

BGP Extensions for BIER

draft-ietf-bier-idr-extensions-08

Xiaohu Xu, Mach Chen, Keyur Patel,
Ice Wijnands, Tony Przygienda, Jeffrey Zhang
IETF 115, London

History

- draft-ietf-bier-idr-extensions had passed WGLC
 - Jeffrey did shepherd review
- The authors had lost steam
- It was only intended for DC scenario (with BGP underlay)
 - If the path to a BFER has a non-BFR, then ingress replication tunneling is done all the way to BFER
 - draft-ietf-bier-multicast-as-a-service has a solution for non-BFRs
 - A general solution cover DC and non-DC scenarios



- With the pre-08 approach:
 - All BFRs/BFERs only advertise their own BIER prefix and encapsulation information
 - When BFR1 calculates BIFT entries for BFER1/2/3, the Nexthop is the non-BFR, so it has to tunnel BIER traffic to BFER1/2/3 directly
 - Not just to BFR2 – since BFR1 does not know BFR2 is in the path
- With the -08 approach:
 - All BFRs will update BIER encapsulation and Nexthop in BIER TLV when re-advertising BIER prefixes
 - This also allows better packing since they all have the same BIER attribute
 - BFRs do not need to advertise their own BIER prefixes
 - When BFR1 receives the BFER1/2/3 prefixes, they all have BFR2's encapsulation information and BFR2 as nexthop
 - As a result, BFR1 only need to tunnel to BFR2 – one copy to reach all BFERs beyond BFR2

Procedures

- BFERs advertise BIER prefixes with their own BIER information
 - Encapsulation sub-TLV , which includes BIFT-ID ranges, etc.
 - Both MPLS and non-MPLS encapsulations are supported
 - Optionally with a BIER Nexthop sub-TLV in the BIER TLV
- When a BFR re-advertises a BIER prefix
 - It replaces/adds the BIER Nexthop with its own address
 - It also replace encapsulation info (e.g., BIFT-ID ranges) with its own
- When a non-BFR re-advertise a BIER prefix, nothing is changed
- When a BFR calculates the BIRT/BIFT entry for a BFER
 - The BIER Nexthop is the BFR neighbor to send/tunnel traffic to
 - Using the encapsulation info in the BIER TLV

Proxy Range sub-TLV

- A BFR/BFER may include a Proxy Range sub-TLV in the BIER TLV of its own BIER prefix
 - To advertise BFR-IDs that are reachable via itself
 - draft-ietf-bier-prefix-redistribute
- In the previous example, BFR2 MAY include BFER1/2/3's BFR-IDs in the Proxy Range sub-TLV in the BIER TLV of its own BIER prefix
 - No need to re-advertise BFER1/2/3 anymore
 - BFR1 can still calculate BIFT entries for BFER1/2/3

Summary of -08 revision

- Extended scope for general applicability
 - Non-MPLS encapsulation as well as MPLS
 - General/brownfield deployment vs. BGP-DC only
 - Handles non-BFR – only tunnel to the next BFR (vs. all the way to BFERs)
 - Moved the solution in draft-ietf-bier-multicast-as-a-service here and simplified it
- Updated signaling
 - Transit BFRs do not need to advertise BIER info with their own BIER prefix
 - Added Proxy Range sub-TLV support
- Added section about BIFT calculation
- Updated text for handling at BIER domain boundary

Next Steps

- Ready for another round of WGLC
- New shepherd needed
 - Jeffrey joined as editor
 - Sandy joined as contributor