Data Plane Handling in NVO3

Zu Qiang

draft-zu-nvo3-user-plane-signalling-01.txt

DATA FORWARDING

- TS data forwarding is based on the inner-outer address table in NVE configured by the NVA
- No additional handling is required in NVE except one security consideration:
  - NVA query procedure may be triggered if a packet with an unknown destination address is received
  - Potential DOS attack if the query procedure is triggered too often

L2CP

- For a L2 NVE, the VAP is an emulation of a physical Ethernet port. It shall have the capability to handle any L2CP. More L2CP protocols will be added in this list.
- STP/RSTP/MSTP
  - In NVO3 network, the NVE does not need to propagate any STP messages to the remote NVEs. But, the NVE may need to learn the Root Bridge MAC address and Bridge Priority of the root of the Internal Spanning Tree (IST) of the attached layer 2 segment by listening to the BPDUs.
- LACP
  - A L2 NVE does not have to be involved in the Link Aggregation procedure. It only needs to encapsulate and forward any Link Aggregation Control Protocol Data and data packets between the participated TSs.

ARP AND ND

- There is no need for NVE to support any specific handling of ARP and ND.
- For optimization (avoiding mcast), the NVE does not have to forward the ARP/ND, if it can response to it based on the inner-outer address mapping table. However, the followings shall be considered:
  - The NVE shall have some kind ARP and/or ND filtering functions to avoid any security issues when query NVA for unknown addresses.
  - At multi-homing, one NVE per network segment of a TS for ARP / ND proxy of given destination IP addresses to avoid the racing issue
  - At VM mobility, the same MAC address may be used at the NVE for a given TS. However, at Multi-Homing scenarios, duplicated MAC address shall be avoided

DYNAMIC ROUTING

- It is possible to have user plane router installed in a TS, e.g. vR function enabled in a VM.
- The centralized NVO3 GW function and the user plane router appears as router adjacencies to each other. At user plane routing rules updates (by the vR), the NVO3 Centralized GW updates its routing distribution polices and forward TS data traffic accordingly.
- In this case, L3 is handled by the GW function. Only L2 NVE functions are needed in NVO3.

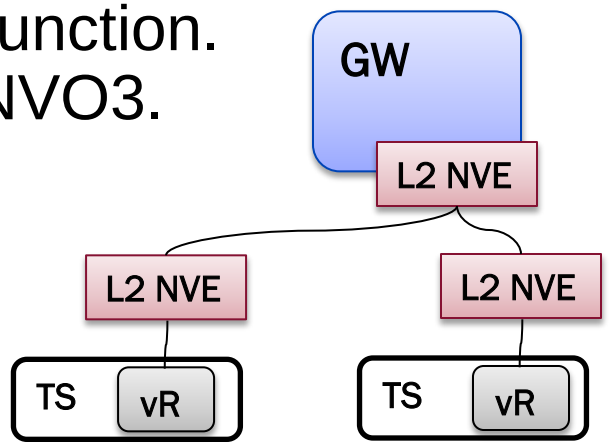

DYNAMIC ROUTING

- With Distributed GW function embedded in the L3 NVE, how to handle the user plane routing signaling is a question.
  - Multiple routes may be available for a destination address. And the related routing policies may be updated by the vR
  - Routing policies to the subnets behind the vR may be updated by the vR, e.g. install, update, remove a route using routing protocol.
- When L3NVE with embedded dGW receives an routing update from the vR, the inner-outer address mapping table needs to be updated, not only at the attached NVE, but also the remote NVEs

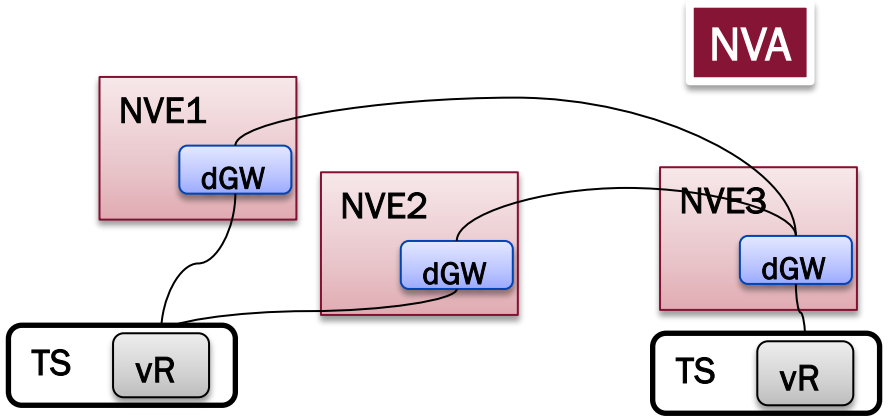

DYNAMIC ROUTING

- Alternatives:
  - A. Disallow dynamic routing at user plane.
  - B. Using NVE-NVE interaction messages to update the peer L3 NVEs.
  - C. Using the NVA-NVE signaling to update the peer L3 NVEs.
  - D. Collocated NVA and GW function.

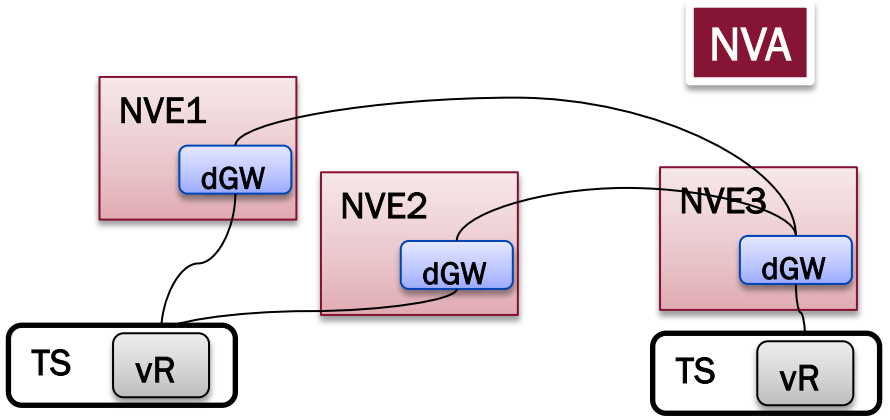

NEXT STEP

- Comments?

- Thanks