

BFR Tethering

IETF 103

Jeffrey Zhang, Juniper



Nils Warnke, DT

Ice Wijnands, Cisco

Brownfield Deployment

How to handle BIER incapable routers?

- Get around them and/or tunnel through them
 - X in the below diagram does not support BIER




Tunnel Through Incapable Routers

And follow unicast topology

- Section 6.9 of BIER Architecture spec
 - At the end of SPF, examine each immediate child node on the SPF tree
 - If it's not BIER capable, replace it with its immediate children
 - If a new child (that just replaced its parent) itself is incapable, replace it with its immediate children
 - The process is repeated until all immediate children are BIER capable
 - If a child is not directly connected, then a tunnel must be used to send BIER traffic to the child
- Any tunnel to the child can be used
 - Static or dynamic (e.g. LDP/SR/GRE)
 - ***No need to announce the tunnels***

Tunneling Alone May Not Be Good Enough



- If the number of BFRs that X connects to is large, then BFR1 needs to tunnel many copies through that BFR1-X link
 - Not good if the BFR1-X connection is long distance and/or BW constrained
- A solution is to tether a BFRx to X with a fat local pipe
 - BFR1 tunnels a single copy to BFRx, who in turn tunnels to other BFRs

Making Tethering Easier

- Tunnels would need to be announced in IGP to make tethering work
 - Otherwise the Section 6.9 method would never put BRFx onto its SPF tree
 - They would need to be announced a BIER specific MT or use FlexAlgo
 - Don't want to send unicast traffic through BFRx
- Unless with the following trick:
 - X advertises that it supports BIER so it will receive BIER packets natively
 - The BIER packets come with a BIER label that normally directs to BIER forwarding
 - X label switches (based on the BIER label) the BIER packets to BFRx
 - BFRx uses the Section 6.9 method to tunnel incoming BIER packets from X to other BFRs through X
 - BFRx knows that X does not really support BIER
 - Through provisioning or additional signaling from X
 - X could also have other connections as long as the X-BFRx metric is 1

Additional Signaling

- X signals “BFRx is my helper”
 - BFRx will know X is actually not capable
 - Other BFRs may use Section 6.9 method to tunnel over X
 - Instead of relying on X to use BIER labels to switch to BIER packets to its helper
- Alternatively, BFRx signals “I am X’s helper”

Summary

- Tethering a BFR helper reduces the copies of tunneled packets over
- Two options to make tethering easier
 - X pretends it supports BIER but label switches incoming BIER packets to its helper BFRx
 - Requires software upgrade on X and its helper only
 - Only works for MPLS
 - BFRx advertises it is X's helper and others will tunnel over X
 - Upgrade needed on BFRs but not on X
- Next Steps
 - Seek comments
 - Request adoption after polishing