

RBridges: ARP/ND Optimization

Radia Perlman

Radia@alum.mit.edu

Issue

- ARP / ND multi-destination frames are a substantial burden on LANs

Layer 2 ½ a natural place

- ARP/ND operates in a layer 2 cloud
- With smart components inside the cloud, and no modification of IP node behavior, can cut down on ARP/ND traffic inside the cloud

What types of things can R Bridges do?

- Proxy ARP/ND (if first R Bridge thinks it knows the answer), learning from
 - previous ARP/ND replies
 - “deep inspection” of data packets
 - ESADI advertisements (layer 3, layer 2) pairs
- Block the query (e.g., another pending to D, “recent” nonreply from D)
- Unicast the query to the expected D’s RB

ESADI review

- End System Address Distribution Information
- Per-VLAN advertisements
- Multicast like a data packet, only received (and explicitly processed) on links with that VLAN
- Optional to transmit, optional to receive
- Currently spec talks about advertising L2 addresses, but the original design allowed also advertising (layer 3, layer 2) pairs

ESADI

- DRB R1 can closely monitor its attached IP nodes without bothering the rest of the cloud, e.g., through
 - Do ARP/ND queries to ensure the endnode is still there
 - Static configuration
- Advertise “definitive” information via ESADI

Blocking ARP/ND Queries

- Defense against obnoxious endnode querying too often
- If multiple sources query for the same destination within some amount of time
- At decapsulating RB R2: If some ports statically configured, R2 need not forward query to ports known not to have D, or if R2 knows (D,x) exists, R2 can proxy respond


What does “unicasting” a query mean?

- The query, as transmitted by the endnode, is not modified in any way
- The TRILL header
 - Flag showing it is not a multidestination frame
 - First RB = (obvious)
 - Last RB = suspected RB attached to D
- Inside, the ARP/ND query, with Ethernet header = broadcast or multicast dest address
- If destination link is multidestination cloud, ARP/ND query is broadcast/multicast **on that one link**

Why Unicast a query?


- Perhaps have timer
- If “small” time has elapsed since (layer 3, layer 2) info learned, confidently proxy reply
- If “medium” time, can verify the cache entry through unicast ARP

RBridge Campus




In the following slides


- Color = VLAN
- Orange and Green use ESADI
- Blue does not



Orange ESDAI Virtual Link



Green ESDAI Virtual Link



END

Radia Perlman

Radia@alum.mit.edu