

Schemes for TRILL's Directory Assisted RBridge Edge


Linda Dunbar
Donald Eastlake
Radia Perlman
Igor Gashinsky
Yizhou Li

Goal

- Reduction in Multi-destination Frames in the Data Center Environment where all (or most) of the configuration is centrally orchestrated.
 - Reduce or eliminate flooding of unknown unicast messages.
 - Reduce or eliminate flooding of ARP/ND messages by responding from the ingress RBridge when possible.
- The mechanism is VLAN scoped.

Terminologies


- **TRILL-DIR**: an RBridge which has Directory Server either directly attached or embedded .
 - Standalone Directory Server: has Nickname, appear as RBridge to others, but not processing native Ethernet data frames to/from end stations.



- The communication between pure directory server and RBridge (real or dummy) is beyond the scope of this draft.

Push Model

Use ESADI base protocol with some minor changes

- Original ESADI: only advertise local attached MACs
 - TRILL-DIR-PUSH: using message format defined by ESADI to send out global mapping entries
 - IPv4, IPv6, MAC, **Nickname**
-  **Doesn't have to be yourself**
- ESADI consumer listens
 - ESADI provider advertises
 - Local information
 - TRILL-DIR: Global information (attachment nickname can be other than yourself)

Push Model

Use ESADI base protocol with some minor changes

- TRILL-DIR announces (in its LSP) its supported VLANs to all Rbridges:
 - One bit for “I am a Push Directory” (Uses ESADI)
 - One bit for “I am a Pull Directory” (Uses Query/Response)
 - Potentially another bit to indicate that it doesn’t want to receive user data frames in the interested VLANs
- ESADI Consumer: no change


Push Model

suggested changes

- TRILL-DIR announces (in its LSP) its supported VLANs to all Rbridges:

- Section 2.3.6 of RFC 6326 (TRILL uses of IS-IS): Interested VLANs Sub-TLV

- Interested VLANs: The Interested VLANs field is formatted as shown below.



- Using BIT#3 to indicate that this is the Push Directory for those VLANs.
- Using BIT #2 to indicate that it is the Pull Directory for those VLANs
- (potentially use another bit to indicate that it doesn't want to receive user data frames in the interested VLANs)

Push Model: Directory server Sending mapping entries

- “Designated TRILL-DIR”
 - When there are multiple TRILL-DIRs for any given VLAN, the TRILL-DIR with the highest system ID is “Designated TRILL-DIR”
 - Only the “Designated TRILL-DIR” sends the full mapping for the VLAN via ESADI.
 - (Same as ESADI, VLAN scoped.)

Pull Model

- **Request and Reply use Rbridge Channel syntax**
 - **TRILL doesn't require TRILL-DIR to have a MAC address**

PULL Request/Response

- **Pull Request:**
 - Target address family (IPv4/IPv6, [MAC/VLAN])
 - Optional: Source Address family (IPv4/IPv6, MAC, VLAN, Source RBridge (Nickname))

- **Pull Response:**
 - Target address family (IP, Target MAC, VLAN) + Egress RBridge (Nickname)
 - Optional “Timeout information”
 - Or Reject:
 - Need Source address family for validation
 - Not allowed

Push-Pull Hybrid Model

- Push model are used for some VIDs, and pull model are used for other VIDs.
 - It can be operator's decision (i.e. by configuration) on which VIDs' mapping entries are pushed down from directory (e.g. frequently used) and which VIDs' mapping entries are pulled (e.g. rarely used).
 - Useful for Gateway nodes where great number of VLANs are enabled.