

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: April 26, 2015

A. Przygienda
Ericsson
L. Ginsberg
Cisco Systems
S. Aldrin
Huawei
J. Zhang
Juniper Networks, Inc.
October 23, 2014

BIER support via ISIS
draft-przygienda-bier-isis-ranges-01

Abstract

Specification of an ISIS extension to support BIER 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 April 26, 2015.

Copyright Notice

Copyright (c) 2014 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 in Extended Reachability TLVs	4
4.2.	BIER Domains	4
5.	Procedures	4
5.1.	Enabling a BIER Domain	4
5.2.	Encapsulation	5
5.3.	Label Advertisements for MPLS encapsulated BIER domains .	5
5.3.1.	Special Consideration	5
5.4.	BFR-id Advertisements	5
5.5.	Flooding	6
6.	Packet Formats	6
6.1.	BIER Info sub-TLV	6
6.2.	BIER MPLS Encapsulation sub-sub-TLV	7
7.	Security Considerations	7
8.	Acknowledgements	8
9.	Normative References	8
	Authors' Addresses	9

[1.](#) Introduction

Bit Index Explicit Replication (BIER)

[I-D.[draft-wijnands-bier-architecture-00](#)] 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-wijnands-mpls-bier-encapsulation-01](#)]. 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 first attempt at necessary extensions to the currently deployed ISIS for IP [[RFC1195](#)] protocol to support distribution of information necessary for operation of BIER domains. This document defines a new TLV to be advertised by every router participating in such BIER domains.

[2.](#) Terminology

Some of the terminology specified in [I-D.[draft-wijnands-bier-architecture-00](#)] 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).

BM: Bit Mask (A bit stream of a certain fixed length. Each Bit represents a receiver).

P-BM: Packet Bit Mask (A Bit Mask included in the Multicast Packet).

BP: Bit Position (A single Bit from the Bit Mask that represents a receiver).

BFR: Bit Forwarding Router (A router that participates in Bit Index Multipoint Forwarding).

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.

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

BIFT: Bit Index Forwarding Table (A 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.

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

IGP signalled BIER domain: A BIER domain information carried in IGP and identified by its multi-topology and bitmask length.

[3.](#) IANA Considerations

This document adds the following new sub-TLVs 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

[4.](#) Concepts

[4.1.](#) BIER Domains in Extended Reachability TLVs

This draft introduces a sub-TLV in the extended reachability TLVs to distribute information about BIER domains and services they carry. To satisfy the requirements for BIER prefixes per [I-D.[draft-wijnands-bier-architecture-00](#)] additional information may be carried in [I-D.[draft-ginsberg-isis-prefix-attributes](#)].

[4.2.](#) BIER Domains

ISIS extensions are capable of carrying BIER information not only for a single BIER domains but for multiple ones. A BIER domain in ISIS is currently always uniquely identified by the tuple of multi-topology MT and bitmask length ML it belongs to denoted as <MT,ML>.

Each such domain itself has as its unique attributes the encapsulation used and the type of tree it is using to forward BIER

frames (currently always SPF).

[5.](#) Procedures

[5.1.](#) Enabling a BIER Domain

A given domain with masklength ML in a multi-topology MT [[RFC5120](#)] (denoted as <MT,ML>) is normally not advertised to preserve the scaling of the protocol (i.e. ISIS carries no TLVs containing any of the elements related to <MT,ML>) and is enabled by a first BIER sub-TLV ([Section 6.1](#)) containing <MT,ML> being advertised into the area. The trigger itself is outside the scope of this draft but can be for example a VPN desiring to initiate a BIER layer as MI-PMSI [[RFC6513](#)] tree. It is outside the scope of this document to describe what trigger for a router capable of participating <MT,ML> is used to start the origination of the necessary information to join into it.

[5.2.](#) Encapsulation

All routers in the flooding scope of the BIER TLVs SHOULD advertise the same encapsulation for a given <MT,ML>. A router discovering encapsulation advertised that is different from its own MUST report a misconfiguration of a specific <MT,ML>. Each router MUST compute BFTs for <MT,ML> using only routers having the same encapsulation as its own advertised encapsulation in BIER sub-TLV for <MT,ML>.

[5.3.](#) Label Advertisements for MPLS encapsulated BIER domains

Each router MAY advertise within the BIER MPLS Encapsulation sub-sub-TLV ([Section 6.2](#)) of a BIER Info sub-TLV ([Section 6.1](#)) for <MT,ML> (denoted further as TLV<MT,ML>) a valid starting label value and a non-zero range length. It MUST advertise a valid label value and a non-zero range length in case it has computed itself as being on the BFT rooted at any of the BFRs with valid BFR-ids (except itself if it does NOT have a valid BFR-id) participating in <MT,ML>.

A router CAN decide to not advertise its TLV<MT,ML> if it does not want to participate in the domain due to resource constraints, label space optimization, administrative configuration or any other

reasons.

[5.3.1.](#) Special Consideration

A router MUST advertise for <MT,ML> label range size that guarantees to cover the maximum BFR-id injected into <MT,ML> (which implies a certain set id as described in [I-D.[draft-wijnands-bier-architecture-00](#)]). Any router that violates this condition MUST be excluded from BIER BFTs for <MT,ML>.

[5.4.](#) BFR-id Advertisements

Each BFER MAY advertise with its TLV<MT,ML> the BFR-id that it has administratively chosen.

If a router discovers that two BFRs it can reach advertise the same value for BFR-id for <MT,ML>, it MUST report a misconfiguration and disregard those routers for all BIER calculations and procedures for <MT,ML> to align with [I-D.[draft-wijnands-bier-architecture-00](#)]. It is worth observing that based on this procedure routers with colliding BFR-id assignments in <MT,ML> MAY still act as BFIRs in <MT,ML> but will be never able to receive traffic from other BFRs in <MT,ML>.

[5.5.](#) Flooding

BIER domain information SHOULD change and force flooding infrequently. Further discussion TBD.

[6.](#) Packet Formats


All ISIS BIER information is carried within the TLVs 235, 237 [[RFC5120](#)] and TLVs 135,236 [[RFC5305](#)],[[RFC5308](#)].

[6.1.](#) BIER Info sub-TLV

This sub-TLV carries the information for the BIER domains that the router participates in as BFR. It can repeat multiple times for different domain <MT,ML> combinations. If the same <MT,ML> domain is

advertised multiple times with different encapsulations, the result is unspecified.

The sub-TLV carries a single <MT,ML> combination followed by optional sub-sub-TLVs specified within its context such as e.g. BIER MPLS Encapsulation per [Section 6.2](#).



Type: as indicated in IANA section.

Length: 1 octet.

Local BitMask Length (BM Len): Bitmask length for this BIER domain that this router is advertising per [I-D.[draft-wijnands-mpls-bier-encapsulation-01](#)]. 4 bits.


Reserved	reserved, must be 0 on transmission, ignored on reception.
4 bits	

BFR-id A 2 octet field encoding the BFR-id, as documented in [I-D.[draft-wijnands-bier-architecture-00](#)]. If set to the invalid BFR-id advertising router is not owning any BFR-id.

6.2. BIER MPLS Encapsulation sub-sub-TLV

This sub-sub-TLV carries the information for the BIER MPLS encapsulations for a certain <MT,ML> and is carried within the BIER

Info sub-TLV ([Section 6.1](#)) that the router participates in as BFR. It can repeat only once within it. If this sub-sub-TLV is included more than once, the result is unspecified.



Type: value of 0 indicating MPLS encapsulation.

Length: 1 octet.

Label Range Size: Number of labels in the range used on encapsulation for this BIER domain, 1 octet. This MUST never be advertise as 0 (zero) and otherwise, this sub-sub-TLV must be treated as if not present for BFT calculations and a misconfiguration SHOULD be reported by the receiving router.

Label: First label of the range used on encapsulation for this BIER domain and service, 20 bits. The label is used for example by [I-D.[draft-wijnands-mpls-bier-encapsulation-01](#)] to forward traffic to sets of BFERs.

```
Reserved    reserved, must be 0 on transmission, ignored on reception.
    4 bits
```

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 draft is aligned with the [I-D.[draft-psenak-ospf-bier-extension-01](#)] draft as far as the protocol mechanisms overlap.

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

9. Normative References

- [I-D.[draft-ginsberg-isis-prefix-attributes](#)]
Ginsberg et al., U., "IS-IS Prefix Attributes for Extended IP and IPv6 Reachability", internet-draft [draft-ginsberg-isis-prefix-attributes-00.txt](#), October 2014.
- [I-D.[draft-psenak-ospf-bier-extension-01](#)]
Psenak, P. and IJ. Wijnands, "OSPF Extension for Bit Index Explicit Replication", internet-draft [draft-ietf-ospf-prefix-link-attr-01.txt](#), October 2014.
- [I-D.[draft-wijnands-bier-architecture-00](#)]
Wijnands, IJ., "Stateless Multicast using Bit Index Explicit Replication Architecture", internet-draft [draft-wijnands-bier-architecture-00.txt](#), February 2014.
- [I-D.[draft-wijnands-mpls-bier-encapsulation-01](#)]
Wijnands et al., IJ., "Bit Index Explicit Replication using MPLS encapsulation", internet-draft [draft-wijnands-mpls-bier-encapsulation-01.txt](#), February 2014.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), December 1990.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4971] Vasseur, JP., Shen, N., and R. Aggarwal, "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", [RFC 4971](#), July 2007.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), February 2008.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), October 2008.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", [RFC 5308](#), October 2008.

[RFC6513] Rosen, E. and R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", [RFC 6513](#), February 2012.

Authors' Addresses

Tony Przygienda
Ericsson
300 Holger Way
San Jose, CA 95134
USA

Email: antoni.przygienda@ericsson.com

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

Email: ginsberg@cisco.com

Sam Aldrin
Huawei
2330 Central Expressway
Santa Clara, CA 95051
USA

Email: aldrin.ietf@gmail.com

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

Email: zzhang@juniper.net

