



BIER Use Case in VXLAN

[draft-wang-bier-vxlan-use-case-00](#)

Linda Wang (Presenting)
Sandy. Zhang & F. Hu



Background


- Motivated by draft-ietf-bier-use-cases, in section 3.6, it proposes there are many advantages introducing BIER in Data Center Virtualization.
- proposes solutions how to extend protocol to take advantage of BIER in Data Center Virtualization.
 - Until now, received several precious comments.

The Problem

- In data center virtualization, there are two most common solutions to forward BUM traffic on the overlay network:
 - Ingress Replication: However, it doesn't provide the optimal forwarding of multicast packets.
 - Enable the multicast capability in the underlay, such as PIM-SM or PIM-BIDIR and so on. However, it requires data centers to run multicast protocol and maintain multicast state in all the edge nodes and intermediate nodes.
- How to eliminate above issues?
- Introducing BIER in Data Center Virtualization
 - Providing the optimal forwarding of multicast packets
 - Doesn't require the data centers to run multicast protocol and maintain in multicast state in the intermediate nodes

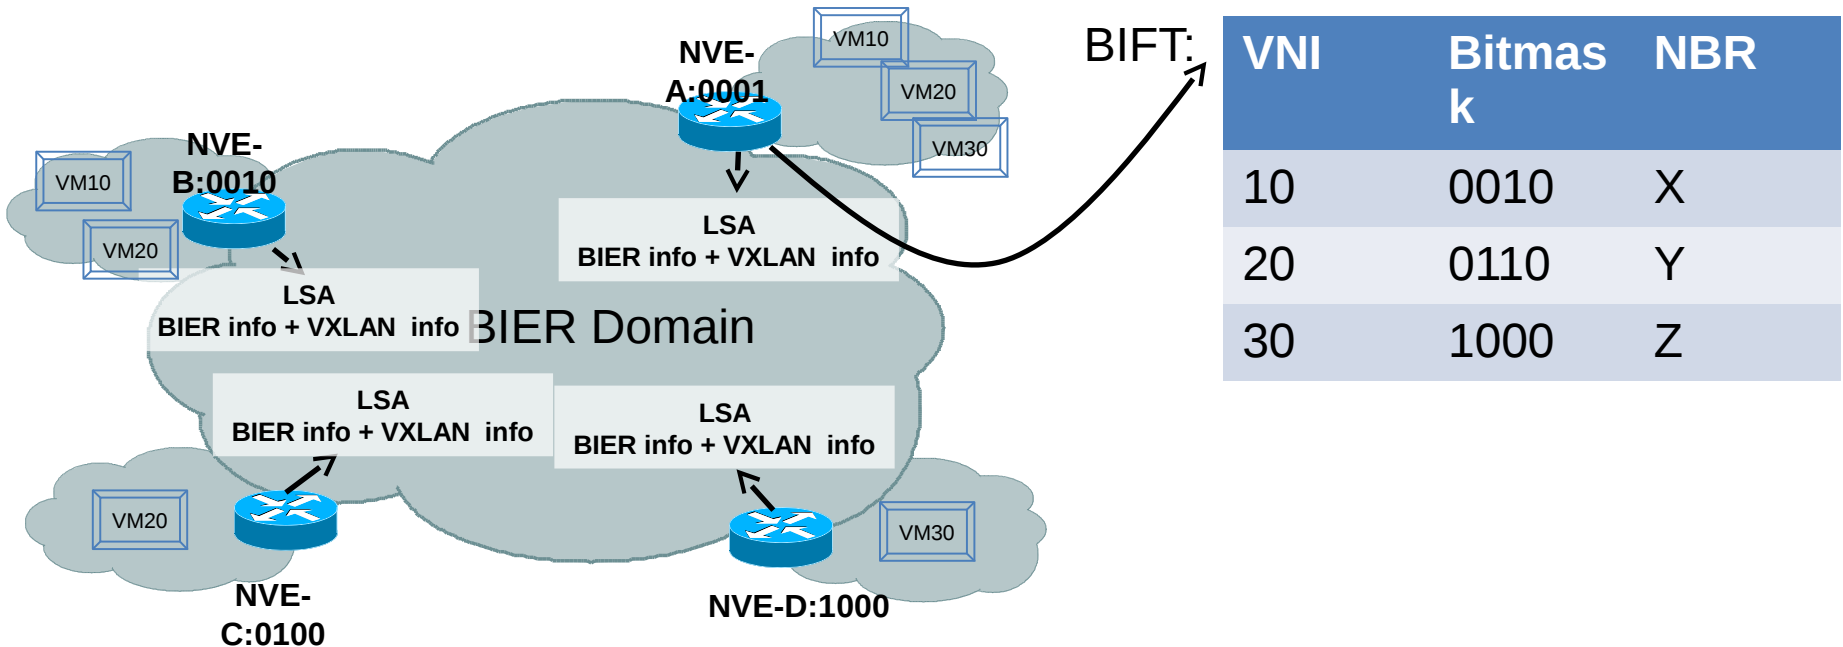
The Solution for Using BIER in DC

- The main idea:
 - Each NVE advertises VXLAN information together with BIER information to other NVEs through IGP/BGP



- Each NVE creates a mapping between the VXLAN Network Identifier and the Bitstring representing remote NVEs in the same virtual network
- Then, forwarding the BUM traffic according to the mapping Bitstring in the BIER Header along the BIER multicast forwarding path.

BIER Data Center Use Case



1. Assign a unique Bit Position to each NVE in the data center.
2. Each NVE floods BIER information + VXLAN information.
3. Each NVE creates the mapping between VXLAN Network Identifier and the Bitstring representing the remote NVEs sending the same VNI information.
4. BUM packets are received from virtual network tenant, NVE lookups the mapping and encapsulates the VXLAN information as well as the BIER information.
5. BUM packets are forwarded and replicated hop-by-hop according to the Bitstring in the BIER header using the Bit Index Forwarding Table.


Precious Comment 1

- Comment From Andrew Qu:


- As VxLAN/NvGRE/.. all data planes are using similar technology to support L2 over L3, how about using one draft to take care of all similar data planes.

Suggesting change the specific VxLAN network Identifier field into two sections:

- 1) protocol [vxLAN/vxlanGPE/nvGRE etc]
- 2) BUM domain like VNI/TSI, but may be more general term can be used.



VXLAN Info sub-sub-TLV



NVO3 Info sub-sub-TLV

Accepted. I will update this draft to cover differentiated NVO3 technologies, rather than just VXLAN.

Precious Comment 2

- Comment From Tony:
 - VXLAN-IGP extension will melt down the IGP flooding, what about using MLD/IGMP extension described in draft-pfister-bier-mld to carry the VXLAN information?

Accepted. I will update this draft to cover MLD/IGMP extension for BIER-VXLAN technologies, rather than using IGP extension.



More comments are welcome!

Thank you!

