

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: March 31, 2015

A. Przygienda
Ericsson
September 27, 2014

BIER support via ISIS
draft-przygienda-bier-isis-ranges-00

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 March 31, 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

Internet-Draft [draft-przygienda-bier-isis-ranges-00](#) September 2014

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	3
4.	Concepts	3
4.1.	BIER as Capability	4
4.2.	BIER Domain Identifier	4
5.	Procedures	4
5.1.	Enabling a BIER Domain	4
5.2.	Length of Bitmasks	4
5.2.1.	Special Consideration	5
5.3.	Encapsulation	5
5.4.	Label Advertisements	5
5.4.1.	Special Consideration	5
5.5.	BFR-id Advertisements	5
6.	Packet Formats	6
6.1.	BIER BFR sub-TLV	6
7.	Security Considerations	7
8.	Acknowledgements	7
9.	Normative References	7
	Author's Address	8

[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-00](#)]. 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 [[RFC7142](#)] protocol to support distribution of information necessary for operation of BIER domains. This document defines a new TLV to be distributed by every router participating in such BIER domains.

Internet-Draft [draft-przygienda-bier-isis-ranges-00](#) September 2014

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.

[3.](#) IANA Considerations

This document requests IANA to assign sub-TLV type values from the ISIS router capability TLV [[RFC4971](#)] registry.

[4.](#) Concepts

Przygienda

Expires March 31, 2015

[Page 3]

Internet-Draft

[draft-przygienda-bier-isis-ranges-00](#)

September 2014

[4.1.](#) BIER as Capability

This draft introduces a sub-TLV in the router capabilities TLV [[RFC4971](#)] to distribute the information. Any of the router's loopback addresses that it originates are considered BFR prefixes as required by [I-D.[draft-wijnands-bier-architecture-00](#)]. The question whether a particular loopback address is routable in a specific topology [[RFC5120](#)] can be resolved by [I-D.[draft-xu-isis-routable-ip-address-01](#)].

[4.2.](#) BIER Domain Identifier

ISIS can carry BIER information not only for a single BIER domain but for multiple, distinct domains. This allows to run many disjoint BIER layers within the same Multi-Topology [[RFC5120](#)] easily instead of always forcing different multicast overlays to share the exactly same set of BFRs and resources. Moreover, multi topology [[RFC5120](#)] can be used for the purpose of restricting links that certain set of BIER domains can use or change metrics of such links. A BIER set is therefore always uniquely identified by the tuple of topology T, domain D it belongs to and its number S, denoted as <T,D,S>. The domain itself has as its unique attributes the encapsulation, bitmask length 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 D in a multi-topology T [[RFC5120](#)] (denoted as $\langle T, D \rangle$ from now on) is normally not advertised to preserve the scaling of the protocol (i.e. ISIS carries no TLVs containing any of the elements related to $\langle T, D \rangle$) and is enabled by a first BIER sub-TLV ([Section 6.1](#)) containing $\langle T, D \rangle$ being advertised into the area. The trigger itself is outside the scope of this draft but can be .e.g. a VPN desiring to initiate a domain as MP2MP or P2MP tree or a BMP being administratively assigned to a BFER and advertised via BIER TLV into the area or any other means within Multicast BIER Overlay(s) using BIER domains.

[5.2.](#) Length of Bitmasks

All routers in the flooding scope of the BIER TLVs SHOULD advertise the same bit mask length for a given $\langle T, D \rangle$. A router discovering bitmask lengths advertised that are shorter than its own MUST report a misconfiguration of a specific $\langle T, D \rangle$. Each router MUST compute BFTs for $\langle T, D \rangle$ using only routers having the same mask length as its own advertised Bit Mask Length in BIER sub-TLV for $\langle T, D \rangle$.

[5.2.1.](#) Special Consideration

The same router MAY advertise for different $\langle T, D \rangle$ combinations two different mask lengths. This allows to cleanly delineate domains crossing the same router but using different mask lengths in the encoding, even within the same topology.

[5.3.](#) Encapsulation

Since encapsulation is an attribute of a domain $\langle T, D \rangle$ just like bitmask length, all rules that apply to Bitmask Length per [Section 5.2](#) apply to it well.

[5.4.](#) Label Advertisements

Each router MAY advertise within the sub-TLV of an according $\langle T, D \rangle$ (denoted further as $\text{TLV}\langle T, D \rangle$) a valid starting label value and a non-zero range length. It MUST advertise a valid label value and a non-zero range length IF it has computed itself as being on the BFT rooted at any of the BFRs with valid BFR-ids (except itself) participating in $\langle T, D \rangle$.

A router CAN withdraw its TLV<T,D> if it does not want to participate in the domain due to resource constraints, label space optimization, administrative configuration or any other reasons. In case a router advertises a label range size of 0 for <T,D> it MUST be excluded from the BIER BFTs for <T,D>.

[5.4.1.](#) Special Consideration

A router MUST advertise a for <T,D> label range size that guarantees to cover the maximum BFR-id injected into <T,D> (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 <T,D>.

[5.5.](#) BFR-id Advertisements

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

If two BFRs advertise in their TLV<T,D> the same value for BFR-id, all routers MUST report it as misconfiguration and disregard those routers for all BIER calculations and procedures to align with [I-D.[draft-wijnands-bier-architecture-00](#)]. Such routers with colliding assignments MAY still act as BFIRs but will be never able to receive traffic.

[6.](#) Packet Formats

All ISIS BIER information is carried within the router capability TLV [[RFC4971](#)] with S bit clear.

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

This sub-TLV carries the information for the BIER domains that the router participates in as BFR. It can repeat multiple times. If the same <T,D> is advertised more than once, the first one in the first sub-TLV in the fragment with the lowest ID MUST be used.

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 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|      MT-ID      | Bier Domain ID | Bit Msk Len |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Lbl Range Size | Label |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| BFR-id | A|R|T| Reserv | Encaps |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Type: TBD1.

Length: 2 octets.

MT-ID: Multi-Topology [[RFC5120](#)], 1 octet.

BIER Domain ID: Unique identifier for a BIER domain, 2 octets.

Label Range Size: Number of labest in the range used on encapsulation for this BIER domain, 1 octet.

Label: First label of the range used on encapsulation for this BIER domain, 20 bits. The label is used by e.g. [I-D.[draft-wijnands-mpls-bier-encapsulation-00](#)] to forward traffic to sets of BFERs.

Local BitMask Length: Bitmask length for this BFR per [I-D.[draft-wijnands-bier-architecture-00](#)].

Encapsulation Type: The BIER encapsulation type, 1 octet. Allowed values are:

0 MPLS per [I-D.[draft-wijnands-mpls-bier-encapsulation-00](#)].

A Indicates administratively set value if set, otherwise the BFR-id value MUST be considered as not assigned in this TLV.

R Reserved for future use. MUST be 0.

T Reserved for future use. MUST be 0.

Reserved MUST be 0 on send, ignored on receive.

7. Security Considerations

The extension does not introduce any known new protocol vulnerabilities.

8. Acknowledgements

The draft is aligned with the [I-D.[draft-kumar-ospf-bier-extension-00](#)] draft as far as the protocol mechanisms overlap.

9. Normative References

- [I-D.[draft-kumar-ospf-bier-extension-00](#)]
Psenak, P. and IJ. Wijnands, "OSPF Extension for Bit Index Explicit Replication", internet-draft [draft-ietf-ospf-prefix-link-attr-00.txt](#), September 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-00](#)]
Wijnands et al., IJ., "Bit Index Explicit Replication using MPLS encapsulation", internet-draft [draft-wijnands-mpls-bier-encapsulation-00.txt](#), February 2014.
- [I-D.[draft-xu-isis-routable-ip-address-01](#)]
Chunduri et al., U., "Carrying Routable IP Addresses in IS-IS Router Capability TLV", internet-draft [draft-xu-isis-routable-ip-address-01.txt](#), September 2014.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

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.

[RFC7142] Shand, M. and L. Ginsberg, "Reclassification of [RFC 1142](#) to Historic", [RFC 7142](#), February 2014.

Author's Address

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

Email: antoni.przygienda@ericsson.com