

bier
Internet-Draft
Intended status: Standards Track
Expires: April 2, 2022

Z. Zhang
A. Przygienda
Juniper Networks
September 29, 2021

BIER with Network Slicing and Flow Differentiation
draft-zzhang-bier-slicing-and-differentiation-00

Abstract

This document specifies how BIER works in the context of IETF Network slicing, with or without fined-grained traffic differentiation.

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 April 2, 2022.

Copyright Notice

Copyright (c) 2021 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 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.

Internet-Draft BIER with Slicing and Differentiation September 2021

Table of Contents

1.	Introduction	2
2.	BIER with IETF Network Slicing	3
3.	BIER with Slice Aggregates	4
4.	Specifications	4
4.1.	ISIS Signaling	4
4.1.1.	OSPF Signaling	5
4.1.2.	BGP Signaling	5
4.2.	BIER Extension Header	5
5.	Security Considerations	5
6.	IANA Considerations	6
7.	References	6
7.1.	Normative References	6
7.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Introduction

Network slicing has been a topic widely discussed in and beyond IETF. According to [[I-D.ietf-teas-ietf-network-slices](#)]:

"An IETF Network Slice is a logical network topology connecting a number of endpoints using a set of shared or dedicated network resources that are used to satisfy specific Service Level Objectives (SLOs).

An IETF Network Slice combines the connectivity resource requirements and associated network behaviors such as bandwidth, latency, jitter, and network functions with other resource behaviors such as compute and storage availability."

It is expected that traffic associated with an IETF network slice is identified with a slice identifier (e.g. an MPLS label) and each node in the path uses the slice identifier to identify the slice in which the traffic is forwarded.

[I-D.bestbar-teas-ns-packet] introduces the notion of Slice Aggregate which comprises of one or more IETF network slice traffic streams. A Slice Aggregate is identified by a Slice Selector (SS), and packets carry the SS so that associated forwarding treatment or S-PHB (Slice policy Per Hop Behavior – the externally observable forwarding behavior applied to a specific packet belonging to a slice aggregate)

- can be applied along the path.

[I-D.li-apn-problem-statement-usecases] describes challenges faced by network operators when attempting to provide fine-grained traffic operations to satisfy the various requirements demanded by new

applications that require differentiated service treatment and [I-D.li-apn-framework] proposes a framework for solution:

"... proposes a new framework, named Application-aware Networking (APN), where application-aware information (i.e. APN attribute) including APN identification (ID) and/or APN parameters (e.g. network performance requirements) is encapsulated at network edge devices and carried in packets traversing an APN domain in order to facilitate service provisioning, perform fine-granularity traffic steering and network resource adjustment."

The authors of this document believe that the IETF Network Slicing framework, when augmented by the Slice Aggregate, addresses the APN problem domain very well. This document describes how BIER [RFC8279] works together with IETF network slicing, with or without Slice Aggregate to provide fine granularity traffic differentiation (e.g. down to per-flow level) that is demanded in the APN problem statement.

[2.](#) BIER with IETF Network Slicing

Since an IETF Network Slice is a logical network topology, each slice may have its BIRT (which maps to a set of BIFTs when BitStringLength and SetID are considered). While it is tempting and seems logical to map a slice to a BIER sub-domain, and it is straightforward to do so when the number of slices is smaller than 256 (the max number of sub-domains), this document allows to map a slice directly to a BIRT instead of a sub-domain.

Now a BIRT corresponds to a <sub-domain, slice> tuple, and each BIFT corresponds to a <subdomain-id, slice-id, bitstring length, set-id> tuple. In forwarding plane a BIFT is only identified by a 20-bit opaque number locally on a BFR, which could be an MPLS label or just a plain number in case of non-MPLS data plane. Therefore, it is feasible to have many slices in the same sub-domain - each slice will have its own BIRT so that the same BFER in the same sub-domain can be

reached via different nexthop BFRs according to different BIRTs (i.e. different set of corresponding BIFTs) for different slices.

With this, up to 2^{20} slices could be supported in theory - the only limit is the number of BIFT entries that a BFR can hold.

Mapping a slice directly to a BIRT instead of a sub-domain not only allows more than 256 slices but also reduces the burden of sub-domain related provisioning (e.g. a BFR-ID is needed for each <sub-domain, BFIR/BFER>). Of course, as mentioned earlier, if the number of slices is smaller than 256 then a slice can map to a sub-domain as well.

3. BIER with Slice Aggregates

Per [[I-D.bestbar-teas-ns-packet](#)], a Slice Aggregate may be the aggregation of several entire slices, or just a particular flow in a slice. With a Slice Aggregate for several entire slices, the different slices (of the same Slice Aggregate) also map to the same BIRT. In that case, for the same destination BFER, traffic in those different slices are forwarded to the same (set of ECMP) nexthop BFER according to the shared BIRT, yet other forwarding treatment (e.g. queuing) could still be different.

In [[RFC8279](#)], a sub-domain is associated with only one topology and each sub-domain has its own BIRT calculated using the topology information. When multiple slices are associated with a single sub-domain, each slice (or a set of slices) also has its own BIRT calculated based on the slice's (or the set of slices') topology information. Therefore, having a sub-domain with multiple slices does not violate the underlying principle of BIER architecture, i.e., a BIRT is calculated on a corresponding topology, whether the topology is for a sub-domain as in [[RFC8279](#)] or for a <sub-domain, slice or set of slices> tuple as in this document.

The BIER header has a 6-bit DSCP field. If that is not enough to identify different slices or slice aggregates that share the same BIRT, an explicit Slice Selector can be carried in "BIER Extension Header" [[I-D.zzhang-intarea-generic-delivery-functions](#)].

This means that, even for a transit BFR, if provisioned to support slice aggregates identified by a Slice Selector in the extension

header, it must check if the "Proto" field is set to a value for BIER Extension Header.

Note: while the concept of "BIER Extension Header" is first brought up in that Generic Delivery Functions draft [[I-D.zzhang-intarea-generic-delivery-functions](#)] in intarea WG, it is expected that BIER specific work will be brought to the BIER WG.

4. Specifications

BIER signaling for OSPF/ISIS/BGP is extended to include slice information so that slice-specific BIRTs can be built.

4.1. ISIS Signaling

A BIER MPLS Encapsulation Extended Sub-sub-TLV is defined with a new type to allow sub-sub-sub-TLVs in it. Besides the new type and additional sub-sub-sub-TLVs, the rest are the same as original BIER MPLS Encapsulation Sub-sub-TLV [RFC8401].

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																													
Max SI										BS Len										Label																			
sub-sub-sub-TLVs (variable)																																							

Type: Value of TBD indicating Extended sub-sub-TLV for MPLS

Length: Variable

Sub-sub-sub-TLVs: for information like Slice Selector

Sub-sub-sub-TLVs will be defined to include Slice Selector information [[I-D.bestbar-teas-ns-packet](#)] that identifies a slice or a Slice Aggregate, and potentially other information. Note that the Slice Aggregate here is for a set of slices instead of a flow in a slice. Future revisions will have more details.

Similar encoding will be defined for non-MPLS encapsulation in future revisions.

[4.1.1.](#) OSPF Signaling

Similar encoding will be defined for OSPF signaling in future revisions.

[4.1.2.](#) BGP Signaling

Similar encoding will be defined for BGP signaling in future revisions.

[4.2.](#) BIER Extension Header

This will be tracked by a separate BIER draft. For now, please refer to [[I-D.zzhang-intarea-generic-delivery-functions](#)].

[5.](#) Security Considerations

To be provided.

[6.](#) IANA Considerations

To be provided.

[7.](#) References

[7.1.](#) Normative References

[I-D.bestbar-teas-ns-packet]

Saad, T., Beeram, V. P., Wen, B., Ceccarelli, D., Halpern, J., Peng, S., Chen, R., Liu, X., Contreras, L. M., and R. Rokui, "Realizing Network Slices in IP/MPLS Networks", [draft-bestbar-teas-ns-packet-03](#) (work in progress), July 2021.

- [I-D.zzhang-intarea-generic-delivery-functions]
Zhang, Z., Bonica, R., Kompella, K., and G. Mirsky,
"Generic Delivery Functions", [draft-zzhang-intarea-generic-delivery-functions-02](#) (work in progress), August 2021.
- [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>>.

7.2. Informative References

- [I-D.ietf-teas-ietf-network-slices]
Farrel, A., Gray, E., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura,
"Framework for IETF Network Slices", [draft-ietf-teas-ietf-network-slices-04](#) (work in progress), August 2021.
- [I-D.li-apn-framework]
Li, Z., Peng, S., Voyer, D., Li, C., Liu, P., Cao, C., Mishra, G., Ebisawa, K., Previdi, S., and J. N. Guichard,
"Application-aware Networking (APN) Framework", [draft-li-apn-framework-03](#) (work in progress), May 2021.

- [I-D.li-apn-problem-statement-usecases]
Li, Z., Peng, S., Voyer, D., Xie, C., Liu, P., Qin, Z., Mishra, G., Ebisawa, K., Previdi, S., and J. N. Guichard,
"Problem Statement and Use Cases of Application-aware Networking (APN)", [draft-li-apn-problem-statement-usecases-04](#) (work in progress), June 2021.

Zhaohui Zhang
Juniper Networks

Email: zzhang@juniper.net

Antoni Przygienda
Juniper Networks

Email: prz@juniper.net