

BIER
Internet-Draft
Intended status: Standards Track
Expires: September 6, 2018

Zhaohui. Zhang
Shaowen. Ma
Juniper Networks
Zheng. Zhang
ZTE Corporation
March 5, 2018

Supporting BIER with RIFT
draft-zzhang-bier-rift-00

Abstract

This document specifies extensions to RIFT protocol to support BIER.

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 [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 <https://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 September 6, 2018.

Copyright Notice

Copyright (c) 2018 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

Internet-Draft

bier-rift

March 2018

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.	Terminologies	2
2.	Introduction	2
3.	Advertising BIER Information For non-MPLS Encapsulation . . .	3
4.	Advertising BIER Information Northbound	4
5.	Advertising BIER Information Southbound	4
5.1.	Local BIER Information	4
5.2.	Proxied BFR-ID Ranges	4
6.	Information Elements Schema	5
6.1.	bier.thrift	5
6.2.	Additions to encoding.thrift	6
7.	IANA Considerations	6
8.	Acknowledgements	6
9.	References	6
9.1.	Normative References	6
9.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Terminologies

Familiarity with BIER and RIFT protocols and procedures is assumed. Some terminologies are listed below for convenience.

[To be added.]

[2.](#) Introduction

BIER [[RFC8279](#)] ... (to be expanded)

RIFT [[I-D.przygienda-rift](#)] is a new protocol specifically designed for CLOS and fat-tree network topologies. As a hybrid between Link State Routing and Distance Vector Routing, it does LSR in northbound (towards the spine) and DVR in southbound (towards the leaves).

[[I-D.ietf-bier-isis-extensions](#)] and [[I-D.ietf-bier-ospf-bier-extensions](#)] specify ISIS/OSPF extensions to support BIER in an ISIS/OSPF domain. The same approach applies to

RIFT in the northbound LSR.

[I-D.zwzw-bier-prefix-redistribute] specifies methods to advertise BIER information via default or summary/aggregate routes advertised

Zhang, et al.

Expires September 6, 2018

[Page 2]

Internet-Draft

bier-rift

March 2018

from one IGP area/domain to another. Similar approach applies to RIFT in the southbound DVR.

BIER encapsulation, whether it is based on MPLS or not, is covered in [\[RFC8296\]](#). However, the OSPF/ISIS extensions for BIER only covers signaling needed for MPLS encapsulation. RIFT is targeted at DC deployments, where MPLS may not be used. This document covers signaling for both BIER MPLS and non-MPLS encapsulation with RIFT.

The details are provided in following sections.

3. Advertising BIER Information For non-MPLS Encapsulation

In the BIER architecture, a BIER sub-domain may have multiple BitStringLengths (BSLs) and multiple Encapsulations (Encaps). A single multicast packet coming from outside the BIER sub-domain may be sent as multiple BIER packets, one for each set that is identified by a SetID (SI). An incoming BIER packet is forwarded according to a BIFT for the <SD,Encap,BSL,SI> tuple. Each BIFT is identified by a 20-bit opaque number (BIFT-ID) in the packet.

With MPLS encapsulation, the BIFT-ID for an incoming BIER packet is simply an MPLS label allocated by the receiving BFR for the BIFT. For each <SD,BSL> tuple, OSPF/ISIS advertises a block of contiguous labels, one label for each SI needed for the tuple, in the MPLS Encapsulation sub-sub-TLV as part of the BIER sub-TLV, which is attached to the Extended Reachability TLV (ISIS case) or the Extended Prefix TLV (OSPF case) for the BFR's BIER Prefix.

With non-MPLS encapsulation, the BIFT-ID in the packet is at the same position as the label in MPLS encapsulation case. Its semantics is no different from the MPLS case in that as an 20-bit opaque value, it leads to the BIFT according to which the BIER packet is forwarded. Beyond the semantics, there are two differences from the MPLS case though:

- o MPLS infrastructure is not needed.
- o While each BFR could allocate local BIFT-IDs independently and advertise them just like in MPLS case, for the same $\langle \text{SD}, \text{Encap}, \text{BSL}, \text{SI} \rangle$ tuple all BFRs could optionally auto-derive or be provisioned with the same BIFT-ID and no signaling is needed in that case.

One may consider that if MPLS would allow to use consistently provisioned BIER labels on all BFRs, then the second difference listed above does not exist anymore.

In this specification, if locally significant BIFT-IDs are to be used with non-MPLS encapsulation, the BIFT-IDs are advertised the same way as in the MPLS case - by a BIFT-ID block, which is a block of contiguous labels in MPLS case or a block of contiguous opaque 20-bit values in non-MPLS case. The only difference is the type of encapsulation.

If consistently provisioned or auto-derived BIFT-IDs are used with non-MPLS encapsulation, then no BIFT-ID block is signaled. Just the encapsulation type is signaled.

[4.](#) Advertising BIER Information Northbound

Nothing special here compared to OSPF/ISIS. A node's local BIER information as described in the previous section is attached to a local BIER Prefix. Details to be added.

[5.](#) Advertising BIER Information Southbound

[5.1.](#) Local BIER Information

Similar to the northbound case, a node's local BIER information is attached to a local BIER prefix that is advertised southbound.

[5.2.](#) Proxied BFR-ID Ranges

On the southbound, a node advertises a default route, plus certain prefixes to prevent blackholing or suboptimal routing upon link failures. Those prefixes and default route are like the summary

routes and default route in [[I-D.zwzw-bier-prefix-redistribute](#)], and similarly they carry BFR-IDs corresponding to the covered BIER Prefixes.

Consider a RIFT network with a BIER sub-domain of 200 BFIR/BFERS. Each non-leaf node is provisioned that BFR-ID 1-200 are used. Suppose a node X advertise southbound a default route RT1 and disaggregation routes RT2/RT3. RT2 and RT3 MUST advertise BFR-IDs covered by them (e.g. BFR-ID 100/102/150 covered by RT2 and BFR-ID 101/103 covered by RT3), while the default route RT1 can always advertise that all BFR-ID 1~200 are covered by it and does not need to exclude BFR-ID 100/102/150 and 101/103 that are covered by RT2/RT3. When a southern node receives RT1 and RT2/RT3, it installs BFR-ID 100/102/150 in its BIFT according to RT2, 101/103 in its BIFT according to RT3, and installs other BFR-IDs (or just a default route) in its BIFT according to RT1.

[6.](#) Information Elements Schema

This document introduces a bier.thrift schema with definitions to be used in RIFT encoding.thrift.

[6.1.](#) bier.thrift

```
typedef i8      SubdomainIdType
typedef i16     BfrIdType
typedef i8      BARType
typedef i8      IPAType
typedef i16     BSLType      /* Number of bits */
typedef i32     BiftIdType   /* Only the most significant 20 bits are used

enum EncapsulationType {
    mpls          = 0;
    non-mpls      = 1;
}

/* Similar to the label range in OSPF/ISIS extensions for BIER */
struct BiftIdBlock {
    1: required BiftIdType      bift_id_base;
```

```

    2: required i8                bift_id_range;
}

/* Similar to the MPLS Encapsulation sub-sub-TLV in OSPF/ISIS */
struct EncapStruct {
    1: required EncapsulationType    encap_type;
    2: required BSLType              bsl;
    3: optional BiftIdBlock          bift_id_block;
}

/*BIER node information. Similar to BIER sub-TLV in OSPF/ISIS. */
struct BierInfo {
    1: required SubdomainIdType      subdomain_id;
    2: required BfrIdType            bfr_id;
    3: required BARType              bar;
    4: required IPAType              ipa;
    5: required EncapStruct          encaps;          /* one or more */
}

struct ProxyBfrIdRange {
    1: required SubdomainIdType      subdomain_id;
    2: required BfrIdType            bfr_id_base;
    3: required BSLType              bfr_id_range;
}

```

[6.2.](#) Additions to encoding.thrift

The PrefixAttributes in encoding.rift now has two optional elements:

```

struct PrefixAttributes {
    ...
    2: optional BierInfo          bier_info;    /* BIER info for a
                                                * local BIER Prefix */
    3: optional ProxyBfrIdRange  proxy_bfr_id; /* one or more proxy
                                                * BFR-ID ranges covered
                                                * by this prefix */
}

```

[7.](#) IANA Considerations

8. Acknowledgements

9. References

9.1. Normative References

- [I-D.przygienda-rift]
Przygienda, T., Sharma, A., Atlas, A., and J. Drake,
"RIFT: Routing in Fat Trees", [draft-przygienda-rift-05](#)
(work in progress), March 2018.
- [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>>.
- [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>>.

9.2. Informative References

Zhang, et al. Expires September 6, 2018 [Page 6]

Internet-Draft

bier-rift

March 2018

- [I-D.ietf-bier-isis-extensions]
Ginsberg, L., Przygienda, T., Aldrin, S., and Z. Zhang,
"BIER support via ISIS", [draft-ietf-bier-isis-extensions-09](#)
(work in progress), February 2018.
- [I-D.ietf-bier-ospf-bier-extensions]
Psenak, P., Kumar, N., Wijnands, I., Dolganow, A.,
Przygienda, T., Zhang, Z., and S. Aldrin, "OSPF Extensions

for BIER", [draft-ietf-bier-ospf-bier-extensions-15](#) (work in progress), February 2018.

[I-D.zwzw-bier-prefix-redistribute]

Zhang, Z., Bo, W., Zhang, Z., and I. Wijnands, "BIER Prefix Redistribute", [draft-zwzw-bier-prefix-redistribute-00](#) (work in progress), January 2018.

Authors' Addresses

Zhaohui Zhang
Juniper Networks

EMail: zzhang@juniper.net

Shaowen Ma
Juniper Networks

EMail: mashao@juniper.net

Zheng Zhang
ZTE Corporation

EMail: zhang.zheng@zte.com.cn