

Framework for DC Network Virtualization

`draft-ietf-nvo3-framework-00`

Authors

Marc Lasserre

Florin Balus

Thomas Morin

Nabil Bitar

Yakov Rekhter

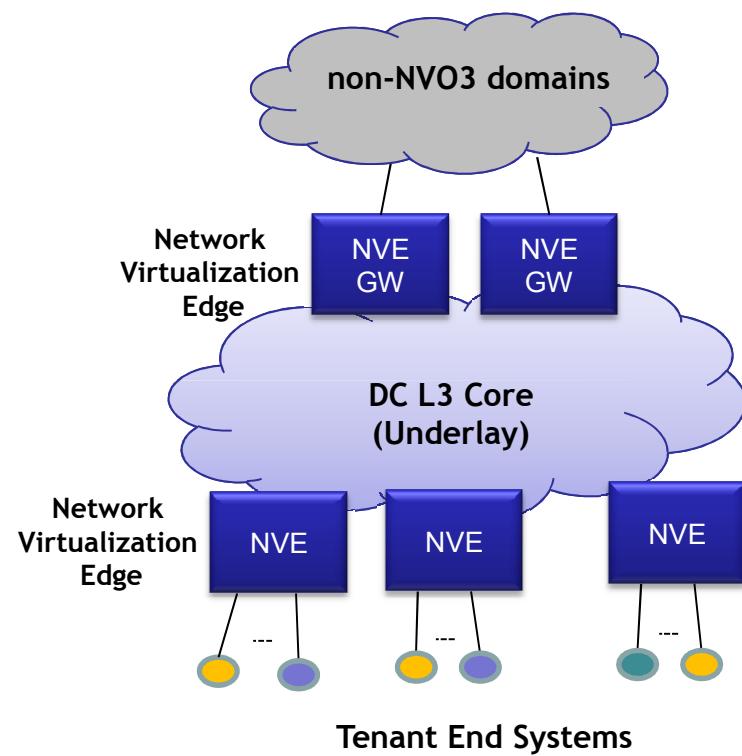
Purpose of the draft

- This document provides a framework for Data Center Network Virtualization over L3 tunnels:
 - Terminology
 - Reference model & functional components
 - Overlay issues to consider

Changes in -00

- Submitted as WG document
- Minor text cleanup

DC Network Virtualization Framework



DC L3 Core

Utilize L3 networking to interconnect NVEs

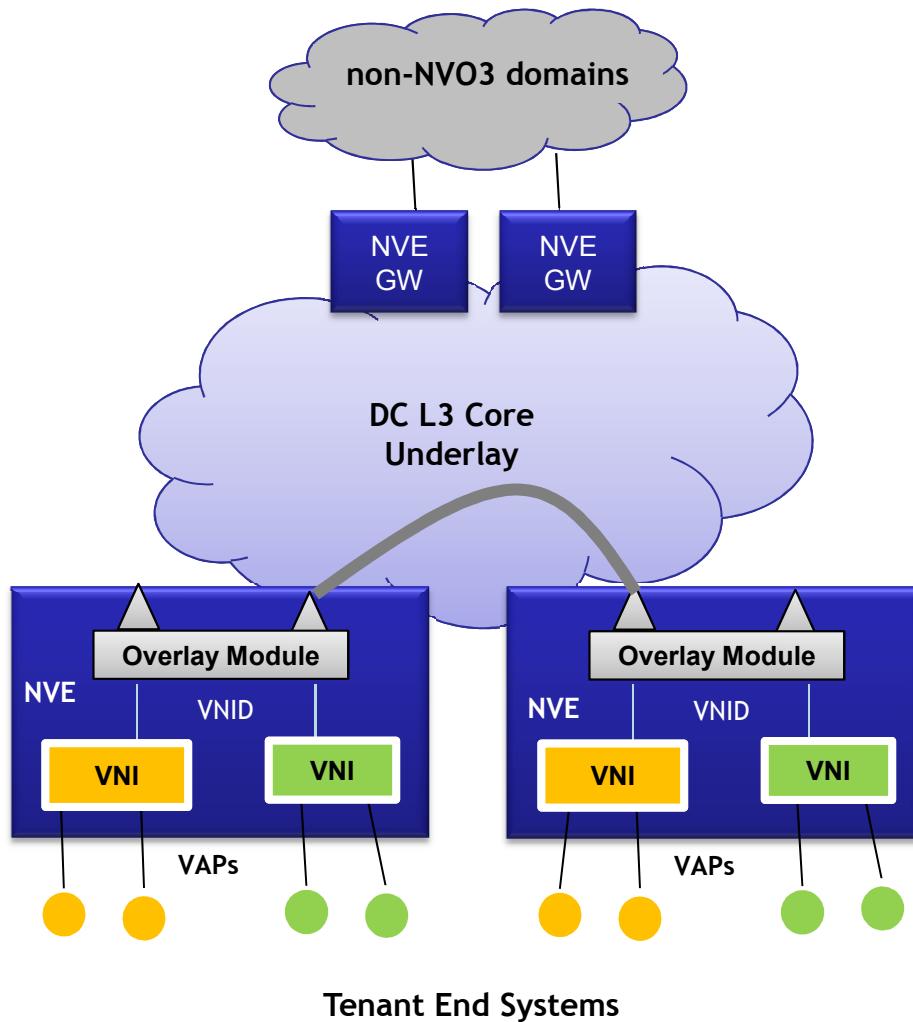
Network Virtualization Edge (NVE)

Handle L2/L3 service instances
Hide tenant addressing information
Map traffic into some tunnel,
Perform NVO3 tunnel encapsulation

Tenant End Systems (TES)

Attach to NVE directly or through L2 network
TES examples: VMs, Bare Metal servers,
Appliance Instances

Generic Reference Model for NVE



Overlay Modules

Provide tunneling capabilities
Encode TNI in tunnel header
Can use any encapsulation option
VXLAN, NVGRE, L2TP, IPSec

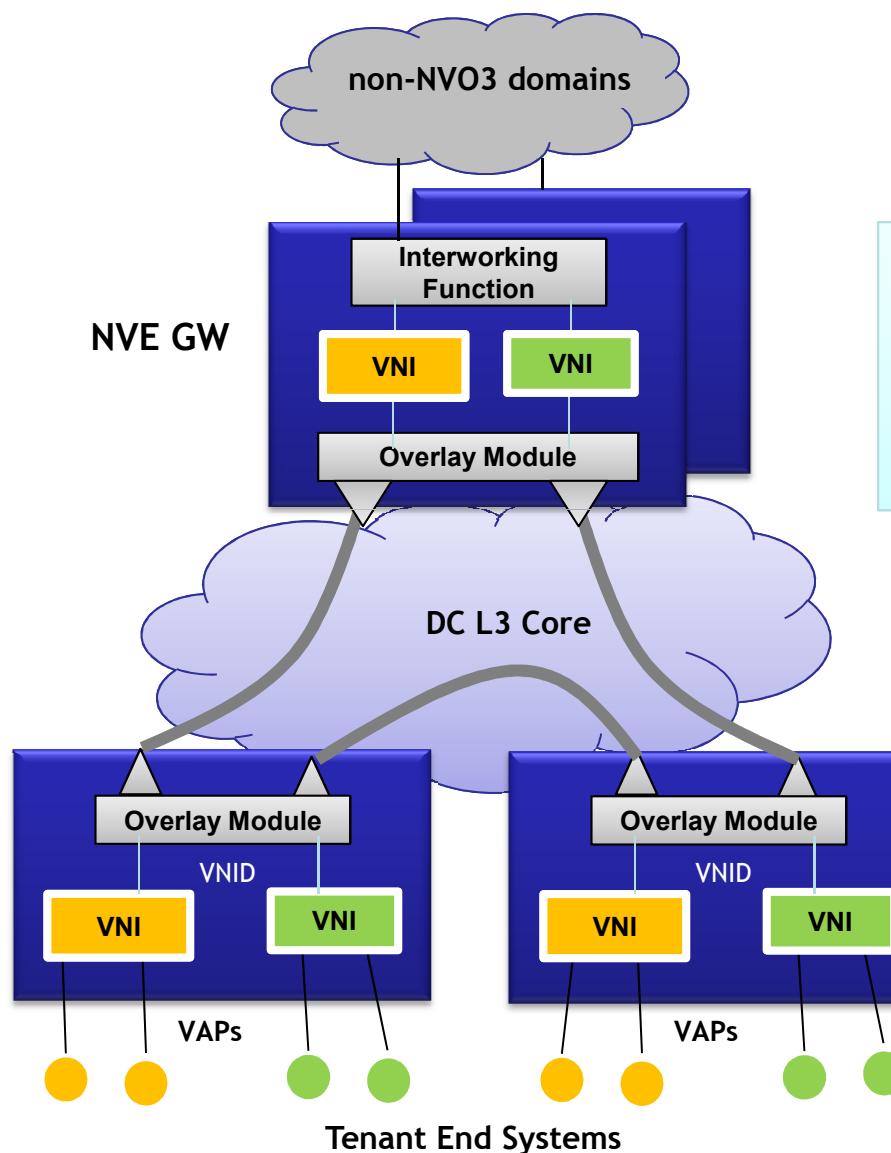
Virtual Network Instance (VNI)

Identified by a VNID
Represents tenant specific configuration
Per-tenant L2 or L3 FIBs

Virtual Access Point (VAP)

Attach tenant end-systems to VNIs
Physical or virtual ports

NVE Gateway Function



NVE GW details (to be added in a future revision)
Provides service handoff to non-NVO3 domains
Could be software and/or hardware based
Multi-homing using existing mechanisms
Examples of non-NVO3 domain: VPNs, VLANs

Solicit WG input on open issues

- Multi-homing discussion
 - Request to add text
- Request to add a service OAM section
- Solicit additional input on security comments

Next Steps

- Terminology cleanup:
 - End station -> Tenant End System
 - Tenant Service Instance -> Virtual Network Instance
- Clarification on Control Plane driven procedure
- Clarification on VN context to VNI mapping
- Add discussion on NVE Gateway