BIER-TE-ARCH
IETF105 Prague

draft-ietf-bier-te-arch-03

Toerless Eckert (tte@cs.fau.de)

Gregory Cauchie (GCAUCHIE@bouyguestelecom.fr)

Michael Menth (menth@uni-tuebingen.de)

Since IETF104

- Entered WGLC after IETF104
- Authors reviewed text.
  - No technical changes
  - But hard to understand document because of circular dependencies
  - And missing explanation of relationship between traffic engineering and
    - Native and/or overlay topologies
  - Solution: Add more introductory text to resolve those issues

- Rev'ed to draft-ietf-bier-te-arch-03

Overview

Abstract improved

- This memo introduces ***per-packet stateless strict and loose path engineered replication and forwarding*** for Bit Index Explicit Replication packets ([RFC8279]). This is called BIER-TE.
- …(more improvements)… pls. Re-read

1.1 New section: Basic Examples

- One topology, full set of bits, explaining different packets
- “native” – adjacencies are L2 (forward_connected)
- “overlay” – adjacencies are remote (forward_routed)

2. Layering

- Add topology and explains in 3 sentences how topology comes into play.

8. Comparison with Segment Routing

- Extended: if you want to think of bits as segments, what does it mean ?
- Hopefully instructive to users of SR

“Native” BIER-TE example

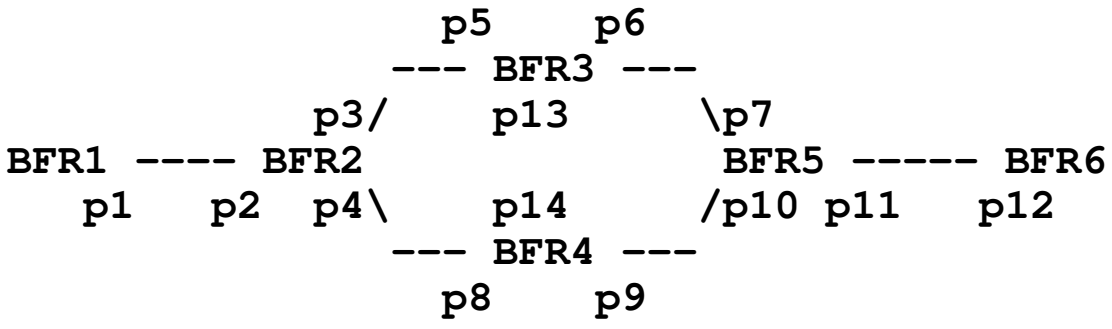

(simplified)

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
        p12 -> local_decap
```

“Overlay” BIER-TE example

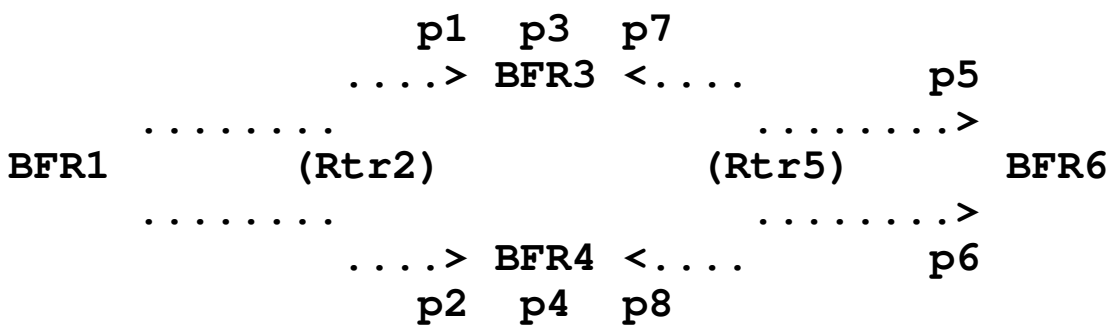

(simplified)
BIER-TE Bit Index Forwarding Tables (BIFT):

BFR1:	p1	-> forward_routed to BFR3
	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

Rtr2/Rtr5 do not speak BIER-TE

BIER-TE Topology in the architecture

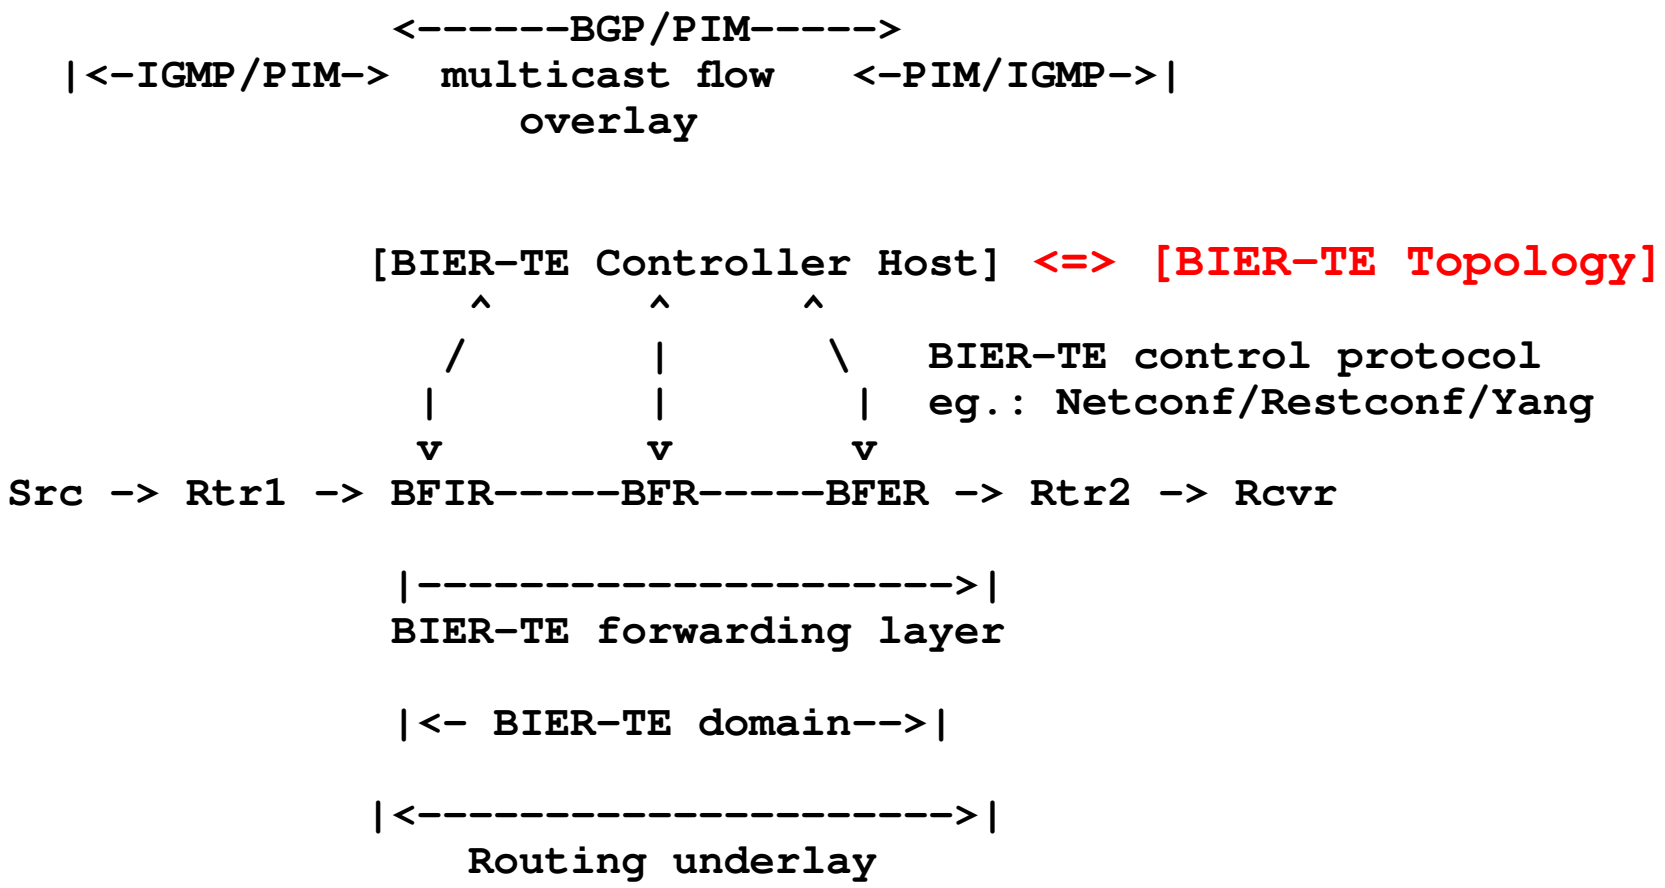

During bring-up or modifications of the network topology, **the controller discovers the network topology and creates the BIER-TE topology from it: determine which adjacencies are required/desired and assign BitPositions to them**. Then it signals the resulting of BitPositions and their adjacencies to each BFR to set up their BIER- TE BIFTs.

BIER-TE Topology in the architecture

- **BIER-TE BitPosition (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 ingres 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-ingres BFR**.
- **SR can naturally be combined with BIER-TE and help to optimize it**. For example, instead of defining BitPositions for non-replicating hops, it is equally possible to use segment routing encapsulations (eg: 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-by- hop 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.

THE END