BIER Tethering Update

Draft-ietf-bier-tether-01

Z. Zhang, N. Warnke, I. Wijnands, D. Awduche
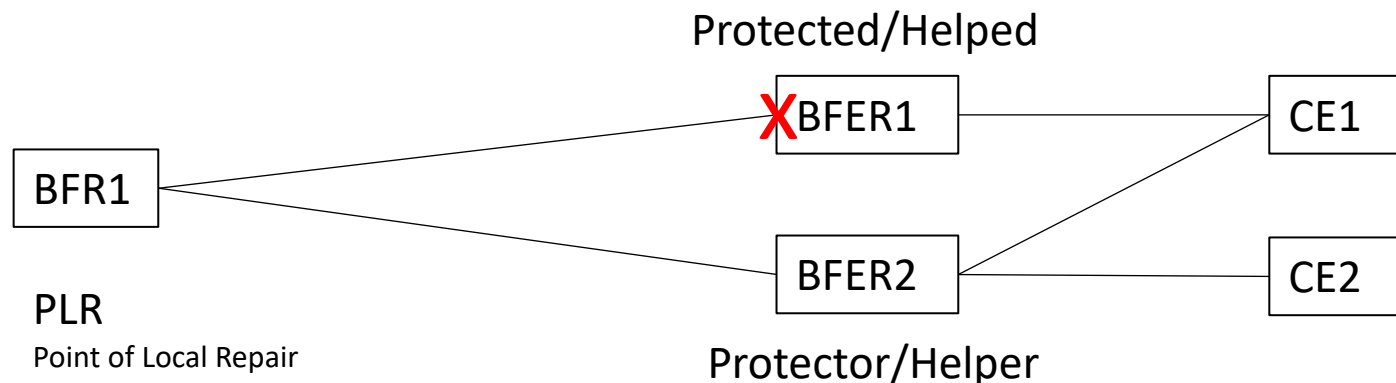
IETF 110

Overview of Tethering

- A non-BFR may have one or more designated helper "tethered to it"
  - Could be via a tunnel
  - A helper advertises that it is the helper of a helped node
- Update to Section 6.9 procedure of RFC 8296, for post-SPF handling of non-BFRs:
  - Instead of replacing a non-BFR on the tree with its BFR children, replacing it with its helper – if that does not introduce loop
  - Traffic tunneled to the helper, who then tunnel to downstream BFRs
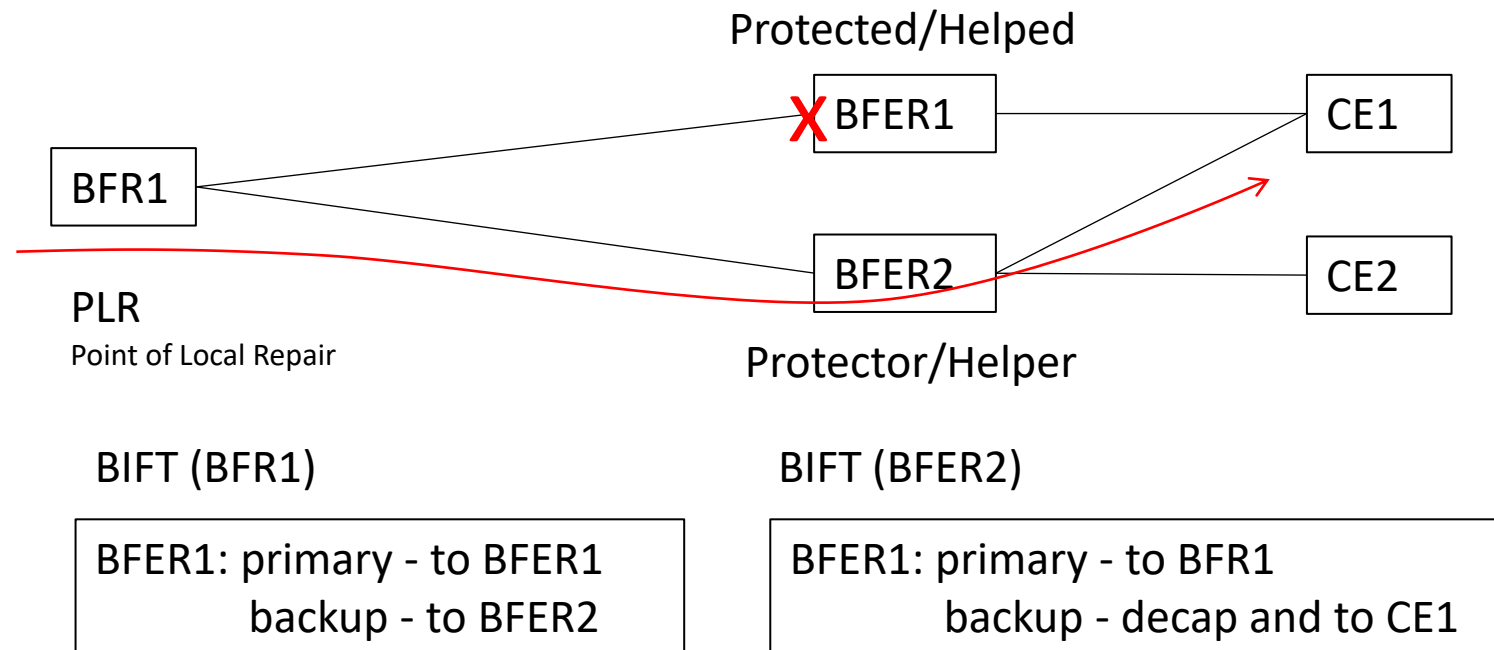- -01 revision introduces egress protection

Tethering Applied to Egress Protection

- For egress protection, a BFER may have a designated protector BFER

- A failed BFER (BFER1) can be considered as a non-BFR, its protector BFER (BFER2) considered as the helper

- No additional signaling besides existing tethering signaling
  - The fact that the helped node is capable of BIER indicates it is for egress protection

Protected/Helped

```
                                    X BFER1 ————————— CE1
                                                     ╲
BFR1 ————————————————————                      ╲
                          ╲                     ╲
                           BFER2 ———————————————————— CE2
PLR
Point of Local Repair                  Protector/Helper
```

Tethering Applied to Egress Protection

- On PLR, for the BIFT entry for a protected BFER, if the primary path is not through another BFR, its node-protection path is to tunnel to the protector/helper BFER

- On protector/helper BFER, for the BIFT entry for the protected BFER, the node-protection path is to decap and send to protected BFER's flow overlay receivers

Protected/Helped

```
┌────────┐              ┌──────┐
X│ BFER1  │──────────────│ CE1  │
 └────────┘              └──────┘
┌────────┐
│  BFR1  │
└────────┘
              ┌────────┐          ┌──────┐
              │ BFER2  │──────────│ CE2  │
              └────────┘          └──────┘
```

PLR
Point of Local Repair

Protector/Helper

BIFT (BFR1)

┌─────────────────────────────────┐
│ BFER1: primary - to BFER1   │
│   backup - to BFER2   │
└─────────────────────────────────┘

BIFT (BFER2)

┌─────────────────────────────────┐
│ BFER1: primary - to BFR1   │
│   backup - decap and to CE1 │
└─────────────────────────────────┘

A Few Notes

- Even if both CE1 and CE2 needs to receive traffic, if the FRR'ed packet only has BFER1 bit set, it is only sent to CE1 after decap
  - This is why PLR should not flip the BFER1-bit to BFER2-bit
- For the backup path of the BFER1 entry on BFER2, the F-BM only includes BFER1 itself
  - This is actually not a special rule – any decap branch's F-BM is like that
  - This ensures that if PLR sends one copy to BFER2 with BFER1+BFER2 bits set, one decap'ed copy is sent to CE2 and one decap'ed copy is sent to CE1
- The primary path for the BFER1 entry on BFER2 is through BFR1
  - This may cause loop if PLR detects BFER1 failure before BFER2 does
  - This can be avoided if BFER2 always uses the backup path for BFER1

Comments?