

BIER Anycast and Overlay Considerations

[draft-zzhang-bier-anycast](#)


[draft-zzhang-bier-egress-protection-overlay](#)

Zhaohui (Jeffrey) Zhang, IJsbrand Wijnands, Zheng Zhang, Mankamana Prasad Mishra, Huaimo Chen

IETF 116, Yokohama

Anycast for Service Egress Protection

- It's well-accepted practice to use anycast addresses for multi-homing PEs for service egress protection (RFC8679) - at least for unicast
- BIER (RFC8279) did not consider anycast



BIER Can Easily Support Anycast

- A non-zero BFR-ID can be advertised from different BFERs with the same anycast BFR-prefix
 - The rest works out automagically – it's just like ECMP to a single BFER
- Different BFR-prefixes must have different BFR-IDs (if none-zero)
- A BFER may have multiple BFR-prefixes (anycast or not)
 - Each with its own BFR-IDs
 - This is to have different multi-homing protection groups
 - Traffic must be sent appropriate BFR-ID in the header
 - Number of BFR-prefixes for a BFER should be minimized
 - To reduce the number of bits needed

Backward Compatibility?


- A “legacy” BFR might treat same BFR-IDs advertised from different BFERs’ anycast BFR-prefix as duplicate
 - The BFR-ID will be treated as invalid and relevant BFERs won’t receive traffic
 - An error message will be logged
- To make use of anycast based egress protection, all BFRs must support it
 - The error log message from a “legacy” BFR will help identify the situation
- Signaling a BFR’s support of anycast?
 - It does not help
 - It is not needed anyway

MVPN Flow Overlay

- The CE connecting to the PE/BFERs must signal multicast interest to all multi-homing PEs of the same anycast protection groups
 - PIM/mLDP MoFRR
 - IGMP/MLD proxy with multiple upstream interface
- All those PEs will signal C-Multicast routes and Leaf-AD routes
 - Anycast BFR-prefix and corresponding BFR-ID used in the Leaf AD route's PTA
- The CE must be prepared to receive traffic from all of them
 - Only one PE/BFER will receive traffic and deliver to the CE
 - This happens naturally – no special handling is needed
- This does not work for EVPN

MVPN/EVPN Tunnel Segmentation

- Segmentation points can be anycast BFERs
 - BFERs in the upstream sub-domain
 - BFIRs in the downstream sub-domain
 - PMSI routes are sent with anycast BFR-prefix/ID
- PEs (PE4/5) and segmentation points (SEG3/4) in the downstream region:
 - are BFERs in the downstream sub-domain
 - Leaf AD routes are sent to upstream PE or segmentation points (SEG1/2) for flow overlay signaling w/ anycast BFR-prefix/ID in PTA
 - are like CEs to the anycast PEs wrt the upstream sub-domain
 - Leaf AD routes are like the PIM joins in the PE-CE case
 - They trigger Leaf AD routes for the upstream sub-domain
- This does work with EVPN
 - If the segmentation points are not also MH PEs



Other Tunnel Types?

- Anycast segmentation points work with Ingress Replication the same way as with BIER
- Other tunnel types could be made work with additional procedures
 - Some details in the draft

Two Drafts

- [draft-zhang-bier-anycast](#) - extension to base spec
 - Explicitly allow anycast BFERs
 - Explicitly tighten the duplication detection rule
 - Standards track
- [draft-zhang-bier-egress-protection-overlay](#)
 - Describe how overlay signaling works, including tunnel segmentation case
 - Clarify the use of anycast BFR-prefix/ID in PMSI/Leaf AD routes
 - Considerations for PIM/mLDP/MLD overlay may be added
 - Standards/Informational