



BIER Use Case in NVO3

[draft-wang-bier-vxlan-use-case-01](#)

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

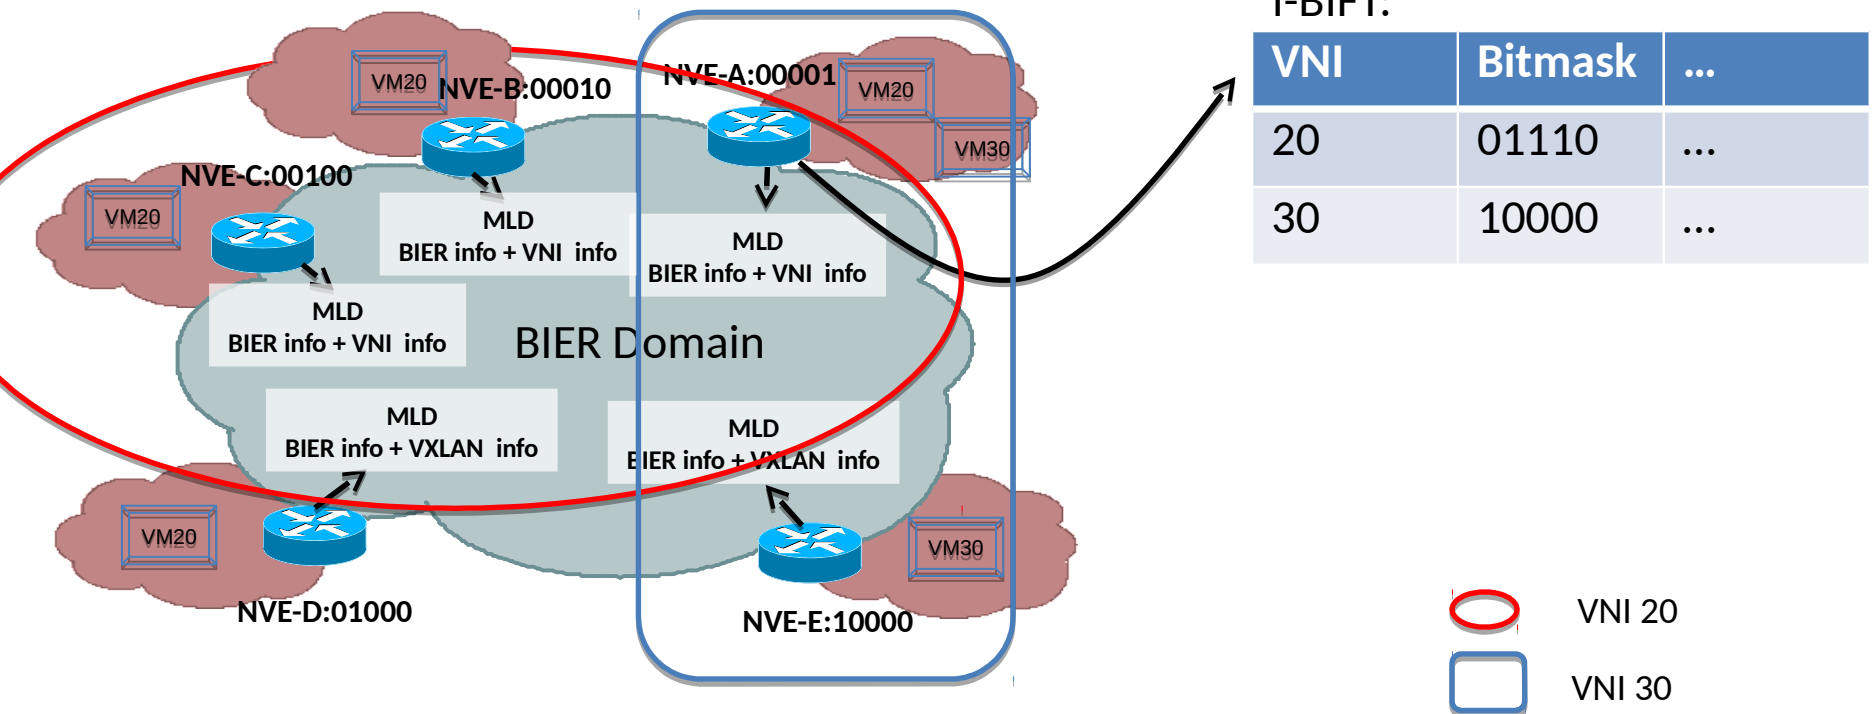
- In data center virtualization, there are two most common solutions to forward BUM traffic on the overlay network:
 - Ingress Replication: However, it doesn't provide the optimal forwarding of multicast packets.
 - Enable the multicast capability in the underlay, such as PIM-SM or PIM-BIDIR and so on. However, it requires data centers to run multicast protocol and maintain multicast state in all the edge nodes and intermediate nodes.
- Proposed EVPN solution
 - Mapping the VXLAN VNI or NVGRE VSID to an EVI
 - Advertising inclusive Multicast Ethernet Tag routes to carry the VNI/VSID field through EVPN control plane
 - Advertising BGP Encapsulation extended community to carry the specific data plane encapsulation for the BUM traffic, such as VXLAN/NVGRE/MPLS/MPLS in GRE/VXLAN GPE and so on.
 - However, it requires NVE/VTEP to run EVPN

The Solution

- Introducing BIER at the transport layer in Data Center Virtualization
- Furthermore, introducing MLD as control plane to construct the overlay mapping.
 - Providing the optimal forwarding of multicast packets
 - Doesn't require the data centers to run multicast protocol and maintain multicast state in the intermediate nodes
 - Doesn't require the data centers to run EVPN


The Sketch of this Solution

- ✓ Specifically, using MLD as an overlay control plane to collect the VTEP/NVE belonging to the same VNI/VSID (like constructing I-PMSI tunnel), through carrying VNI/VSID information in the MLD extension.



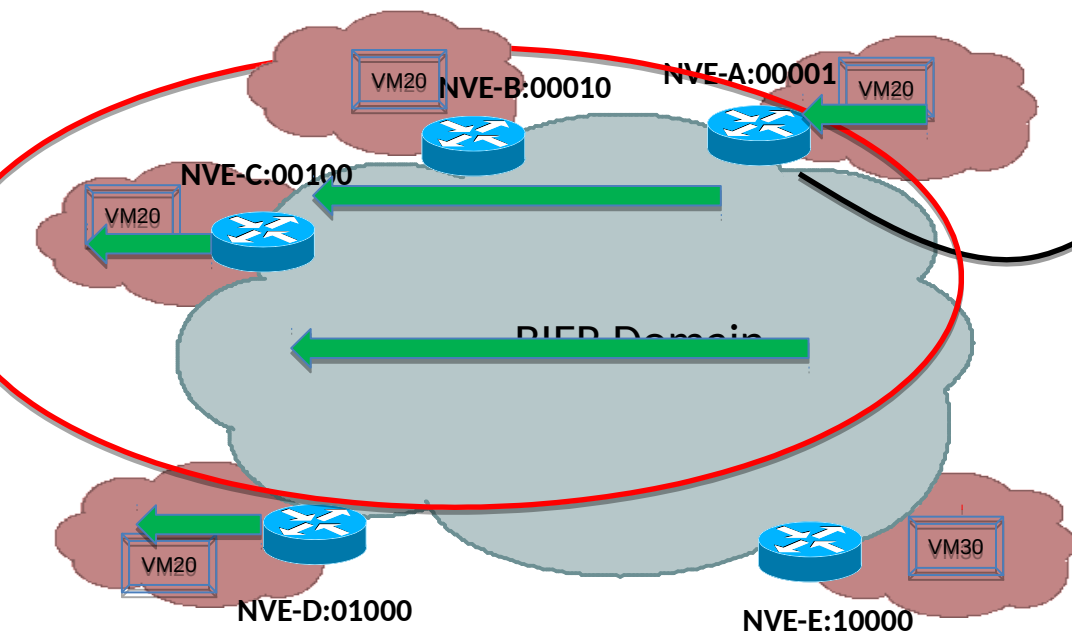
The Sketch of this Solution

- ✓ Then, using MLD as an overlay control plane as well to collect the VTEP/NVE attaching to the same tenant who join the same multicast group (like constructing S-PMSI tunnel), through carrying $\langle S,G \rangle$ information in the MLD extension.



The Sketch of this Solution

- ✓ Using the $\langle S,G \rangle$ information in the multicast traffic to lookup S-BIFT, encapsulating the multicast traffic with the Bitmask in the table, and forwarding the multicast traffic using BIER forwarding.





S-BIFT:

$\langle S,G \rangle$	VNI	Bitmask	...
$\langle S,G \rangle$	20	00110	...

MLD extension

- The main idea:
 - Each NVE advertises Virtual Network Identifier information with BIER information to other NVEs through MLD, including VXLAN network identifier, VSID and so on.



Virtual Network Identifier Info sub-sub-TLV



More comments are welcome!

Thank you!

