

Directory Assistance Mechanisms

draft-dunbar-trill-scheme-for-directory-assist-04.txt

Linda Dunbar

Donald E. Eastlake, 3rd

Yizhou Li

Huawei Technologies

Radia Perlman

Intel Labs

Igor Gashinsky

Yahoo

Directory Assistance

- Goal:
 - Reduce Multi-Destination Traffic, in particular
 - ARP, ND, RARP, ...
 - Unknown destination unicast
- Method:
 - Directory information to answer ARP, ND, and RARP and decrease unknown unicast.
 - Discard unknown unicast if Directory is complete.

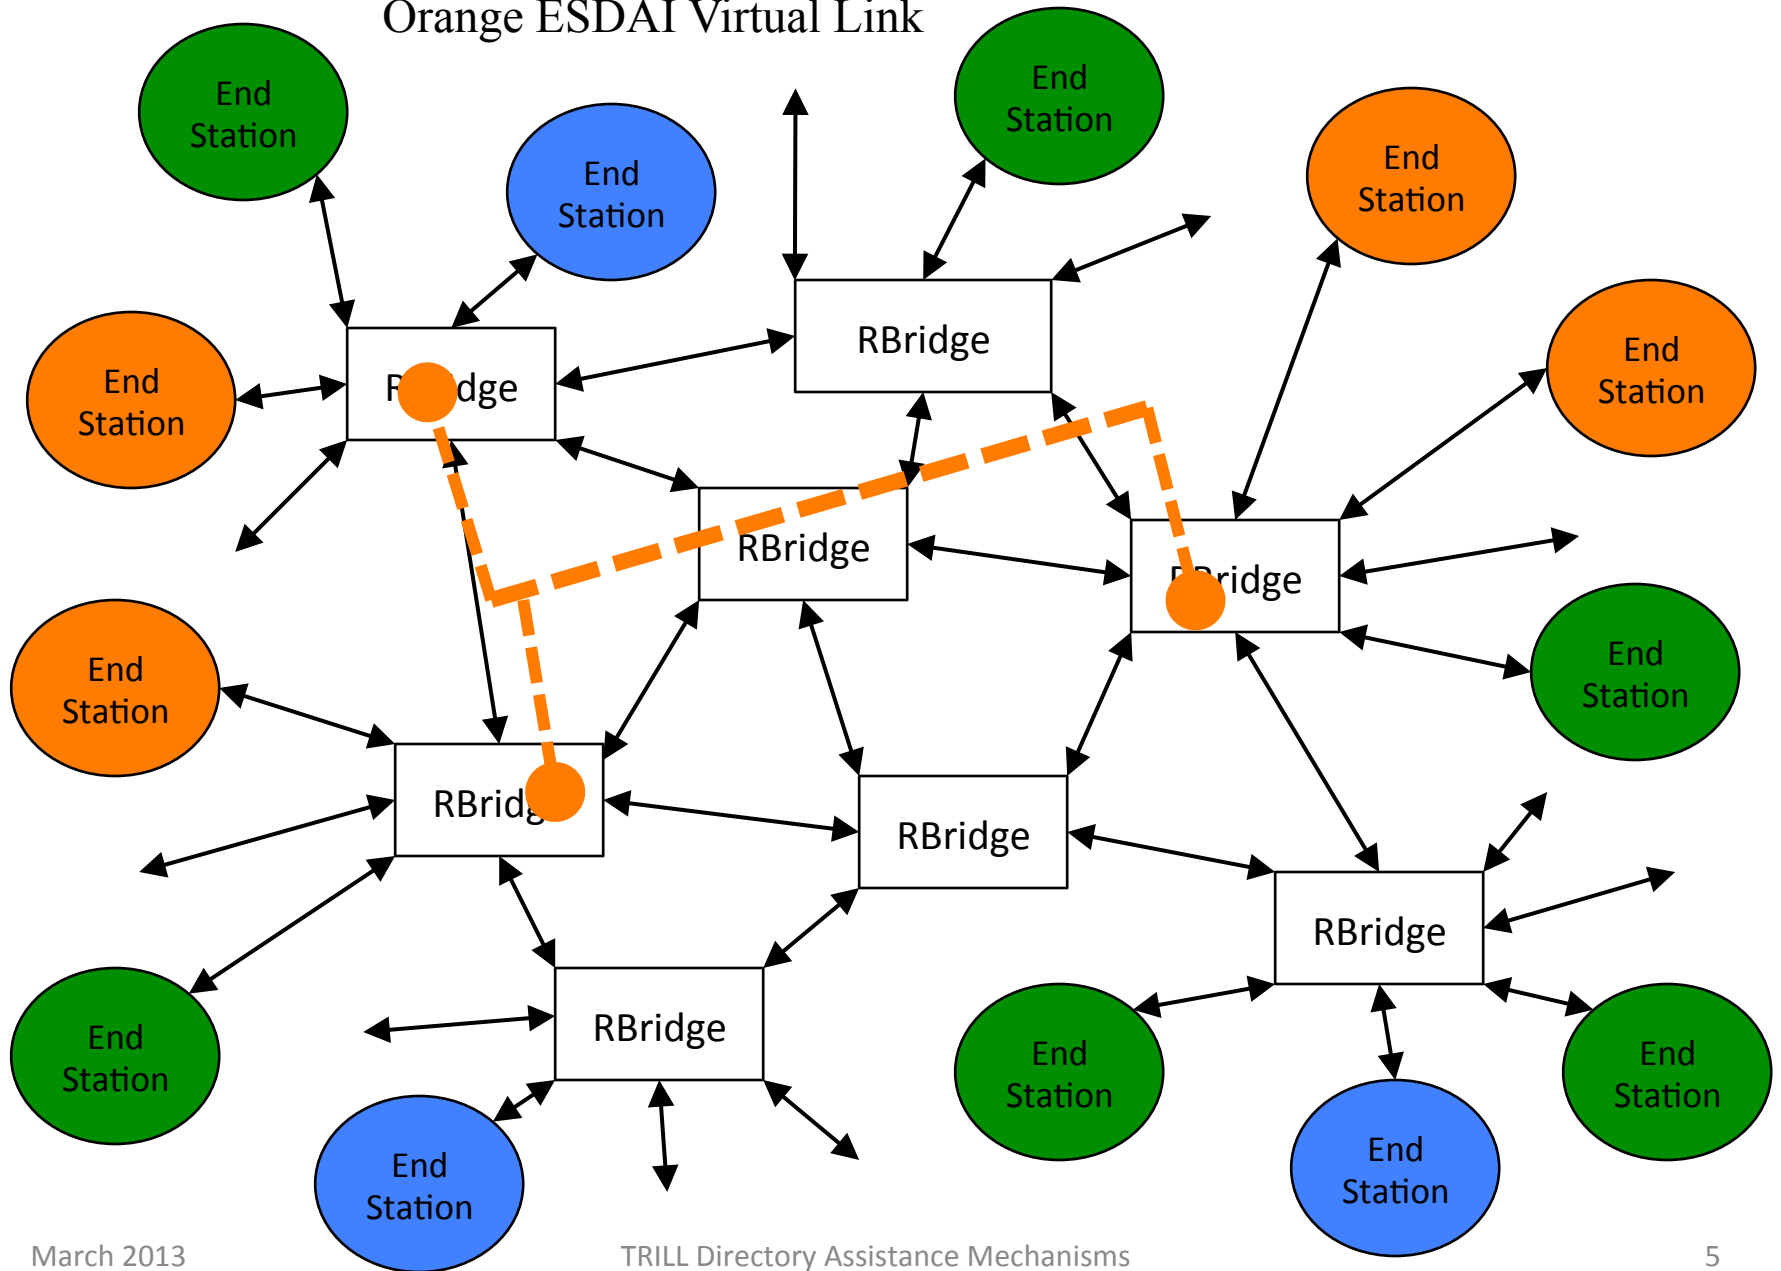
Directory Assistance Mechanisms

- Push Directories
 - Uses ESADI: Client subscribes by data label (VLAN or fine-grained label) and gets information pushed to it by Push Directory servers.
- Pull Directories
 - Uses RBridge Channel messages to query RBridges that advertise themselves as a Pull Directory for one or more data labels.

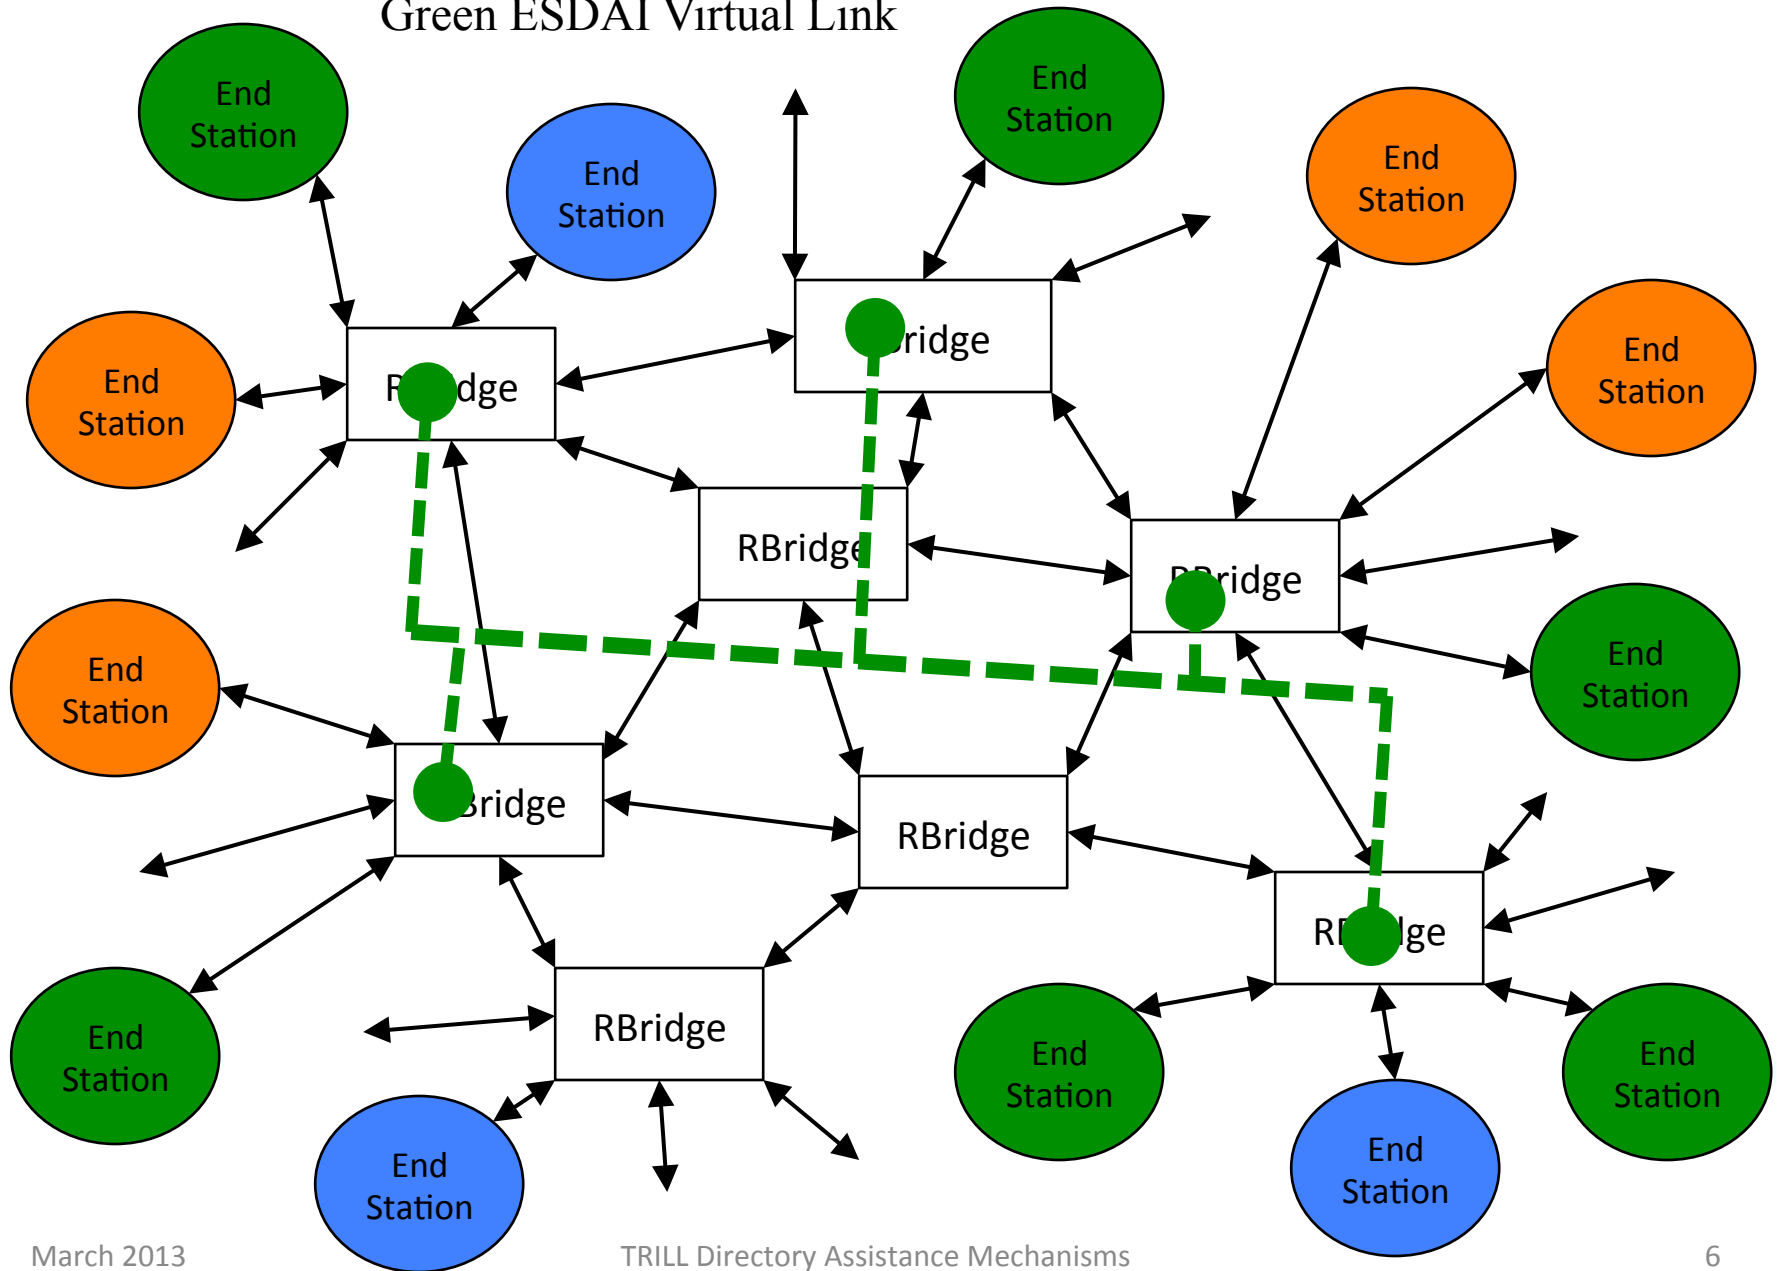
Digression on ESADI

- ESADI is a data label scoped method of flooding information.
 - All RBridges participating in ESADI for the same data label appear to be directly connected to the same virtual link. Information is tunneled through transit RBridges.
- Used in current TRILL protocol as an optional way for RBridges to report locally learned attached MAC addresses using the MAC Reachability TLV (RFC 6165).

Color = VLAN
Orange ESDAI Virtual Link



Color = VLAN
Green ESDAI Virtual Link



New IA Data Structure

- Push and Pull Directories use a new Interface Addresses APPsub-TLV (IA) data structure.
- Each IA contains lists of addresses that all identify a single interface (port) and the RBridge by which that interface is reachable.
 - For example, a 48-bit MAC address, an IPv4 address, and an IPv6 address.
- IA is very flexible and currently specified in the Directory Assistance Mechanisms draft but could be a separate document.

Push Directories

- When a client participates in ESADI for a data label, directory information for that label is pushed to them.
- A Push Directory advertises its existence in ESADI and says whether its information is complete.
 - So Push Directories can see each other. For robustness, normally the two highest priority Push Directories are Active while any others are Passive and do not push their copy of the database.

Pull Directories

- Pull Directories advertise themselves in the Interested VLANs or Interested Labels sub-TLV in the core IS-IS database. ESADI is not used.
- Queries for information related to one or more addresses, usually MAC or IP addresses, is sent to a Pull Directory, by default the one that is least cost from the client.
- Replies provide either a negative response for the address (unknown or administratively denied) or an IA with the other addresses associated and the RBridge by which they are reachable.

Pull Directories (cont.)

- Both positive and negative results are cached for a time set by the Directory.
 - Caching of negative results is important. Most ARPs and NDs are for non-existent addresses.
- Pull Directories **MUST** track outstanding cached data at some level of granularity and send unsolicited updates if the information changes.
- Pull Directory can operate from an end station – proxying RBridge just forwards queries and responses as appropriate.

Modes of Client Operation

- Modes:
 - On a per data label basis: No Directory Use, Use Push Only, Use Pull Only, Use both.
 - Do/Don't discard unicast if you have Push Directory information that claims to be complete.
 - If not answered by cached information:
 - Hold frame while doing a Pull.
 - Flood frame in parallel with doing a Pull.

Directory Conflicts

- It is possible to have conflicting data from Push and/or Pull Directories and/or local learning.
 - Higher confidence data always wins.
 - Data from Push Directory that is not data reachable is ignored.
 - For conflicts between Push Directories with equal confidence, the highest priority Push Directory wins.

The No Data Bit

- Existence of an Interested VLANs or Interested Labels sub-TLV originated by an RBridge has indicated interest in that RBridge receiving multi-destination data for the VLAN(s) or Label(s) listed.
- But Push Directories use this sub-TLV to indicate ESADI participation and Pull Directories use this sub-TLV to advertise themselves.
- RBridges advertising those sub-TLVs as a Directory may not want to receive multi-destination user data for the label(s) involved. So a “No Data” bit is provided for them to indicate this.

Directory Assistance Mechanisms

- Some Possible Enhancements:
 - Method to declare Push Directory information is complete for less than an entire data label. For example, for a range of addresses.
 - Specification of how Push Directory could work from an end station to the RBridge proxying for it.

Next Step

- Current draft is a bit rough.
- Comments welcome.
- Call for WG adoption after one more revision by authors.

END