

draft-ietf-nvo3-ts-address-migration-00

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.
- 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, ...
- “VM” -> “TS”; “mobility” -> “Migration”
- Added L3 address migration issue & solution (Thanks to Tom Herbert for his suggestions & comments)

Solution Highlight – VLAN-ID conflict

- **NVA based solution to address Same VN using different provider administered VLAN-IDs in L2 access domains**
 - 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 – Layer 2 Extension

- **Dynamically interconnect (add/drop) NVEs to which TSs of one specific VN are attached:**
 - **NVA based solution:**
 - NVA informs all the NVEs to which the TSs of the given VN are attached
 - When the last TS of a VN is moved out of a NVE, NVE can either confirm with 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
 - **E-VPN based solution:**
 - VN scoped Route distribution to each NVE (NVE-PE)
 - When a given TS moves to a new L2 site, if in the new site this TS is the only TS from its L2-based VN, then the PE-NVE(s) connected to the new site need to be provisioned with the E-VPN Instances (EVI) of the E-VPN associated with this L2-based VN
 - if after the move the old site no longer has any TSs that are in the same L2-based VN as the TS that moved, the PE-NVE(s) connected to the old site need to be de-provisioned with the EVI of the E-VPN.

Solution Highlight – Optimal IP Routing (Outbound)

- Optimal routing of a TS's outbound traffic (avoid tri-angular routing to default gateway when TS is moved)
 - E-VPN based TS default gateway solutions
 - Option 1: Requiring all TSs to have Anycast default gateway IP address and Anycast default gateway MAC address
 - Option 2:
 - Each default gateway for a given L2-based VN advertises in the E-VPN control plane its default gateway IP and MAC address using the MAC advertisement route, and indicates that such route is associated with the default gateway
 - Each PE-NVE that receives this route and imports it as per procedures of [E-VPN] MUST create MAC forwarding state that enables it to apply IP forwarding to the packets destined to the MAC address carried in the route. The PE-NVE that receives this E-VPN route follows procedures in Section 12 of [E-VPN] when replying to ARP/ND Requests that it receives if such Requests are for the IP address in the received E-VPN route.
 - 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 IP Routing (Inbound)

- Optimal routing of TS's inbound traffic (proper routes announcement to intermediate nodes in underlay to avoid/minimize triangular routing)
 - NVA Based Intra Data Center Triangular Routing
 - NVA distribute all routes to NVEs.
 - E-VPN Based
 - To avoid triangular routing among TSs within the same data center, E-VPN routes originated by one of the EVIs within such set should be imported by all other EVIs in that set, irrespective of whether these other EVIs belong to the same E-VPN as the EVI that originates the routes

Solution Highlight – L3 Address Migration

- TS migration cause fragmented IP addresses
- Main Issue: Inbound IP optimal routing
- 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.