

# BIER Brownfield Deployment Options

draft-zhang-bier-brownfield-options-00  
draft-przygienda-bier-migration-options-00

Jeffrey Zhang, Tony Przygienda – Juniper

Sandy Zhang – ZTE

IETF120, Vancouver

# Brownfield Deployment

- Mixed BFRs and BIER-incapable routers (non-BFRs) in a BIER domain
- Informational Draft on brownfield deployment options
  - “Naked” MT
  - Section 6.9 of RFC8279
  - Skip non-BFRs in BIRT/BIFT calculation
  - Controller-programmed BIFTs
  - BIER Tethering
  - BIER PHP

# “Naked” MT

- A separate topology that only includes BIER-capable routers and connections between them
  - The connections can be native links or tunnels that can be viewed as IGP adjacencies
- No protocol extensions needed

# Section 6.9 of RFC8279

- At the end of SPF calculation, examine all the children of the root node of SPF tree
  - A child could be a reparented one as in the next bullet
- If a child is BIER-incapable, remove it from the tree, and reparent its children to the root (the calculating router) directly
- Eventually all children of the root are BFRs
  - Some are reparented ones and they must be reached via tunnels
  - Native forwarding or tunneling to the children
- The concept of “tunnel to downstream BFRs” can be achieved with BGP signaling as well
  - With the “BIER Nexthop” in the BIER attribute that is only changed by BFRs

# Skip non-BFRs in BIER-specific SPF Calculations

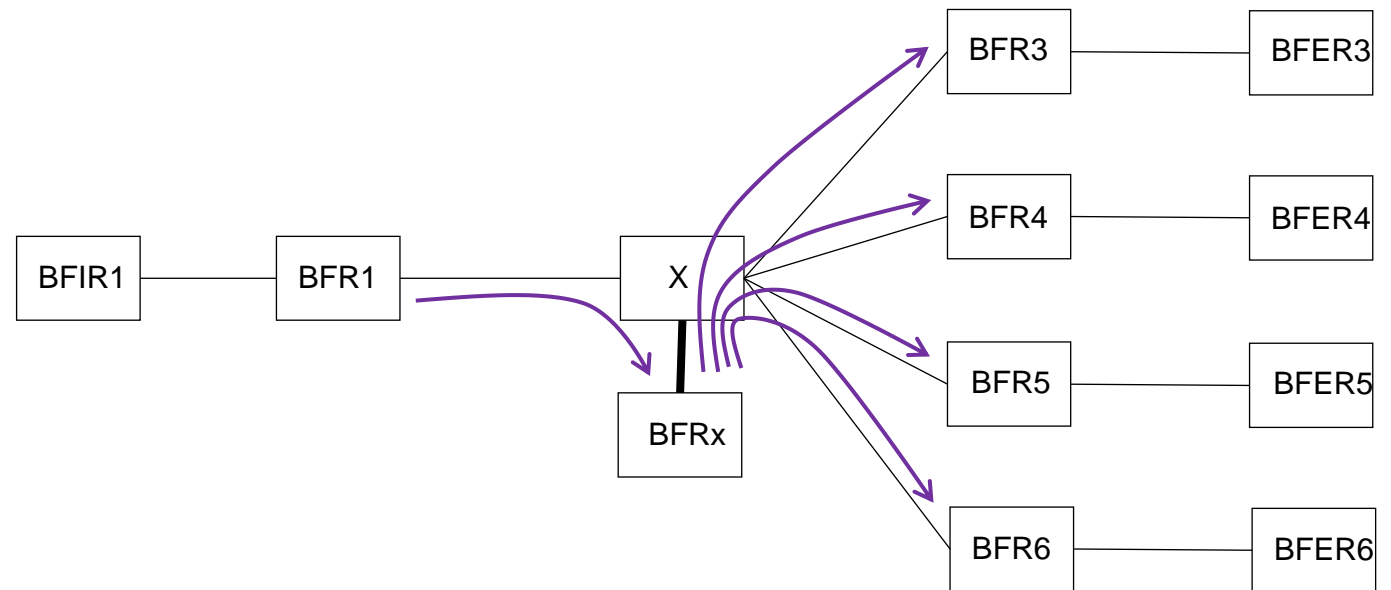

- Separate SPF for BIRT/BIFT calculations
- Signal a BAR value that indicates a BIER-specific calculation algorithm that skip non-BFRs
  - During the calculation the non-BFRs will not be considered
  - There must be a complete topology that connects all BFRs
    - With native links or tunnels as IGP adjacencies

# Controller-programmed BIRTs/BIFTs

- An omniscient/omnipotent controller figures out everything and programs the BIRTs/BIFTs
  - Forwarding over native links or tunnels between BFRs

# Tethering

- All previous options may lead to ingress-replication-like tunneling
- Tethering avoids that
  - It's like turning a non-BFR into a BFR by tethering a helper to it
  - W/ very simple IGP signaling
  - Works with BGP as well



# BIER PHP

- All previous options assume that the flow overlay edge routers support BIER
  - E.g., MVPN PEs are BFIRs/BFERs and capable of BIER forwarding
- BFERs may request its upstream to do PHP so that they do not need to handle BIER forwarding
  - Complete BIER control plane plus PHP request
  - No BIER data plane
- Ingress PEs need to be real BFIRs



# Summary

- Brownfield deployment options
  - “Naked” MT
  - Section 6.9 of RFC8279 and its BGP equivalent
  - Skip non-BFRs in BIRT/BIFT calculation
  - Controller-programmed BIRTs/BIFTs
  - BIER Tethering
  - BIER PHP
- The most practical solutions may be:
  - Section 6.9 of RFC8279 and its BGP equivalent
  - Tethering (IGP & BGP)
  - PHP

# Next Steps

- Discussions and comments
- WG adoption?