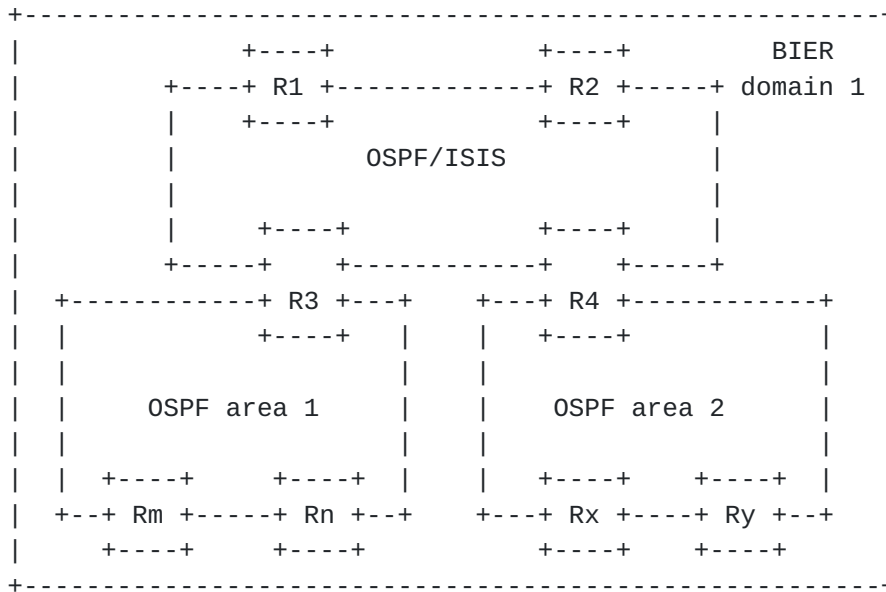


Figure 3

It is more efficient to deploy BIER by creating one BIER domain for the hybrid network to achieve forwarding benefit.

Since the limitation of the BIER routing protocol scope, BFR-id is confined to only one routing area. A BIER proxy function is introduced to transport BIER BFR-id information in a BIER domain across multiple routing protocol areas. So BIER forwarding tables can be built across multiple underlay routing protocols to replace encapsulation/ decapsulation processing. In the current deployment, border router (ABR) has a similar role, ABR summaries unicast routing information from one routing protocol area and sends it to another routing area by new routing protocol messages. So ABR can implement BIER proxy function to summarize BIER BFR-id information from one routing protocol area and sends it to another routing area.

In figure 3, R3/R4 connects two areas which running same or different routing protocols, they can be used as BIER proxies to transport BIER information. For example, after R3 receives BFR-ids information from OSPF area 1 and sends it to ISIS routing area (the upper area), the routers in ISIS routing area can generate BIER forwarding items toward the BFR-ids in OSPF area 1 through R3. Similarly, R3 receives BFR-ids information from the upper area and sends them into OSPF area 1, the routers in OSPF area 1 can build BIER forwarding items toward

the BFR-ids in ISIS area. R4 does the same function, the BIER forwarding plane is constructed accordingly.

For example, in this network, suppose that Rm and Rn have the prefix of 201.1.1.1/32, 201.1.1.2/32. In order to build one BIER domain which includes these three IGP areas, R3 advertises the BFR-ids of Rm/Rn with associated prefixes (201.1.1.1/32, 201.1.1.2/32) into the upper area. Similarly, R4 advertises the BFR-ids of Rx/Ry with associated prefixes (202.1.1.1/32, 202.1.1.2/32) into the upper area too.

And R3/R4 advertises the prefixes of R1 and R2 (suppose that the prefixes are 200.1.1.1/32 and 200.1.1.2/32) with associated BFR-ids into IGP area 1 and area 2. Also, R3 advertises the prefixes learned from R4 (202.1.1.1/32, 202.1.1.2/32) with associated BFR-ids into IGP area 1. R4 also advertises the prefixes (201.1.1.1/32, 201.1.1.2/32) with associated BFR-ids into IGP area 2.

After the path calculation, the BIER forwarding plane is built. When R1/R2 receives multicast packet which should be sent to Rm/Rn/Rx/Ry, R1/R2 encapsulates the packet with BIER header and send it into the upper area. When R3/R4 receives the packet, R3/R4 forwards the packet into IGP area 1 and area 2 according to the BIER forwarding table without doing the decapsulation and re-encapsulation actions.

Obviously, in order to build the large BIER domain, the BFR-id of edge router in each IGP area MUST NOT be overlapped.

3. Advertisement

According to [[RFC8279](#)], each BFER needs to have a unique (in each sub-domain) BFR-id, and each BFR and BFER floods itself BIER info sub-TLV and associated sub-sub-TLVs in the BIER domain. To keep consistent with the definition in [[RFC8444](#)], [[I-D.ietf-bier-ospfv3-extensions](#)], and [[RFC8401](#)], BIER info sub-TLV defined in [[RFC8401](#)] and BIER sub-TLV defined in [[RFC8444](#)], [[I-D.ietf-bier-ospfv3-extensions](#)] is reused to convey the BFR-id information. OSPF extended Prefix Opaque LSA [[RFC7684](#)] and TLVs 235, 237 defined in [[RFC5120](#)], or TLVs 135 [[RFC5305](#)], or TLV 236 [[RFC5308](#)] are still used to carry the BFR-id/ BFR-prefix information, etc.

The key parameters got from the original routing protocol should be adapted to the format of next routing protocol, such as BFR-prefix, BFR-id, subdomain-id, and so on. Some parameters like BAR, MT-ID has local significance, So they should be set to same values with BIER proxy own advertisement when BIER proxy advertise them to the next routing area.

And as the two BIER info sub-sub-TLVs (sub-TLVs) including MPLS encapsulation and BSL conversion also have local significance. The information carried in these two sub-sub-TLV need not, but MAY, be advertised to next routing area.

In the same example, when R3 advertises the prefixes of Rm and Rn into the upper area, R3 may advertise the prefixes one by one, that is R3 advertises 201.1.1.1/32 with associated BFR-id, and R3 advertises 201.1.1.2/32 with associated BFR-id. If there is dozens of edge routers in area 1, R3 advertises dozens of prefixes and associated BFR-ids into the upper area. When R3 advertises the prefixes from the upper area into area 1, R3 advertises the prefixes of R1 and R2 with associated BFR-ids separately, and R3 advertises the prefixes of Rx and Ry which come from R4 into area 1 one by one. R4 does it as well.

3.1. BIER proxy range sub-TLV

In some cases unicast default route and aggregated/ summarized routes are used in some routing areas and routers in next area can not see the specific BFR-prefix routes from original area. Like in figure 3, in case R3/R4 does not advertise specific ISIS unicast routes to OSPF area and only advertises unicast default route or aggregated/ summarized route to OSPF area 1/2, when R3/R4 advertises BIER info sub-TLV to OSPF area 1/2, R3 MUST advertise the prefix with default route or aggregated/ summarized route. In that case, multiple BFR-ids will be mapped to one prefix. In order to advertise BFR-ids optimally, we define a new BIER proxy range sub-TLV to advertise the information of BFR-ids.

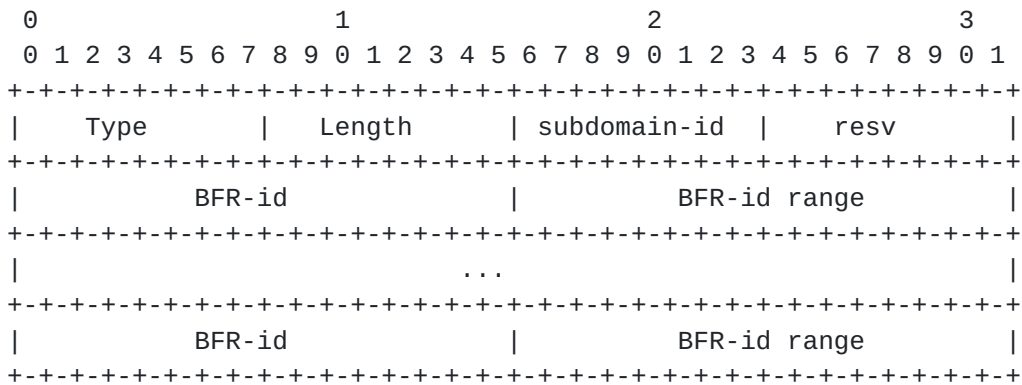


figure 4

- o Type: 8-bit unsigned integer. TBD to indicate the BIER proxy range sub-TLV.
- o Length: 8-bit unsigned integer. Length of the BIER proxy range sub-TLV in 4-octet units, not including the first 4 octets.

- o Subdomain-id: 8-bit unsigned integer. The subdomain-id from original advertisement.
- o resv: 8-bit unsigned integer. The reserved field.
- o BFR-id: 16-bit unsigned integer. The first BFR-id from original advertisement.
- o BFR-id range: 16-bit unsigned integer. The range of BFR-ids with one subdomain-id.

The BIER proxy range sub-TLV is attached to the aggregated/ summarized route prefix or default route prefix. The summarized/ aggregated/ default prefix may need multiple BIER proxy range sub-TLVs when the BFR-ids covered by the prefix are allocated from different ranges (even if they're from a single range but some BFR-ids in the range map to the BIER prefixes that are covered by a different summarized/ aggregated prefix, then that single large range needs to be broken into smaller ranges).

When a BFR receives a default/summary route with a BIER proxy range sub-TLV, it builds a BIRT route with that default/summary prefix. It also builds multiple BIFT entries for each BFR-IDs covered in the proxy range sub-TLV, using the same forwarding information as in the BIRT route. It is possible that a BFR-ID is covered in the proxy range sub-TLV of multiple default/summary routes. In that case, ECMP can be used and longest match SHOULD be used. For example, if ABR1 advertises default/summary route P1 while ABR2 advertises a more specific summary route P2 and both have a proxy range sub-TLV that covers BFR-ID 100, then the BIFT entry for BFR-ID 100 only follows the P2 route in BIRT.

The proxy range sub-TLV can also be attached to a host BIER prefix itself. Consider the situation where BGP-LU [[RFC8277](#)] is used for a seamless MPLS [[I-D.ietf-mpls-seamless-mpls](#)] environment. An ABR/ASBR advertises BIER prefixes via BGP over an area/AS to other ABR/ASBRs but the BIER prefixes are not advertised into the IGP for the area/AS. The ABR/ASBR does advertise the BIER prefix for itself into the IGP for the area/AS, with a BIER proxy range sub-TLV to cover the BFR-IDs for BFRs that the ABR/ASBR has learned (either through IGP or through BGP signaling). When an internal BFR receives it, it treats as if a default summary route had been received with the proxy range sub-TLV. Note that this imaginary default summary route is only for the purpose of building BIRT/BIFT entries and not used for unicast forwarding.

With this scheme, even though the BIER prefixes are not advertised into the IGP for the area/AS and unicast traffic for those BIER

prefixes are tunneled through, corresponding BIFT entries are maintained inside the area/AS for the purpose of efficient BIER forwarding. Otherwise, BIER forwarding through the area/AS would be tunneled just like unicast case.

The range in the BIER proxy range sub-TLV can be as granular as to advertise individual BFR-ids. Though a larger range can increase advertisement efficiency, it requires careful planning for BFR-id assignment.

When the proxy range sub-TLV is used, the mapping between a BIER prefix and its BFR-id is no longer conveyed in the routing underlay. As a result, the mapping must be provided by other means, e.g. in the multicast overlay.

3.2. Proxy range sub-TLV usage

In the same example of figure 3, in case there are 40 edge routers in area 1, the BFR-ids of area 1 start from 51 to 90, and the prefixes of these routers start from 201.1.1.1/32 to 201.1.1.40/32. These prefixes are not overlapped with the prefixes in any other area.

When R3 advertises these prefixes into the upper area, the proxy range sub-TLV can be used to optimize the advertisement. That is R3 can advertise only one prefix 201.1.1.0/24, with the BFR-id set to 51, the BFR-id range set to 40, into the upper area. Then the BIER overlay is built among R1, R2, Rm, Rn, Rx and Ry. R3 and R4 need not maintain the multicast overlay states.

When R3 advertises the prefixes from the upper area and area 2 into area 1, R3 may advertise only one default route (0.0.0.0/0) into area 1 if one or more continuous BFR-id range can be attached. Suppose that the BFR-id in the upper area starts from 1001 to 1050, the BFR-id in area 2 starts from 201 to 250, and these ranges are not overlapped with the ranges in any other area. Suppose that the sub-domain ID is 1, the BIER proxy range sub-TLV may be advertised like this:




figure 5

The optimized summary/ aggregated or default prefix can be generated by the operation policy which is configured by the network administrator.

In case the range of BFR-ids in one area is overlapped with the BFR-ids in any other area, the proxy range sub-TLV can not be used. In the same example above, if the BFR-ids in area 1 are 21, 31, 41, etc., the BFR-ids in area 2 are 22, 32, 42, etc., even if the summarized prefixes are not overlapped with the prefixes in any other area, when R3 advertises the summarized prefixes in area 1 into the upper area, the proxy range sub-TLV may not optimize the advertisement.

After the forwarding plane is built, when R1 receives multicast packet, and the receivers of this flow are connected by Rm and Rx, R1 encapsulates BIER header in front of the flow packet with BFR-ids set to Rm and Rx. When R3/R4 receives the packet, R3/R4 need not decapsulate and re-encapsulate the packet. R3/R4 just forwards the packet according to the BIER forwarding table. When the packet reaches Rm and Rx, Rm and Rx remove the BIER header and forward it to receivers.

4. IANA Considerations

IANA is requested to set up a new types of sub-TLV (TLV) registry value for BIER proxy range advertisement in OSPF, ISIS, BGP, etc.

5. Security Considerations

Implementations must assure that malformed TLV and Sub-TLV permutations do not result in errors which cause hard protocol failures.

6. Acknowledgements

The authors would like to thank Stig Venaas for his valuable comments and suggestions.

7. References

7.1. Normative References

[I-D.ietf-bier-idr-extensions]

Xu, X., Chen, M., Patel, K., Wijnands, I., and T. Przygienda, "BGP Extensions for BIER", [draft-ietf-bier-idr-extensions-07](#) (work in progress), September 2019.

- [I-D.ietf-bier-ospfv3-extensions]
Psenak, P., Nainar, N., and I. Wijnands, "OSPFv3 Extensions for BIER", [draft-ietf-bier-ospfv3-extensions-02](#) (work in progress), May 2020.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", [RFC 5308](#), DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.
- [RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", [RFC 7684](#), DOI 10.17487/RFC7684, November 2015, <<https://www.rfc-editor.org/info/rfc7684>>.
- [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>>.
- [RFC8401] Ginsberg, L., Ed., Przygienda, T., Aldrin, S., and Z. Zhang, "Bit Index Explicit Replication (BIER) Support via IS-IS", [RFC 8401](#), DOI 10.17487/RFC8401, June 2018, <<https://www.rfc-editor.org/info/rfc8401>>.
- [RFC8444] Psenak, P., Ed., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPFv2 Extensions for Bit Index Explicit Replication (BIER)", [RFC 8444](#), DOI 10.17487/RFC8444, November 2018, <<https://www.rfc-editor.org/info/rfc8444>>.

7.2. Informative References

- [I-D.ietf-mpls-seamless-mpls]
Leymann, N., Decraene, B., Filsfils, C., Konstantynowicz, M., and D. Steinberg, "Seamless MPLS Architecture", [draft-ietf-mpls-seamless-mpls-07](#) (work in progress), June 2014.

- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, [RFC 7761](https://www.rfc-editor.org/info/rfc7761), DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", [RFC 8277](https://www.rfc-editor.org/info/rfc8277), DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.

Authors' Addresses

Zheng(Sandy) Zhang
ZTE Corporation

E-Mail: zzhang_ietf@hotmail.com

Bo Wu
Individual

E-Mail: w1973941761@163.com

Zhaohui Zhang
Juniper Networks

E-Mail: zzhang@juniper.net

IJsbrand Wijnands
Cisco Systems, Inc.

E-Mail: ice@cisco.com

Yisong Liu
China Mobile

E-Mail: liuyisong.ietf@gmail.com

