NVO3 Split-NVE CP Protocol Requirements and VDP extensions

Yizhou Li
Lucy Yong
Lawrence Kreeger
Thomas Narten
David Black

Pat Thaler
Joe Pelissier
Eric Gray

Split-NVE structure

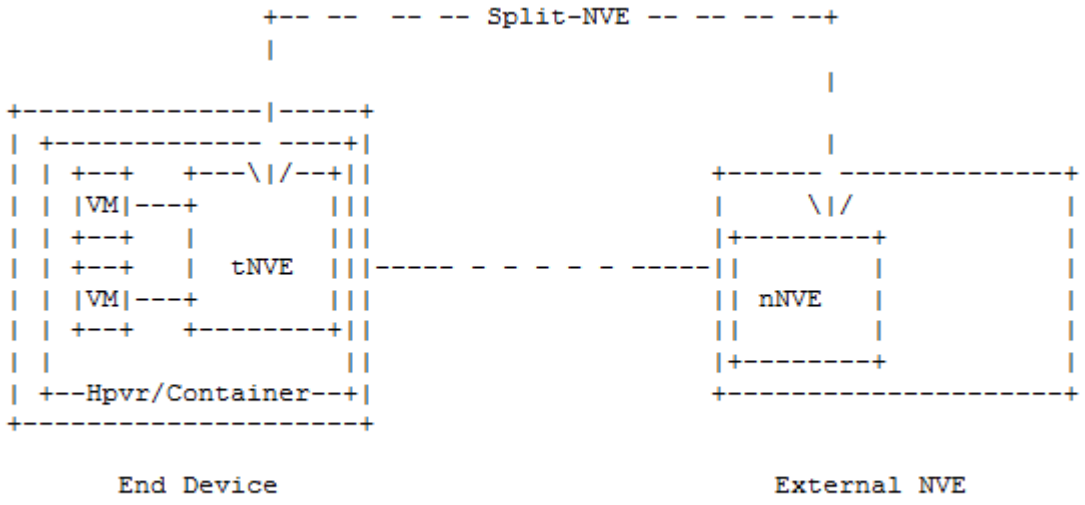

Figure 1 Split-NVE structure

Split-NVE: a type of NVE that the functionalities of it are split across an end device supporting virtualization and an external network device.
tNVE: the portion of Split-NVE functionalities located on the end device supporting virtualization.
nNVE: the portion of Split-NVE functionalities located on the network device which is directly or indirectly connects to the end device holding the corresponding tNVE.
External NVE: the physical network device holding nNVE

State Transition of a VAP Instance on an External NVE

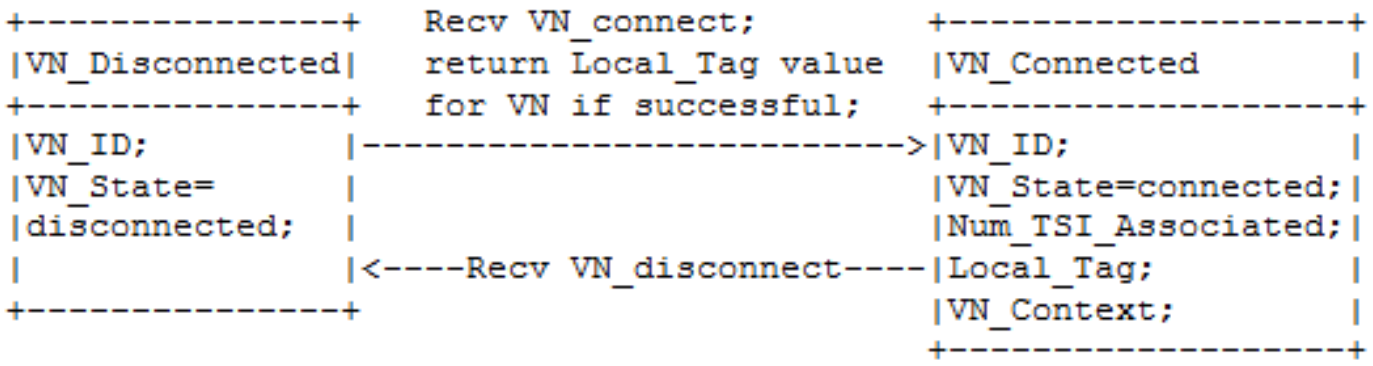

State Transition of a TSI Instance on an External NVE

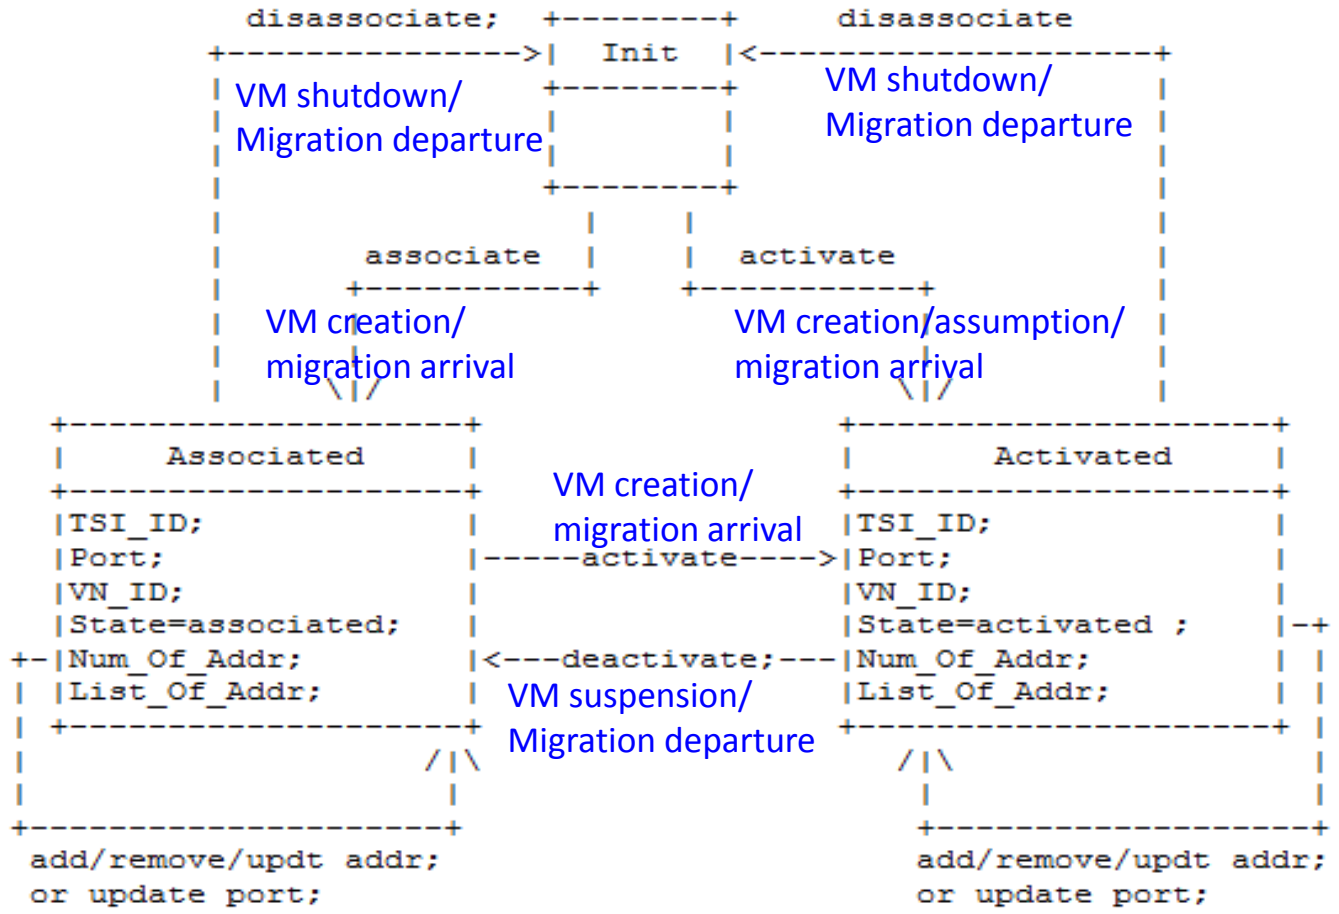

TSI: Tenant System Identifier Interface to a Virtual Network as presented to a Tenant System. To the Tenant System, the TSI is like a NIC.

VDP Background

- VDP (VSI discovery and configuration protocol) carries the virtual machine networking state between station and bridge in EVB structure defined by IEEE802.1Qbg.
- It is a candidate control plane protocol to be used between hypervisor and the external NVE in NVO3 context. draft-pt-nvo3-vdp-vm2nve-gap-analysis-00 describes the gap between current VDP and the earlier version of requirement doc.

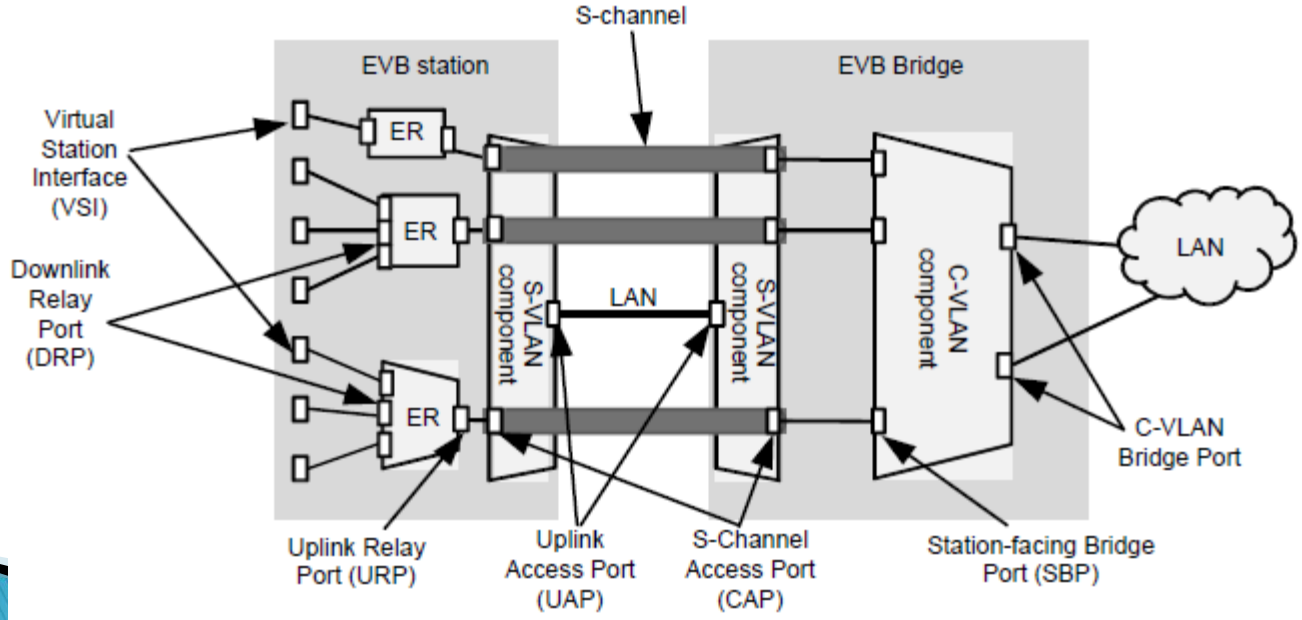

Requirements – 1

#	req	VDP Extension required?
1	support a bridged network connecting End Devices to External NVE.	Yes. Dest MAC can be a specific unicast MAC besides Nearest Customer Bridge group MAC
2	support multiple End Devices sharing the same External NVE via the same physical port across a bridged network.	
3	MAY support an End Device using multiple external NVEs simultaneously, but only one external NVE for each VN	Yes. Dest MAC needs to be a unicast MAC for both direct and indirect connection case.
4	MAY support an End Device using multiple external NVEs simultaneously for the same VN	Yes. Dest MAC needs to be a unicast MAC in indirect connection case.
5	allow the End Device initiating a request to its associated External NVE to be connected/disconnected to a given VN.	No. VN is indicated by GroupID
6	allow an External NVE initiating a request to its connected End Devices to be disconnected to a given VN	No. Bridge may send De-Associate

Requirements – 2

#	req	VDP Extension required?	
7	When a TS attaches to a VN, the protocol MUST allow for an End Device and its external NVE to negotiate a locally-significant tag for carrying traffic associated with a specific VN (e.g., 802.1Q tags).	No. VID==0 in request and bridge returns the assigned value in response Or specify GroupID in request and get VID assigned in returning response	
8	allow an End Device initiating a request to associate/disassociate and/or activate/deactivate address(es) of a TSI instance to a VN on an NVE port.	requirements	VDP equivalence
		associate/disassociate activate/deactivate	pre-asso/de-associate associate/de-associate
		No. Needs clarification to allow transition from associate->pre-assoc	
9	allow the External NVE initiating a request to disassociate and/or deactivate address(es) of a TSI instance to a VN on an NVE port.	No. VDP bridge may initiate de-associate	

Requirements – 3

#	req	VDP Extension required?
10	allow an End Device initiating a request to add, remove or update address(es) associated with a TSI instance on the external NVE. Addresses can be expressed in different formats, for example, MAC, IP or pair of IP and MAC.	Yes. Needs extension for IPv4/IPv6 address association. Add a new “filter info format” type
11	MUST allow the External NVE to authenticate the End Device connected.	Yes. Needs a new TLV for integrity check. May reuse current 802.1x or MacSec for direct connection. May need IETF work (IPSec?) for indirect connection.
12	be able to run over L2 links between the End Device and its External NVE.	No. L2 protocol naturally
13	SHOULD support the End Device indicating if an associate or activate request from it results from a VM hot migration event.	Yes. New bits for migration indication in new “filter info format” type (Current VDP has M bit for migrated VM on destination hypervisor and S bit for that on source hypervisor. It is indistinguishable when M/S is 0 between no guidance and events not caused by migration where NVE may act differently.)

Summary

- Extensions include:
  - Specific unicast destination MAC other than nearest customer bridge group
  - Authentication: TLV for integrity check.
    - Direct connection: using existing mechanism like .1x or macsec.
    - Indirect connection : IP sec? needs ietf work
  - IP address binding: add new filter info format type
  - Clearer migration indicator: put bits into new filter info format
  - Clarify assoc->pre-assoc : state machine in current VDP allows it but not intentionally

Status and Next Steps

- Discussed in IEEE 802.1 plenary two weeks ago. 802.1 has preliminary intention to extend VDP upon NVO3 request.
- Another round of draft editing before LC.
- Liaison from IETF to IEEE802.1 to request:
  - Amendments on VDP spec about unicast dest MAC
  - Extensions based on IETF requirement document about new filter info and TLV types