

TRILL issue: Mixed links

Radia Perlman

Radia.Pperlman@sun.com

Issue

- RBridges could work over lots of types of layer 2 links
- The proposed encapsulation format was written for Ethernet
 - Add TTL, “next hop”
- But no reason why it can’t support other link types (MPLS, Infiniband, Fiberchannel)
- Question: how ambitious should we be?

Other link types, homogeneous environment

- Should be easy to define encapsulations for other link types, assuming a homogeneous campus (all links the same type)
- MPLS is probably trivial (it already has a hop count, you don't have to worry about shared links)

Mixing link types (type A and B)

- Less ambitious goal
 - A nodes talk to A nodes, B nodes talk to B nodes. (A does not talk to B)
 - A nodes can talk to A nodes across B links
- More ambitious goal
 - A nodes talk to B nodes by having RBridges translate

Two types of “more ambitious”

- Supporting IP nodes on link types A and B
- Supporting non-IP nodes on heterogeneous link types talking to each other
- My guess: supporting IP nodes won't be too hard (other than finding an IP router—see later slide), but mixing non-IP incompatible link types is too hard and hopefully not important

Translation issues

- What if packets sizes on A and B are different?
- What if information (such as priority) gets lost?
- What if (layer 2) addresses are incompatible?

IP nodes

- Let's say S is on link type A, D is on link type B, and S's RBridge is R1, D's is R2
- S does ARP query, which gets flooded
- R2 does ARP for layer 3 address "D". Gets reply (layer 3=D, layer 2=d)
- R2 sends (type B, D, d) in LSP
- R1 says "whoops. A won't understand "d""
- So R1 replies with either R2's layer 2 address or R1's layer 2 address

Forwarding for an IP node of type A to an IP node of type B

- A transmits to “r” (a type A MAC of R1)
- A’s frame gets forwarded from RBridge to RBridge, sometimes with a type A encapsulation hdr, sometimes with a type B encapsulation hdr
- When R2 forwards to D, it replaces the layer 2 header in the original frame with a type B header

Summary

- Just think of it as forwarding an IP packet, but not using the hop count in the IP packet (so S and D will think they are neighbors)
- No problem with forwarding IP packets across incompatible link types

But an issue:

- Suppose the IP routers on the link have type B addresses, and S is type A
- S needs to have a type A layer 2 address for an IP router in order to transmit to a “non-neighbor” IP node
- Various solutions possible
 - Require all IP routers to have both a type A and a type B address, and have RBridges know how to route to it
 - Have address translation table in RBridges for on-link IP routers, perhaps even learned by listening to router control traffic

Conclusions

- Hopefully we don't need to bridge non-IP heterogeneous links, though bridging across heterogeneous links is possible
- IP over heterogeneous links is not very hard
- Which link types should we support?