

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 26 April 2022

J. Xie  
F. Duan  
Huawei Technologies  
23 October 2021

BIER Deployment and Operation: Challenges and Solution Approaches  
draft-xie-mboned-bier-deploy-00

## Abstract

As a new multicast architecture, BIER [[RFC8279](#)] has been an IETF standard for years. It has been evaluated in some networks for some scenarios. Some challenges related to its deployment, operation, maintenance, and extensibility are raised. This document reviews and describes the challenges related to its deployment, and try to figure out the potential solution approaches.

## 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.

## 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 26 April 2022.

## Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

BIER Deployment Challenges

October 2021

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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Challenges for Inter-domain Deployment . . . . .	<a href="#">2</a>
<a href="#">2.1.</a>	Protocol ambiguity for BIER advertisement in BGP . . . . .	<a href="#">2</a>
<a href="#">2.2.</a>	Multi-hop BFR-NBR Support as an Inherent Requirement . . . . .	<a href="#">4</a>
<a href="#">2.3.</a>	Anycast BIER-Label for Redundant ASBR deployment . . . . .	<a href="#">5</a>
<a href="#">2.4.</a>	Overlapped BFR-id Assignment in Different Domains . . . . .	<a href="#">7</a>
<a href="#">2.5.</a>	Summary of Challenges in Inter-domain Deployment . . . . .	<a href="#">10</a>
<a href="#">3.</a>	Challenges for Brownfield Deployment . . . . .	<a href="#">11</a>
<a href="#">3.1.</a>	Adding Bypass tunnel for Intermediate Nodes . . . . .	<a href="#">11</a>
<a href="#">3.2.</a>	Removing (Popping) BIER Header for Edge Nodes . . . . .	<a href="#">12</a>
<a href="#">3.3.</a>	Summary of Challenges in Brownfield Deployment . . . . .	<a href="#">13</a>
<a href="#">4.</a>	Security Considerations . . . . .	<a href="#">13</a>
<a href="#">5.</a>	IANA Considerations . . . . .	<a href="#">14</a>
<a href="#">6.</a>	Acknowledgements . . . . .	<a href="#">14</a>
<a href="#">7.</a>	Normative References . . . . .	<a href="#">14</a>
	Authors' Addresses . . . . .	<a href="#">16</a>

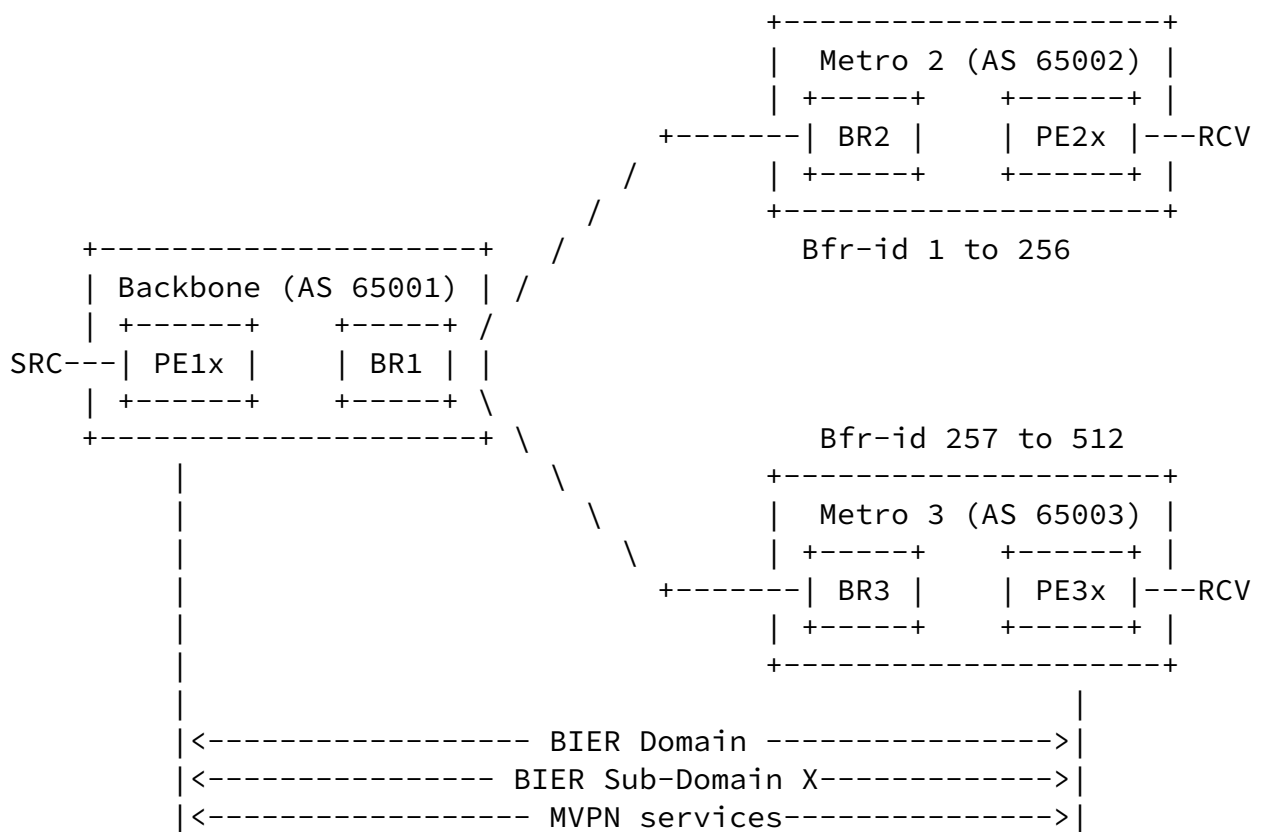
## [1.](#) Introduction

As a new multicast architecture, BIER [[RFC8279](#)] has been an IETF standard for years. It has been evaluated in some networks for some scenarios. Some challenges related to its deployment, operation, maintenance, and extensibility are raised. This document reviews and describes the challenges related to its deployment, and try to figure out the potential solution approaches.

## [2.](#) Challenges for Inter-domain Deployment

### [2.1.](#) Protocol ambiguity for BIER advertisement in BGP

The following figure demonstrates an inter-domain network where a single BIER sub-domain is deployed across the whole BIER domain (multiple domains of an administrative entity, e.g., a Service Provider's network), and where MVPN service(s) deployed on the Edge PE1x/PE2x/PE3x.



BR = Border Router  
 SRC = Multicast Source  
 RCV = Multicast Receiver

Figure 1: BIER Inter Domain Deployment

In this figure, router BR2 needs to receive BIER information advertisement from PE2x and other routers in Metro 2 by IGP (e.g., IS-IS or OSPF), and re-advertise these BIER information to BR1 using eBGP. This means that, BR2 needs to use mixed protocols for BIER information advertisement in a single sub-domain.

Such kind of mixed underlay-protocols usage for a single BIER sub-

domain could also happen in intra-domain case.

[[I-D.ietf-bier-prefix-redistribute](#)] describes such a case where an Area-Border-Router(ABR) uses different IGP protocols in different interfaces, and these interfaces are belonging to a single BIER sub-domain.

However, BIER architecture [[RFC8279](#)] requires that only one routing underlay could be used for a single Sub-domain.

Draft [[I-D.ietf-bier-prefix-redistribute](#)] defines a new TLV structure named BIER proxy range sub-TLV for mixed protocols, including BGP to use. However, it does not clearly specify how to use this TLV in BGP, e.g, what BGP attribute to carry this TLV, and what AFI/SAFI used for the advertisement.

Draft [[I-D.ietf-bier-idr-extensions](#)] defines a new BGP Path named BIER Path in [section 3](#), with a sub-TLV named BIER MPLS Encapsulation sub-TLV in the same section, but lacks description what SAFI/AFI used to carry the BIER Path attribute explicitly.

This situation leads to the challenge of BIER deployment in Inter-domain network as illustrated above. This document suggests clarifications are made in the above document by IETF. For example, it may be considered to use BGP SAFI-1/2 routes to carry BIER Path, and use BIER Path to carry the BIER proxy range sub-TLV for inter-domain advertisement.

## [2.2](#). Multi-hop BFR-NBR Support as an Inherent Requirement

In above figure, router BR2 re-advertise BIER information to BR1 using eBGP. Accordingly, BR1 needs to receive the BIER information advertisement from BR2 using eBGP, and re-advertise these BIER information to PE1x, either through IGP or iBGP.

A common practice for inter-domain routing is to separate the routing procedure into two layers, 1st layer is BGP routing to determine non-direct BGP next-hop, 2nd layer is IGP routing to determine direct IGP

next-hop based on the BGP/non-direct next-hop. For BIER inter-domain deployment as illustrated in the above figure, the preferred solution is to use iBGP on BR1 to re-advertise the BIER information to PE1x, and PE1x set BR1 (the BGP/non-direct nexthop) as the non-direct BFR-NBR to BFRs in Metro-2 and Metro-3. In another word, non-direct BFR-NBR support is an inherent requirement for BIER inter-domain deployment.

Unfortunately the BIER architecture [[RFC8279](#)] is built on L2 and the non-direct BFR-NBR or Multi-hop BFR-NBR support is optional. This means that, every time a Non-direct BFR-NBR is used, another layer of "bypass tunnel" needs to be used on the top of the BIER header for multi-hop BFR-NBR reaching. The function seems fine, but there are implications to the operation and maintenance aspects.

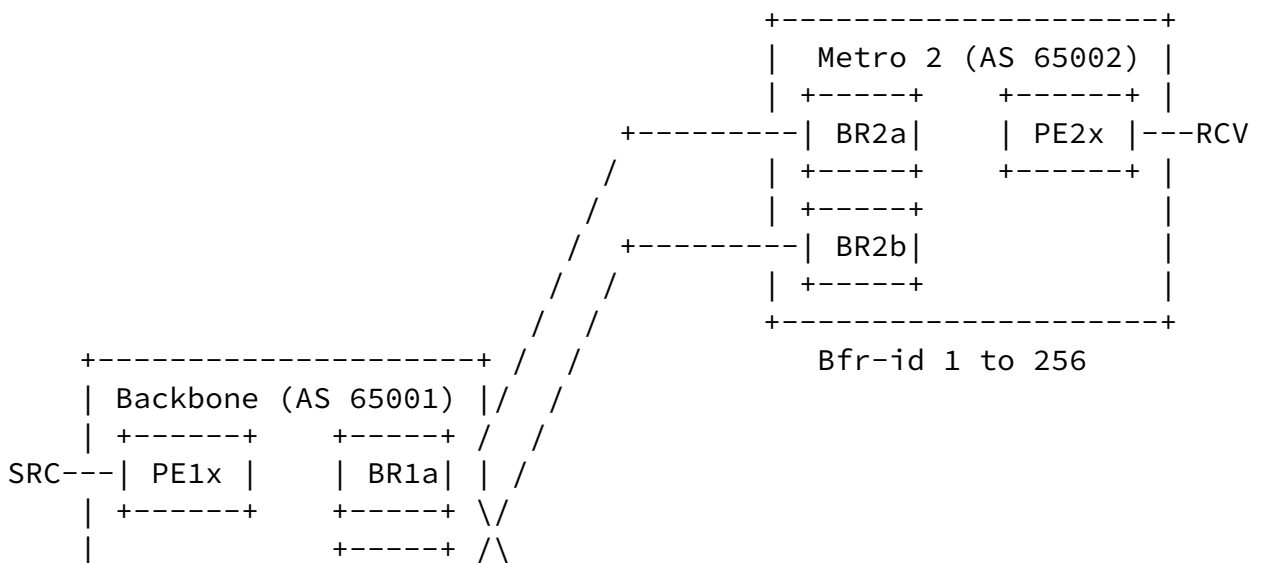
Firstly a policy should be configured to select what type(s) of tunnel are preferred. Some network operator may prefer to use "MPLS LSP" as the "bypass tunnel", and then there are multiple options "LDP LSP", "RSVP-TE LSP", "SR-MPLS LSP" for the selection. Some network operator may prefer to use IP, GRE, or UDP tunnel.

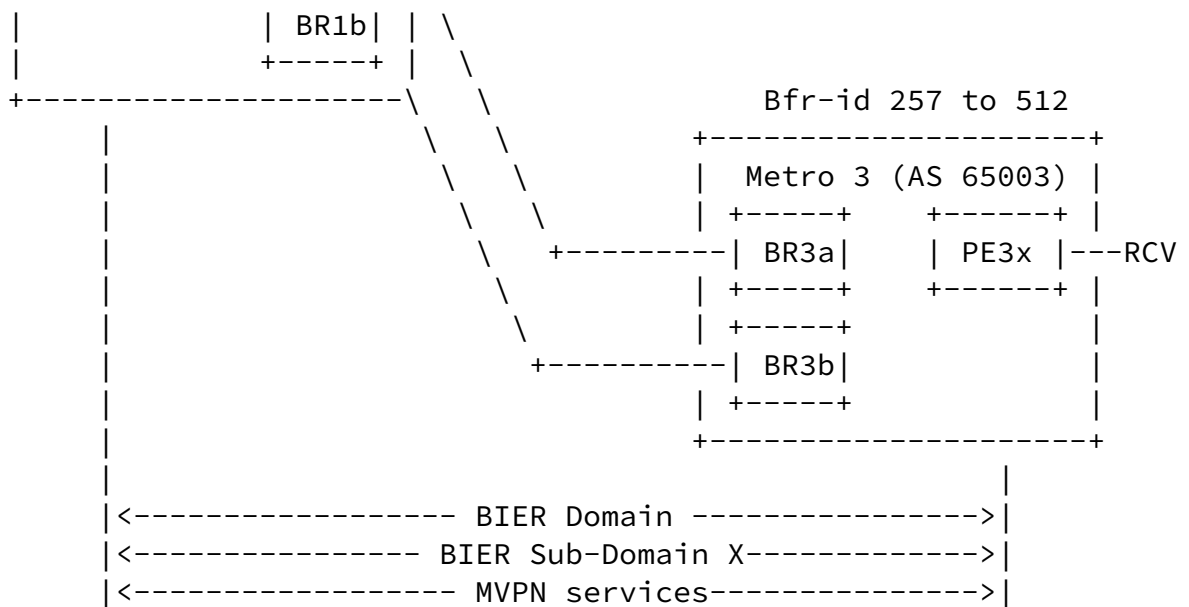
Secondly the protocols and identifiers bound to these "bypass tunnel" have to be taken into the BIER routing and forwarding information. Different tunnel type means different tunnel identifier in control plane for operation and maintenance. For example, LDP tunnel means FEC object [[RFC5036](#)], RSVP-TE tunnel means Session Object [[RFC3209](#)], SR tunnel means SRGB block and the index object [[RFC8660](#)], and IP/UDP/GRE tunnel means the IP Endpoint/UDP port/GRE key object. Accordingly, network administrators need to debug these protocols and the various identifiers additionally in operation and troubleshooting.

### [2.3](#). Anycast BIER-Label for Redundant ASBR deployment

The following figure demonstrates an inter-domain network (of an administrative entity, e.g., a Service Provider's network), where

redundant ASBRs is deployed.





BR = Border Router  
 SRC = Multicast Source  
 RCV = Multicast Receiver

Figure 2: BIER Inter Domain Deployment Redundant

In this figure, a common practice is to use anycast mechanism for service protection. For example, BR1a and BR1b share a same identifier called anycast ID, where the anycast ID could be an IP address or an SRGB label [RFC8402]. Take unicast IP address as an example, PE1x send a BIER packet to Metro-2 or Metro-3 through the backbone border using the anycast IP address without awareness of the two nodes and its state.

Usually such an anycast ID is a per node-pair allocation policy.

In BIER-MPLS [RFC8296] design, the first 4 octets of BIER header is an MPLS label with the S bit set to 1 to indicate the bottom of the MPLS label stack. It applies to a per-SD/BSL/SI allocation policy and thus defined as "locally significant" in [section 2.1.1.1 of \[RFC8296\]](#).

If BR1a and BR1b want to deploy anycast mechanism for service protection, then SRGB label need to be used for BIER-MPLS label

allocation on a Per-SD/BSL/SI base. It is needed to use manual configuration on each node due to the shortage of SRGB label for automatic allocation.

In addition, the BR1a and BR1b need to have the same (at least overlapped) SRGB label space to ensure the anycast BIER-MPLS value is absolutely equal. Basically it does not require the SRGB label space to be absolutely equal in Segment Routing architecture [[RFC8402](#)], but in anycast Label case, it needs the absolute equivalence. If BR1a and BR1b have different SRGB label space, the deployment of anycast BIER-MPLS scheme is still challengeable.

Note that, Non-MPLS BIER encapsulation uses a different L2 protocol indication (typically Ethertype 0xAB37) to indicate the BIER header following the L2 header. In such case, the first 4 octets of the BIER header is not an MPLS label encoding. The 20-bit BIFT-id field of BIER header is wide enough for automatically mapping from SD/BSL/SI by using the method in [[I-D.ietf-bier-non-mpls-bift-encoding](#)].

#### [2.4.](#) Overlapped BFR-id Assignment in Different Domains

Intra-domain usually needs to consider a router to be added without much impact to existing routers. Given this, the BFR-id assignment in Intra-domain scenario need to reserve some BFR-id space (holes) in the earlier range.

Following is quoted from [RFC8279 section 2](#) verbatim.

The procedure for assigning a particular BFR-id to a particular BFR is outside the scope of this document. However, it is RECOMMENDED that the BFR-ids for each sub-domain be assigned "densely" from the numbering space, as this will result in a more efficient encoding (see [Section 3](#)). That is, if there are 256 or fewer BFRs, it is RECOMMENDED to assign all the BFR-ids from the range [1,256]. If there are more than 256 BFRs but less than 512, it is RECOMMENDED to assign all the BFR-ids from the range [1,512], with as few "holes" as possible in the earlier range. However, in some deployments, it may be advantageous to depart from this recommendation; this is discussed further in [Section 3](#).

Following is quoted from [RFC8279 section 3](#) verbatim.



In order to minimize the number of copies that must be made of a given multicast packet, it is RECOMMENDED that the BFR-ids used in a given sub-domain be assigned "densely" (see [Section 2](#)) from the numbering space. This will minimize the number of SIs that have to be used in that sub-domain. However, depending upon the details of a particular deployment, other assignment methods may be more advantageous. Suppose, for example, that in a certain deployment, every multicast flow is intended either for the "east coast" or for the "west coast", but not for both coasts. In such a deployment, it would be advantageous to assign BFR-ids so that all the "west coast" BFR-ids fall into the same SI-subset and so that all the "east coast" BFR-ids fall into the same SI-subset.

The BFR-id assignment recommendation above can be summarized as a hierarchical rule: A whole set/block (1-256 are a set for example) is assigned to a network-region (east coast or west coast) of a network, where some "holes" in the block are reserved to allow nodes adding in the future without crossing the border of a set.

Inter-domain allows a domain to be added without much impact to existing domains. It adds another level of constraints when considering the BFR-id assignment. Extending the hierarchical rule, an overlapped BFR-id assignment scheme may be considered. The following figure illustrates the scheme:




Figure 3: Overlapped BFR-id Assignment

In this scheme, each domain (Backbone, Metro-2, Metro-3) has a bigger BFR-id space (1 to 1024 for example), and within each domain, the hierarchical rule is further used to assign block of BFR-id (1 to 256 for example) to a network-region, and holes in the block are reserved for each network-region to expend its nodes in the future.

When a packet is transmitted from PE1x to PE2x, PE1x forwards it through a multi-hop crossing-domain tunnel from PE1x to BR2, and BR2 does the rest BIER forwarding inside the domain. Again, the multi-hop BFR-NBR support is an Inherent Requirement in this scheme.

This scheme will provide a clear BFR-id assignment for inter-domain deployment, and extends the multicast deployment to beyond the 256 set limitation by using the combination of Ingress-Replication and BIER. Note the multiple Set is a combination of Ingress-Replication and BIER too. Below is an example of such extensibility:

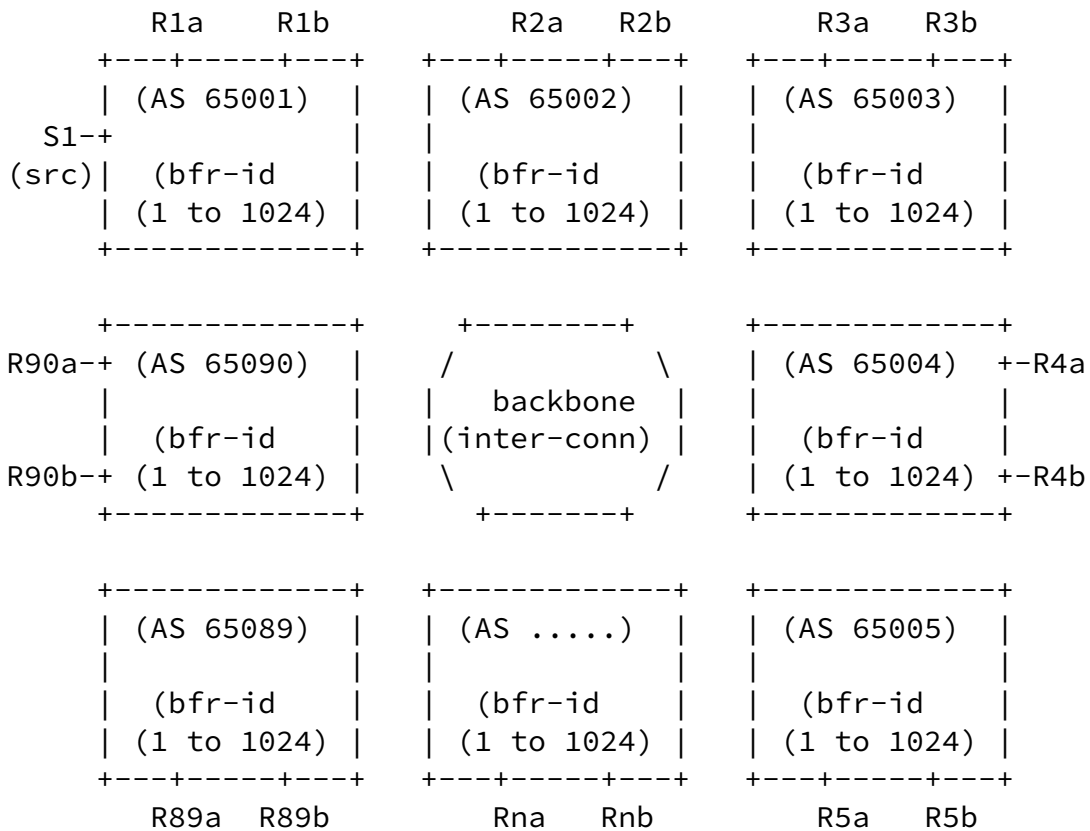


Figure 4: Overlapped BFR-id Assignment 2

In this example, AS 65001, 65002, 65003, ..., 65090 each has an

initial BFR-id scope 1 to 1024, totally 90K BFERs could be supported. Source PE S1 could send packet to R1a/R1b, R2a/R2b, R3a/R3b, ...,

R90a/R90b using Ingress-Replication (e.g., to each domain) and BIER BitString encapsulation (e.g., to multiple nodes in a domain). When the inter-connected ASs is increased to AS 65091, it could continue its extension linerly and the limitation is the replication capability on S1 (which could be hundreds or thousands on modern routers).

However, BIER is usually deployed with Overlay service and procedure like MVPN [[RFC8556](#)]. Such service depends on the packet to have the "Source" information for Reverse-Path-Forwarding (RPF) check to ensure the packet is coming from the correct Upstream-Multicast-Hop (UMH, [[RFC6513](#)]). In BIER architecture, the "Source" information in the packet is the "BFIR-id" field within the Context of BIER Sub-domain-id. The BIER Sub-domain-id is further mapped from the BIFT-id field. The BFIR-id (in the context of Sub-domain-id) need to be uniquely bound to an IP address of the BFIR node. This makes the above assignment scheme unavailable due to the overlapped BFR-id assignment among multiple domains. Note all the BFR-id demonstrated in the above diagram is in a single Sub-domain-id.

The root cause of the challenge in this case is the use of BFR-id in both the "destinations" and "source" of a packet. If these two things are separated, and the unique IP address of the BFIR node is used in the BIER encoding, this problem will be solved.

## [2.5](#). Summary of Challenges in Inter-domain Deployment

These are some of the challenges for BIER deployment in Inter-domain environment. This section also reviews the BIER architecture for each challenge to try to figure out the gap that may need to consider in the future work.

### Challenge-1:

Description: Protocol ambiguity for BIER advertisement in BGP.

Gap: BIER architecture does not support multiple/mixed routing-underlay protocols for a single sub-domain. BGP protocol extension for inter-domain BIER deployment is still unclear.

Challenge-2:

Description: Multi-hop BFR-NBR Support as an Inherent Requirement.

Gap: BIER architecture is built on L2 and Multi-hop BFR-NBR support is optional.

Challenge-3:

Description: Anycast BIER-Label for Redundant ASBR deployment.

Gap: The locally-significant per-SD/BSL/SI BIER-label is opposite to the per-Node Anycast ID assignment.

Challenge-4:

Description: Overlapped BFR-id Assignment in Different Domains.

Gap: BFR-id is coupled for both "destinations" and "source" in BIER architecture. It makes BFR-id assignment constrained and lacking of extensibility.

### [3.](#) Challenges for Brownfield Deployment

#### [3.1.](#) Adding Bypass tunnel for Intermediate Nodes

Bypass tunnel is a mechanism in BIER for deployment in brownfield network where an intermediate router does not support BIER forwarding. This is the same as Multi-hop BFR-NBR as previously described. The function could be implemented by additional encapsulation of a "bypass tunnel" of any type. The impact is in the "operation and maintenance" aspects, as shown in previously in this document. The improvements to minimize the impact is to select a "default" type of bypass tunnel and mandate it in a standard, thus different implementations could interop at the bottom.

For BIER-MPLS, the preferred bypass tunnel is highly likely to be MPLS LSP. But still need the further step to mandate an option as "default" from SR, LDP or RSVP-TE.

For Non-MPLS BIER, the MPLS LSP and the functions built on it (like TI-LFA) is no longer available. The preferred bypass tunnel is highly likely to be an IP based tunnel. For IPv4, an IPv4+UDP tunnel may be preferred due to the lack of ECMP in IP itself. For IPv6, an IPv6 tunnel may be preferred.

Such diversity of options make it challengeable toward implementation and operation, including the selection between BIER-MPLS and Non-MPLS BIER first, and further interop test, deployment, troubleshooting and so on in each case.

### [3.2.](#) Removing (Popping) BIER Header for Edge Nodes

Another typical challenge in brownfield network is the Edge router(s) not supporting BIER forwarding. Some networks may have many edge routers connected to a few core routers, and it is highly possible there are some of these edge routers not supporting BIER. Using Penultimate Hop Pop (PHP) or Penultimate Segment Pop (PSP, [[RFC8986](#)]) on the upstream node of an edge router can increase the deployability of BIER greatly.

Draft [[I-D.ietf-bier-php](#)] defines a method for BIER PHP. An edge router not supporting BIER forwarding acts as a Pseudo-BFER node. It has a valid BFR-id assignment, and it signals other router in a BIER domain (typical an IGP domain using IS-IS or OSPF) a "PHP request". When an upstream router has a BIER packet with the bit corresponding to the BFR-id of this pseudo-BFER set to 1, the BIER header of the BIER packet is removed, and the packet is then unicast to this pseudo-BFER.

However, as is pointed in the document, penultimate hop popping the BIER header ahead of the pseudo-BFER means that, the BFR-id or the source identifier in the BIER header is also lost. RPF function depending on the source identifier is no longer available. Note that, RPF function is a basic function in multicast, to solve the problem in UMH [[RFC6513](#)] changing scenario, Source Redundancy

[[RFC9026](#)] scenario and so on.

Another problem is that, BIER header has a "proto" field to indicate the payload that follows the BIER header. It is a "BIER specific" Next header indication. When the BIER header is popped, the next header indication will have to scatter to the preceding header. E.g., if the preceding header is a link-level Ethernet header, each "BIER proto" value need to have a Ether-Type. If the preceding header is a "bypass tunnel" of IP/GRE/UDP type, each "BIER Proto" value need to have an IP Proto/GRE Proto/UDP Port. A typical case is the "Echo Request" currently defined in [[I-D.ietf-bier-ping](#)] as BIER payload (using BIER Proto 5). Using the BIER PHP, the BIER ping/trace will no longer be available. Note that, some other functions like the BFD bootstrap may depend on the BIER ping/trace, and will suffer the same as a consequence.

A possible way for solving these problems in PHP deployment is to expend another layer over BIER header. For example, an IP header is followed as a shim layer in the lifecycle of a BIER packet from BFIR to BFER. Thus, when the BIER header is popped, the BFER could still get the "source identifier" from the IP header. Accordingly, the BIER ping, BIER bfd also need to build the upper-layer body after the IP header.

An alternative way is to use use an additional layer of header when popping the BIER header. For example, the upstream router A get the "source identifier" from the BIER packet, and encapsulate the BIER payload with an additional IP header whose IP source is the "source identifier" of the BFIR node. Thus when the pseudo-BFER receives the packet without BIER header can still get the source identifier necessary for RPF function and the like. Of course, the IP proto should be able to identify the BIER payload as mentioned above.

### [3.3.](#) Summary of Challenges in Brownfield Deployment

These are some of the challenges for BIER deployment in Brown-field network.

Challenge-5:

Description: Adding Bypass tunnel on top of BIER header for Intermediate nodes.

Gap: Select a "default" type of bypass tunnel to help the interop between different implementations.

Challenge-6:

Description: Removing (Popping) the BIER header for Edge nodes.

Gap: Not losing basic functions like MVPN RPF/UMH, Source Redundancy, Ping/Trace when possible, using additional encapsulation either from BFIR or from the popping BFR node.

It can be seen that, the BIER architecture [[RFC8279](#)] is built on L2 and thus is dependent completely on additional mechanisms for brownfield deployment. The mechanisms include: Some kind of routing mechanism (IP or LSP are both included in this terminology) is needed for multi-hop BIER neighboring. Additional identifier of a BFIR node is needed for MVPN RPF and UMH-Redundancy functions, and additional Next Header identifier is needed for BIER payload indicating when popping BIER header.

#### [4.](#) Security Considerations

This document introduces no new security considerations beyond those already specified in [[RFC8279](#)], [[RFC8296](#)], [[RFC8556](#)].

#### [5.](#) IANA Considerations

This document contains no actions for IANA.

#### [6.](#) Acknowledgements

Thanks Xuesong Geng for the review of this document.

#### [7.](#) Normative References

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

Xu, X., Chen, M., Patel, K., Wijnands, I., and A. Przygienda, "BGP Extensions for BIER", Work in Progress, Internet-Draft, [draft-ietf-bier-idr-extensions-07](https://www.ietf.org/archive/id/draft-ietf-bier-idr-extensions-07), 6 September 2019, <<https://www.ietf.org/archive/id/draft-ietf-bier-idr-extensions-07.txt>>.

[I-D.ietf-bier-non-mpls-bift-encoding]

Wijnands, I., Mishra, M., Xu, X., and H. Bidgoli, "An Optional Encoding of the BIFT-id Field in the non-MPLS BIER Encapsulation", Work in Progress, Internet-Draft, [draft-ietf-bier-non-mpls-bift-encoding-04](https://www.ietf.org/archive/id/draft-ietf-bier-non-mpls-bift-encoding-04), 30 May 2021, <<https://www.ietf.org/archive/id/draft-ietf-bier-non-mpls-bift-encoding-04.txt>>.

[I-D.ietf-bier-php]

Zhang, Z., "BIER Penultimate Hop Popping", Work in Progress, Internet-Draft, [draft-ietf-bier-php-06](https://www.ietf.org/archive/id/draft-ietf-bier-php-06), 24 August 2020, <<https://www.ietf.org/archive/id/draft-ietf-bier-php-06.txt>>.

[I-D.ietf-bier-ping]

Kumar, N., Pignataro, C., Akiya, N., Zheng, L., Chen, M., and G. Mirsky, "BIER Ping and Trace", Work in Progress, Internet-Draft, [draft-ietf-bier-ping-07](https://www.ietf.org/archive/id/draft-ietf-bier-ping-07), 11 May 2020, <<https://www.ietf.org/archive/id/draft-ietf-bier-ping-07.txt>>.

[I-D.ietf-bier-prefix-redistribute]

Zhang, Z., Wu, B., Zhang, Z., Wijnands, I., and Y. Liu, "BIER Prefix Redistribute", Work in Progress, Internet-Draft, [draft-ietf-bier-prefix-redistribute-00](https://www.ietf.org/archive/id/draft-ietf-bier-prefix-redistribute-00), 4 August 2020, <<https://www.ietf.org/archive/id/draft-ietf-bier-prefix-redistribute-00.txt>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](https://www.rfc-editor.org/info/rfc2119), [RFC 2119](https://www.rfc-editor.org/info/rfc2119), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.



- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", [RFC 5036](#), DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", [RFC 8986](#), DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9026] Morin, T., Ed., Kebler, R., Ed., and G. Mirsky, Ed., "Multicast VPN Fast Upstream Failover", [RFC 9026](#), DOI 10.17487/RFC9026, April 2021, <<https://www.rfc-editor.org/info/rfc9026>>.

#### Authors' Addresses

Jingrong Xie  
Huawei Technologies

Email: [xiejingrong@huawei.com](mailto:xiejingrong@huawei.com)

Fanghong Duan  
Huawei Technologies

Email: [duanfanghong@huawei.com](mailto:duanfanghong@huawei.com)

