TRILL Directory Assistance Mechanisms

draft-dunbar-trill-scheme-for-directory-assist-05
draft-eastlake-trill-ia-appsubtlv-00

Linda Dunbar, Donald Eastlake
Radia Perlman, Igor Gashinsky, Yizhou Li

Goal

- To reduce multi-destination traffic by
  - reducing or eliminating unknown unicast flooding
  - when appropriate, locally responding to ARP, ND, and RARP requests or converting them to unicast requests

  using address mapping directories.

- An appropriate source of directory information is a Data Center orchestration system.

Two Directory Types

- ## Push Directories

  - Data and updates pushed out on a per Data Label (VLAN or Fine Grained Label) basis to all subscribing candidates.

  - Redundant push directories supported that can be configured as to how many should be active at once.

  - Data conflicts and which push directories should be active arbitrated by priority.

- ## Pull Directories

  - Each pull directory responds to requests in a set of Data Labels it advertises.

  - Data and negative responses include an expiry time for caching.

  - Unsolicited updates sent for unexpired cached information.

Out of Scope

- How directories are populated with information and updated as it changes.

  – But orchestration systems seem like a good source.

- How data in multiple directories are kept synchronized.

  – But updating clients when directory data they are holding changes is in scope.

Push Directory Mechanisms

- Uses the TRILL ESADI protocol for reliable Data Label scoped data flooding services.

    - draft-ietf-trill-esadi

- Client simply advertises (in core IS-IS) participation in ESADI for a Data Label and will be sent the data and updates.

- Push Directories for a Data Label can see each other in that ESADI instance to arbitrate which should be active.

Pull Directory Mechanisms

- Uses the RBridge Channel mechanism to encode requests and responses.

  - draft-ietf-trill-rbridge-channel

- Can optionally include the actual frame that caused the pull in pull request if it is small enough.

- Can be hosted on an end stations, in which case the TRILL switch by which it is reachable proxies for it.

Push Pull Policies

- Clients can have a wide variety of policies:
  - Type of directory use:
    - Just use pushed data.
    - Just use pulled data.
    - Consult pushed and cached pull data and do a pull if no match found.
  - Behavior if no match
    - Either immediately or after failed pull:
      - Discard packet
      - Flood packet

Push Pull Policies

- Push Directories may be the most appropriate for Data Labels with a smaller number of end stations that mostly all talk to each other.

- Pull Directories may be the most appropriate for Data Labels with larger numbers of end stations with sparse intercommunication.

- If a Data Label has a few end stations everyone talks to, but otherwise has sparse intercommunication, you could push information for the few and use pull for the rest.

- A client pulling could just pull from the nearest relevant pull directory or could pull from all revenant pull directories and use the first response it gets, etc.

TRILL Directory Assist Mechanisms

Address Mapping Data Representation

- Both push and pull directories use the "Interface Addresses" APPsub-TLV

  – draft-eastlake-ia-apsubtlv-00.txt

- Provides for encoding a very flexible set of addresses that all represent the same Interface (port).

  – For example { an IPv6 address, a 48-bit MAC address, a Data Label, a TRILL switch nickname }.

    - Could be used to look up the IPv6 address or the MAC address within the Data Label to get the other addresses.

Next Steps

- Working Group consensus for the Directory Assistance Framework document indicates that is a general direction that the WG has decided to pursue.

- Plan is to polish the directory assistance mechanisms draft next week and issue a call for WG adoption the week after next.

- Solicit reviewers at this meeting to review after WG adoption.