

Potential uses of TRILL/LISP mechanisms for an NVE<->NVA Interface


draft-dunbar-nvo3-nva-gap-analysis-01

Linda Dunbar
Donald Eastlake

TRILL Directory & LISP Mapping System Applicability to NVO3

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

Various ways of Directories (NVAs) connecting to Edge nodes (NVEs)



Locations:

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

Contents:

- Each Directory has mapping for a subset of VNs
- multiple Directories have mapping entries for a VN

Pull Model Used by LISP and TRILL

- **LISP: Map-Request/Map-Reply/Map-Register/Map-Notify**
 - Requests are sent to Map-Resolver (Portal to Mapping system)
 - Map-Resolver has to find the Map-Server(s) which are responsible for the block of address space
 - Map-Requests are then sent to the appropriate Map-Server via LISP-DDT, which forward it to registered ETRs
 - The messages are encapsulated in a LISP control message called a "Encapsulated Control Message".
- **TRILL: request/reply**
 - Requests can be sent to any of the valid directories
 - TRILL directory 100% depends on other authoritative sources. Therefore, doesn't need mechanism to get input from Edge nodes
 - use Rbridge Channel syntax for those messages.
- **Pull Request 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 Responses by TRILL/LISP

- **When the mapping entry is available in the Directory/MappingSystem**
 - Valid Response (TRILL/LISP)
- **When the mapping is not available:**
 - LISP: policy to ITR : “native-forward”, or “drop”
 - TRILL: policy to Edge: “drop” or “native-forward” (i.e. flooding)
- **TRILL extra:**
 - cache timer

Push Model

- **Only in TRILL. LISP is PULL only**
- **Push:**
 - Directory servers (NVAs) pushed VN scoped entries to Edges.
- **Incremental Push Service Update**
 - Achieved by Link State Update to distribute the incremental updates.
- **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
- **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)

Push-Pull Hybrid Model

- Specific to TRILL
- 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.

Comparison Summary

	Types of Content	Content Source	Push /Pull?	Incremental Update
TRILL Directory	<ul style="list-style-type: none"> •Inner-outer mapping for target in the same virtual network as the Source. 	100% from authoritative system (e.g. VM orchestration)	Per VN based Push Pull, or combination	Yes
LISP MappingSys	<ul style="list-style-type: none"> •Inner-outer mapping for targets in any Virtual Networks 	100% populated via control plane	Pull only	Not applicable
NVO3' NVA	<p><u>When attached VNs are L2:</u></p> <ul style="list-style-type: none"> -Inner-outer mapping for target in the same VN as Source <p><u>When the attached VNs are L3:</u></p> <ul style="list-style-type: none"> -Inner-outer mapping for targets in any Virtual Networks •Gateway's forwarding table for inter-subnet (or inter-VN) communication • Potentially the Inter-VN communication policies (??) 	<p>coming from authoritative systems</p> <p>But also allow NVEs to announce their directly attached TSes.</p> <p>(conflict?)</p>	Push, Pull, Hybrid	???