

BIER NVO3 Analysis

draft-zhangwu-bier-nvo3-analysis-00

BIER WG
IETF99# Prague

Zheng Zhang
Bo Wu

NVO3 multicast background

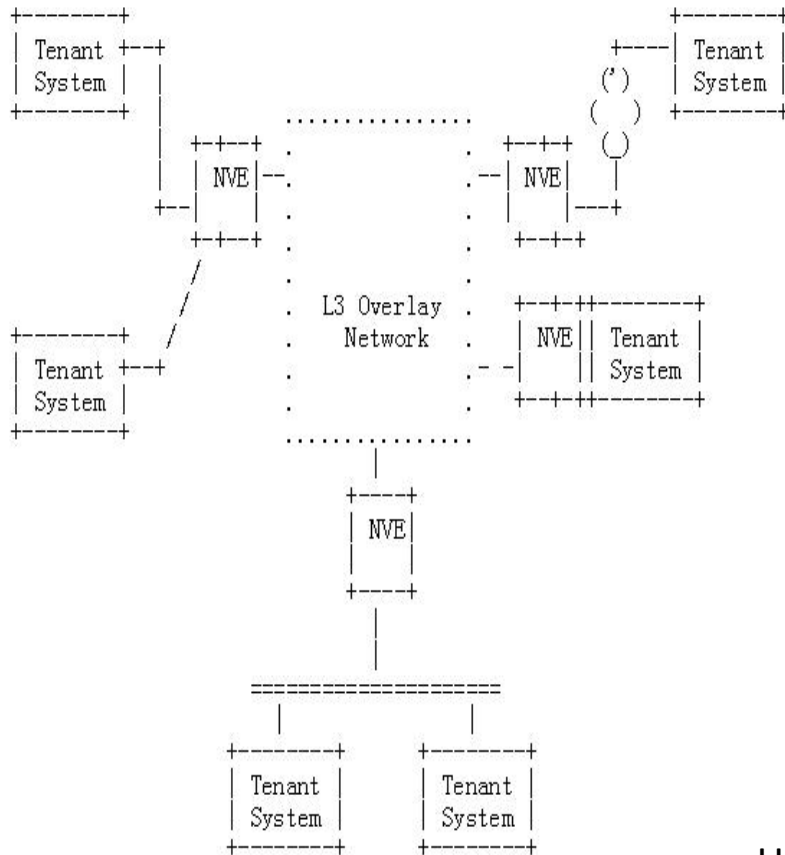


Figure 1: NVO3 Generic Reference Model

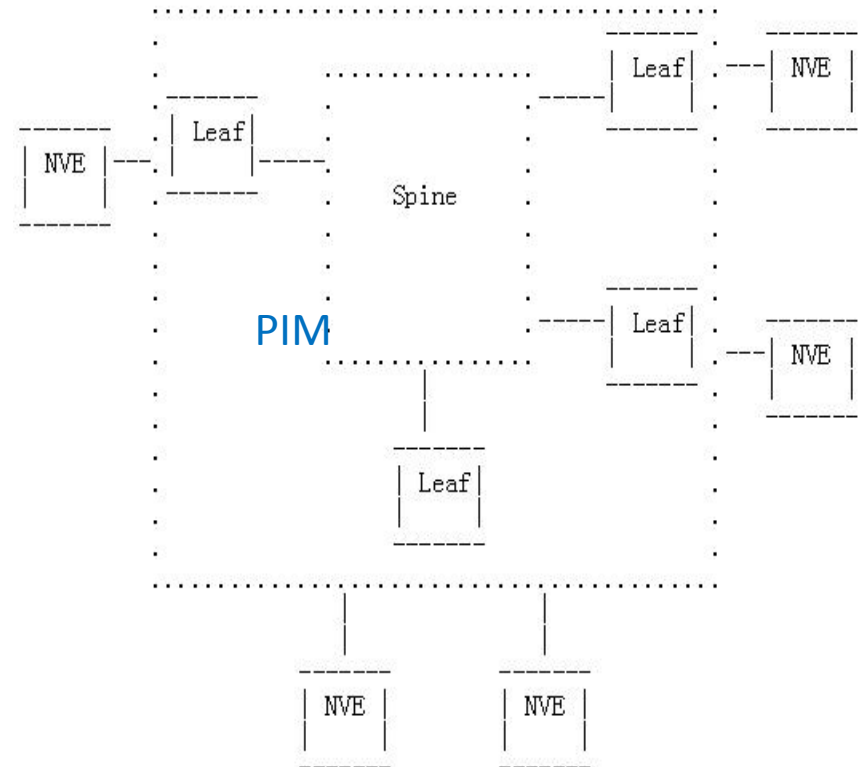


Figure 2: Example topology of NVO3 underlay network

- There is multicast requirement in NVO3 network. PIM and BIER can be used as IP multicast underlay technology.

Using PIM as IP multicast underlay:

- PIM needs multicast tree building in leaf and spine devices.
- Every leaf and spine devices along the tree must maintain multicast state for every multicast flow.
- In case one of the source leaf or receiver leaf changed, the multicast states along the tree needs to be changed.

BIER technology

BIER multicast technology:

- BIER is another choice for the forwarding of multicast data packets. It does not require a protocol for explicitly building multicast distribution trees, nor does it require intermediate devices to maintain any per-flow state.
- So use BIER as multicast underlay will eliminate the per-flow state and the explicit tree-building protocols, decrease the burden of intermediate devices.

BIER requirements:

- Every edge node (BFIR, BFER) need to be assigned a unique BFR-ID.
- Edge nodes should exchange BFR-ID and BFR-Prefix by IGP/BGP extension.
- BIER Ingress node for a multicast flow encapsulates the multicast packet in a BIER header which destination is the BFR-ID set of BIER egress nodes.

Two possible ways to handle multicast packets are identified for using BIER in NVO3 network:

- Placing BIER boundaries at leaf switches.
- Putting the BIER boundary at NVE to achieve better efficiency.

Leaf devices as the edge BIER nodes

- NVE needs to support IGMP/MLD protocol.
- The BUM flows are encapsulated corresponding multicast group address.
- BMLD protocol is used among all the leaf devices to exchange multicast group information.
- Leaf device should be allocated with one unique BFR-ID.
- Leaf device encapsulates the packet with according BIER header.

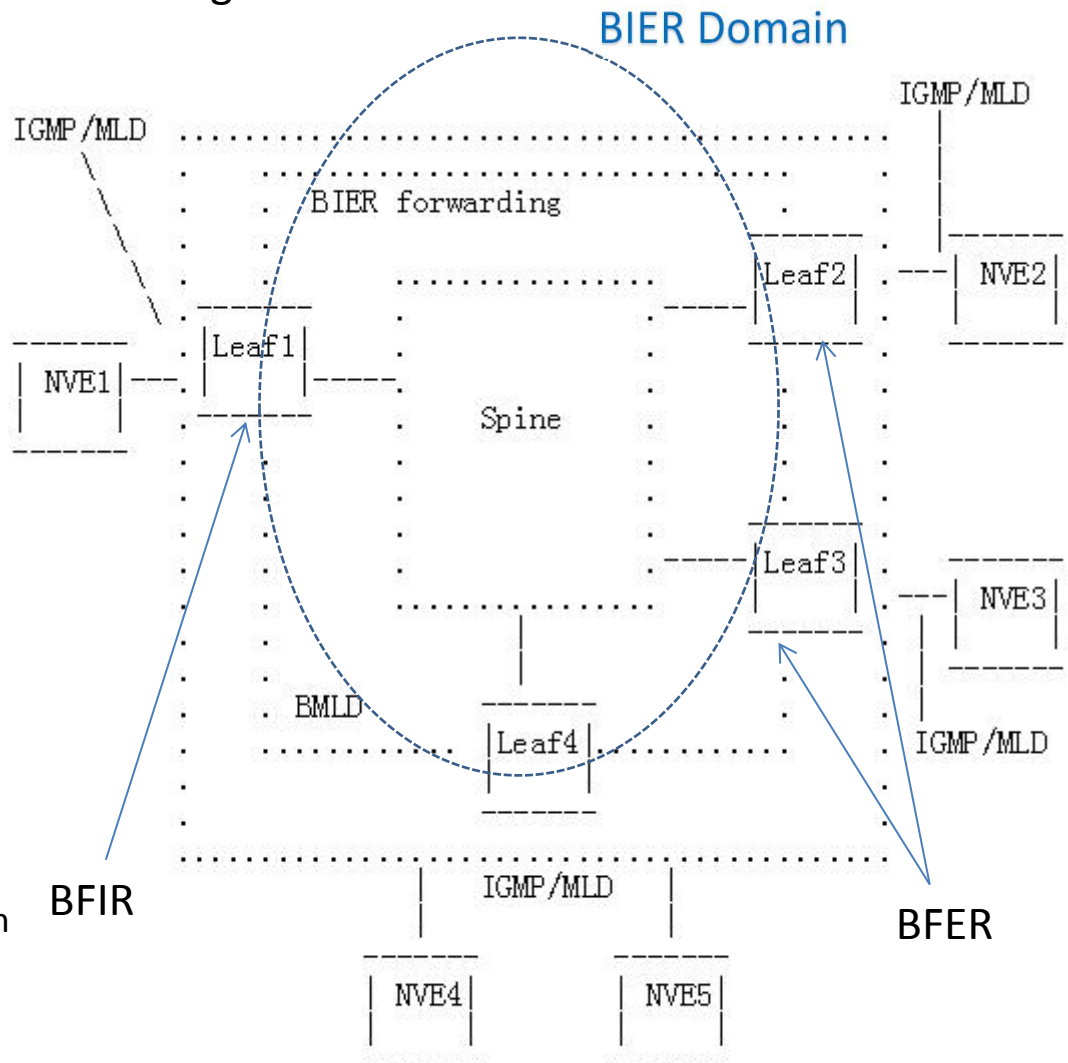


Figure 3: BIER in VN03

Consideration:

- IGMP/MLD is still needed to run between NVE and leaf devices.
- The multicast groups for Tenants and specific multicast applications are needed.
- BitStringLength limitation.

NVE as edge BIER nodes

- Every NVE should be allocated with one unique BFR-id.
- IGMP/MLD protocol does not need to run between NVE and leaf device. NVE encapsulates the BIER header with the BFR-ids of destination NVEs straightly and send it to leaf devices.
- Leaf and spine devices build BIER forwarding plane normally.
- NVE gets the destination NVEs' BFR-id from SDN controller or other method.

Consideration:

- If NVE does not support IGP or BGP; NVE can not take part in the BIER information exchange. So the BFR-id of NVE should be advertised by some other method.
- BitStringLength limitation.

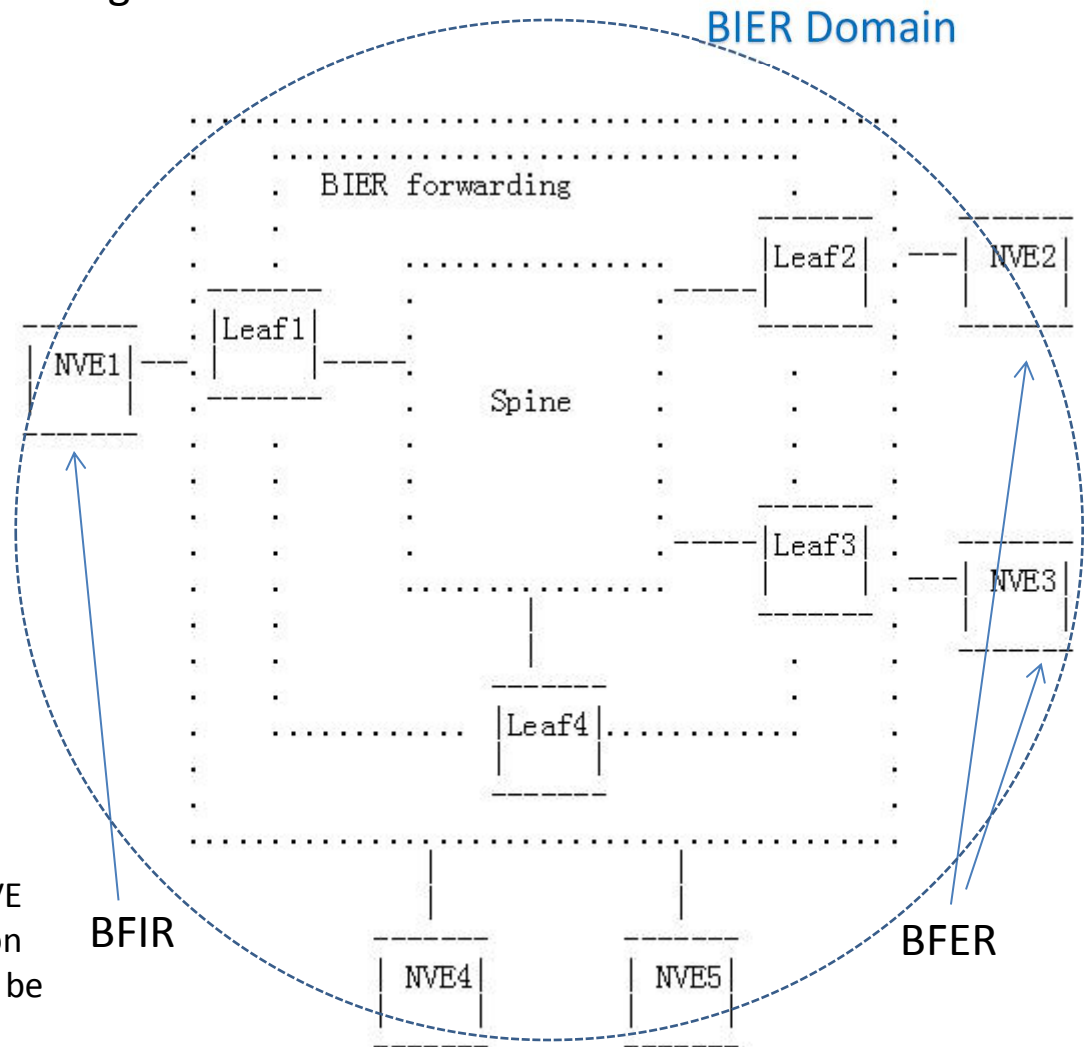


Figure 4: BIER in VN03

- Any comments are welcome 😊

Thanks!