Bit Indexed Explicit Replication
Internet-Draft
Intended status: Standards Track
Expires: 25 July 2026

P. Psenak
Cisco Systems, Inc.
N. Nainar, Ed.
Oracle
IJ. Wijnands
Individual Contributor
Z. Zhang, Ed.
ZTE Corporation
21 January 2026

OSPFv3 Extensions for BIER
draft-ietf-bier-ospfv3-extensions-09

Abstract

Bit Index Explicit Replication (BIER) is an architecture that provides multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain multicast related per-flow state. The BIER architecture uses MPLS or other encapsulations to steer the multicast traffic towards the receivers.

This document describes the OSPFv3 protocol extensions required for BIER with MPLS encapsulation. Support for other encapsulation types is outside the scope of this document.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Bit Index Explicit Replication (BIER) [RFC8279] is an architecture that provides optimal multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain any multicast related per-flow state. BIER also does not explicitly require a tree-building protocol for its operation. A multicast data packet enters a BIER domain at a "Bit-Forwarding Ingress Router" (BFIR), and leaves the BIER domain at one or more "Bit-Forwarding Egress Routers" (BFERs). The BFIR router adds a BIER header to the packet. The BIER header contains a bit-string in which each bit represents exactly one BFER to which the packet could be forwarded. The set of BFERs to which the multicast packet needs to be forwarded is expressed by setting the bits that correspond to those routers in the BIER header.

The BIER architecture requires routers participating in BIER to exchange BIER related information within a given domain. The BIER architecture permits link-state routing protocols to perform distribution of such information. [RFC8444] defines the OSPFv2 protocol extensions to distribute BIER specific information. This document describes extensions to OSPFv3 to enable it to advertise BIER specific information in the case where BIER uses MPLS encapsulation as described in [RFC8296].

1.1. Requirements Language

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

2. Flooding of BIER Information in OSPFv3

All BIER specific information that a Bit-Forwarding Router (BFR) needs to advertise to other BFRs is associated with a BFR-Prefix. A BFR prefix is a unique (within a given BIER domain) routable IPv4 or IPv6 address that is assigned to each BFR as described in more detail in [RFC8279].

[RFC8362] defines the format of TLV that allows additional information to be carried in OSPFv3 LSAs. This section defines the required Sub-TLVs to carry BIER information that is associated with the BFR-Prefix. The Sub-TLV defined in this section can be carried in the OSPFv3 Extended LSA TLVs [RFC8362] listed below:

Inter-Area-Prefix TLV

External-Prefix TLV

Intra-Area-Prefix TLV

2.1. BIER Sub-TLV

A Sub-TLV of the above mentioned Prefix TLVs is defined for distributing BIER information. The Sub-TLV is called the BIER Sub-TLV. Multiple BIER Sub-TLVs may be included in any of the above mentioned Prefix TLV.

The format is the same with the definition in [RFC8444]:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Sub-domain-ID |     MT-ID     |              BFR-id           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     BAR       |     IPA       |         Reserved              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Sub-TLVs (variable)                      |
+-                                                             -+
|                                                               |
```

Type: TBD1

Length: The length in octets of the BIER Sub-TLV’s value part.

Sub-domain-ID: Unique value identifying the BIER sub-domain within the BIER domain, as described in [RFC8279]

MT-ID: Multi-Topology ID (as defined in [RFC4915])that identifies the topology that is associated with the BIER sub-domain.

BFR-id: A 2-octet field encoding the BFR-id, as documented in Section 2 of [RFC8279]. Typically, a BFR as an intermediate node only is not assigned a BFR-id, but still needs to announce and flood the BIER Sub-TLV and related MPLS encapsulation Sub-TLV for BIER BIFT construction.

BAR: Single octet BIER specific algorithm used to calculate underlay paths to reach other BFRs. Values are allocated from the "BIER Algorithm" registry which is defined in [RFC8401].

IPA: Single octet IGP algorithm to either modify, enhance or replace the calculation of underlay paths to reach other BFRs as defined by the BAR value. Values are defined in the "IGP Algorithm Types" registry in [RFC8665].

Reserved: A 2-octet field, MUST be set to 0 on transmission and MUST be ignored by the receiver.

Each BIER sub-domain MUST be associated with one and only one OSPF topology that is identified by the MT-ID. If the association between BIER sub-domain and value of the MT-ID field advertised in the BIER Sub-TLV by other BFRs is in conflict with the association locally configured on the receiving router, the received BIER Sub-TLV MUST be ignored.

If the MT-ID value is outside of the values specified in [RFC4915], the BIER Sub-TLV MUST be ignored by the receiver.

A Prefix-TLV can be used for multiple BIER sub-domains. For each sub-domain indicated by Sub-domain-ID, there is only one BIER Sub-TLV. If a BFR advertises the same Sub-domain-ID in multiple BIER Sub-TLVs, the BFR MUST be treated as if it did not advertise a BIER Sub-TLV for such sub-domain.

All BFRs MUST detect advertisement of duplicate valid BFR-IDs for a given Sub-domain-ID. When such duplication is detected by the BFR, it MUST behave as described in section 5 of [RFC8279].

The supported BAR and IPA algorithms MUST be consistent for all routers supporting a given BFR sub-domain. A router receiving BIER Sub-TLV advertisement with a value in BAR or IPA fields which does not match the locally configured value for a given BFR sub-domain, MUST report a misconfiguration for such BIER sub-domain and MUST ignore such BIER Sub-TLV.

Implementations should set the BAR and IPA fields to zero by default. Other values may be carried in these fields, but the processing is outside the scope of this document.

When the BIER Sub-TLV is ignored due to any of the reasons specified in this section, the flooding of the TLV is not affected.

2.2. BIER MPLS Encapsulation Sub-TLV

The BIER MPLS Encapsulation Sub-TLV is a Sub-TLV of the BIER Sub-TLV defined in Section 2.1. The BIER MPLS Encapsulation Sub-TLV is used in order to advertise MPLS specific information used for BIER. It MAY appear multiple times in the BIER Sub-TLV.

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

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Max SI    |                     Label                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|BS Len |                     Reserved                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Type: Set to TBD2.

Length: 8 octets

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 the bitstring length indicated by the BS Len field.

Label: A 3-octet field, where the 20 rightmost bits represent the first label in the label range. The 4 leftmost bits MUST be ignored by the receiver.

Bit String Length: A 4 bits field indicating the supported BitString length associated with this BFR-prefix using the exponential encoding defined in section 2.1.2 [RFC8296]. The set of values allowed in this field are specified in that section.

Reserved: A 28 bits field, MUST be set to 0 on transmission and MUST be ignored by the receiver.

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 Set Identifiers (SI) (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 Set Identifier exceeds the 20 bit range, the BIER MPLS Encapsulation Sub-TLV MUST be ignored by the receiver.

If the BS length is set to a value that does not match any of the allowed values specified in [RFC8296], the BIER MPLS Encapsulation Sub-TLV MUST be ignored.

If same BS length is repeated in multiple BIER MPLS Encapsulation Sub-TLV inside the same BIER Sub-TLV, all MPLS encapsulation Sub-TLVs MUST be ignored by the receiver.

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 by the receiver.

When the BIER Sub-TLV is ignored due to any of the reasons specified in this section, the flooding of the TLV is not affected..

2.3. Flooding scope of BIER Information

The flooding scope of the Extended LSAs [RFC8362] that is used for advertising the BIER Sub-TLV is area-local. To allow BIER deployment in a multi-area environment, OSPFv3 must propagate BIER information between areas.

```
                 (  )         (  )         (  )
               (      )     (      )     (      )
            R1  Area 1   R2  Area 0   R3  Area 2  R4
               (      )     (      )     (      )
                 (  )         (  )         (  )
```

Figure 1: BIER propagation between areas

The following procedure is used in order to propagate BIER related information between areas:

When an OSPFv3 Area Border Router (ABR) advertises E-Inter-Area-Prefix-LSA from an intra-area or inter-area prefix to all its attached areas, it determines whether a BIER Sub-TLV should be included in this LSA. To achieve this, an OSPFv3 ABR will:

- Examine its best path to the prefix in the source area and find the advertising router associated with the best path to that prefix.

- Determine if such advertising router advertised a BIER Sub-TLV for the prefix. If yes, the ABR will copy the information from such BIER Sub-TLV when advertising BIER Sub-TLV to each attached area.

In the Figure 1, R1 advertises a prefix 2001:db8:b1e6::1/128 in Area 1. It also includes BIER Sub-TLV in E-Intra-Area-Prefix-LSA. ABR R2 calculates the reachability for prefix 2001:bdb8:b1e6::1/128 inside Area 1 and propagates it to Area 0 using E-Inter-Area-Prefix-LSA. When doing so, it copies the entire BIER Sub-TLV (including all its Sub-TLVs) it received from R1 in Area 1 and includes it in the E-Inter-Area-Prefix-LSA it generates for the prefix in Area 0. ABR R3 calculates the reachability for prefix 2001:bdb8:b1e6::1/128 inside Area 0 and propagates it to Area 2. When doing so, it copies the entire BIER Sub-TLV (including all its Sub-TLVs) it received from R2 in Area 0 and includes it in E-Inter-Area-Prefix-LSA it generates for 2001:bdb8:b1e6::1/128 in Area 2.

3. Security Considerations

This document introduces new Sub-TLVs for OSPFv3 Extended-LSAs. It does not introduce any new security risks to OSPFv3. Existing security concerns documented in [RFC8362] is applicable for the Sub-TLVs defined in this document.

It is assumed that both BIER and OSPF layer is under a single administrative domain. There can be deployments where potential attackers have access to one or more networks in the OSPFv3 routing domain. In these deployments, stronger authentication mechanisms such as those specified in [RFC4552] SHOULD be used.

The Security Considerations section of [RFC8279] discusses the possibility of performing a Denial of Service (DoS) attack by setting too many bits in the BitString of a BIER-encapsulated packet. However, this sort of DoS attack cannot be initiated by modifying the OSPF BIER advertisements specified in this document. A BFIR decides which systems are to receive a BIER-encapsulated packet. In making this decision, it is not influenced by the OSPF control messages. When creating the encapsulation, the BFIR sets one bit in the encapsulation for each destination system. The information in the OSPF BIER advertisements is used to construct the forwarding tables that map each bit in the encapsulation into a set of next hops for the host that is identified by that bit, but is not used by the BFIR to decide which bits to set. Hence an attack on the OSPF control plane cannot be used to cause this sort of DoS attack.

While a BIER-encapsulated packet is traversing the network, a BFR that receives a BIER-encapsulated packet with n bits set in its BitString may have to replicate the packet and forward multiple copies. However, a given bit will only be set in one copy of the packet. That means that each transmitted replica of a received packet has fewer bits set (i.e., is targeted to fewer destinations) than the received packet. This is an essential property of the BIER forwarding process as defined in [RFC8279]. While a failure of this process might cause a DoS attack (as discussed in the Security Considerations of [RFC8279]), such a failure cannot be caused by an attack on the OSPF control plane.

Implementations MUST assure that malformed TLV and Sub-TLV defined in this document are detected and do not provide a vulnerability for attackers to crash the OSPFv3 router or routing process. Reception of malformed TLV or Sub-TLV SHOULD be counted and/or logged for further analysis. Logging of malformed TLVs and Sub-TLVs SHOULD be rate-limited to prevent a Denial of Service (DoS) attack (distributed or otherwise) from overloading the OSPFv3 control plane.

4. IANA Considerations

The document requests two new allocations from the OSPFv3 Extended-LSA Sub-TLVs registry as defined in [RFC8362] with the range: 4-32767.

BIER Sub-TLV: TBD1

BIER MPLS Encapsulation Sub-TLV: TBD2

5. Acknowledgements

The authors would like to thank Mankamana Mishra, Tony Przygienda, Huaimo Chen and Greg Shepherd for their review comments.

6. 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>.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <https://www.rfc-editor.org/info/rfc4552>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <https://www.rfc-editor.org/info/rfc4915>.

[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>.

[RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <https://www.rfc-editor.org/info/rfc8362>.

[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>.

[RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <https://www.rfc-editor.org/info/rfc8665>.

Authors’ Addresses

Peter Psenak
Cisco Systems, Inc.
Apollo Business Center
821 09 Mlynske nivy 43
Slovakia
Email: ppsenak@cisco.com

Nagendra Kumar Nainar (editor)
Oracle
Research Triangle Park, NC 27709
United States of America

Email: nagendrakumar.nainar@gmail.com

IJsbrand Wijnands
Individual Contributor
1831 Diegem
Belgium
Email: ice@braindump.be

Zheng Zhang (editor)
ZTE Corporation
China
Email: zhang.zheng@zte.com.cn