draft-sajassi-bess-evpn-fabric-migration-02.txt

A. Sajassi (Cisco), C. Wang, N. Malhotra

IETF 123, July 2025

Madrid

History

- Rev 01 (Presented in IETF 120)
  - Problem Statement of EVPN underlay network migration
  - Fabric Migration Procedures
  - BGP Procedures
  - Multihoming Operations
- Rev 02 (Current)
  - Clarification on “Originating Router’s IP Address”
    - Referencing to RFC7432bis about the behaviors
  - Operational considerations:
    - Reverting before migration completion
    - Reverting after migration completion

Fabric Migration (Rev 01)

- Problem Statement
  - To migrate EVPN from IPv4 to IPv6 (or vise versa) SEAMLESSLY!
  - Migration is independent of BGP control-plane session – i.e., user data traffic is migrated to use IPv6 while BGP can remain IPv4
- Solution Architecture
  - Enable IPv4/IPv6 dual-stack underlay on the network
    - Continue Using NLRI NH field to carry primary IP address
    - Use optional BGP Tunnel Encapsulation Attribute to carry secondary IP address
  - Originating Router's IP Address
    - IPv4 or IPv6, independent of the BGP Next Hop (VTEP) address type for that NLRI

Fabric Migration from V4 to V6 (Rev 01)

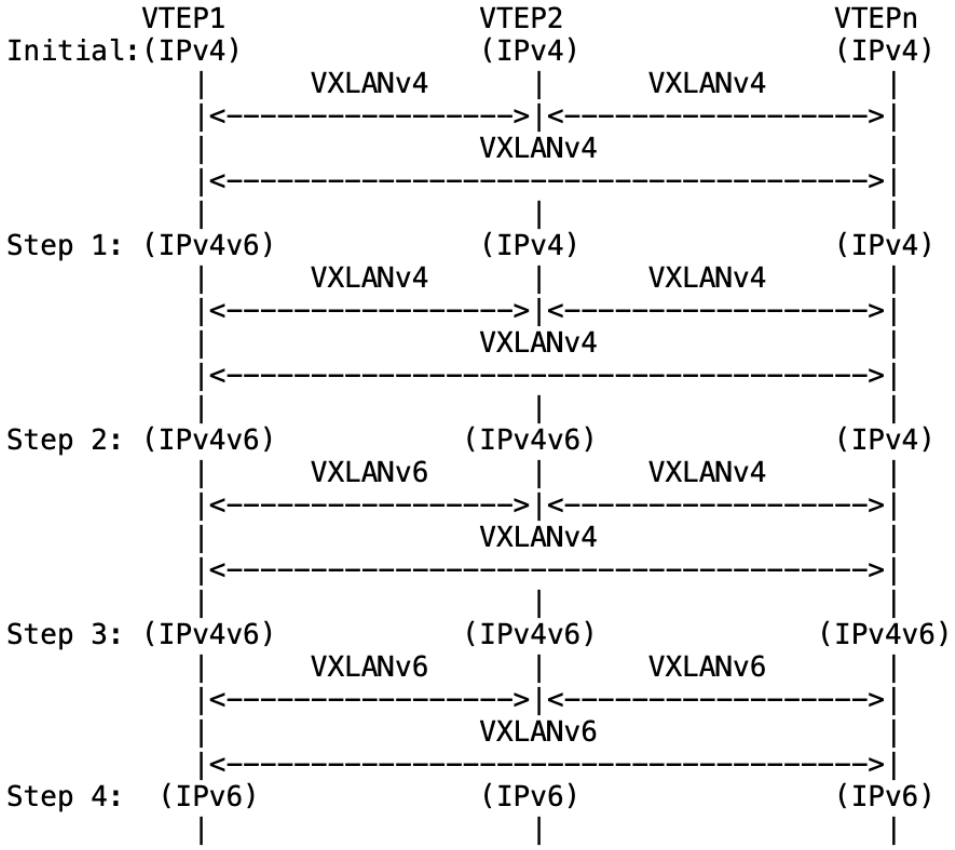

Figure 1: Fabric Migration from VXLANv4 to VXLANv6

Clarification on “Originating Router’s IP Address” (Rev 02)

- draft-ietf-bess-rfc7432bis-13

Section 8.1.1 (ES)/Section 11.1 (IMET)

“The Originating Router's IP Address field value MUST be set to an IP address of the PE (e.g., this address may be the PE's loopback address). The IP Address Length field is in bits. The Originating Router's IP address does not need to be a routable address and its purpose is to identify the originator of that EVPN route uniquely. **It can be either IPv4 or IPv6 address independent of the BGP next hop address type for that NLRI and it must remain the same for all EVPN routes advertised by that PE (across all EVIs).**”

Revert before migration completion (Rev 02)

- Reverting the migration process **back to VxLANv4** before its completion requires careful handling to ensure network stability and avoid traffic disruption. In this scenario, some PEs are dual-stack and some PEs are still single-stack IPv4.
- Reverting steps
  - Find all the migrated VTEPs (dual-stack mode)
  - Start with the first one, revert it back to single-stack “ipv4”, but do NOT unconfigure local IPv6 VTEP address to ensure seamless migration
  - Confirm the revert is successful and there is no connectivity issue
  - Repeat and revert the rest
  - Once all done, the IPv6 VTEP addresses can be unconfigured

Revert after migration completion (Rev 02)

- This is equivalent to fabric migration from IPv6 to IPv4 underlay. In this scenario, all PEs are single-stack IPv6.
- Reverting steps
  - Confirm the fabric is operating in single- stack "ipv6" mode
  - Start with the first VTEP, migrate it to dual-stack mode and ensure that the EVPN routes from this VTEP carry IPv6 VTEP address in the Next Hop of MP_REACH_NLRI attribute and IPv4 VTEP address in the Tunnel Encapsulation attribute.
  - Confirm the migration is successful and there is no connectivity issue
  - Repeat and migrate the rest
  - Once all the VTEPs are in dual-stack mode, follow the same procedure in previous slide to revert to single-stack (VXLANv4)

Next Step

- Solicit Input from Working Group

THANK YOU!