

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 28, 2018

L. Ginsberg, Ed.
Cisco Systems
A. Przygienda
Juniper Networks
S. Aldrin
Google
J. Zhang
Juniper Networks, Inc.
July 27, 2017

BIER support via ISIS
draft-ietf-bier-isis-extensions-05

Abstract

Specification of an ISIS extension to support BIER domains and sub-domains.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)] .

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 January 28, 2018.

Copyright Notice

Copyright (c) 2017 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	2
2.	Terminology	3
3.	IANA Considerations	4
4.	Concepts	4
4.1.	BIER Domains and Sub-Domains	4
4.2.	Advertising BIER Information	5
5.	Procedures	5
5.1.	Enabling a BIER Sub-Domain	5
5.2.	Multi Topology and Sub-Domain	6
5.3.	Encapsulation	6
5.4.	Tree Type	6
5.5.	Label advertisements for MPLS Encapsulation	6
5.6.	BFR-id Advertisements	6
5.7.	Reporting Misconfiguration	7
5.8.	Flooding Reduction	7
6.	Packet Formats	7
6.1.	BIER Info sub-TLV	7
6.2.	BIER MPLS Encapsulation sub-sub-TLV	8
6.3.	Optional BIER sub-domain Tree Type sub-sub-TLV	9
6.4.	Optional BIER sub-domain BSL conversion sub-sub-TLV	9
7.	Security Considerations	10
8.	Acknowledgements	10
9.	References	10
9.1.	Normative References	10
9.2.	Informative References	11
	Authors' Addresses	11

[1.](#) Introduction

Bit Index Explicit Replication (BIER)

[I-D.[draft-ietf-bier-architecture-07](#)] defines an architecture where all intended multicast receivers are encoded as bitmask in the Multicast packet header within different encapsulations such as [I-D.[draft-ietf-bier-mpls-encapsulation-07](#)]. A router that receives such a packet will forward the packet based on the Bit Position in the packet header towards the receiver(s), following a precomputed

tree for each of the bits in the packet. Each receiver is represented by a unique bit in the bitmask.

This document presents necessary extensions to the currently deployed ISIS for IP [[RFC1195](#)] protocol to support distribution of information necessary for operation of BIER domains and sub-domains. This document defines a new TLV to be advertised by every router participating in BIER signaling.

2. Terminology

Some of the terminology specified in [I-D.[draft-ietf-bier-architecture-07](#)] 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). Each BFIR must have a valid BFR-id assigned.

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 signalled 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-TLV to the registry of sub-TLVs for TLVs 235, 237 [[RFC5120](#)] and TLVs 135, 236 [[RFC5305](#)], [[RFC5308](#)].

Value: 32 (suggested - to be assigned by IANA)

Name: BIER Info

This document also introduces a new registry for sub-sub-TLVs for the BIER Info sub-TLV added above. The registration policy is Expert Review as defined in [[RFC8126](#)]. This registry is part of the "IS-IS TLV Codepoints" registry. The name of the registry is "sub-sub-TLVs for BIER Info sub-TLV". The defined values are:

Type	Name
----	----
1	BIER MPLS Encapsulation
2	BIER sub-domain Tree Type
3	BIER sub-domain BSL conversion

4. Concepts

4.1. BIER Domains and Sub-Domains

An ISIS signalled BIER domain is aligned with the scope of distribution of BFR-prefixes that identify the BFRs within ISIS. ISIS acts in such a case as the supporting BIER underlay.

Within such a domain, the extensions defined in this document advertise BIER information for one or more BIER sub-domains. Each sub-domain is uniquely identified by a subdomain-id. Each subdomain is associated with a single ISIS topology [[RFC5120](#)], which may be any of the topologies supported by ISIS. Local configuration controls which <MT,SD> pairs are supported by a router. The mapping of sub-domains to topologies MUST be consistent within a BIER flooding domain.

Each BIER sub-domain has as its unique attributes the encapsulation used and the type of tree it is using to forward BIER frames (currently always SPF). Additionally, per supported bitstring length in the sub-domain, each router will advertise the necessary label ranges to support it.

At present, IS-IS support for a given BIER domain/sub-domain is limited to a single area - or to the IS-IS L2 sub-domain.

4.2. Advertising BIER Information

BIER information advertisements are associated with a new sub-TLV in the extended reachability TLVs. BIER information is always associated with a host prefix which MUST be a node address for the advertising node. The following restrictions apply:

- o Prefix length MUST be 32 for an IPv4 prefix or 128 for an IPv6 prefix
- o When the Prefix Attributes Flags sub-TLV is present N flag MUST be set and X and R flags MUST NOT be set. [[RFC7794](#)]
- o BIER sub-TLVs MUST NOT be included when a prefix reachability advertisement is leaked between levels.

5. Procedures

5.1. Enabling a BIER Sub-Domain

A given sub-domain with identifier SD with supported bitstring lengths MLs in a multi-topology MT [[RFC5120](#)] is denoted further as <MT,SD,MLs> and does not have to be advertised by default by BFRs to preserve the scaling of the protocol (i.e. ISIS carries no TLVs containing any of the elements related to <MT,SD>). The advertisement may be triggered e.g. by a first BIER sub-TLV ([Section 6.1](#)) containing <MT,SD> advertised into the Level 1 area (or the Level 2 IS-IS sub-domain). The specific trigger itself is outside the scope of this RFC but can be for example a VPN desiring to initiate a BIER sub-domain as MI-PMSI [[RFC6513](#)] tree or a pre-configured BFER (since BFERs will always advertise the BIER sub-TLV to make sure they can be reached). It is outside the scope of this document to describe what trigger for a router capable of participating in <MT,SD> is used to start the origination of the necessary information to join into it.

[5.2.](#) Multi Topology and Sub-Domain

A given sub-domain is supported within one and only one topology. All routers in the flooding scope of the BIER sub-TLVs MUST advertise the same sub-domain within the same multi-topology. A router receiving an <MT,SD> advertisement which does not match the locally configured pair MUST report a misconfiguration of the received <MT, SD> pair. All received BIER advertisements associated with the conflicting <MT, SD> pair MUST be ignored.

[5.3.](#) Encapsulation

All routers in the flooding scope of the BIER TLVs MUST advertise the same encapsulation for a given <MT,SD>. A router discovering encapsulation advertised that is different from its own MUST report a misconfiguration of a specific <MT,SD>. All received BIER advertisements associated with the conflicting <MT, SD> pair MUST be ignored.

[5.4.](#) Tree Type

All routers in the flooding scope of the BIER TLVs MAY advertise a supported tree type for a given <MT,SD>. Tree type indicates the algorithm used when calculating the optimal path. Currently only the default algorithm "SPF" is defined - which has a tree type of 0. If no tree type is advertised tree type 0 is assumed. The supported tree type MUST be consistent for all routers supporting a given <MT,SD>.

[5.5.](#) Label advertisements for MPLS Encapsulation

A router that desires to participate in <MT,SD> MUST advertise for each bitstring length it supports in <MT,SD> a label range size that guarantees to cover the maximum BFR-id injected into <MT,SD> (which implies a certain maximum set id per bitstring length as described in [I-D.[draft-ietf-bier-architecture-07](#)]). Any router that violates this condition MUST be excluded from BIER BFTs for <MT,SD>.

[5.6.](#) BFR-id Advertisements

Each BFER/BFIR MAY advertise with its TLV<MT,SD> the BFR-id that it has administratively chosen. A valid BFR-id MUST be unique within the flooding scope of the BIER advertisements. All BFERs/BFIRs MUST detect advertisement of duplicate valid BFR-IDs for a given <MT, SD>. When such duplication is detected all of the routers advertising duplicates MUST be treated as if they did not advertise a valid BFR-id. This implies they cannot act as BFER or BFIR in that <MT,SD>.

5.7. Reporting Misconfiguration

Whenever an advertisement is received which violates any of the constraints defined in this document the receiving router **MUST** report the misconfiguration.

5.8. Flooding Reduction

BIER domain information SHOULD change infrequently. Frequent changes will increase the number of Link State PDU (LSP) updates and negatively impact performance in the network.

6. Packet Formats

All ISIS BIER information is carried within the TLVs 235, 237 [RFC5120] or TLVs 135 [RFC5305], or TLV 236 [RFC5308].

6.1. BIER Info sub-TLV

This sub-TLV carries the information for the BIER sub-domains that the router participates in as BFR. This sub-TLV MAY appear multiple times in a given prefix-reachability TLV - once for each sub-domain supported in the associated topology.

The sub-TLV advertises a single <MT,SD> combination followed by optional sub-sub-TLVs as described in the following sections.

```

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           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Reserved       | subdomain-id       | BFR-id                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: as indicated in IANA section.

Length: 1 octet.

Reserved: MUST be 0 on transmission, ignored on reception. May be used in future versions. 8 bits

subdomain-id: Unique value identifying the BIER sub-domain. 1 octet

BFR-id: A 2 octet field encoding the BFR-id, as documented in [I-D.[draft-ietf-bier-architecture-07](#)]. If no BFR-id has been assigned this field is set to the invalid BFR-id.

6.2. BIER MPLS Encapsulation sub-sub-TLV

This sub-sub-TLV carries the information for the BIER MPLS encapsulation including the label range for a specific bitstring length for a certain <MT,SD>. It is advertised within the BIER Info sub-TLV ([Section 6.1](#)). This sub-sub-TLV MAY appear multiple times within a single BIER info sub-TLV.

On violation of any of the following conditions, the receiving router MUST ignore the encapsulating BIER Info sub-TLV.

- o Label ranges in multiple sub-sub-TLV MUST NOT overlap.
- o Bitstring lengths in multiple sub-sub-TLVs MUST NOT be identical.
- o The sub-sub-TLV MUST include the required bitstring lengths encoded in precisely the same way as in [I-D.[draft-ietf-bier-mpls-encapsulation-07](#)].
- o The label range size MUST be greater than 0.
- o All labels in the range MUST represent valid label values

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	2	3	4	5	6	7	8	9
+-+-+-----+-+-+-----+-+-+-----+-+-+-----+																																							
Type										Length																													
+-+-+-----+-+-+-----+-+-+-----+-+-+-----+																																							
Lbl Range Size										BS Len										Label																			
+-+-+-----+-+-+-----+-+-+-----+-+-+-----+																																							

Type: value of 1 indicating MPLS encapsulation.

Length: 1 octet.

Local BitString Length (BS Len): Encoded bitstring length as per [I-D.[draft-ietf-bier-mpls-encapsulation-07](#)]. 4 bits.

Label Range Size: Number of labels in the range used on encapsulation for this BIER sub-domain for this bitstring length, 1 octet.

Label: First label of the range, 20 bits. The labels are as defined in [I-D.[draft-ietf-bier-mpls-encapsulation-07](#)].

6.3. Optional BIER sub-domain Tree Type sub-sub-TLV

This sub-sub-TLV carries the information associated with the supported BIER tree type for a <MT,SD> combination. It is carried within the BIER Info sub-TLV ([Section 6.1](#)) that the router participates in as BFR. This sub-sub-TLV is optional and its absence has the same semantics as its presence with Tree Type value 0 (SPF). When Tree Type 0 is used it is recommended that this sub-sub-TLV be omitted in order to reduce the space consumed in the parent TLV.

This sub-sub-TLV MUST NOT occur more than once in a BIER Info sub-TLV. If multiple occurrences of this sub-sub-TLV are present in a single BIER Info sub-TLV the encapsulating BIER Info sub-TLV MUST be ignored.

If the tree type (implied or explicitly advertised) does not match the locally configured tree type associated with the matching <MT, SD> pair the encapsulating sub-TLV MUST be ignored.

```

      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           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Tree Type        |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Type: value of 1 indicating BIER Tree Type.

Length: 1 octet.

Tree Type: 1 octet

6.4. Optional BIER sub-domain BSL conversion sub-sub-TLV

This sub-sub-TLV indicates whether the BFR is capable of imposing a different Bit String Length (BSL) than the one it received in a BIER encapsulated packet. Such a capability may allow future, advanced tree types which ensure simple migration procedures from one BSL to another in a given <MT,SD> or prevent stable blackholes in scenarios where not all routers support the same set of BSLs in a given <MT,SD>. It is carried within the BIER Info sub-TLV ([Section 6.1](#)). This sub-sub-TLV is optional and its absence indicates that the

router is NOT capable of imposing different BSLs but will always forward the packet with the BSL unchanged. This sub-sub-TLV MAY occur at most once in a given BIER info sub-TLV. If multiple occurrences of this sub-sub-TLV are received in a given BIER info sub-TLV the encapsulating sub-TLV MUST be ignored.

```

      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           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

7. Security Considerations

Implementations must assure that malformed TLV and Sub-TLV permutations do not result in errors which cause hard protocol failures.

8. Acknowledgements

The RFC is aligned with the [I-D.[draft-ietf-bier-ospf-bier-extensions-07](#)] draft as far as the protocol mechanisms overlap.

Many thanks for comments from (in no particular order) Hannes Gredler, Ijsbrand Wijnands, Peter Psenak and Chris Bowers.

9. References

9.1. Normative References

- [I-D.[draft-ietf-bier-mpls-encapsulation-07](#)]
 - Wijnands et al., IJ., "Bit Index Explicit Replication using MPLS encapsulation", internet-draft [draft-ietf-bier-mpls-encapsulation-07.txt](#), Jun 2017.
- [I-D.[draft-ietf-bier-ospf-bier-extensions-07](#)]
 - Psenak et al., P., "OSPF Extension for Bit Index Explicit Replication", internet-draft [draft-ietf-bier-ospf-bier-extensions-07.txt](#), Jul 2017.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), DOI 10.17487/RFC1195, December 1990, <<http://www.rfc-editor.org/info/rfc1195>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), DOI 10.17487/RFC5120, February 2008, <<http://www.rfc-editor.org/info/rfc5120>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", [RFC 5308](#), DOI 10.17487/RFC5308, October 2008, <<http://www.rfc-editor.org/info/rfc5308>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", [RFC 6513](#), DOI 10.17487/RFC6513, February 2012, <<http://www.rfc-editor.org/info/rfc6513>>.
- [RFC7794] Ginsberg, L., Ed., Decraene, B., Previdi, S., Xu, X., and U. Chunduri, "IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability", [RFC 7794](#), DOI 10.17487/RFC7794, March 2016, <<http://www.rfc-editor.org/info/rfc7794>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 8126](#), DOI 10.17487/RFC8126, June 2017, <<http://www.rfc-editor.org/info/rfc8126>>.

[9.2. Informative References](#)

- [I-D.[draft-ietf-bier-architecture-07](#)]
Wijnands et al., IJ., "Stateless Multicast using Bit Index Explicit Replication Architecture", internet-draft [draft-ietf-bier-architecture-07.txt](#), Jun 2017.

Authors' Addresses

Les Ginsberg (editor)
Cisco Systems
510 McCarthy Blvd.
Milpitas, CA 95035
USA

Email: ginsberg@cisco.com

Tony Przygienda
Juniper Networks

Email: prz@juniper.net

Sam Aldrin
Google
1600 Amphitheatre Parkway
Mountain View, CA
USA

Email: aldrin.ietf@gmail.com

Jeffrey (Zhaohui) Zhang
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886
USA

Email: zzhang@juniper.net

