

BIER WG
Internet-Draft
Intended status: Standards Track
Expires: June 23, 2016

Z. Zhang
C. Wang
ZTE Corporation
December 21, 2015

Automatic Assignment of BIER BFR-ids in OSPF
draft-zhang-bier-bfrid-assignment-00

Abstract

[I-D.ietf-bier-architecture] has introduced a new method to forward multicast flow, without explicit multicast states storing in every node along the multicast paths. This document introduces a method to allocate BFR-id automatically through OSPF extension.

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 <http://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 June 23, 2016.

Copyright Notice

Copyright (c) 2015 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	4
3.	IANA considerations	5
4.	Procedures	6
4.1.	D-BFR Election Algorithm	6
4.2.	D-BFR Procedures	8
4.2.1.	Assignment of BMPs to BFRs	8
4.3.	BD-BFR Procedures	8
4.4.	BFR Procedures	8
5.	Special Considerations	9
5.1.	BD-BFR to D-BFR Transition	9
5.2.	Election FSM for BFR	9
5.2.1.	States	10
5.2.2.	Events	11
5.2.3.	Rules	12
5.2.4.	Warning and Logging	12
6.	Packet Formats	13
6.1.	BIER-PE BIER Protocol Election Sub-sub-TLV	13
6.2.	Reuse of the Reserved Bits in BIER Info sub-TLV	13
6.3.	BIER-PE-BMP: BIER PE BMP Assignments TLV	14
7.	Security Considerations	17
8.	Normative References	18
	Authors' Addresses	19

1. Introduction

[I-D.ietf-bier-architecture] defines a new efficient forwarding way for multicast flow. The crucial object in this method is BFR-id for every BFERs. All nodes in the BIER domain learn the BFR-ids of BFERs, and forward the packet according to the BIER header that are composed by the BFR-ids.

Although the BFR-id can be acquired by central controllers or statically, it will be more convenient if there is a way to allocate the BFR-id automatically.

This document introduces a new method to allocate BFR-id for BFERs in BIER domain. And this document also defines the format of OSPF extension for BFR-id auto assignment.

This document gets benefit from the DR election algorithm that is defined in [RFC2328](#). And the main idea of this document is the same as that is defined in "[draft-prz-bier-bfrid-assignment](#)". The only difference between the two documents is the protocol format.

2. Terminology

Some of the terminology specified in [[I-D.ietf-bier-architecture](#)] is replicated here and extended by necessary definitions:

BIER: Bit Index Explicit Replication (The overall architecture of forwarding multicast using a Bit Position).

BIER-OL: BIER Overlay Signaling. (The method for the BFIR to learn about BFER's).

BFR: Bit Forwarding Router (A router that participates in Bit Index Multipoint Forwarding). A BFR is identified by a unique BFR-prefix in a BIER domain.

BFIR: Bit Forwarding Ingress Router (The ingress border router that inserts the BM into the packet).

BFER: Bit Forwarding Egress Router. A router that participates in Bit Index Forwarding as leaf. Each BFER must be a BFR. Each BFER must have a valid BFR-id assigned.

BFT: Bit Forwarding Tree used to reach all BFERs in a domain.

BIFT: Bit Index Forwarding Table.

BMS: Bit Mask Set. Set containing bit positions of all BFER participating in a set.

BMP: Bit Mask Position, a given bit in a BMS.

Invalid BMP: Unassigned Bit Mask Position, consisting of all 0s.

IGP signaled BIER domain: A BIER underlay where the BIER synchronization information is carried in IGP. Observe that a multi-topology is NOT a separate BIER domain in IGP.

BIER sub-domain: A further distinction within a BIER domain identified by its unique sub-domain identifier. A BIER sub-domain can support multiple BitString Lengths.

BFR-id: An optional, unique identifier for a BFR within a BIER sub-domain.

Invalid BFR-id: Unassigned BFR-id, consisting of all 0s.

3. IANA considerations

This document adds the following new sub-sub-TLVs to the registry of sub-TLVs for BIER Info sub-TLV.

BIER Protocol Election sub-sub-TLV Value: TBD (suggested - to be assigned by IANA)

This document adds the following new TLV to the registry of OSPF TLVs.

BIER PE BMP Assignments TLV Value: TBD (suggested - to be assigned by IANA)

4. Procedures

At first, all the BIER nodes collect the information of D-BFR candidates and BD-BFR candidates through the advertisements of BIER protocol election sub-sub-TLVs. All the BFRs flood the sub-sub-TLVs per sub-domain or per BMS to all other nodes.

The D-BFR election algorithm is most like the DR elect function in OSPF protocol. And the FSM is also like the function in OSPF protocol. The algorithm described below is most from [RFC2328](#).

OSPF floods the DR/BDR information through OSPF hello packets. BIER nodes flood the BIER protocol election sub-sub-TLVs along with BIER information sub-TLV.

ALL the BIER nodes elect the D-BFR and BD-BFR through the Designated Router BFR function. And the D-BFR assigns BFR-ids according to the received BIER information sub-TLV which request to allocate a BFR-id. After D-BFR assigns all the BFR-ids and floods the assignment to all the BIER nodes, BD-BFR mirrors the assignment as its base assignment. If there are some collisions existing, the BFRs that are not allocated BFR-id negotiate the BFR-id assignment procedure with D-BFR again.

4.1. D-BFR Election Algorithm

The Designated Router BFR election algorithm proceeds as follows:

- o Call the router doing the calculation Router X. The list of neighbors attached to the network and having established bidirectional communication with Router X is examined.
- o The state of BFRs that may be D-BFR or BD-BFR should be examined by SPF computation.
- o Router X itself is also considered to be on the list.
- o Discard all routers from the list that are ineligible to become DR-BDR. (Routers having Router Priority of 0 are ineligible to become D-BFR.) The following steps are then executed, considering only those routers that remain on the list:

(1) Note the current values for the network's D-BFR and BD-BFR. This is used later for comparison purposes.

(2) Calculate the new BD-BFR for the network as follows. Only those routers on the list that have not declared themselves to be D-BFR are eligible to become BD-BFR. If one or more of these routers

have declared themselves BD-BFR (i.e., they are currently listing themselves as BD-BFR, but not as D-BFR, in their sub-sub-TLVs) the one having highest Router Priority is declared to be BD-BFR. In case of a tie, the one having the highest Router ID is chosen. If no routers have declared themselves BD-BFR, choose the router having highest Router Priority, (again excluding those routers who have declared themselves D-BFR), and again use the Router ID to break ties.

(3) Calculate the new Designated Router for the network as follows. If one or more of the routers have declared themselves D-BFR (i.e., they are currently listing themselves as D-BFR in their BIER PE sub-sub-TLV) the one having highest Router Priority is declared to be D-BFR. In case of a tie, the one having the highest Router ID is chosen. If no routers have declared themselves D-BFR, assign the D-BFR to be the same as the newly elected BD-BFR.

(4) If Router X is now newly the D-BFR or newly the BD-BFR, or is now no longer the D-BFR or no longer the BD-BFR, repeat steps 2 and 3, and then proceed to step 5. For example, if Router X is now the DR-BDR, when step 2 is repeated X will no longer be eligible for BD-BFR election. Among other things, this will ensure that no router will declare itself both BD-BFR and D-BFR.

(5) As a result of these calculations, the router itself may now be D-BFR or BD-BFR. See Section4.2 and Section4.3 for the additional duties this would entail.

(6) If the above calculations have caused the identity of either the D-BFR or BD-BFR to change, all the routers must re-evaluate whether they have been selected D-BFR or BD-BFR and initiate according procedures. In case the new D-BFR is not advertising according bitmask assignment and they are needed, they initiate according procedures in Section4.2.

The reason behind the election algorithm's complexity is the desire for an orderly transition from BD-BFR to D-BFR, when the current D-BFR fails. This orderly transition is ensured through the introduction of hysteresis: no new BD-BFR can be chosen until the old Backup accepts its new D-BFR responsibilities.

The above procedure may elect the same router to be both D-BFR and BD-BFR, although that router will never be the calculating router (Router X) itself. The elected D-BFR may not be the router having the highest Router Priority, nor will the BD-BFR necessarily have the second highest Router Priority. If Router X is not itself eligible to become D-BFR, it is possible that neither a BD-BFR nor a D-BFR will be selected in the above procedure. Note also that if Router X

is the only attached router that is eligible to become D-BFR, it will select itself as D-BFR and there will be no BD-BFR for the network.

4.2. D-BFR Procedures

Similar to the D-BFR and BD-BFR election procedure, the assignment of D-BFR is also base on a sub-domian or a BMS.

4.2.1. Assignment of BMPs to BFRs

The procedure is initiated by a BFER announcing R bit in BIER Info sub-TLV. The D-BFR assigns BMPs to such BFER or announces collisions.

Observe that BFERs can request (or announce) the R bits even before a D-BFR has been chosen so the election and assignment are largely orthogonal sets of procedures.

The BFR-ids in one sub-domain or a BMS should be assigned one by one as far as possible.

4.3. BD-BFR Procedures

BD-BFR mirrors the BIER PE BMP Assignments TLV from the advertisement of D-BFR. And BD-BFR uses the existing assignment as the initial input of probably allocation.

4.4. BFER Procedures

BFER sends the BIER protocol Election sub-sub-TLV at first. If the BFER wants itself to be a D-BFR or BD-BFR, it should adjust the D-BFR priority in advance. After BFER receives the BIER protocol Election sub-sub-TLVs from other BIER nodes, it elects the D-BFR and BD-BFR according to the function defined in Section4.1.

If the BFER finds that itself is the D-BFR, it should do the assignment of D-BFR. If the BFER finds that itself is the BD-BFR, it mirrors the assignment advertisement of D-BFR. If the BFER is neither D-BFR nor BD-BFR, it should only care the interaction between itself and D-BFR.

BFER which need be allocated BFR-id sends the request in BIER info sub-TLV. If one certain BFR-id is pre-configured, BFER sends this BFR-id to D-BFR along with BIER info sub-TLV. And D-BFR takes the certain BFR-id into account preferential. If BFER can't receive the satisfied result from the PE BMP assignments TLV, it should log the error and negotiate with D-BFR again.

5. Special Considerations

5.1. BD-BFR to D-BFR Transition

BD-BFR stores the assignments of D-BFR advertisement. And BD-BFR treats this existing allocation as initial state. When BD-BFR should take charge of D-BFR and continue allocating BFR-ids, it MUST NOT change existing allocation, in other words, BD-BFR should allocate new BFR-ids to the new nodes of the network.

5.2. Election FSM for BFR

This section describes the finite state machine that runs on every BFR.

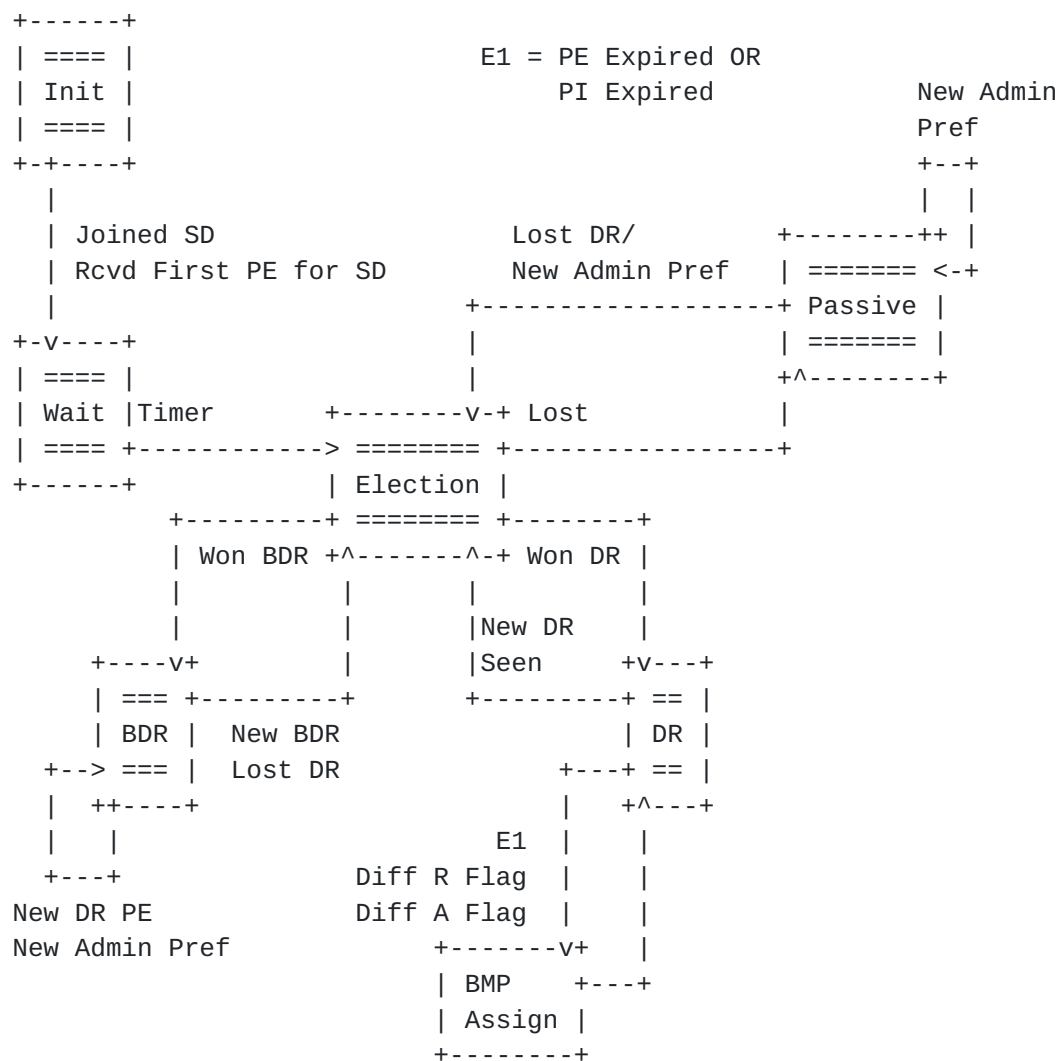


Figure 1: D-BFR/BD-BFR election FSM

5.2.1. States

Init: Initial State of the Machine

Wait: State waiting for routers to update their PEs for the sub-domain on startup

Election: State that runs the election procedures and generates according events that progress it into another state immediately

Passive: State entered when lost both DR and BDR in election.

Elected D-BFR

Elected BD-BFR

BMP Assign: State in which the assignment of bits happens upon requests from BFRs.

5.2.2. Events

Timer: Initial timer waiting for s of other routers before election is triggered.

Signalling/Rcvd First PE: First PE for has been received or signaling enabled for the set S on BFR.

Lost DR: Current D-BFR cannot be reached anymore via SPF computation in standard topology.

Lost: Lost election for D-BFR and BD-BFR.

Won BDR: Won election for BD-BFR.

Won DR: Won election for D-BFR.

New BDR: A new BD-BFR has been elected by the D-BFR.

New DR PE: New BIER-PE Instance from D-BFR.

New Admin Pref: Changed Administrative preference. And it triggers the election of BD-BFR.

Diff R Flag: R flag has been announced by a BFR which was not present before. In case of a new R flag, an assignment should be attempted. In case of R flag being deleted

- if the A flag is set, the validity of the copied BFR-id with the assignment is checked

- if the A flag is clear, the value is assumed non-negotiable and re-assignments may be necessary

Diff A Flag: A flag has been withdrawn or announced.

- If A flag was present before and

- R flag is clear, the value is assumed non-negotiable and re-assignments may be necessary.

- R flag is set, a new assignment is requested.

- If A flag was not present before and

R flag is clear, the validity of the copied BFR-id with the assignment is checked

R flag is set, the client MUST be declared faulty and disregarded.

5.2.3. Rules

When a new BFR originates its BIER protocol election advertisement, and it is one candidate to be D-BFR or BD-BFR, it should announce itself to be BD-BFR instead of D-BFR. If the administrative priority is set to 0, it MUST NOT announce itself to be D-BFR or BD-BFR.

5.2.4. Warning and Logging

Election failure If there is no candidate for D-BFR and BD-BFR after election timer expired, D-BFR and BD-BFR can't be elected, it should trigger a warning or an error log.

D-BFR reachability After a D-BFR is elected, if the D-BFR is unreachable, it should trigger a warning or an error log.

Flag error If the C bit, R bit and A bit are used incorrect, it should trigger a warning or an error log.

6. Packet Formats

The BIER information is advertised in a sub-TLV, and this information is associated with the BFR-prefix, this definition is described in [\[I-D.ietf-bier-ospf-bier-extensions\]](#) .

A new sub-sub-TLV that is defined for BIER DR election algorithm is included in the BIER Info sub-TLV of the according sub-domain as specified by [\[I-D.ietf-bier-ospf-bier-extensions\]](#). It MUST be included in the BIER Info sub-TLV only once, otherwise the first instance is used. As the [\[I-D.ietf-bier-architecture\]](#) said, the middle nodes that will not be BFER do not need the BFR-id. But in some situations, one of the middle node will be used to be treated as D-BFR to allocate BFR-ids, the middle node should also send the sub-sub-TLV with the BIER info sub-TLV to indicate that it should be treated as one of the candidate of D-BFR.

6.1. BIER-PE BIER Protocol Election Sub-sub-TLV

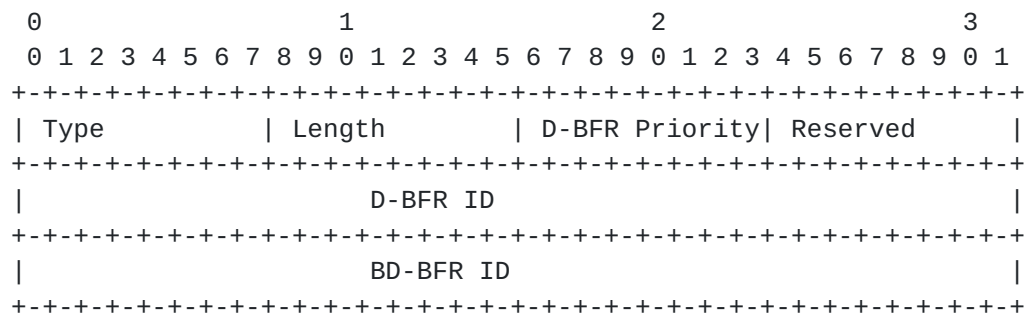


Figure 2 BIER Protocol Election sub-sub-TLV

- o Type: TBD.
- o Length: 1 octet.
- o Priority: Priority at which this router is set to become D-BFR for the sub-domain.
- o D-BFR ID: ID of the router chosen as D-BFR. If the router elected itself as D-BFR it MUST set it to its own ID.
- o BD-BFR ID: ID of the router chosen as BD-BFR. If the router elected itself as BD-BFR it MUST set it to its own ID.

6.2. Reuse of the Reserved Bits in BIER Info sub-TLV

The format listed here may seem more like the format that is defined in [\[I-D.ietf-bier-isis-extensions\]](#) than that is defined in [\[I-D.ietf-bier-ospf-bier-extensions\]](#), because the

[I-D.ietf-bier-isis-extensions] has been discussed more sufficient, and the format of BIER info sub-TLV will be uniform later between ISIS and OSPF.




Figure 3 Reuse of the Reserved Bits in BIER Info sub-TLV

Version: Version of the protocol. It remains at 0.

C: The compatibility bit. It is set according to following rules:

If the R bit is set, C is set to 0, i.e. the TLV is not compatible with version 0 of the BIER information. This will prevent routers not implementing this specification from looking at this advertisement.

If the R bit is clear, C is set to 1. In case the BFR-id has been obtained without an error by requesting it from a D-BFR, the value is copied into BFR-id of this sub-TLV, otherwise it is set to invalid BFR-id.

R: Request Bit. When set, this bit advertises that the BFER is willing to accept another BMP than the one administratively desired from D-BFR. The value of BMP is then determined by the according element in BIER-PE-BMP of the D-BFR.

A: When this bit is set, the BFER advertises that the value indicated in the BFR-id has been copied from the assignment provided by D-BFR. If clear and BFR-id is set, the value is administratively assigned and is non-negotiable.

BFR-id: If set and R bit is clear, it indicates the BFR-id the BFR is occupying to the D-BFR. If the R bit is set, it indicates the desired BFR-id to be assigned or no preference.

6.3. BIER-PE-BMP: BIER PE BMP Assignments TLV

This TLV is advertised only for a sub-domain or a BMS for which the router has been elected to be D-BDR or BD-BDR. It can repeat multiple times.


```
2 error: Out of Bit Mask Positions for the Topology and Set, Bit
Mask Position cannot be used. E-bit MUST be set.
```


all other values reserved, MUST NOT be used.

The assignments SHOULD be sorted on BFER-ID. Assignments MUST NOT repeat when the TLV is advertised multiple times and a router discovering such condition MUST issue an adequate warning. When multiple assignments for the same BFR are found, the first one in first TLV MUST be used and all others disregarded.

The assignments MUST NOT repeat any BIER Info sub-TLVs that have the R and A bit cleared, e.g. purely administrative assignments. A router discovering such condition MUST issue an adequate warning and disregard such assignments.

The assignments MUST repeat all assigned BIER Info sub-TLVs (that have A bit set). When such assignment is not advertised anymore, the according BFER MUST interpret that as loss as assignment, i.e. start with R bit again or set the BFR-id to invalid BFR-id.

7. Security Considerations

For general BIER Security Considerations.

8. Normative References

[I-D.ietf-bier-architecture]

Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", [draft-ietf-bier-architecture-02](#) (work in progress), July 2015.

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

Ginsberg, L., Przygienda, T., Aldrin, S., and J. Zhang, "BIER support via ISIS", [draft-ietf-bier-isis-extensions-01](#) (work in progress), October 2015.

[I-D.ietf-bier-mpls-encapsulation]

Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., and S. Aldrin, "Encapsulation for Bit Index Explicit Replication in MPLS Networks", [draft-ietf-bier-mpls-encapsulation-02](#) (work in progress), August 2015.

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

Psenak, P., Kumar, N., Wijnands, I., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPF Extensions For BIER", [draft-ietf-bier-ospf-bier-extensions-01](#) (work in progress), October 2015.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.

Authors' Addresses

Zheng(Sandy) Zhang
ZTE Corporation
No. 50 Software Ave, Yuhuatai District
Nanjing,
China

Phone:
Email: zhang.zheng@zte.com.cn

Cui(Linda) Wang
ZTE Corporation
No. 50 Software Ave, Yuhuatai District
Nanjing,
China

Phone:
Email: wang.cui1@zte.com.cn

