


NVO3 NVA Gap Analysis

Linda Dunbar
Donald Eastlake

TRILL Directory Applicability to NVO3

- **TRILL: L2 in TRILL**
- **NVO3: L2 in IP; L3 in IP**
- **Goal : Resolve Egress edge addresses**

Variety of Directories (NVAs)



Locations:

- Embedded in routers/switches in the core, or as standalone servers attached to them.
- Embedded in Edges or as Standalone servers connected to Edges via the client side

Contents:

- Each DIR has mapping for a subset of VNs
- multiple DIRs have content for a VN

Push Model in TRILL

- **Policies:** When ingress edge can't find entries for the incoming data frame:
 - simply drop the data frame,
 - flood it to all other edges that are in the same VN, or
 - start the “pull” process to get information from Pull Directory Server(s) (or NVA)
- **Requesting Push Service:**
 - Push Directory servers (NVAs) (or the Edge by which they are attached) use VN scoped reliable Link State flooding to announce their availability to push mapping information.
 - Edge nodes use VN scoped reliable Link State flooding to announce all the Virtual Networks in which they are participating
- **Push:**
 - Directory servers (NVAs) pushed VN scoped entries to Edges.
- **Incremental Push Service Update**
 - Achieved by Link State Update to distribute the incremental updates.

Push Model: Multiple Directories

- When there are multiple Directories for a given VN, all of them can push, and configured number of them actually do.
 - An Edge takes the content from the Directory with the highest system ID

Pull Model

- **Request and Reply use separate control channel (Rbridge Channel syntax)**
- **Pull Request:**
 - Target address family (IPv4/IPv6, [MAC/VLAN])
 - Optional: Source Address family (IPv4/IPv6, MAC, VLAN, Source NVE)
 - Triggered by:
 - An edge node (NVE) receives an ingress data frame with a destination whose attached edge (NVE) is unknown, or
 - The edge node (NVE) receives an ingress ARP/ND request for a target whose link address (MAC) or attached edge (NVE) is unknown.

PULL Response

- Address mapping, coupled with the valid cache timer
- The target being queried is not available, or
- The requestor is administratively prohibited from getting an informative response.
- If no response is received within a configurable timeout, the request should be re-transmitted with the same Sequence Number up to a configurable number of times that defaults to three.

Push-Pull Hybrid Model

- Push model are used for some VNs, and pull model are used for other VNs.
 - It can be operator's decision (i.e. by configuration) on which VNs' mapping entries are pushed down from directory (e.g. frequently used) and which VNs' mapping entries are pulled (e.g. rarely used).
 - Useful for Gateway nodes where great number of VNs are enabled.
- Or, a portion of hosts in a VN is pushed, other portion has to be pulled.

Cache Consistency

- When data in a Pull Directory changes and there may be unexpired stale information at a querying edge (NVE), the Pull Directory must send an unsolicited message to the edge (NVE) to update the pulled entries.
- When data in Push model changes, incremental update is sent to all edges that have pushed entries.

Gap Summary

- TRILL's Push and Pull Policies are applicable to NVO3
- TRILL's Push/Pull requests/reply mechanism are applicable to NVO3
- NVO3's NVA will need one additional attribute: VN context (VN ID and/or VN Name).