

Internet Engineering Task Force (IETF)
Request for Comments: 9272
Updates: 8401, 8444
Category: Standards Track
ISSN: 2070-1721

Z. Zhang
A. Przygienda
Juniper Networks
A. Dolganow
Individual
H. Bidgoli
Nokia
IJ. Wijnands
Individual
A. Gulko
Edward Jones Wealth Management
September 2022

Underlay Path Calculation Algorithm and Constraints for Bit Index
Explicit Replication (BIER)

Abstract

This document specifies general rules for the interaction between the BIER Algorithm (BAR) and the IGP Algorithm (IPA) used for underlay path calculation within the Bit Index Explicit Replication (BIER) architecture. The semantics defined in this document update RFC 8401 and RFC 8444. This document also updates the "BIER Algorithm" registry established in RFC 8401.

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/rfc9272>.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
 - 1.1. Requirements Language
2. Updated Definitions for IPA and BAR Fields
3. General Rules for the BAR and IPA Interaction
 - 3.1. When BAR Is Not Used
 - 3.2. Exceptions or Extensions to the General Rules
4. Examples
5. IANA Considerations
6. Security Considerations
7. Normative References

1. Introduction

In the Bit Index Explicit Replication (BIER) architecture [RFC8279], packets with a BIER encapsulation header are forwarded to the neighbors on the underlay paths towards Bit-Forwarding Egress Routers (BFRs) that are represented by bits set in the BIER header's BitString. The paths are calculated in the underlay topology for each sub-domain following a calculation algorithm specific to the sub-domain. The topology or algorithm may or may not be congruent with unicast. The algorithm could be a BIER-specific algorithm or could be a generic IGP one, e.g., Shortest Path First (SPF).

In [RFC8401] and [RFC8444], an 8-bit BAR (BIER Algorithm) field and 8-bit IPA (IGP Algorithm) field are defined to signal the BIER-specific algorithm and generic IGP Algorithm, respectively, and only value 0 is allowed for both fields in those two documents.

This document specifies general rules for the interaction between the BIER Algorithm (BAR) and the IGP Algorithm (IPA) used for underlay path calculation when other BAR and/or IPA values are used. The semantics defined in this document update [RFC8401] and [RFC8444]. This document also updates the "BIER Algorithm" registry defined in [RFC8401] by renaming the "Experimental Use" range to "Private or Experimental Use".

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. Updated Definitions for IPA and BAR Fields

The definitions for the IPA and BAR fields in Section 6.1 of [RFC8401] and Section 2.1 of [RFC8444] are updated as follows.

IPA: IGP Algorithm. Specifies a generic Routing Algorithm and related Routing Constraints to calculate underlay paths to reach other Bit-Forwarding Routers (BFRs). Values are from the "IGP Algorithm Types" registry. One octet.

BAR: BIER Algorithm. Specifies a BIER-specific Algorithm and BIER-specific Constraints used to either modify, enhance, or replace the calculation of underlay paths to reach other BFRs as defined by the IPA value. Values are allocated from the "BIER Algorithm" registry. One octet.

When a BAR value is defined, the corresponding BIER-specific Algorithm (BA) and BIER-specific Constraint (BC) semantics SHOULD be specified. For an IGP Algorithm to be used as a BIER IPA, its Routing Algorithm (RA) and Routing Constraint (RC) semantics SHOULD be specified. If any of these semantics is not specified, it MUST be interpreted as the "NULL" algorithm or constraint. For example, the IGP Algorithm 0 defined in [RFC8665] is treated as having a NULL RC, i.e., no constraints (see Section 3).

If a specification is not available for a specific BAR value, its value MUST be from the Private or Experimental Use range of the registry.

3. General Rules for the BAR and IPA Interaction

For a particular sub-domain, all BFRs MUST be provisioned with and signal the same BAR and IPA values. If a BFR discovers another BFR advertising a different BAR or IPA value for a sub-domain, it MUST treat the advertising router as incapable of supporting BIER for that

sub-domain. (One way of handling incapable routers is documented in Section 6.9 of [RFC8279], and additional methods may be defined in the future.)

For a particular topology X that a sub-domain is associated with, a router MUST calculate the underlay paths according to its BAR and IPA values in the following way:

1. Apply the BIER constraints, resulting in BC(X). If BC is NULL, then BC(X) is X itself.
2. Apply the routing constraints, resulting in RC(BC(X)). If RC is NULL, then RC(BC(X)) is BC(X).
3. Select the algorithm AG as follows:
 - a. If BA is NULL, AG is set to RA.
 - b. If BA is not NULL, AG is set to BA.
4. Run AG on RC(BC(X)).

It's possible that the resulting AG is not applicable to BIER. In that case, no BIER paths will be calculated, and this is a network design issue that an operator needs to avoid when choosing the BAR or IPA.

3.1. When BAR Is Not Used

BAR value 0 is defined as "No BIER-specific algorithm is used" [RFC8401]. This value indicates NULL BA and BC. Following the rules defined above, the IPA value alone identifies the calculation algorithm and constraints to be used for a particular sub-domain.

3.2. Exceptions or Extensions to the General Rules

Exceptions or extensions to the above general rules may be specified in the future for specific BAR and/or IPA values. When that happens, compatibility with defined BAR and/or IPA values and semantics need to be specified.

4. Examples

As an example, one may define a new BAR with a BIER-specific constraint of "excluding BIER-incapable routers". No BIER-specific algorithm is specified, and the BIER-specific constraint can go with any IPA, i.e., any RC defined by the IPA is augmented with "excluding BIER-incapable routers". (Routers that do not support BIER are not considered when applying the IGP Algorithm.)

If the BC and RC happen to conflict and lead to an empty topology, then no BIER forwarding path will be found. For example, the BC could be "exclude BIER-incapable routers", and the RC could be "include green links only". If all the green links are associated with BIER-incapable routers, it results in an empty topology. This is a network design issue that an operator needs to avoid when choosing the BAR or IPA.

In another example, a BAR value can be specified to use the Steiner tree algorithm and used together with IPA 0 (which uses an SPF algorithm). According to the general rules, the BIER-specific algorithm takes precedence so SPF is not used.

5. IANA Considerations

The "BIER Algorithm" registry has been updated as follows:

1. The "Experimental Use" range has been renamed "Private or Experimental Use".
2. This document has been added as a reference both for the registry

itself and for values 240-254 in the registry.

6. Security Considerations

This document specifies general rules for the interaction between the BIER Algorithm (BAR) and the IGP Algorithm (IPA) used for underlay path calculation. It does not change the security aspects as discussed in [RFC8279], [RFC8401], and [RFC8444].

7. 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>>.
- [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>>.
- [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>>.

Acknowledgements

The authors thank Alia Atlas, Eric Rosen, Senthil Dhanaraj and many others for their suggestions and comments. In particular, the BC/BA/RC/RA representation for the interaction rules is based on Alia's write-up.

Authors' Addresses

Zhaohui Zhang
Juniper Networks
Email: zhang@juniper.net

Antoni Przygienda
Juniper Networks
Email: prz@juniper.net

Andrew Dolganow
Individual
Email: adolgano@gmail.com

Hooman Bidgoli
Nokia
Email: hooman.bidgoli@nokia.com

IJsbrand Wijnands
Individual
Email: ice@braindump.be

Arkadiy Gulko
Edward Jones Wealth Management
Email: arkadiy.gulko@edwardjones.com