Network Working Group Internet-Draft Intended status: Standards Track Expires: January 10, 2022 T. Eckert, Ed. Futurewei G. Cauchie Bouygues Telecom M. Menth University of Tuebingen July 9, 2021

Tree Engineering for Bit Index Explicit Replication (BIER-TE) draft-ietf-bier-te-arch-10

Abstract

This memo describes per-packet stateless strict and loose path steered replication and forwarding for Bit Index Explicit Replication packets (<u>RFC8279</u>). It is called BIER Tree Engineering (BIER-TE) and is intended to be used as the path steering mechanism for Traffic Engineering with BIER.

BIER-TE introduces a new semantic for bit positions (BP) that indicate adjacencies, as opposed to BIER in which BPs indicate Bit-Forwarding Egress Routers (BFER). BIER-TE can leverage BIER forwarding engines with little changes. Co-existence of BIER and BIER-TE forwarding in the same domain is possible, for example by using separate BIER sub-domains (SDs). Except for the optional routed adjacencies, BIER-TE does not require a BIER routing underlay, and can therefore operate without depending on an Interior Gateway Routing protocol (IGP).

As it operates on the same per-packet stateless forwarding principles, BIER-TE can also be a good fit to support multicast path steering in Segment Routing (SR) networks.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

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 <u>https://datatracker.ietf.org/drafts/current/</u>.

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."

Eckert, et al.

Expires January 10, 2022

This Internet-Draft will expire on January 10, 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>https://trustee.ietf.org/license-info</u>) 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

<u>1</u> . Overview	<u>3</u>
<u>1.1</u> . Requirements Language	<u>5</u>
$\underline{2}$. Introduction	<u>5</u>
<u>2.1</u> . Basic Examples	<u>5</u>
2.2. BIER-TE Topology and adjacencies	<u>8</u>
2.3. Relationship to BIER	<u>9</u>
<u>2.4</u> . Accelerated/Hardware forwarding comparison	<u>11</u>
<u>3</u> . Components	<u>11</u>
<u>3.1</u> . The Multicast Flow Overlay	<u>12</u>
3.2. The BIER-TE Control Plane	<u>12</u>
<u>3.2.1</u> . The BIER-TE Controller	<u>13</u>
<u>3.2.1.1</u> . BIER-TE Topology discovery and creation	<u>14</u>
<u>3.2.1.2</u> . Engineered Trees via BitStrings	<u>14</u>
<u>3.2.1.3</u> . Changes in the network topology	<u>15</u>
<u>3.2.1.4</u> . Link/Node Failures and Recovery	<u>15</u>
<u>3.3</u> . The BIER-TE Forwarding Plane	<u>15</u>
<u>3.4</u> . The Routing Underlay	<u>16</u>
<u>3.5</u> . Traffic Engineering Considerations	<u>16</u>
<u>4</u> . BIER-TE Forwarding	<u>17</u>
<u>4.1</u> . The Bit Index Forwarding Table (BIFT)	<u>17</u>
<u>4.2</u> . Adjacency Types	<u>18</u>
<u>4.2.1</u> . Forward Connected	<u>18</u>
<u>4.2.2</u> . Forward_routed	<u>19</u>
<u>4.2.3</u> . ECMP	<u>19</u>
<u>4.2.4</u> . Local Decap(sulation)	<u>19</u>
<u>4.3</u> . Encapsulation / Co-existence with BIER	<u>20</u>
<u>4.4</u> . BIER-TE Forwarding Pseudocode	<u>21</u>
<u>4.5</u> . Basic BIER-TE Forwarding Example	<u>24</u>
<u>4.6</u> . BFR Requirements for BIER-TE forwarding	<u>26</u>

5. BIER-TE	Controller Operational Considerations			<u>26</u>
<u>5.1</u> . Bit	position Assignments			<u>26</u>
<u>5.1.1</u> .	P2P Links			<u>27</u>
<u>5.1.2</u> .	BFER			<u>27</u>
<u>5.1.3</u> .	Leaf BFERs			<u>27</u>
<u>5.1.4</u> .	LANS			<u>28</u>
<u>5.1.5</u> .	Hub and Spoke			<u>28</u>
<u>5.1.6</u> .	Rings			<u>29</u>
<u>5.1.7</u> .	Equal Cost MultiPath (ECMP)			<u>30</u>
<u>5.1.8</u> .	Forward_routed adjacencies			<u>32</u>
<u>5.1.8</u>	<u>1</u> . Reducing bit positions			<u>32</u>
<u>5.1.8</u>	<u>2</u> . Supporting nodes without BIER-TE			<u>33</u>
<u>5.1.9</u> .	Reuse of bit positions (without DNC)			<u>33</u>
<u>5.1.10</u> .	Summary of BP optimizations			<u>35</u>
<u>5.2</u> . Avoi	ding duplicates and loops			<u>36</u>
<u>5.2.1</u> .	Loops			<u>36</u>
<u>5.2.2</u> .	Duplicates			<u>36</u>
<u>5.3</u> . Mana	iging SI, sub-domains and BFR-ids			<u>37</u>
<u>5.3.1</u> .	Why SI and sub-domains			<u>37</u>
<u>5.3.2</u> .	Assigning bits for the BIER-TE topology			<u>38</u>
<u>5.3.3</u> .	Assigning BFR-id with BIER-TE			<u>39</u>
<u>5.3.4</u> .	Mapping from BFR to BitStrings with BIER-TE			<u>40</u>
<u>5.3.5</u> .	Assigning BFR-ids for BIER-TE			<u>41</u>
<u>5.3.6</u> .	Example bit allocations			<u>41</u>
5.3.6.	<u>1</u> . With BIER			<u>41</u>
<u>5.3.6</u>	<u>2</u> . With BIER-TE			<u>42</u>
<u>5.3.7</u> .	Summary			<u>43</u>
<u>6</u> . BIER-TE	and Segment Routing			<u>44</u>
7. Security	Considerations			<u>45</u>
8. IANA Cor	siderations			<u>46</u>
9. Acknowle	edgements			<u>46</u>
<u>10</u> . Change]	.og [RFC Editor: Please remove]			<u>46</u>
<u>11</u> . Referenc	es			<u>56</u>
<u>11.1</u> . Nor	mative References			<u>56</u>
<u>11.2</u> . Inf	ormative References			<u>56</u>
Authors' Add	Iresses			<u>58</u>

1. Overview

BIER-TE is based on architecture, terminology and packet formats with BIER as described in [RFC8279] and [RFC8296]. This document describes BIER-TE in the expectation that the reader is familiar with these two documents.

BIER-TE introduces a new semantic for bit positions (BP) that indicate adjacencies, as opposed to BIER in which BPs indicate Bit-Forwarding Egress Routers (BFER). With BIER-TE, the BIFT of each BFR is only populated with BP that are adjacent to the BFR in the BIER-TE

Topology. Other BPs are empty in the BIFT. The BFR replicate and forwards BIER packets to adjacent BPs that are set in the packet. BPs are normally also cleared upon forwarding to avoid duplicates and loops. This is detailed further below.

BIER-TE can leverage BIER forwarding engines with little or no changes. It can also co-exist with BIER forwarding in the same domain, for example by using separate BIER sub-domains. Except for the optional routed adjacencies, BIER-TE does not require a BIER routing underlay, and can therefore operate without depending on an Interior Gateway Routing protocol (IGP).

As it operates on the same per-packet stateless forwarding principles, BIER-TE can also be a good fit to support multicast path steering in Segment Routing (SR) networks.

This document is structured as follows:

- <u>Section 2</u> introduces BIER-TE with two reference forwarding examples, followed by an introduction of the new concepts of the BIER-TE (overlay) topology and finally a summary of the relationship between BIER and BIER-TE and a discussion of accelerated hardware forwarding.
- <u>Section 3</u> describes the components of the BIER-TE architecture, Flow overlay, BIER-TE layer with the BIER-TE control plane (including the BIER-TE controller) and BIER-TE forwarding plane, and the routing underlay.
- <u>Section 4</u> specifies the behavior of the BIER-TE forwarding plane with the different type of adjacencies and possible variations of BIER-TE forwarding pseudocode, and finally the mandatory and optional requirements.
- o Section 5 describes operational considerations for the BIER-TE controller, foremost how the BIER-TE controller can optimize the use of BP by using specific type of BIER-TE adjacencies for different type of topological situations, but also how to assign bits to avoid loops and duplicates (which in BIER-TE does not come for free), and finally how SI, sub-domains and BFR-ids can be managed by a BIER-TE controller, examples and summary.
- o <u>Section 6</u> concludes the technology specific sections of document by further relating BIER-TE to Segment Routing (SR).

Note that related work, [<u>I-D.ietf-roll-ccast</u>] uses Bloom filters [<u>Bloom70</u>] to represent leaves or edges of the intended delivery tree. Bloom filters in general can support larger trees/topologies with

fewer addressing bits than explicit BitStrings, but they introduce the heuristic risk of false positives and cannot clear bits in the BitString during forwarding to avoid loops. For these reasons, BIER-TE uses explicit BitStrings like BIER. The explicit BitStrings of BIER-TE can also be seen as a special type of Bloom filter, and this is how related work [ICC] describes it.

<u>1.1</u>. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in <u>BCP</u> <u>14</u> [<u>RFC2119</u>], [<u>RFC8174</u>] when, and only when, they appear in all capitals, as shown here.

2. Introduction

2.1. Basic Examples

BIER-TE forwarding is best introduced with simple examples.

BIER-TE Topology: Diagram: p5 p6 --- BFR3 --p3/ p13 \p7 p15 BFR1 ---- BFR2 BFR5 ---- BFR6 p1 p2 p4\ p14 /p10 p11 p12 --- BFR4 --p8 p9 (simplified) BIER-TE Bit Index Forwarding Tables (BIFT): BFR1: p1 -> local_decap p2 -> forward_connected() to BFR2 BFR2: p1 -> forward_connected() to BFR1 p5 -> forward_connected() to BFR3 p8 -> forward_connected() to BFR4 BFR3: p3 -> forward_connected() to BFR2 p7 -> forward_connected() to BFR5 p13 -> local_decap BFR4: p4 -> forward_connected() to BFR2 p10 -> forward_connected() to BFR5 p14 -> local_decap BFR5: p6 -> forward_connected() to BFR3 p9 -> forward_connected() to BFR4 p12 -> forward_connected() to BFR6 BFR6: p11 -> forward_connected() to BFR5 p15 -> local_decap

Figure 1: BIER-TE basic example

Consider the simple network in the above BIER-TE overview example picture with 6 BFRs. p1...p14 are the bit positions (BP) used. All BFRs can act as an ingress BFR (BFIR), BFR1, BFR3, BFR4 and BFR6 can also be egress BFRs (BFERs). Forward_connected() is the name for adjacencies that are representing subnet adjacencies of the network. Local_decap() is the name of the adjacency to decapsulate BIER-TE packets and pass their payload to higher layer processing.

Assume a packet from BFR1 should be sent via BFR4 to BFR6. This requires a BitString (p2,p8,p10,p12,p15). When this packet is examined by BIER-TE on BFR1, the only bit position from the BitString that is also set in the BIFT is p2. This will cause BFR1 to send the only copy of the packet to BFR2. Similarly, BFR2 will forward to BFR4 because of p8, BFR4 to BFR5 because of p10 and BFR5 to BFR6 because of p12. p15 finally makes BFR6 receive and decapsulate the packet.

To send in addition to BFR6 via BFR4 also a copy to BFR3, the BitString needs to be (p2,p5,p8,p10,p12,p13). When this packet is examined by BFR2, p5 causes one copy to be sent to BFR3 and p8 one copy to BFR4. When BFR3 receives the packet, p13 will cause it to receive and decapsulate the packet.

If instead the BitString was (p2,p6,p8,p10,p12,p13,p15), the packet would be copied by BFR5 towards BFR3 because of p6 instead of being copied by BFR2 to BFR3 because of p5 in the prior case. This is showing the ability of the shown BIER-TE Topology to make the traffic pass across any possible path and be replicated where desired.

BIER-TE has various options to minimize BP assignments, many of which are based on assumptions about the required multicast traffic paths and bandwidth consumption in the network.

The following picture shows a modified example, in which Rtr2 and Rtr5 are assumed not to support BIER-TE, so traffic has to be unicast encapsulated across them. To emphasize non-L2, but routed/tunneled forwarding of BIER-TE packets, these adjacencies are called "forward_routed". Otherwise there is no difference in their processing over the aforementioned "forward_connected" adjacencies.

In addition, bits are saved in the following example by assuming that BFR1 only needs to be BFIR but not BFER or transit BFR.

BIER-TE Topology:

p2 -> forward_routed() to BFR4
BFR3: p3 -> local_decap
p5 -> forward_routed() to BFR6

- BFR4: p4 -> local_decap p6 -> forward_routed() to BFR6
- BFR6: p5 -> local_decap p6 -> local_decap p7 -> forward_routed() to BFR3 p8 -> forward_routed() to BFR4

Figure 2: BIER-TE basic overlay example

To send a BIER-TE packet from BFR1 via BFR3 to BFR6, the BitString is (p1,p5). From BFR1 via BFR4 to BFR6 it is (p2,p6). A packet from BFR1 to BFR3,BFR4 and from BFR3 to BFR6 uses (p1,p2,p3,p4,p5). A packet from BFR1 to BFR3,BFR4 and from BFR4 to BFR uses (p1,p2,p3,p4,p6). A packet from BFR1 to BFR4, and from BFR4 to BFR6 and from BFR6 to BFR3 uses (p2,p3,p4,p6,p7). A packet from BFR1 to BFR3, and from BFR3 to BFR6 and from BFR6 to BFR4 uses (p1,p3,p4,p5,p8).

2.2. BIER-TE Topology and adjacencies

The key new component in BIER-TE compared to BIER is the BIER-TE topology as introduced through the two examples in <u>Section 2.1</u>. It is used to control where replication can or should happen and how to minimize the required number of BP for adjacencies.

BIER-TE ARCH

The BIER-TE Topology consists of the BIFTs of all the BFR and can also be expressed as a directed graph where the edges are the adjacencies between the BFR labelled with the BP used for the adjacency. Adjacencies are naturally unidirectional. BP can be reused across multiple adjacencies as long as this does not lead to undesired duplicates or loops as explained further down in the text.

If the BIER-TE topology represents (a subset of) the underlying (layer 2) topology of the network as shown in the first example, this may be called a "native" BIER-TE topology. A topology consisting only of "forward_routed" adjacencies as shown in the second example may be called an "overlay" BIER-TE topology. A BIER-TE topology will both "forward_connected" and "forward_routed" adjacencies may be called a "hybrid" BIER-TE topology.

<u>2.3</u>. Relationship to BIER

BIER-TE is designed so that is forwarding plane is a simple extension to the BIER forwarding plane, hence allowing for it to be added to BIER deployments where it can be beneficial.

BIER-TE is also intended as an option to expand the BIER architecture into deployments where BIER may not be the best fit, such as statically provisioned networks with needs for path steering but without desire for distributed routing protocols.

- 1. BIER-TE inherits the following aspects from BIER unchanged:
 - 1. The fundamental purpose of per-packet signaled packet replication and delivery via a BitString.
 - The overall architecture consisting of three layers, flow overlay, BIER(-TE) layer and routing underlay.
 - 3. The supportable encapsulations, [<u>RFC8296</u>] or other (future) encapsulations.
 - 4. The semantic of all [<u>RFC8296</u>] header elements used by the BIER-TE forwarding plane other than the semantic of the BP in the BitString.
 - 5. The BIER forwarding plane, with the exception of how bits have to be cleared during replication.
- 2. BIER-TE has the following key changes with respect to BIER:
 - 1. In BIER, bits in the BitString of a BIER packet header indicate a BFER and bits in the BIFT indicate the BIER

control plane calculated next-hop toward that BFER. In BIER-TE, bits in the BitString of a BIER packet header indicate an adjacency in the BIER-TE topology, and only the BFRs that are upstream of this adjacency have this bit populated with the adjacency in their BIFT.

- 2. In BIER, the implied reference option for the core part of the BIER layer control plane is the BIER extension to distributed routing protocol, such as standardized in ISIS/ OSPF extensions for BIER, [RFC8401] and [RFC8444]. The reference option for the core part of the BIER-TE control plane is the BIER-TE controller. Nevertheless, both BIER and BIER-TE BIFT forwarding plane state could equally be populated by any mechanism.
- 3. Assuming the reference options for the control plane, BIER-TE replaces in-network autonomous path calculation by explicit paths calculated by the BIER-TE controller.
- 3. The following element/functions described in the BIER architecture are not required by the BIER-TE architecture:
 - 1. BIRTs on BFR for BIER-TE are not required when using a BIER-TE controller because the controller can directly populate the BIFTs. In BIER, BIRTs are populated by the distributed routing protocol support for BIER, allowing BFR to populate their BIFTs locally from their BIRTs. Other BIER-TE control plane or management plane options may introduce requirements for BIRTs for BIER-TE BFR.
 - 2. The BIER-TE layer forwarding plane does not require BFR to have a unique BP and therefore also no unique BFR-id. See for example See <u>Section 5.1.3</u>.
 - 3. Identification of BFR by the BIER-TE control plane is outside the scope of this specification. Whereas the BIER control plane uses BFR-ids in its BFR to BFR signaling, a BIER-TE controller may choose any form of identification deemed appropriate.
 - 4. BIER-TE forwarding does not use the BFR-id field of the BIER packet header.
- 4. Co-existence of BIER and BIER-TE in the same network requires the following:
 - 1. The BIER/BIER-TE packet header needs to allow addressing both BIER and BIER-TE BIFT. Depending on the encapsulation

option, the same SD may or may not be reusable across BIER and BIER-TE. See <u>Section 4.3</u>. In either case, a packet is always only forwarded end-to-end via BIER or via BIER-TE (ships in the nights forwarding).

2. BIER-TE deployments will have to assign BFR-ids to BFR and insert them into the BFR-id field of BIER packet headers as BIER does, whenever the deployment uses (unchanged) components developed for BIER that use BFR-id, such as multicast flow overlays or BIER layer control plane elements. See also <u>Section 5.3.3</u>.

<u>2.4</u>. Accelerated/Hardware forwarding comparison

Forwarding of BIER-TE is designed to easily build/program common forwarding hardware with BIER. The pseudocode in <u>Section 4.4</u> shows how existing BIER/BIFT forwarding can be modified to support the REQUIRED BIER-TE forwarding functionality, by using BIER BIFT's "Forwarding Bit Mask" (F-BM): Only the clearing of bits to avoid duplicate packets to a BFR neighbor is skipped in BIER-TE forwarding because it is not necessary and could not be done when using BIER F-BM.

Whether to use BIER or BIER-TE forwarding is simply a choice of the mode of the BIFT indicated by the packet (BIER or BIER-TE BIFT). This is determined by the BFR configuration for the encapsulation, see <u>Section 4.3</u>.

3. Components

BIER-TE can be thought of being constituted from the same three layers as BIER: The "multicast flow overlay", the "BIER layer" and the "routing underlay". The following picture also shows how the "BIER layer" is constituted from the "BIER-TE forwarding plane" and the "BIER-TE control plane" represent by the "BIER-TE Controller".

```
<---->BGP/PIM---->
 |<-IGMP/PIM-> multicast flow <-PIM/IGMP->|
                 overlay
     BIER-TE [BIER-TE Controller] <=> [BIER-TE Topology]
            Λ
                     \wedge
     control
                          Λ
     plane
              /
                    \ BIER-TE control protocol
             | e.g. YANG/Netconf/RestConf
                                 PCEP/...
             v
                     V
                            v
Src -> Rtr1 -> BFIR-----BFR----BFER -> Rtr2 -> Rcvr
            |<---->|
           BIER-TE forwarding plane
             |<- BIER-TE domain->|
           |<---->|
              Routing underlay
```

Figure 3: BIER-TE architecture

3.1. The Multicast Flow Overlay

The Multicast Flow Overlay has the same role as described for BIER in [RFC8279], Section 4.3. See also Section 3.2.1.2.

3.2. The BIER-TE Control Plane

In the BIER architecture [<u>RFC8279</u>], the BIER control plane is not explicitly separated from the BIER forwarding plane, but instead their functions are summarized together in <u>Section 4.2</u>. Example standardized options for the BIER control plane include ISIS/OSPF extensions for BIER, [<u>RFC8401</u>] and [<u>RFC8444</u>].

For BIER-TE, the control plane includes at minimum the following functionality.

- During initial provisioning of the network and/or during modifications of its topology and/or services: protocols and/or procedures to establish BIER-TE BIFTs:
 - Determine the desired BIER-TE topology for a BIER-TE subdomains: the native and/or overlay adjacencies that are assigned to BPs.
 - 2. Determine the per-BFR BIFT from the BIER-TE topology.

- 3. Optionally assign BFR-id to BFIR for later insertion into BIER-TE headers on BFIR. Alternatively, bfir-id in BIER packet headers may be managed solely by the flow overlay layer and/or be unused.
- 4. Install/update the BIFTs into the BFRs and optionally BFR-id into BFIR.
- 2. During operations of the network: Protocols and/or procedures to support creation/change/removal of overlay flows on BFIR:
 - 1. Process the BIER-TE requirements for the multicast overlay flow: BFIR and BFERs of the flow as well as policies for the path selection of the flow.
 - 2. Determine the BitStrings and optionally Entropy.
 - Install state on the BFIR to imposition the desired BIER packet header(s) for packets of the overlay flow.
 - 4. Install the necessary state on the BFERs to decapsulate the BIER packet header and properly dispatch its payload.

<u>3.2.1</u>. The BIER-TE Controller

Notwithstanding other options, this architecture describes the BIER control plane as shown in Figure 3 to consists of:

- o A single centralized BIER-TE controller.
- Data-models and protocols to communicate between controller and BFR in step 1, such YANG/Netconf/RestConf.
- o Protocols to communicate between controller and BFIR in step 2, such as BIER-TE extensions for [<u>RFC5440</u>].

The BIER control plane could equally be implemented without any active dynamic components by an operator via CLI on the BFRs. In that case, operator configured local policy on the BFIR would have to determine how to set the appropriate BIER header fields. The BIER-TE control plane could also be decentralized and/or distributed, but this document does not consider any additional protocols and/or procedures which would then be necessary to coordinate its entities to achieve the above described functionality.

<u>3.2.1.1</u>. BIER-TE Topology discovery and creation

Step 1.1 includes network topology discovery and BIER-TE topology creation. The latter describes the process by which a Controller determines which routers are to be configured as BFR and the adjacencies between them.

In statically managed networks, such as in industrial environments, both discovery and creation can be a manual/offline process.

In other networks, topology discovery may rely on protocols including extending a Link-State-Protocol (LSP) based IGP into the BIER-TE controller itself, [<u>RFC7752</u>] (BGP-LS) or [<u>RFC8345</u>] (Yang topology) as well as BIER-TE specific methods, for example via [<u>I-D.ietf-bier-te-yang</u>]. These options are non-exhaustive.

Dynamic creation of the BIER-TE topology can be as easy as mapping the network topology 1:1 to the BIER-TE topology by assigning a BP for every network subnet adjacency. In larger networks, it likely involves more complex policy and optimization decisions including how to minimize the number of BP required and how to assign BP across different BitStrings to minimize the number of duplicate packets across links when delivering an overlay flow to BFER using different SIS/BitStrings. These topics are discussed in <u>Section 5</u>.

When the topology is determined, the BIER-TE Controller then pushes the BitPositions/adjacencies to the BIFT of the BFRs, populating only those SI:BitPositions to the BIFT of each BFR to which that BFR should be able to send packets to - adjacencies connecting to this BFR.

Communications between the BIER-TE Controller and BFRs (beside topology discovery) is ideally via standardized protocols and datamodels such as Netconf/RestConf/Yang/PCEP. Vendor-specific CLI on the BFRs is also an option (as in many other SDN solutions lacking definition of standardized data model).

<u>3.2.1.2</u>. Engineered Trees via BitStrings

In BIER, the same set of BFER in a single sub-domain is always encoded as the same BitString. In BIER-TE, the BitString used to reach the same set of BFER in the same sub-domain can be different for different overlay flows because the BitString encodes the paths towards the BFER, so the BitStrings from different BFIR to the same set of BFER will often be different, and the BitString from the same BFIR to the same set of BFER can different for different overlay flows for policy reasons such as shortest path trees, Steiner trees (minimum cost trees), diverse path trees for redundancy and so on.

See also [<u>I-D.ietf-bier-multicast-http-response</u>] for a solution describing this interaction.

<u>3.2.1.3</u>. Changes in the network topology

If the network topology changes (not failure based) so that adjacencies that are assigned to bit positions are no longer needed, the BIER-TE Controller can re-use those bit positions for new adjacencies. First, these bit positions need to be removed from any BFIR flow state and BFR BIFT state, then they can be repopulated, first into BIFT and then into the BFIR.

3.2.1.4. Link/Node Failures and Recovery

When link or nodes fail or recover in the topology, BIER-TE could quickly respond with out-of-scope FRR procedures such as [<u>I-D.eckert-bier-te-frr</u>]. It can also more slowly react by recalculating the BitStrings of affected multicast flows. This reaction is slower than the FRR procedure because the BIER-TE Controller needs to receive link/node up/down indications, recalculate the desired BitStrings and push them down into the BFIRs. With FRR, this is all performed locally on a BFR receiving the adjacency up/down notification.

3.3. The BIER-TE Forwarding Plane

The BIER-TE Forwarding Plane constitutes of the following components:

- On BFIR imposition of BIER header for packets from overlay flows. This is driven by a combination of state established by the BIER-TE control plane and/or the multicast flow overlay as explained in <u>Section 3.1</u>.
- On BFR (including BFIR and BFER), forwarding/replication of BIER packets according to their BitString as explained below and optionally Entropy. Processing of other BIER header fields such as DSCP is outside the scope of this document.
- On BFER removal of BIER header and dispatching of the payload according to state created by the BIER-TE control plane and/or overlay layer.

When the BIER-TE Forwarding Plane receives a packet, it simply looks up the bit positions that are set in the BitString of the packet in the Bit Index Forwarding Table (BIFT) that was populated by the BIER-TE Controller. For every BP that is set in the BitString, and that has one or more adjacencies in the BIFT, a copy is made according to the type of adjacencies for that BP in the BIFT. Before sending any

copy, the BFR clears all BPs in the BitString of the packet for which the BFR has one or more adjacencies in the BIFT, except when the adjacency indicates "DoNotClear" (DNC, see <u>Section 4.2.1</u>). This is done to inhibit that packets can loop.

<u>3.4</u>. The Routing Underlay

For forward_connected() adjacencies, BIER-TE is sending BIER packets to directly connected BIER-TE neighbors as L2 (unicasted) BIER packets without requiring a routing underlay. For forward_routed() adjacencies, BIER-TE forwarding encapsulates a copy of the BIER packet so that it can be delivered by the forwarding plane of the routing underlay to the routable destination address indicated in the adjacency. See <u>Section 4.2.2</u> for the adjacency definition.

BIER relies on the routing underlay to calculate paths towards BFERs and derive next-hop BFR adjacencies for those paths. This commonly relies on BIER specific extensions to the routing protocols of the routing underlay but may also be established by a controller. In BIER-TE, the next-hops of a packet are determined by the BitString through the BIER-TE Controller established adjacencies on the BFR for the BPs of the BitString. There is thus no need for BFER specific routing underlay extensions to forward BIER packets with BIER-TE semantics.

Encapsulation parameters can be provisioned by the BIER-TE controller into the forward_connected() or forward_routed() adjacencies directly without relying on a routing underlay.

If the BFR intends to support FRR for BIER-TE, then the BIER-TE forwarding plane needs to receive fast adjacency up/down notifications: Link up/down or neighbor up/down, e.g. from BFD. Providing these notifications is considered to be part of the routing underlay in this document.

<u>3.5</u>. Traffic Engineering Considerations

Traffic Engineering ([<u>I-D.ietf-teas-rfc3272bis</u>]) provides performance optimization of operational IP networks while utilizing network resources economically and reliably. The key elements needed to effect TE are policy, path steering and resource management. These elements require support at the control/controller level and within the forwarding plane.

Policy decisions are made within the BIER-TE control plane, i.e., within BIER-TE Controllers. Controllers use policy when composing BitStrings and BFR BIFT state. The mapping of user/IP traffic to specific BitStrings/BIER-TE flows is made based on policy. The

specific details of BIER-TE policies and how a controller uses them are out of scope of this document.

Path steering is supported via the definition of a BitString. BitStrings used in BIER-TE are composed based on policy and resource management considerations. For example, when composing BIER-TE BitStrings, a Controller must take into account the resources available at each BFR and for each BP when it is providing congestion loss free services such as Rate Controlled Service Disciplines [RCSD94]. Resource availability could be provided for example via routing protocol information, but may also be obtained via a BIER-TE control protocol such as Netconf or any other protocol commonly used by a PCE to understand the resources of the network it operates on. The resource usage of the BIER-TE traffic admitted by the BIER-TE controller can be solely tracked on the BIER-TE Controller based on local accounting as long as no forward_routed() adjacencies are used (see <u>Section 4.2.1</u> for the definition of forward_routed() adjacencies). When forward_routed() adjacencies are used, the paths selected by the underlying routing protocol need to be tracked as well.

Resource management has implications on the forwarding plane beyond the BIER-TE defined steering of packets. This includes allocation of buffers to guarantee the worst case requirements of admitted RCSD traffic and potential policing and/or rate-shaping mechanisms, typically done via various forms of queuing. This level of resource control, while optional, is important in networks that wish to support congestion management policies to control or regulate the offered traffic to deliver different levels of service and alleviate congestion problems, or those networks that wish to control latencies experienced by specific traffic flows.

<u>4</u>. **BIER-TE Forwarding**

<u>4.1</u>. The Bit Index Forwarding Table (BIFT)

The Bit Index Forwarding Table (BIFT) exists in every BFR. For every sub-domain in use, it is a table indexed by SI:bit position and is populated by the BIER-TE control plane. Each index can be empty or contain a list of one or more adjacencies.

Like BIER, BIER-TE can support multiple sub-domains, each with a separate BIFT.

In [RFC8279], Figure 2, indices into the BIFT are both SI:BitString and BFR-id, where BitString is indicating a BP: BFR-id = SI * 2^BSL + BP. As shown in Figure 4, in BIER-TE, only SI:BP are used as indices into a BIFT because they identify adjacencies and not BFR.

| Index: | Adjacencies: | SI:bit position | <empty> or one or more per entry _____ | 0:1 forward_connected(interface,neighbor{,DNC}) | _____ forward_connected(interface,neighbor{,DNC}) | 0:2 forward_connected(interface,neighbor{,DNC}) | _____ local_decap({VRF}) 0:3 _____ 0:4forward_routed({VRF,}l3-neighbor) _____ | <empty> 0:5 _____ ECMP({adjacency1,...adjacencyN}, seed) 0:6 _____ | BitStringLength | ... _____ Bit Index Forwarding Table

Figure 4: BIFT adjacencies

The BIFT is programmed into the data plane of BFRs by the BIER-TE Controller and used to forward packets, according to the rules specified in the BIER-TE Forwarding Procedures.

Note that a BIFT index (SI:BP) may be populated in the BIFT of more than one BFR. See <u>Section 5.1.6</u> for an example of how a BIER-TE controller could assign BPs to (logical) adjacencies shared across multiple BFRs, <u>Section 5.1.3</u> for an example of assigning the same BP to different adjacencies, and <u>Section 5.1.9</u> for guidelines regarding re-use of BPs across different adjacencies.

{VRF} indicates the Virtual Routing and Forwarding context into which the BIER payload is to be delivered. This is optional and depends on the multicast flow overlay.

<u>4.2</u>. Adjacency Types

4.2.1. Forward Connected

A "forward_connected" adjacency is towards a directly connected BFR neighbor using an interface address of that BFR on the connecting interface. A forward_connected() adjacency does not route packets but only L2 forwards them to the neighbor.

Packets sent to an adjacency with "DoNotClear" (DNC) set in the BIFT MUST NOT have the bit position for that adjacency cleared when the BFR creates a copy for it. The bit position will still be cleared for copies of the packet made towards other adjacencies. This can be used for example in ring topologies as explained below.

4.2.2. Forward_routed

A "forward_routed" adjacency is an adjacency towards a BFR that uses a (tunneling) encapsulation which will cause the packet to be forwarded by the routing underlay toward the adjacent BFR. This can leverage any feasible encapsulation, such as MPLS or tunneling over IP/IPv6, as long as the BIER-TE packet can be identified as a payload. This identification can either rely on the BIER/BIER-TE coexistence mechanisms described in <u>Section 4.3</u>, or by explicit support for a BIER-TE payload type in the tunneling encapsulation.

"forward_routed" adjacencies are necessary to pass BIER-TE traffic across non BIER-TE capable routers or to minimize the number of required BP by tunneling over (BIER-TE capable) routers on which neither replication nor path-steering is desired, or simply to leverage path redundancy and FRR of the routing underlay towards the next BFR. They may also be useful to a multi-subnet adjacent BFR to leverage the routing underlay ECMP independent of BIER-TE ECMP (Section 4.2.3).

4.2.3. ECMP

BIER ECMP is tied to the BIER BIFT processing semantic and are therefore not directly usable with BIER-TE.

A BIER-TE "Equal Cost Multipath" (ECMP) adjacency has a list of two or more non-ECMP adjacencies and a seed parameter. When a BIER-TE packet is copied onto such an ECMP adjacency, an implementation specific so-called hash function will select one out of the lists adjacencies to which the packet is forwarded. This ECMP hash function MUST select the same adjacency from that list for all packets with the same entropy parameter. The seed parameter allows to design hash functions that are easy to implement at high speed without running into polarization issues across multiple consecutive ECMP hops. See <u>Section 5.1.7</u> for more explanations.

<u>4.2.4</u>. Local Decap(sulation)

A "local_decap" adjacency passes a copy of the payload of the BIER-TE packet to the protocol within the BFR (IPv4/IPv6, Ethernet,...) responsible for that payload according to the packet header fields. A local_decap() adjacency turns the BFR into a BFER for matching
packets. Local_decap() adjacencies require the BFER to support routing or switching for NextProto to determine how to further process the packet.

4.3. Encapsulation / Co-existence with BIER

Specifications for BIER-TE encapsulation are outside the scope of this document. This section gives explanations and guidelines.

Because a BFR needs to interpret the BitString of a BIER-TE packet differently from a BIER packet, it is necessary to distinguish BIER from BIER-TE packets. In the BIER encapsulation [RFC8296], the BIFTid field of the packet indicates the BIFT of the packet. BIER and BIER-TE can therefore be run simultaneously, when the BIFT-id address space is shared across BIER BIFT and BIER-TE BIFT. Partitioning the BIFT-id address space is subject to BIER-TE/BIER control plane procedures.

When [RFC8296] is used for BIER with MPLS, BIFT-id address ranges can be dynamically allocated from MPLS label space only for the set of actually used SD:BSL BIFT. This allows to also allocate nonoverlapping label ranges for BIFT-ids that are to be used with BIER-TE BIFTs.

With MPLS, it is also possible to reuse the same SD space for both BIER-TE and BIER, so that the same SD has both a BIER BIFT and according range of BIFT-ids and a disjoint BIER-TE BIFT and non-overlapping range of BIFT-ids.

When a fixed mapping from BSL, SD, SI is used without specifically distinguishing BIER and BIER-TE, such as proposed for non-MPLS forwarding with [RFC8296] in [I-D.ietf-bier-non-mpls-bift-encoding] revision 04, section 5., then it is necessary to allocate disjoint SDs to BIER and BIER-TE BIFT so that both can be addressed by the BIFT-ids. The encoding proposed in section 6. of the same document does not statically encode BSL or SD into the BIFT-id, but allows for a mapping, and hence could provide for the same freedom as when MPLS is being used (same or different SD for BIER/BIER-TE).

"forward_routed" requires an encapsulation permitting to unicast BIER-TE packets to a specific interface address on a target BFR. With MPLS encapsulation, this can simply be done via a label stack with that addresses label as the top label - followed by the label assigned to the (BSL,SD,SI) BitString. With non-MPLS encapsulation, some form of IP encapsulation would be required (for example IP/GRE).

The encapsulation used for "forward_routed" adjacencies can equally support existing advanced adjacency information such as "loose source

routes" via e.g. MPLS label stacks or appropriate header extensions (e.g. for IPv6).

4.4. BIER-TE Forwarding Pseudocode

The following pseudocode, Figure 5, for BIER-TE forwarding is based on the BIER forwarding pseudocode of [RFC8279], section 6.5 with one modification.

```
void ForwardBitMaskPacket_withTE (Packet)
{
    SI=GetPacketSI(Packet);
    Offset=SI*BitStringLength;
    for (Index = GetFirstbit position(Packet->BitString); Index ;
         Index = GetNextbit position(Packet->BitString, Index)) {
        F-BM = BIFT[Index+Offset]->F-BM;
        if (!F-BM) continue;
                                                         [3]
        BFR-NBR = BIFT[Index+Offset]->BFR-NBR;
        PacketCopy = Copy(Packet);
        PacketCopy->BitString &= F-BM;
                                                         [2]
        PacketSend(PacketCopy, BFR-NBR);
        // The following must not be done for BIER-TE:
        // Packet->BitString &= ~F-BM;
                                                         [1]
    }
}
```

Figure 5: BIER-TE Forwarding Pseudocode for required functions, based on BIER Pseudocode

In step [2], the F-BM is used to clear bit(s) in PacketCopy. This step exists in both BIER and BIER-TE, but the F-BMs need to be populated differently for BIER-TE than for BIER for the desired clearing.

In BIER, multiple bits of a BitString can have the same BFR-NBR. When a received packets BitString has more than one of those bits set, the BIER replication logic has to avoid that more than one PacketCopy is sent to that BFR-NBR ([1]). Likewise, the PacketCopy sent to a BFR-NBR must clear all bits in its BitString that are not routed across BFR-NBR. This protects against BIER replication on any possible further BFR to create duplicates ([2]).

To solve both [1] and [2] for BIER, the F-BM of each bit needs to have all bits set that this BFR wants to route across BFR-NBR. [2] clears all other bits in PacketCopy->BitString, and [1] clears those bits from Packet->BitString after the first PacketCopy.

Internet-Draft

In BIER-TE, a BFR-NBR is an adjacency, forward_connected, forward_routed or local_decap. There is no need for [2] to suppress duplicates in the way BIER does because in general, different BP would never have the same adjacency. If a BIER-TE controller actually finds some optimization in which this would be desirable, then the controller is also responsible to ensure that only one of those bits is set in any Packet->BitString, unless the controller explicitly wants for duplicates to be created.

For BIER-TE, F-BM is handled as follows:

- The F-BM of all bits without an adjacency has all bits clear. This will cause [3] to skip further processing of such a bit.
- All bits with an adjacency (with DNC flag clear) have an F-BM that has only those bits set for which this BFR does not have an adjacency. This causes [2] to clear all bits from PacketCopy->BitString for which this BFR does have an adjacency.
- 3. [1] is not performed for BIER-TE. All bit clearing required by BIER-TE is performed by [2].

This Forwarding Pseudocode can support the REQUIRED BIER-TE forwarding functions (see <u>Section 4.6</u>), forward_connected, forward_routed() and local decap, but not the RECOMMENDED functions DNC flag and multiple adjacencies per bit nor the OPTIONAL function, ECMP adjacencies. The DNC flag cannot be supported when using only [1] to mask bits.

The modified and expanded Forwarding Pseudocode in Figure 6 specifies how to support all BIER-TE forwarding functions (required, recommended and optional):

- o This pseudocode eliminates per-bit F-BM, therefore reducing the size of BIFT state by BitStringLength^2*SI and eliminating the need for per-packet-copy masking operation except for adjacencies with the DNC flag set:
 - * AdjacentBits[SI] are bits with a non-empty list of adjacencies. This can be computed whenever the BIER-TE Controller updates the adjacencies.
 - * Only the AdjacentBits need to be examined in the loop for packet copies.
 - * The packets BitString is masked with those AdjacentBits before the loop to avoid doing this repeatedly for every PacketCopy.

o The code loops over the adjacencies because there may be more than

```
one adjacency for a bit.
o When an adjacency has the DNC bit, the bit is set in the packet
  copy (to save bits in rings for example).
o The ECMP adjacency is shown. Its parameters are a
  ListOfAdjacencies from which one is picked.
o The forward_local, forward_routed, local_decap() adjacencies are
  shown with their parameters.
void ForwardBitMaskPacket_withTE (Packet)
 {
     SI=GetPacketSI(Packet);
     Offset=SI*BitStringLength;
     AdjacentBits = Packet->BitString &= ~AdjacentBits[SI];
     Packet->BitString &= AdjacentBits[SI];
     for (Index = GetFirstbit position(AdjacentBits); Index ;
          Index = GetNextbit position(AdjacentBits, Index)) {
         foreach adjacency BIFT[Index+Offset] {
             if(adjacency == ECMP(ListOfAdjacencies, seed) ) {
                 I = ECMP_hash(sizeof(ListOfAdjacencies),
                               Packet->Entropy, seed);
                 adjacency = ListOfAdjacencies[I];
             }
             PacketCopy = Copy(Packet);
             switch(adjacency) {
                 case forward_connected(interface, neighbor, DNC):
                     if(DNC)
                         PacketCopy->BitString |= 1<<(Index-1);</pre>
                     SendToL2Unicast(PacketCopy, interface, neighbor);
                 case forward_routed({VRF}, neighbor):
                     SendToL3(PacketCopy, {VRF, }13-neighbor);
                 case local_decap({VRF}, neighbor):
                     DecapBierHeader(PacketCopy);
                     PassTo(PacketCopy, {VRF, }Packet->NextProto);
             }
         }
     }
 }
```

Figure 6: Complete BIER-TE Forwarding Pseudocode for required, recommended and optional functions

<u>4.5</u>. Basic BIER-TE Forwarding Example

[RFC Editor: remove this section.]

THIS SECTION TO BE REMOVED IN RFC BECAUSE IT WAS SUPERCEEDED BY <u>SECTION 1.1</u> EXAMPLE - UNLESS REVIEWERS CHIME IN AND EXPRESS DESIRE TO KEEP THIS ADDITIONAL EXAMPLE SECTION. ALVARO RETANA DID NOT MIND ANOTHER EXAMPLE.

Step by step example of basic BIER-TE forwarding. This example does not use ECMP or forward_routed() adjacencies nor does it try to minimize the number of required BitPositions for the topology.

[BIER-TE Controller]					
/	Ι				
V	v v				
p13 p1					
+- BFIR2+	I				
. p2	p6	. LAN2			
. + B	FR3+	.			
.	p7	p11			
Src -+ .	+ B	FER1+			
. p3	p8	.			
. + B	FR4+	. + Rcv1			
.	I	.			
.					
p14 p4					
+- BFIR1+	I				
. + B	FR5+ p10	p12			
LAN1. p5	р9 + В	FER2+			
	I	. + Rcv2			
		.			
		. LAN3			
IP BIER-T	E network	IP			

Figure 7: BIER-TE Forwarding Example

pXX indicate the BitPositions number assigned by the BIER-TE Controller to adjacencies in the BIER-TE topology. For example, p9 is the adjacency towards BFR5 on the LAN connecting to BFER2.

BIFT BFIR2: p13: local_decap p2: forward_connected(BFR3) BIFT BFR3: p1: forward_connected(BFIR2) p7: forward_connected(BFER1) p8: forward_connected(BFR4) BIFT BFER1: p11: local_decap p6: forward_connected(BFR3) p8: forward_connected(BFR4) Figure 8: BIER-TE Forwarding Example Adjacencies ...and so on. For example, we assume that some multicast traffic seen on LAN1 needs to be sent via BIER-TE by BFIR2 towards Rcv1 and Rcv2. The BIER-TE Controller determines it wants it to pass this traffic across the following paths: -> BFER1 ----> Rcv1 BFIR2 -> BFR3 -> BFR4 -> BFR5 -> BFER2 -> Rcv2 Figure 9: BIER-TE Forwarding Example Paths These paths equal to the following BitString: p2, p5, p7, p8, p10, p11, p12. This BitString is assigned by BFIR2 to the example multicast traffic received from LAN1. Then BFIR2 forwards this multicast traffic with BIER-TE based on that BitString. The BIFT of BFIR2 has only p2 and p13 populated. Only p2 is in the BitString and this is an adjacency towards BFR3. BFIR2 therefore clears p2 in the BitString and sends a copy towards BFR2.

BFR3 sees a BitString of p5,p7,p8,p10,p11,p12. For those BPs, it has only adjacencies for p7,p8. It creates a copy of the packet to BFER1 (due to p7) and one to BFR4 (due to p8). It clears p7, p8 before sending.

BFER1 sees a BitString of p5,p10,p11,p12. For those BPs, it only has an adjacency for p11. p11 is a "local_decap" adjacency installed by the BIER-TE Controller to receive a copy of the BIER packet - dispose

of the BIER header and pass the payload to IP multicast. IP multicast will then forward the packet out to LAN2 because it did receive PIM or IGMP joins on LAN2 for the traffic.

Further processing of the packet in BFR4, BFR5 and BFER2 accordingly.

<u>4.6</u>. BFR Requirements for BIER-TE forwarding

BFR MUST support to configure the BIFT of sub-domains so that they use BIER-TE forwarding rules instead of BIER forwarding rules. Every BP in the BIFT MUST support to have zero or one adjacency. Forwarding MUST support the adjacency types forward_connected() with clear DNC flag, forward_routed() and local_decap. As explained in <u>Section 4.4</u>, these REQUIRED BIER-TE forwarding functions can be implement via the same Forwarding Pseudocode as BIER forwarding except for one modification (skipping one masking with F-BM).

BIER-TE forwarding SHOULD support forward_connected() adjacencies with a set DNC flag, as this is highly useful to save bits in rings (see <u>Section 5.1.6</u>).

BIER-TE forwarding SHOULD support more than one adjacency on a bit. This allows to save bits in hub&spoke scenarios (see <u>Section 5.1.5</u>).

BIER-TE forwarding MAY support ECMP adjacencies to save bits in ECMP scenarios, see <u>Section 5.1.7</u> for an example. This is a MAY requirement, because the deployment importance of ECMP adjacencies for BIER-TE is unclear as one can also leverage ECMP of the routing underlay via forwarded_routed adjacencies and/or might prefer to have more explicit control of the path chosen via explicit BP/adjacencies for each ECMP path alternative.

5. BIER-TE Controller Operational Considerations

<u>5.1</u>. Bit position Assignments

This section describes how the BIER-TE Controller can use the different BIER-TE adjacency types to define the bit positions of a BIER-TE domain.

Because the size of the BitString limits the size of the BIER-TE domain, many of the options described exist to support larger topologies with fewer bit positions (4.1, 4.3, 4.4, 4.5, 4.6, 4.7, 4.8).

BIER-TE ARCH

5.1.1. P2P Links

On a P2P link that connects two BFR, the same bit position can be used on both BFR for the adjacency to the neighboring BFR. A P2P link requires therefore only one bit position.

5.1.2. BFER

Every non-Leaf BFER is given a unique bit position with a local_decap adjacency.

5.1.3. Leaf BFERs

BFR1(P) BFR2(P)	BFR1(P)	BFR2(P)
X		
BFER1(PE) BFER2(PE)	BFER1(PE)	BFER2(PE)

^ U-turn link

Leaf BFER /	Non-Leaf BFER /		
PE-router	PE-router		

Figure 10: Leaf vs. non-Leaf BFER Example

A leaf BFERs is one where incoming BIER-TE packets never need to be forwarded to another BFR but are only sent to the BFER to exit the BIER-TE domain. For example, in networks where Provider Edge (PE) router are spokes connected to Provider (P) routers, those PEs are Leaf BFERs unless there is a U-turn between two PEs.

Consider how redundant disjoint traffic can reach BFER1/BFER2 in Figure 10: When BFER1/BFER2 are Non-Leaf BFER as shown on the right hand side, one traffic copy would be forwarded to BFER1 from BFR1, but the other one could only reach BFER1 via BFER2, which makes BFER2 a non-Leaf BFER. Likewise BFER1 is a non-Leaf BFER when forwarding traffic to BFER2. Note that the BFERs in the left hand picture are only guaranteed to be leaf-BFER by fitting routing configuration that prohibits transit traffic to pass through the BFERs, which is commonly applied in these topologies.

All leaf-BFERs in a BIER-TE domain can share a single bit position. This is possible because the bit position for the adjacency to reach the BFER can be used to distinguish whether or not packets should reach the BFER.

This optimization will not work if an upstream interface of the BFER is using a bit position optimized as described in the following two sections (LAN, Hub and Spoke).

<u>5.1.4</u>. LANs

In a LAN, the adjacency to each neighboring BFR is given a unique bit position. The adjacency of this bit position is a forward_connected() adjacency towards the BFR and this bit position is populated into the BIFT of all the other BFRs on that LAN.

BFR1 |p1 LAN1-+-+---+ p3| p4| p2| BFR3 BFR4 BFR7

Figure 11: LAN Example

If Bandwidth on the LAN is not an issue and most BIER-TE traffic should be copied to all neighbors on a LAN, then bit positions can be saved by assigning just a single bit position to the LAN and populating the bit position of the BIFTs of each BFRs on the LAN with a list of forward_connected() adjacencies to all other neighbors on the LAN.

This optimization does not work in the case of BFRs redundantly connected to more than one LAN with this optimization because these BFRs would receive duplicates and forward those duplicates into the opposite LANs. Adjacencies of such BFRs into their LAN still need a separate bit position.

5.1.5. Hub and Spoke

In a setup with a hub and multiple spokes connected via separate p2p links to the hub, all p2p links can share the same bit position. The bit position on the hub's BIFT is set up with a list of forward_connected() adjacencies, one for each Spoke.

This option is similar to the bit position optimization in LANs: Redundantly connected spokes need their own bit positions, unless they are themselves Leaf-BFER.

This type of optimized BP could be used for example when all traffic is "broadcast" traffic (very dense receiver set) such as live-TV or situation-awareness (SA). This BP optimization can then be used to explicitly steer different traffic flows across different ECMP paths

in Data-Center or broadband-aggregation networks with minimal use of BPs.

5.1.6. Rings

In L3 rings, instead of assigning a single bit position for every p2p link in the ring, it is possible to save bit positions by setting the "DoNotClear" (DNC) flag on forward_connected() adjacencies.

For the rings shown in Figure 12, a single bit position will suffice to forward traffic entering the ring at BFRa or BFRb all the way up to BFR1:

On BFRa, BFRb, BFR30,... BFR3, the bit position is populated with a forward_connected() adjacency pointing to the clockwise neighbor on the ring and with DNC set. On BFR2, the adjacency also points to the clockwise neighbor BFR1, but without DNC set.

Handling DNC this way ensures that copies forwarded from any BFR in the ring to a BFR outside the ring will not have the ring bit position set, therefore minimizing the chance to create loops.



Figure 12: Ring Example

Note that this example only permits for packets intended to make it all the way around the ring to enter it at BFRa and BFRb, and that packets will always travel clockwise. If packets should be allowed to enter the ring at any ring BFR, then one would have to use two ring bit positions. One for each direction: clockwise and counterclockwise.

Both would be set up to stop rotating on the same link, e.g. L1. When the ingress ring BFR creates the clockwise copy, it will clear the counterclockwise bit position because the DNC bit only applies to the bit for which the replication is done. Likewise for the clockwise bit position for the counterclockwise copy. As a result, the ring ingress BFR will send a copy in both directions, serving BFRs on either side of the ring up to L1.

5.1.7. Equal Cost MultiPath (ECMP)

The ECMP adjacency allows to use just one BP per link bundle between two BFRs instead of one BP for each p2p member link of that link bundle. In Figure 13, one BP is used across L1,L2,L3.

> --L1-----BFR1 --L2---- BFR2 --L3----

BIFT entry in BFR1:

Index	A	djacencies		
0:6 	E	CMP({forward_connected(L1, forward_connected(L2, forward_connected(L3,	BFR2), BFR2), BFR2)},	 seed)

```
BIFT entry in BFR2:

Index | Adjacencies |

0:6 | ECMP({forward_connected(L1, BFR1), |

forward_connected(L2, BFR1), |

forward_connected(L3, BFR1)}, seed)
```

Figure 13: ECMP Example

This document does not standardize any ECMP algorithm because it is sufficient for implementations to document their freely chosen ECMP algorithm. This allows the BIER-TE Controller to calculate ECMP paths and seeds. Figure 14 shows an example ECMP algorithm:

forward(packet, ECMP(adj(0), adj(1),... adj(N-1), seed)): i = (packet(bier-header-entropy) XOR seed) % N forward packet to adj(i)

Figure 14: ECMP algorithm Example

In the following example, all traffic from BFR1 towards BFR10 is intended to be ECMP load split equally across the topology. This example is not meant as a likely setup, but to illustrate that ECMP can be used to share BPs not only across link bundles, but also

across alternative paths across different transit BFR, and it explains the use of the seed parameter.



Figure 15: Polarization Example

Note that for the following discussion of ECMP, only the BIFT ECMP adjacencies on BFR1, BFR2, BFR3 are relevant. The re-use of BP across BFR in this example is further explained in <u>Section 5.1.9</u> below.

With the setup of ECMP in the topology above, traffic would not be equally load-split. Instead, links L22 and L31 would see no traffic at all: BFR2 will only see traffic from BFR1 for which the ECMP hash in BFR1 selected the first adjacency in the list of 2 adjacencies given as parameters to the ECMP. It is link L11-to-BFR2. BFR2 performs again ECMP with two adjacencies on that subset of traffic using the same seed1, and will therefore again select the first of its two adjacencies: L21-to-BFR4. And therefore L22 and BFR5 sees no traffic. Likewise for L31 and BFR6.

This issue in BFR2/BFR3 is called polarization. It results from the re-use of the same hash function across multiple consecutive hops in topologies like these. To resolve this issue, the ECMP adjacency on BFR1 can be set up with a different seed2 than the ECMP adjacencies on BFR2/BFR3. BFR2/BFR3 can use the same hash because packets will not sequentially pass across both of them. Therefore, they can also use the same BP 0:7.

Note that ECMP solutions outside of BIER often hide the seed by autoselecting it from local entropy such as unique local or next-hop identifiers. Allowing the BIER-TE Controller to explicitly set the seed gives the ability for it to control same/different path selection across multiple consecutive ECMP hops.

5.1.8. Forward_routed adjacencies

<u>5.1.8.1</u>. Reducing bit positions

Forward_routed() adjacencies can reduce the number of bit positions required when the path steering requirement is not hop-by-hop explicit path selection, but loose-hop selection. Forward_routed() adjacencies can also allow to operate BIER-TE across intermediate hop routers that do not support BIER-TE.

....BFR1--... ...-L1-- BFR2...Routers. ...-L2--/BFR4--... BFR3... | L0 Network Area 1

Figure 16: Forward_routed Adjacencies Example

Assume the requirement in Figure 16 is to explicitly steer traffic flows that have arrived at BFR1 or BFR4 via a shortest path in the routing underlay "Network Area 1" to one of the following three next segments: (1) BFR2 via link L1, (2) BFR2 via link L2, or (3) via BFR3.

To enable this, both BFR1 and BFR4 are set up with a forward_routed adjacency bit position towards an address of BFR2 on link L1, another forward_routed() bit position towards an address of BFR2 on link L2 and a third forward_routed() bit position towards a node address L0 of BFR3.

<u>5.1.8.2</u>. Supporting nodes without BIER-TE

Forward_routed() adjacencies also enable incremental deployment of BIER-TE. Only the nodes through which BIER-TE traffic needs to be steered - with or without replication - need to support BIER-TE. Where they are not directly connected to each other, forward_routed adjacencies are used to pass over non BIER-TE enabled nodes.

<u>5.1.9</u>. Reuse of bit positions (without DNC)

bit positions can be re-used across multiple BFR to minimize the number of BP needed. This happens when adjacencies on multiple BFR use the DNC flag as described above, but it can also be done for non-DNC adjacencies. This section only discusses this non-DNC case.

Because BP are cleared when passing a BFR with an adjacency for that BP, reuse of BP across multiple BFR does not introduce any problems with duplicates or loops that do not also exist when every adjacency has a unique BP. Instead, the challenge when reusing BP is whether it allows to still achieve the desired Tree Engineering goals.

BP cannot be reused across two BFR that would need to be passed sequentially for some path: The first BFR will clear the BP, so those paths cannot be built. BP can be set across BFR that would (A) only occur across different paths or (B) across different branches of the same tree.

An example of (A) was given in Figure 15, where BP 0:7, BP 0:8 and BP 0:9 are each reused across multiple BFRs because a single packet/path would never be able to reach more than one BFR sharing the same BP.

Assume the example was changed: BFR1 has no ECMP adjacency for BP 0:6, but instead BP 0:5 with forward_connected() to BFR2 and BP 0:6 with forward_connected() to BFR3. Packets with both BP 0:5 and BP 0:6 would now be able to reach both BFR2 and BFR3 and the still existing re-use of BP 0:7 between BFR2 and BFR3 is a case of (B) where reuse of BP is perfect because it does not limit the set of useful path choices:

If instead of reusing BP 0:7, BFR3 used a separate BP 0:10 for its ECMP adjacency, no useful additional path steering options would be enabled. If duplicates at BFR10 where undesirable, this would be done by not setting BP 0:5 and BP 0:6 for the same packet. If the duplicates where desirable (e.g.: resilient transmission), the additional BP 0:10 would also not render additional value.



Figure 17: Reuse of BP

Reuse may also save BPs in larger topologies. Consider the topology shown in Figure 20. A BFIR/sender (e.g.: video headend) is attached to area 1, and area 2...6 contain receivers/BFER. Assume each area had a distribution ring, each with two BPs to indicate the direction (as explained before). These two BPs could be reused across the 5 areas. Packets would be replicated through other BPs for the Core to the desired subset of areas, and once a packet copy reaches the ring of the area, the two ring BPs come into play. This reuse is a case of (B), but it limits the topology choices: Packets can only flow around the same direction in the rings of all areas. This may or may not be acceptable based on the desired path steering options: If resilient transmission is the path engineering goal, then it is

likely a good optimization, if the bandwidth of each ring was to be optimized separately, it would not be a good limitation.

5.1.10. Summary of BP optimizations

This section reviewed a range of techniques by which a BIER-TE Controller can create a BIER-TE topology in a way that minimizes the number of necessary BPs.

Without any optimization, a BIER-TE Controller would attempt to map the network subnet topology 1:1 into the BIER-TE topology and every subnet adjacent neighbor requires a forward_connected() BP and every BFER requires a local_decap() BP.

The optimizations described are then as follows:

- o P2P links require only one BP (<u>Section 5.1.1</u>).
- o All leaf-BFER can share a single local_decap() BP (<u>Section 5.1.3</u>).
- o A LAN with N BFR needs at most N BP (one for each BFR). It only needs one BP for all those BFR that are not redundantly connected to multiple LANs (Section 5.1.4).
- A hub with p2p connections to multiple non-leaf-BFER spokes can share one BP to all spokes if traffic can be flooded to all spokes, e.g.: because of no bandwidth concerns or dense receiver sets (Section 5.1.5).
- o Rings of BFR can be built with just two BP (one for each direction) except for BFR with multiple ring connections similar to LANs (Section 5.1.6).
- ECMP adjacencies to N neighbors can replace N BP with 1 BP.
 Multihop ECMP can avoid polarization through different seeds of the ECMP algorithm (<u>Section 5.1.7</u>).
- o Forward_routed() adjacencies allow to "tunnel" across non-BIER-TE capable routers and across BIER-TE capable routers where no traffic-steering or replications are required (<u>Section 5.1.8</u>).
- o BP can generally be reused across nodes that do not need to be consecutive in paths, but depending on scenario, this may limit the feasible path steering options (<u>Section 5.1.9</u>).

Note that the described list of optimizations is not exhaustive. Especially when the set of required path steering choices is limited and the set of possible subsets of BFERs that should be able to

receive traffic is limited, further optimizations of BP are possible. The hub & spoke optimization is a simple example of such traffic pattern dependent optimizations.

5.2. Avoiding duplicates and loops

<u>5.2.1</u>. Loops

Whenever BIER-TE creates a copy of a packet, the BitString of that copy will have all bit positions cleared that are associated with adjacencies on the BFR. This inhibits looping of packets. The only exception are adjacencies with DNC set.



Figure 18: Miswired Ring Example

With DNC set, looping can happen. Consider in Figure 18 that link L4 from BFR3 is (inadvertently) plugged into the L1 interface of BFRa (instead of BFR2). This creates a loop where the rings clockwise bit position is never cleared for copies of the packets traveling clockwise around the ring.

To inhibit looping in the face of such physical misconfiguration, only forward_connected() adjacencies are permitted to have DNC set, and the link layer port unique unicast destination address of the adjacency (e.g. MAC address) protects against closing the loop. Link layers without port unique link layer addresses should not be used with the DNC flag set.

5.2.2. Duplicates

BFIR1 / \ / p2 \ p3 BFR2 BFR3 \ p4 / p5 \ / BFER4

Figure 19: Duplicates Example

Duplicates happen when the graph expressed by a BitString is not a tree but redundantly connecting BFRs with each other. In Figure 19, a BitString of p2,p3,p4,p5 would result in duplicate packets to arrive on BFER4. The BIER-TE Controller must therefore ensure to only create BitStrings that are trees.

When links are incorrectly physically re-connected before the BIER-TE Controller updates BitStrings in BFIRs, duplicates can happen. Like loops, these can be inhibited by link layer addressing in forward_connected() adjacencies.

If interface or loopback addresses used in forward_routed() adjacencies are moved from one BFR to another, duplicates can equally happen. Such re-addressing operations must be coordinated with the BIER-TE Controller.

5.3. Managing SI, sub-domains and BFR-ids

When the number of bits required to represent the necessary hops in the topology and BFER exceeds the supported BitStringLength (BSL), multiple SIs and/or sub-domains must be used. This section discusses how.

BIER-TE forwarding does not require the concept of BFR-id, but routing underlay, flow overlay and BIER headers may. This section also discusses how BFR-ids can be assigned to BFIR/BFER for BIER-TE.

5.3.1. Why SI and sub-domains

For BIER and BIER-TE forwarding, the most important result of using multiple SI and/or sub-domains is the same: Packets that need to be sent to BFERs in different SIs or sub-domains require different BIER packets: each one with a BitString for a different (SI,sub-domain) combination. Each such BitString uses one BSL sized SI block in the BIFT of the sub-domain. We call this a BIFT:SI (block).

For BIER and BIER-TE forwarding themselves there is also no difference whether different SIs and/or sub-domains are chosen, but
SI and sub-domain have different purposes in the BIER architecture shared by BIER-TE. This impacts how operators are managing them and how especially flow overlays will likely use them.

By default, every possible BFIR/BFER in a BIER network would likely be given a BFR-id in sub-domain 0 (unless there are > 64k BFIR/BFER).

If there are different flow services (or service instances) requiring replication to different subsets of BFERs, then it will likely not be possible to achieve the best replication efficiency for all of these service instances via sub-domain 0. Ideal replication efficiency for N BFER exists in a sub-domain if they are split over not more than ceiling(N/BitStringLength) SI.

If service instances justify additional BIER:SI state in the network, additional sub-domains will be used: BFIR/BFER are assigned BFR-id in those sub-domains and each service instance is configured to use the most appropriate sub-domain. This results in improved replication efficiency for different services.

Even if creation of sub-domains and assignment of BFR-id to BFIR/BFER in those sub-domains is automated, it is not expected that individual service instances can deal with BFER in different sub-domains. A service instance may only support configuration of a single subdomain it should rely on.

To be able to easily reuse (and modify as little as possible) existing BIER procedures including flow-overlay and routing underlay, when BIER-TE forwarding is added, we therefore reuse SI and subdomain logically in the same way as they are used in BIER: All necessary BFIR/BFER for a service use a single BIER-TE BIFT and are split across as many SIs as necessary (see <u>Section 5.3.2</u>). Different services may use different sub-domains that primarily exist to provide more efficient replication (and for BIER-TE desirable path steering) for different subsets of BFIR/BFER.

<u>5.3.2</u>. Assigning bits for the BIER-TE topology

In BIER, BitStrings only need to carry bits for BFERs, which leads to the model that BFR-ids map 1:1 to each bit in a BitString.

In BIER-TE, BitStrings need to carry bits to indicate not only the receiving BFER but also the intermediate hops/links across which the packet must be sent. The maximum number of BFER that can be supported in a single BitString or BIFT:SI depends on the number of bits necessary to represent the desired topology between them.

"Desired" topology because it depends on the physical topology, and on the desire of the operator to allow for explicit path steering across every single hop (which requires more bits), or reducing the number of required bits by exploiting optimizations such as unicast (forward_routed), ECMP or flood (DNC) over "uninteresting" sub-parts of the topology - e.g. parts where different trees do not need to take different paths due to path steering reasons.

The total number of bits to describe the topology vs. the number of BFERs in a BIFT:SI can range widely based on the size of the topology and the amount of alternative paths in it. The higher the percentage of non-BFER bits, the higher the likelihood, that those topology bits are not just BIER-TE overhead without additional benefit, but instead that they will allow to express desirable path steering alternatives.

5.3.3. Assigning BFR-id with BIER-TE

BIER-TE forwarding does not use the BFR-id, not does it require for the BFR-id field of the BIER header to be set to a particular value. However, other parts of a BIER-TE deployment may need a BFR-id, specifically overlay signaling, and in that case BFR need to also have BFR-ids for BIER-TE SDs.

For example, for BIER overlay signaling, BFIR need to have a BFR-id, because this BFIR BFR-id is carried in the BFR-id field of the BIER header to indicate to the overlay signaling on the receiving BFER which BFIR originated the packet.

In BIER, BFR-id = BSL * SI + BP, such that the SI and BP of a BFER can be calculated from the BFR-id and vice versa. This also means that every BFR with a BFR-id has a reserved BP in an SI, even if that is not necessary for BIER forwarding, because the BFR may never be a BFER but only a BFIR.

In BIER-TE, for a non-leaf BFER, there is usually a single BP for that BFER with a local_decap() adjacency on the BFER. The BFR-id for such a BFER can therefore equally be determined as in BIER: BFR-id = SI * BitStringLength + BP.

As explained in <u>Section 5.1.3</u>, leaf BFERs do not need such a unique local_decap() adjacency, likewise, BFIR who are not also BFER may not have a unique local_decap() adjacency either. For all those BFIR and (leaf) BFER, the controller needs to determine unique BFR-ids that do not collide with the BFR-ids derived from the non-leaf BFER local_decap() BPs.

While this document defines no requirements how to allocate such BFRid, a simple option is to derive it from the (SI,BP) of an adjacency

that is unique to the BFR in question. For a BFIR this can be he first adjacency only populated on this BFIR, for a leaf-BFER, this could be the first BP with an adjacency towards that BFER.

5.3.4. Mapping from BFR to BitStrings with BIER-TE

In BIER, applications of the flow overlay on a BFIR can calculate the (SI,BP) of a BFER from the BFR-id of the BFER and can therefore easily determine the BitStrings for a BIER packet to a set of BFER with known BFR-ids.

In BIER-TE this mapping needs to be equally supported for flow overlays. This section outlines two core options, based on how "complex" the Tree Engineering is that the BIER-TE controller performs for a particular application.

"Independent branches": For a given flow overlay instance, the branches from a BFIR to every BFER are calculated by the BIER-TE controller to be independent of the branches to any other BFER. Shortest path trees are the most common examples of trees with independent branches.

"Interdependent branches": When a BFER is added or deleted from a particular distribution tree, the BIER-TE controller has to recalculate the branches to other BFER, because they may need to change. Steiner trees are examples of interdependent branch trees.

If "independent branches" are used, the BIER-TE Controller can signal to the BFIR flow overlay for every BFER an SI:BitString that represents the branch to that BFER. The flow overlay on the BIFR can then independently of the controller calculate the SI:BitString for all desired BFER by OR'ing their BitStrings. This allows for flow overlay applications to operate independently from the controller whenever it needs to determine which subset of BFERs need to receive a particular packet.

If "interdependent branches" are required, the application would need to inquire the SI:BitString for a given set of BFER whenever the set changes.

Note that in either case (unlike in BIER), the bits may need to change upon link/node failure/recovery, network expansion and network resource consumption by other traffic as part of traffic engineering goals (e.g.: re-optimization of lower priority traffic flows). Interactions between such BFIR applications and the BIER-TE Controller do therefore need to support dynamic updates to the SI:BitStrings.

Communications between BFIR flow overlay and BIER-TE controller requires some way to identify BFER. If BFR-ids are used in the deployment, as outlined in <u>Section 5.3.3</u>, then those are the natural BFR identifier. If BFR-ids are not used, then any other unique identifier, such as the BFR-prefix of the BFR as of [<u>RFC8279</u>] could be used.

<u>5.3.5</u>. Assigning BFR-ids for BIER-TE

It is not currently determined if a single sub-domain could or should be allowed to forward both BIER and BIER-TE packets. If this should be supported, there are two options:

A. BIER and BIER-TE have different BFR-id in the same sub-domain. This allows higher replication efficiency for BIER because their BFRid can be assigned sequentially, while the BitStrings for BIER-TE will have also the additional bits for the topology. There is no relationship between a BFR BIER BFR-id and BIER-TE BFR-id.

B. BIER and BIER-TE share the same BFR-id. The BFR-ids are assigned as explained above for BIER-TE and simply reused for BIER. The replication efficiency for BIER will be as low as that for BIER-TE in this approach. Depending on topology, only the same 20%..80% of bits as possible for BIER-TE can be used for BIER.

<u>5.3.6</u>. Example bit allocations

5.3.6.1. With BIER

Consider a network setup with a BSL of 256 for a network topology as shown in Figure 20. The network has 6 areas, each with 170 BFRs, connecting via a core with 4 (core) BFRs. To address all BFERs with BIER, 4 SIs are required. To send a BIER packet to all BFER in the network, 4 copies need to be sent by the BFIR. On the BFIR it does not make a difference how the BFR-ids are allocated to BFER in the network, but for efficiency further down in the network it does make a difference.

area1 area2 area3 BFR1a BFR1b BFR2a BFR2b BFR3a BFR3b | \ / \ / | Core . | / \ / \ | BFR4a BFR4b BFR5a BFR5b BFR6a BFR6b area4 area5 area6

Figure 20: Scaling BIER-TE bits by reuse

With random allocation of BFR-id to BFER, each receiving area would (most likely) have to receive all 4 copies of the BIER packet because there would be BFR-id for each of the 4 SIs in each of the areas. Only further towards each BFER would this duplication subside - when each of the 4 trees runs out of branches.

If BFR-ids are allocated intelligently, then all the BFER in an area would be given BFR-id with as few as possible different SIs. Each area would only have to forward one or two packets instead of 4.

Given how networks can grow over time, replication efficiency in an area will also easily go down over time when BFR-ids are network wide allocated sequentially over time. An area that initially only has BFR-id in one SI might end up with many SIs over a longer period of growth. Allocating SIs to areas with initially sufficiently many spare bits for growths can help to alleviate this issue. Or renumber BFERs after network expansion. In this example one may consider to use 6 SIs and assign one to each area.

This example shows that intelligent BFR-id allocation within at least sub-domain 0 can even be helpful or even necessary in BIER.

5.3.6.2. With BIER-TE

In BIER-TE one needs to determine a subset of the physical topology and attached BFERs so that the "desired" representation of this topology and the BFER fit into a single BitString. This process needs to be repeated until the whole topology is covered.

Once bits/SIs are assigned to topology and BFERs, BFR-id is just a derived set of identifiers from the operator/BIER-TE Controller as explained above.

Every time that different sub-topologies have overlap, bits need to be repeated across the BitStrings, increasing the overall amount of bits required across all BitString/SIs. In the worst case, random

subsets of BFERs are assigned to different SIs. This is much worse than in BIER because it not only reduces replication efficiency with the same number of overall bits, but even further - because more bits are required due to duplication of bits for topology across multiple SIs. Intelligent BFER to SI assignment and selecting specific "desired" subtopologies can minimize this problem.

To set up BIER-TE efficiently for the above topology, the following bit allocation method can be used. This method can easily be expanded to other, similarly structured larger topologies.

Each area is allocated one or more SIs depending on the number of future expected BFERs and number of bits required for the topology in the area. In this example, 6 SIs, one per area.

In addition, we use 4 bits in each SI: bia, bib, bea, beb: bit ingress a, bit ingress b, bit egress a, bit egress b. These bits will be used to pass BIER packets from any BFIR via any combination of ingress area a/b BFR and egress area a/b BFR into a specific target area. These bits are then set up with the right forward_routed() adjacencies on the BFIR and area edge BFR:

On all BFIRs in an area j|j=2...6, bia in each BIFT:SI is populated with the same forward_routed(BFRja), and bib with forward_routed(BFRjb). On all area edge BFR, bea in BIFT:SI=k|k=2...6 is populated with forward_routed(BFRka) and beb in BIFT:SI=k with forward_routed(BFRkb).

For BIER-TE forwarding of a packet to a subset of BFERs across all areas, a BFIR would create at most 6 copies, with SI=1...SI=6, In each packet, the bits indicate bits for topology and BFER in that topology plus the four bits to indicate whether to pass this packet via the ingress area a or b border BFR and the egress area a or b border BFR, therefore allowing path steering for those two "unicast" legs: 1) BFIR to ingress are edge and 2) core to egress area edge. Replication only happens inside the egress areas. For BFER in the same area as in the BFIR, these four bits are not used.

5.3.7. Summary

BIER-TE can, like BIER, support multiple SIs within a sub-domain to allow re-using the concept of BFR-id and therefore minimize BIER-TE specific functions in any possible BIER layer control plane used in conjunction with BIER-TE, flow overlay methods and BIER headers.

The number of BFIR/BFER possible in a sub-domain is smaller than in BIER because BIER-TE uses additional bits for topology.

Sub-domains (SDs) in BIER-TE can be used like in BIER to create more efficient replication to known subsets of BFERs.

Assigning bits for BFERs intelligently into the right SI is more important in BIER-TE than in BIER because of replication efficiency and overall amount of bits required.

<u>6</u>. BIER-TE and Segment Routing

SR aims to enable lightweight path steering via loose source routing. Compared to its more heavy-weight predecessor RSVP-TE, SR does for example not require per-path signaling to each of these hops.

BIER-TE supports the same design philosophy for multicast. Like in SR, it relies on source-routing - via the definition of a BitString. Like SR, it only requires to consider the "hops" on which either replication has to happen, or across which the traffic should be steered (even without replication). Any other hops can be skipped via the use of routed adjacencies.

BIER-TE bit position (BP) can be understood as the BIER-TE equivalent of "forwarding segments" in SR, but they have a different scope than SR forwarding segments. Whereas forwarding segments in SR are global or local, BPs in BIER-TE have a scope that is the group of BFR(s) that have adjacencies for this BP in their BIFT. This can be called "adjacency" scoped forwarding segments.

Adjacency scope could be global, but then every BFR would need an adjacency for this BP, for example a forward_routed() adjacency with encapsulation to the global SR SID of the destination. Such a BP would always result in ingress replication though. The first BFR encountering this BP would directly replicate to it. Only by using non-global adjacency scope for BPs can traffic be steered and replicated on non-ingress BFR.

SR can naturally be combined with BIER-TE and help to optimize it. For example, instead of defining bit positions for non-replicating hops, it is equally possible to use segment routing encapsulations (e.g. SR-MPLS label stacks) for the encapsulation of "forward_routed" adjacencies.

Note that BIER itself can also be seen to be similar to SR. BIER BPs act as global destination Node-SIDs and the BIER BitString is simply a highly optimized mechanism to indicate multiple such SIDs and let the network take care of effectively replicating the packet hop-byhop to each destination Node-SID. What BIER does not allow is to indicate intermediate hops, or terms of SR the ability to indicate a

sequence of SID to reach the destination. This is what BIER-TE and its adjacency scoped BP enables.

Both BIER and BIER-TE allow BFIR to "opportunistically" copy packets to a set of desired BFER on a packet-by-packet basis. In BIER, this is done by OR'ing the BP for the desired BFER. In BIER-TE this can be done by OR'ing for each desired BFER a BitString using the "independent branches" approach described in <u>Section 5.3.3</u> and therefore also indicating the engineered path towards each desired BFER. This is the approach that [<u>I-D.ietf-bier-multicast-http-response</u>] relies on.

7. Security Considerations

If [<u>RFC8296</u>] is used, BIER-TE shares its security considerations.

BIER-TE shares the security considerations of BIER, [<u>RFC8279</u>], with the following overriding or additional considerations.

In BIER, the standardized methods for the routing underlays as well as to distribute BFR-ids and BFR-prefixes are IGPs such as specified in [<u>RFC8401</u>] for IS-IS and in [<u>RFC8444</u>] for OSPF. Attacking the protocols for the BIER routing underlay or BIER layer control plane, or impairment of any BFR in a domain may lead to successful attacks against the results of the routing protocol, enabling DoS attacks against paths or addressing (BFR-id, BFR-prefixes) used by BIER.

The reference model for the BIER-TE layer control plane is a BIER-TE controller. When such a controller is used, impairment of individual BFR in a domain causes no impairment of the BIER-TE control plane on other BFR. If a routing protocol is used to support forward_routed() adjacencies, then this is still an attack vector as in BIER, but only for BIER-TE forward_routed() adjacencies, and no other adjacencies.

Whereas IGP routing protocols are most often not well secured through cryptographic authentication and confidentiality, communications between controllers and routers such as those to be considered for the BIER-TE controller/control-plane can and are much more commonly secured with those security properties, for example by using Secure SHell (SSH), [RFC4253] for NetConf ([RFC6241]), or via Transport Layer Security (TLS), such as [RFC8253] for PCEP, [RFC5440], or [RFC7589] for NetConf.

For additional, BIER-TE independent security considerations for the use of a central BIER-TE controller, the security section of the protocols and security options in the previous paragraph apply. In addition, the security considerations of [<u>RFC4655</u>] apply.

The most important attack vector in BIER-TE is misconfiguration, either on the BFR themselves or via the BIER-TE controller. Forwarding entries with DNC could be set up to create persistent loops, in which packets only expire because of TTL. To minimize the impact of such attacks (or more likely unintentional misconfiguration by operators and/or bad BIER-TE controller software), the BIER-TE forwarding rules are defined to be as strict in clearing bits as they are. The clearing of all bits with an adjacency on a BFR prohibits that a looping packet creates additional packet amplification through the misconfigured loop on the packets second or further times around the loop, because all relevant adjacency bits would have been cleared on the first round through the loop. In result, BIER-TE has the same degree of looping packets as possible with unintentional or malicious loops in the routing underlay with BIER or even with unicast traffic.

Deployments especially where BIER-TE would likely be beneficial may include operational models where actual configuration changes from the controller are only required during non-productive phases of the networks life-cycle, such as in embedded networks or in manufacturing networks during e.g. plant reworking/repairs. In these type of deployments, configuration changes could be locked out when the network is in production state and could only be (re-)enabled through reverting the network/installation into non-productive state. Such security designs would not only allow to provide additional layers of protection against configuration attacks, but would foremost protect the active production process from such configuration attacks.

8. IANA Considerations

This document requests no action by IANA.

9. Acknowledgements

The authors would like to thank Greg Shepherd, Ijsbrand Wijnands, Neale Ranns, Dirk Trossen, Sandy Zheng, Lou Berger, Jeffrey Zhang, Alvaro Retana and Wolfgang Braun for their reviews and suggestions.

<u>10</u>. Change log [RFC Editor: Please remove]

draft-ietf-bier-te-arch:

10: AD review Alvaro Retana, summary:

Note: rfcdiff shows more changes than actually exist because text moved around.

Summary:

- 1. restructuring: merged all controller sections under common controller ops main section, moved unfitting stuff out to other parts of doc. Split Intro section into Overview and Intro. Shortened Abstract, moved text into Overview, added sections overview.
- enhanced/rewrote: 2.3 Comparison with -> Relationship to BIER-TE
- 3. enhanced/rewrote: 3.2 BIER-TE controller -> BIER-TE control
 plane, 3.2.1 BIER-TE controller, for consistency with rfc8279
- 4. additional subsections for Alvaros asks
- added to: 3.3 BIER-TE forwarding plane (consistency with rfc8279)
- 6. Enhanced description of 4.3/encap considerations to better explain how BIER/BIER-TE can run together.

Notation: Markers (a),(b),... at end of points are references from the review discussion with Alvaro to the changes made.

Details:.

Throughout text: changed term spelling to <u>rfc8279</u> - bit positions, sub-domain, ... (i).

Reset changed to clear, also DNR changed to DNC (Do Not Clear) (q).

Abstract: Shortened. Removed name explanation note (Tree Engineering), (a).

1. Introduction -> Overview: Moved important explanation paragraph from abstract to Introduction. Fixed text, (a).

Added bullet point list explanation of structure of document (e).

Renamed to Overview because that is now more factually correct.

1.1. Fixed bug in example adding bit p15.(1).

2. (New - Introduction): Moved <u>section 1.1</u> - 1.3 (examples, comparison with BIER-TE) from Introduction into new "Overview" section. Primarily so that "requirements language" section (at end of Introduction) is not in middle of document after all the Introduction.

2.1 Removed discussion of encap, moved to 4.2.2 (m).

2.2 enhanced paragraph suggesting native/overlay topology types, also sugest type hybrid (n).

2.3 Overhauled comparison text BIER/BIER-TE, structured into common, different, not-required-by-te, integration-bier-bier-te. Changed title to "Relationship" to allow including last point. (f).

2.4 moved Hardware forwarding comparison section into <u>section 2</u> to allow coalescing of sections into <u>section 5</u> about the controller operations (hardware forwarding was in the middle of it, wrong place). Shortened/improved third paragraph by pointing to BIFT as deciding element for selection between BIER/BIER-TE. Removed notion of experimentation (this now targets standard) (g).

3. (Components): Aligned component name and descriptions better with <u>RFC8279</u>. Now describe exactly same three layers. BIER layer constituted from BIER-TE control plane and BIER-TE forwarding plane. BIER-TE controller is now simply component of BIER-TE control plane. (b).

3.1. shortened/improved paragraph explaining use of SI:BP instead of also bfr-id as index into BIFT, rewrote paragraph talking about reuse of BPs(o).

3.2. rewrote explanation of BIER-TE control plane in the style of RFC8729 Section 4.2 (BIER layer) with numbered points. Note that RFC8729 mixes control and forwarding plane bullet points (this doc does not). Merged text from old sections 2.2.1 and 2.2.3 into list. (b).

3.2.1. Expanded/improved explanation of BIER-TE Controller (b).

3.2.1.1. Added subsection for topology discovery and creation (d).

3.2.1.2. Added subsection for engineered BitStrings as key novel aspect not found in BIER. (X).

3.3. Added numbered list for components of BIER-TE forwarding plane (completing the comparable text from <u>RFC8729 Section 4.2</u>).

3.4 Alvaro does not mind additional example, fixed bugs.

3.5 Removed notion about using IGP BIER extensions for BIER-TE, such as BIFT address ranges. After -10 making use of BIFT

clearer, it now looks to authors as if use of IGP extensions would not be beneficial, as long as we do need to use the BIER-TE controller, e.g. unlike in BIER, a BFR could not learn from the IGP information what traffic to send towards a particular BIFT-ID, but instead that is the core of what the controller needs to provide.

4.2.2 Improved text to explain requirement to identify BIER-TE in the tunnel encap and compress description of use-cases (m).

4.2.3 enhanced ECMP text (p).

4.3. rewrote most of Encapsulation Considerations to better explain to Alvaros question re sharing or not sharing SD via BIER/ BIER-TE. Added reference to I-D.ietf-bier-non-mpls-bift-encoding as a very helpful example. (f).

4.3 Renamed title to "...Co-Existence with BIER" as this is what it is about and to help finding it from abstract/intro ("co-exist") (j).

4.4. Moved BIER-TE Forwarding Pseudocode here to coalesce text logically. Changed text to better compare with BIER pseudo forwarding code. Numerical list of how F-BM works for BIER-TE. Removed efficiency comparison with BIER (too difficult to provide sufficient justification, derails from focus of section) (j).

4.6. (Requirements) Restructured: Removed notion of "basic" BIER-TE forwarding, simply referring to it now as "mandatory" BIER-TE forwarding. Cleaned up text to have requirements for different adjacencies in different paragraphs. (c).

5. Created new main section "BIER-TE Controller operational considerations", coalesced old sections 4., 5., 7. into this new main section. No text changes. (k).

5.1.9 Added new separate picture instead of referring to a picture later in text, adjusted text (r).

5.3.2 Changed title to not include word "comparison" to avoid this being accounted against Alvaros concern about scattering comparison (IMHO text already has little comparison, so title was misleading) (h).

co-authors internal review:

4.4 Added xref to Figure 5.

5.2.1 Duplicated ring picture, added visuals for described miswiring (s).

5.2.2 replace "topology" with graph (wrong word).

5.3.3 rewrote explanation of how to map BFR-id to SI:BP and assign them, clarified BFR-id is option. Retitled to better explain scope of section.

5.3.4 Removed considerations in 5.3.4 for sharing BFR-id across BIER/BIER-TE (t), changed title to explain how BFIR/BIER-TE controller interactions need some form of identifying BFR but this does not have to be BFR-id.

7. Added new security considerations (u).

09: Incorporated fixes for feedback from Shepherd (Xuesong Geng).

Added references for Bloom Filters and Rate Controlled Service Disciplines.

1.1 Fixed numbering of example 1 topology explanation. Improved language on second example (less abbreviating to avoid confusion about meaning).

1.2 Improved explanation of BIER-TE topology, fixed terminology of graphs (BIER-TE topology is a directed graph where the edges are the adjacencies).

2.4 Fixed and amended routing underlay explanations: detailed why no need for BFER routing underlay routing protocol extensions, but potential to re-use BIER routing underlay routing protocol extensions for non-BFER related extensions.

3.1 Added explanation for VRF and its use in adjacencies.

08: Incorporated (with hopefully acceptable fixes) for Lou suggested <u>section 2.5</u>, TE considerations.

Fixes are primarily to the point to a) emphasize that BIER-TE does not depend on the routing underlay unless forward_routed() adjacencies are used, and b) that the allocation and tracking of resources does not explicitly have to be tied to BPs, because they are just steering labels. Instead, it would ideally come from per-hop resource management that can be maintained only via local accounting in the controller.

07: Further reworking text for Lou.

Renamed BIER-PE to BIER-TE standing for "Tree Engineering" after votes from BIER WG.

Removed <u>section 1.1</u> (introduced by version 06) because not considered necessary in this doc by Lou (for framework doc).

Added [RFC editor pls. remove] Section to explain name change to future reviewers.

06: Concern by Lou Berger re. BIER-TE as full traffic engineering solution.

Changed title "Traffic Engineering" to "Path Engineering"

Added intro section of relationship BIER-PE to traffic engineering.

Changed "traffic engineering" term in text" to "path engineering", where appropriate

Other:

Shortened "BIER-TE Controller Host" to "BIER-TE Controller". Fixed up all instances of controller to do this.

05: Review Jeffrey Zhang.

Part 2:

4.3 added note about leaf-BFER being also a propery of routing setup.

4.7 Added missing details from example to avoid confusion with routed adjacencies, also compressed explanatory text and better justification why seed is explicitly configured by controller.

4.9 added section discussing generic reuse of BP methods.

4.10 added section summarizing BP optimizations of section 4.

6. Rewrote/compressed explanation of comparison BIER/BIER-TE forwarding difference. Explained benefit of BIER-TE per-BP forwarding being independent of forwarding for other BPs.

Part 1:

Explicitly ue forwarded_connected adjcency in ECMP adjcency examples to avoid confusion.

4.3 Add picture as example for leav vs. non-leaf BFR in topology. Improved description.

4.5 Exampe for traffic that can be broadcast -> for single BP in hub&spoke.

4.8.1 Simplified example picture for routed adjacency, explanatory text.

Review from Dirk Trossen:

Fixed up explanation of ICC paper vs. bloom filter.

04: spell check run.

Addded remaining fixes for Sandys (Zhang Zheng) review:

4.7 Enhance ECMP explanations:

example ECMP algorithm, highlight that doc does not standardize ECMP algorithm.

Review from Dirk Trossen:

1. Added mentioning of prior work for traffic engineered paths with bloom filters.

2. Changed title from layers to components and added "BIER-TE control plane" to "BIER-TE Controller" to make it clearer, what it does.

2.2.3. Added reference to I-D.ietf-bier-multicast-http-response as an example solution.

2.3. clarified sentence about resetting BPs before sending copies (also forgot to mention DNR here).

3.4. Added text saying this section will be removed unless IESG review finds enough redeeming value in this example given how -03 introduced <u>section 1.1</u> with basic examples.

7.2. Removed explicit numbers 20%/80% for number of topology bits in BIER-TE, replaced with more vague (high/low) description, because we do not have good reference material Added text saying this section will be removed unless IESG review finds enough

redeeming value in this example given how -03 introduced section 1.1 with basic examples.

many typos fixed. Thanks a lot.

03: Last call textual changes by authors to improve readability:

removed Wolfgang Braun as co-authors (as requested).

Improved abstract to be more explanatory. Removed mentioning of FRR (not concluded on so far).

Added new text into Introduction section because the text was too difficult to jump into (too many forward pointers). This primarily consists of examples and the early introduction of the BIER-TE Topology concept enabled by these examples.

Amended comparison to SR.

Changed syntax from [VRF] to {VRF} to indicate its optional and to make idnits happy.

Split references into normative / informative, added references.

02: Refresh after IETF104 discussion: changed intended status back to standard. Reasoning:

Tighter review of standards document == ensures arch will be better prepared for possible adoption by other WGs (e.g. DetNet) or std. bodies.

Requirement against the degree of existing implementations is self defined by the WG. BIER WG seems to think it is not necessary to apply multiple interoperating implementations against an architecture level document at this time to make it qualify to go to standards track. Also, the levels of support introduced in -O1 rev. should allow all BIER forwarding engines to also be able to support the base level BIER-TE forwarding.

01: Added note comparing BIER and SR to also hopefully clarify BIER-TE vs. BIER comparison re. SR.

- added requirements section mandating only most basic BIER-TE forwarding features as MUST.

- reworked comparison with BIER forwarding section to only summarize and point to pseudocode section.

- reworked pseudocode section to have one pseudocode that mirrors the BIER forwarding pseudocode to make comparison easier and a second pseudocode that shows the complete set of BIER-TE forwarding options and simplification/optimization possible vs. BIER forwarding. Removed MyBitsOfInterest (was pure optimization).

- Added captions to pictures.

- Part of review feedback from Sandy (Zhang Zheng) integrated.

00: Changed target state to experimental (WG conclusion), updated references, mod auth association.

- Source now on http://www.github.com/toerless/bier-te-arch

- Please open issues on the github for change/improvement requests to the document - in addition to posting them on the list (bier@ietf.). Thanks!.

draft-eckert-bier-te-arch:

06: Added overview of forwarding differences between BIER, BIER-TE.

05: Author affiliation change only.

04: Added comparison to Live-Live and BFIR to FRR section (Eckert).

04: Removed FRR content into the new FRR draft [I-D.eckert-bier-te-frr] (Braun).

- Linked FRR information to new draft in Overview/Introduction
- Removed BTAFT/FRR from "Changes in the network topology"
- Linked new draft in "Link/Node Failures and Recovery"
- Removed FRR from "The BIER-TE Forwarding Layer"
- Moved FRR section to new draft
- Moved FRR parts of Pseudocode into new draft
- Left only non FRR parts

- removed FrrUpDown(..) and //FRR operations in ForwardBierTePacket(..)

- New draft contains FrrUpDown(..) and ForwardBierTePacket(Packet) from bier-arch-03

- Moved "BIER-TE and existing FRR to new draft

- Moved "BIER-TE and Segment Routing" section one level up

- Thus, removed "Further considerations" that only contained this section

- Added Changes for version 04

03: Updated the FRR section. Added examples for FRR key concepts. Added BIER-in-BIER tunneling as option for tunnels in backup paths. BIFT structure is expanded and contains an additional match field to support full node protection with BIER-TE FRR.

03: Updated FRR section. Explanation how BIER-in-BIER encapsulation provides P2MP protection for node failures even though the routing underlay does not provide P2MP.

02: Changed the definition of BIFT to be more inline with BIER. In revs. up to -01, the idea was that a BIFT has only entries for a single BitString, and every SI and sub-domain would be a separate BIFT. In BIER, each BIFT covers all SI. This is now also how we define it in BIER-TE.

02: Added <u>Section 5.3</u> to explain the use of SI, sub-domains and BFR-id in BIER-TE and to give an example how to efficiently assign bits for a large topology requiring multiple SI.

02: Added further detailed for rings - how to support input from all ring nodes.

01: Fixed BFIR -> BFER for section 4.3.

01: Added explanation of SI, difference to BIER ECMP, consideration for Segment Routing, unicast FRR, considerations for encapsulation, explanations of BIER-TE Controller and CLI.

00: Initial version.
<u>11</u>. References

<u>**11.1</u>**. Normative References</u>

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in <u>RFC</u> 2119 Key Words", <u>BCP 14</u>, <u>RFC 8174</u>, DOI 10.17487/RFC8174, May 2017, <<u>https://www.rfc-editor.org/info/rfc8174</u>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", <u>RFC 8279</u>, DOI 10.17487/RFC8279, November 2017, <<u>https://www.rfc-editor.org/info/rfc8279</u>>.

<u>11.2</u>. Informative References

- [Bloom70] Bloom, B., "Space/time trade-offs in hash coding with allowable errors", Comm. ACM 13(7):422-6, July 1970.
- [I-D.eckert-bier-te-frr]

Eckert, T., Cauchie, G., Braun, W., and M. Menth, "Protection Methods for BIER-TE", <u>draft-eckert-bier-te-</u><u>frr-03</u> (work in progress), March 2018.

[I-D.ietf-bier-multicast-http-response]

Trossen, D., Rahman, A., Wang, C., and T. Eckert, "Applicability of BIER Multicast Overlay for Adaptive Streaming Services", <u>draft-ietf-bier-multicast-http-</u> <u>response-05</u> (work in progress), January 2021.

[I-D.ietf-bier-non-mpls-bift-encoding]

Wijnands, I., Mishra, M., Xu, X., and H. Bidgoli, "An Optional Encoding of the BIFT-id Field in the non-MPLS BIER Encapsulation", <u>draft-ietf-bier-non-mpls-bift-</u> <u>encoding-03</u> (work in progress), November 2020.

[I-D.ietf-bier-te-yang]

Zhang, Z., Wang, C., Chen, R., Hu, F., Sivakumar, M., and H. Chen, "A YANG data model for Traffic Engineering for Bit Index Explicit Replication (BIER-TE)", <u>draft-ietf-</u> <u>bier-te-yang-02</u> (work in progress), August 2020.

[I-D.ietf-roll-ccast]

Bergmann, O., Bormann, C., Gerdes, S., and H. Chen, "Constrained-Cast: Source-Routed Multicast for RPL", <u>draft-ietf-roll-ccast-01</u> (work in progress), October 2017.

[I-D.ietf-teas-rfc3272bis]

Farrel, A., "Overview and Principles of Internet Traffic Engineering", <u>draft-ietf-teas-rfc3272bis-11</u> (work in progress), April 2021.

- [ICC] Reed, M., Al-Naday, M., Thomos, N., Trossen, D., Petropoulos, G., and S. Spirou, "Stateless multicast switching in software defined networks", IEEE International Conference on Communications (ICC), Kuala Lumpur, Malaysia, 2016, May 2016, <<u>https://ieeexplore.ieee.org/document/7511036</u>>.
- [RCSD94] Zhang, H. and D. Domenico, "Rate-Controlled Service Disciplines", Journal of High-Speed Networks, 1994, May 1994, <<u>https://dl.acm.org/doi/10.5555/2692227.2692232</u>>.
- [RFC4253] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Transport Layer Protocol", <u>RFC 4253</u>, DOI 10.17487/RFC4253, January 2006, <<u>https://www.rfc-editor.org/info/rfc4253</u>>.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", <u>RFC 4655</u>, DOI 10.17487/RFC4655, August 2006, <<u>https://www.rfc-editor.org/info/rfc4655</u>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", <u>RFC 5440</u>, DOI 10.17487/RFC5440, March 2009, <<u>https://www.rfc-editor.org/info/rfc5440</u>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", <u>RFC 6241</u>, DOI 10.17487/RFC6241, June 2011, <<u>https://www.rfc-editor.org/info/rfc6241</u>>.
- [RFC7589] Badra, M., Luchuk, A., and J. Schoenwaelder, "Using the NETCONF Protocol over Transport Layer Security (TLS) with Mutual X.509 Authentication", <u>RFC 7589</u>, DOI 10.17487/RFC7589, June 2015, <<u>https://www.rfc-editor.org/info/rfc7589</u>>.

- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", <u>RFC 7752</u>, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.
- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", <u>RFC 8253</u>, DOI 10.17487/RFC8253, October 2017, <<u>https://www.rfc-editor.org/info/rfc8253</u>>.
- [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", <u>RFC 8296</u>, DOI 10.17487/RFC8296, January 2018, <<u>https://www.rfc-editor.org/info/rfc8296</u>>.
- [RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", <u>RFC 8345</u>, DOI 10.17487/RFC8345, March 2018, <<u>https://www.rfc-editor.org/info/rfc8345</u>>.
- [RFC8401] Ginsberg, L., Ed., Przygienda, T., Aldrin, S., and Z. Zhang, "Bit Index Explicit Replication (BIER) Support via IS-IS", <u>RFC 8401</u>, DOI 10.17487/RFC8401, June 2018, <https://www.rfc-editor.org/info/rfc8401>.
- [RFC8444] Psenak, P., Ed., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPFv2 Extensions for Bit Index Explicit Replication (BIER)", <u>RFC 8444</u>, DOI 10.17487/RFC8444, November 2018, <<u>https://www.rfc-editor.org/info/rfc8444</u>>.

Authors' Addresses

Toerless Eckert (editor) Futurewei Technologies Inc. 2330 Central Expy Santa Clara 95050 USA

Email: tte+ietf@cs.fau.de

Gregory Cauchie Bouygues Telecom

Email: GCAUCHIE@bouyguestelecom.fr

Michael Menth University of Tuebingen

Email: menth@uni-tuebingen.de