

# **draft-merged-nvo3-ts-address-migration-01**

Yakov Rekhter  
Linda Dunbar  
Rahul Aggarwal  
Ravi Shekhar  
Wim Henderickx  
LuYuan Fang  
Ali Sajassi

# Status Update

- The draft added solutions to the issues identified by “draft-ietf-nv03-vm-mobility-issues” that was adopted by WG in Dec, 2012.
- Removed E-VPN based solutions
- Added “Solution with Anycast for TS Default Gateways”
- Made changes to L3 address migration issue & solution (Thanks to Tom Herbert for his suggestions & comments)
- Updated Optimal Routing description based on discussion on the list.
- Many people have commented and made suggestions to the proposed solutions
  - Many thanks to David Black, Larry Kreeger, Dave Allen, Tom Herbert, Dino Farinacci, Sri Gundavelli, Lucy Yong, Tom Narten, Andy Malis, Sharon Barkai, Behcet Sarikaya, Reith Lothar, ZuQiang, ...

# Solution Highlight – VLAN-ID conflict

- **NVA based solution to address Same VN using different provider administered VLAN-IDs in L2 access domains**
- **Added those definitions as suggested by Dave Allen:**
  - Customer administered VLAN-IDs (usually hard coded in a TS's Guest OS and can't be changed when the TS move from one NVE to another). Some TSs may not have VLAN-ID attached.
  - Provider administered VLAN-IDs of local significance, and
  - Provider administered VN-IDs of global significance.
- **In a nutshell:**
  - NVA manages unused VLAN-IDs pool in each access L2 domain
  - NVE reports to NVA when first local TS of a VN is reachable, or none of TS in a VN is reachable by the NVE
  - NVA can push the global VN ID <-> locally administered VID mapping to NVE, or NVE can pull upon detecting a newly attached VN.
  - NVA manages the first switch to which TS is attached on mapping between TS's own VLAN-ID and "locally administered VID".

## Solution Highlight – NVA based Layer 2 Extension

- **Dynamically interconnect (add/drop) NVEs to which TSs of one specific VN are attached:**
  - NVA can offer services in a Push, Pull mode, or the combination of the two
  - When the last TS of a VN is moved out of a NVE, NVE can either send query to the NVA or the NVA notifies the NVE for it to remove its connectivity to the VN
  - When an NVE needs to support connectivity to a VN not currently supported (as a result of TS turn up, or TS migration), the NVA will push the necessary VN information into the NVE

# Solution Highlight – Optimal IP Routing

- Drive: TS migration can cause a VN spread across multiple DCBRs and fragmented IP addresses
  - When addresses can't be properly aggregated
  - Optimal routing : Inbound & Outbound
- Preserving Policies at NVEs
  - policies guarding TSs across different VNs, with some being enforced by Firewall, some enforced by NAT/AntiDDOS/IPS/IDS, etc.
  - dynamically changing policies associated with the “middleware” boxes attached to NVEs (if those middle boxes are distributed)
- Distributed Proxy Default Gateway Solution
  - NVEs perform the function of the default gateway for all the TSs attached. Those NVEs are called “Proxy Default Gateway” in this document because those NVEs might not be the Default Gateways explicitly configured on TSs attaches. Some of those proxy default gateway NVEs might not have the complete inter-subnet communications policies for the attached VNs.

## **Solution Highlight – Optimal routing (Inbound)**

- Solution 1: host routing
- Solution 2: designated NVEs
  - For scenarios that small level of triangular routing can be tolerated and underlay routers don't want too many routes when VNs are spread across multiple NVEs.

# Next Step

- This draft is ready for WG Adoption
  - Solutions for issues associated with VM mobility.