Internet Engineering Task Force (IETF)
Request for Comments: 8444
Category: Standards Track
ISSN: 2070-1721

P. Psenak, Ed.
N. Kumar
IJ. Wijnands
Cisco
A. Dolganow
Nokia
T. Przygienda
J. Zhang
Juniper Networks, Inc.
S. Aldrin
Google, Inc.

November 2018

OSPFv2 Extensions for Bit Index Explicit Replication (BIER)

Abstract

Bit Index Explicit Replication (BIER) is an architecture that provides optimal multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain multicast-related, per-flow state. BIER also does not require an explicit 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 adds a BIER packet header to the packet. The BIER packet header contains a BitString in which each bit represents exactly one BFER to forward the packet to. The set of BFERs to which the multicast packet needs to be forwarded is expressed by the set of bits in the BIER packet header.

This document describes the OSPF protocol extension (from RFC 2328) that is required for BIER with MPLS encapsulation (which is defined in RFC 8296). Support for other encapsulation types and the use of multiple encapsulation types are outside the scope of this document.

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/rfc8444.

Table of Contents

1. Introduction 3
2. Flooding of the BIER Information in OSPF 4
2.1. BIER Sub-TLV 4
2.2. BIER MPLS Encapsulation Sub-TLV 5
2.3. Flooding Scope of BIER Information 7
3. Security Considerations 8
4. IANA Considerations 9
5. References 9
5.1. Normative References 9
5.2. Informative References 10
Acknowledgments 11
Authors' Addresses 11

1. Introduction

Bit Index Explicit Replication (BIER) is an architecture that provides optimal multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain any multicast-related, per-flow state. Neither does BIER 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 packet header to the packet. The BIER packet header contains a BitString in which each bit represents exactly one BFER to forward the packet to. The set of BFERs to which the multicast packet needs to be forwarded is expressed by the set of bits in the BIER packet header.

The BIER architecture requires routers participating in BIER to exchange BIER-related information within a given domain and permits link-state routing protocols to perform distribution of such information. This document describes extensions to OSPF necessary to advertise BIER-specific information in the case where BIER uses MPLS encapsulation as described in [RFC8296].

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 the BIER Information in OSPF

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 IP address that is assigned to each BFR as described in detail in Section 2 of [RFC8279].

Given that BIER information must be associated with a BFR-prefix, the OSPFv2 Extended Prefix Opaque LSA [RFC7684] has been chosen for advertisement.

2.1. BIER Sub-TLV

A sub-TLV of the OSPFv2 Extended Prefix TLV (defined in [RFC7684]) is defined for distributing BIER information. The sub-TLV is called the BIER Sub-TLV. Multiple BIER Sub-TLVs may be included in the OSPFv2 Extended Prefix TLV.

The BIER Sub-TLV has the following format:

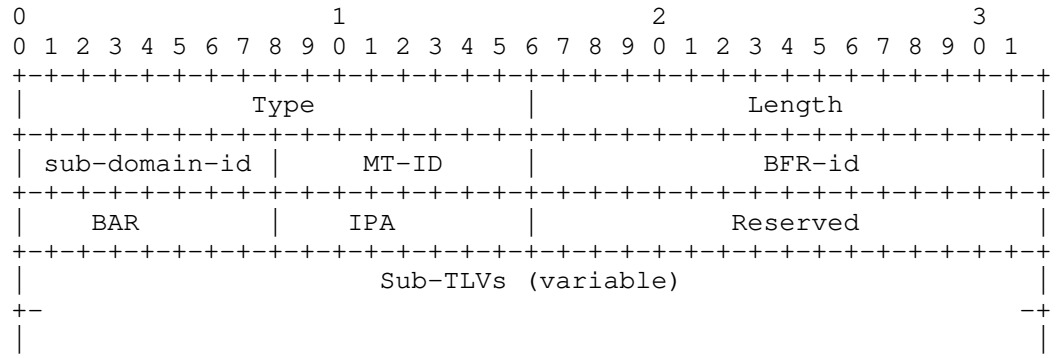

Type: 9

Length: Variable, dependent on sub-TLVs.

sub-domain-id: Unique value identifying the BIER sub-domain within the BIER domain, as described in Section 1 of [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]. If the BFR is not locally configured with a valid BFR-id, the value of this field is set to 0, which is defined as illegal in [RFC8279].

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

IPA: Single-octet IGP Algorithm used 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 [IANA-IGP].

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

If the MT-ID contains an invalid value as specified in [RFC4915], the BIER Sub-TLV for such subdomains with conflict MUST be ignored.

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 MT-ID and 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. If a router receives a BIER Sub-TLV advertisement with a value in the BAR or IPA fields that does not match the locally configured value for a given BFR sub-domain, the router MUST report a misconfiguration for such BIER sub-domain and MUST ignore the BIER Sub-TLV containing the error.

The use of non-zero values in either the BAR field or the IPA field is outside the scope of this document.

2.2. BIER MPLS Encapsulation Sub-TLV

The BIER MPLS Encapsulation Sub-TLV is a sub-TLV of the BIER Sub-TLV. 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:

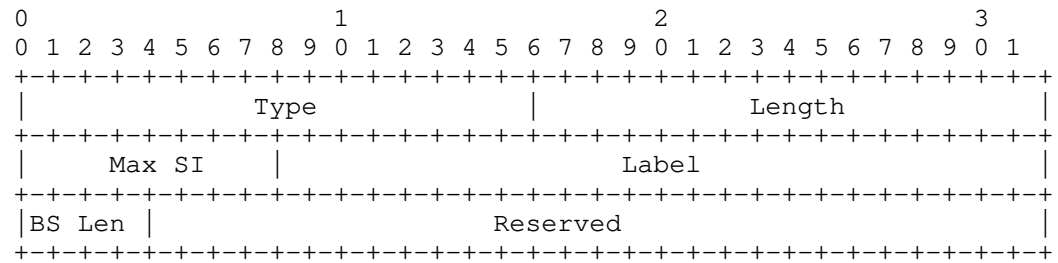

Type: 10

Length: 8 octets

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

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.

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].

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

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 Set Identifier exceeds the 20-bit range, the BIER MPLS Encapsulation Sub-TLV containing the error MUST be ignored.

If the BitString length is set to a value that does not match any of the allowed values specified in [RFC8296], 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 Sub-TLV, the whole BIER Sub-TLV containing the conflicts 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 sub-TLVs advertised by such a router MUST be ignored.

2.3. Flooding Scope of BIER Information

The flooding scope of the OSPFv2 Extended Prefix Opaque LSA [RFC7684] that is used for advertising the BIER Sub-TLV is set to area-local. To allow BIER deployment in a multi-area environment, OSPF must propagate BIER information between areas.

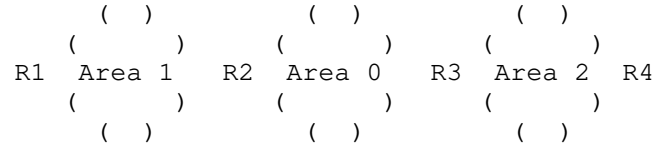

Figure 1: BIER Propagation between Areas

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

When an OSPF Area Border Router (ABR) advertises a Type-3 Summary LSA from an intra-area or inter-area prefix to all its attached areas, it will also originate an OSPFv2 Extended Prefix Opaque LSA, as described in [RFC7684]. The flooding scope of the OSPFv2 Extended Prefix Opaque LSA type will be set to area-local. The route-type in the OSPFv2 Extended Prefix TLV is set to inter-area. When determining whether a BIER Sub-TLV should be included in this LSA, an OSPF 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 the advertising router advertised a BIER Sub-TLV for the prefix. If yes, the ABR will copy the information from that BIER Sub-TLV when advertising the BIER Sub-TLV to each attached area.

In Figure 1, R1 advertises a prefix 192.0.2.1/32 in Area 1. It also advertises an OSPFv2 Extended Prefix Opaque LSA for prefix 192.0.2.1/32 and includes a BIER Sub-TLV in it. ABR R2 calculates the reachability for prefix 192.0.2.1/32 inside Area 1 and

propagates it to Area 0. When doing so, it copies the entire BIER Sub-TLV (including all of its Sub-TLVs) that it received from R1 in Area 1 and includes it in the OSPFv2 Extended Prefix Opaque LSA it generates for 192.0.2.1/32 in Area 0. ABR R3 calculates the reachability for prefix 192.0.2.1/32 inside Area 0 and propagates it to Area 2. When doing so, it copies the entire BIER Sub-TLV (including all of its sub-TLVs) that it received from R2 in Area 0 and includes it in the OSPFv2 Extended Prefix Opaque LSA it generates for 192.0.2.1/32 in Area 2.

3. Security Considerations

This document introduces new sub-TLVs for the existing OSPFv2 Extended Prefix TLV. It does not introduce any new security risks to OSPF. Existing security extensions as described in [RFC2328] and [RFC7684] apply.

It is assumed that both the BIER and OSPF layers are under a single administrative domain. There can be deployments where potential attackers have access to one or more networks in the OSPF routing domain. In these deployments, stronger authentication mechanisms such as those specified in [RFC7474] 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 the information 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. This 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 section of [RFC8279]), such a failure cannot be caused by an attack on the OSPF control plane.

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

4. IANA Considerations

IANA has allocated the following from the "OSPFv2 Extended Prefix TLV Sub-TLVs" registry defined in [RFC7684].

BIER Sub-TLV: 9

BIER MPLS Encapsulation Sub-TLV: 10

5. References

5.1. 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>.

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

[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>.

[RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", RFC 7474, DOI 10.17487/RFC7474, April 2015, <https://www.rfc-editor.org/info/rfc7474>.

[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>.

[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>.

[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>.

5.2. Informative References

[IANA-IGP] IANA, "IGP Algorithm Types", <https://www.iana.org/assignments/igp-parameters/>.

Acknowledgments

The authors would like to thank Rajiv Asati, Christian Martin, Greg Shepherd, and Eric Rosen for their contributions.

Authors' Addresses

Peter Psenak (editor)
Cisco
Apollo Business Center
Mlynske nivy 43
Bratislava 821 09
Slovakia

Email: ppsenak@cisco.com

Nagendra Kumar
Cisco
7200 Kit Creek Road
Research Triangle Park, NC 27709
United States of America

Email: naikumar@cisco.com

IJsbrand Wijnands
Cisco
De Kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com

Andrew Dolganow
Nokia
750 Chai Chee Rd
06-06 Viva Business Park
Singapore 469004
Singapore

Email: andrew.dolganow@nokia.com

Tony Przygienda
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886
United States of America

Email: prz@juniper.net

Jeffrey Zhang
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886
United States of America

Email: zzhang@juniper.net

Sam Aldrin
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA
United States of America

Email: aldrin.ietf@gmail.com