

# TRILL issue: VLANs

Radia Perlman

[Radia.Pperlman@sun.com](mailto:Radia.Pperlman@sun.com)

# Issue

- To support VLANs, we need to keep VLAN separation
- To optimize for VLANs, we would like to optimize bandwidth and memory

# Broadcast packets within VLANs

- ARP/ND queries, or unknown destinations need to be flooded, but only within that VLAN
- Rbridges should announce, in their link state info, which VLAN(s) they support (have potential endnodes for)
- Flood VLAN A packet only to Rbridges that support A (as per multicasts)

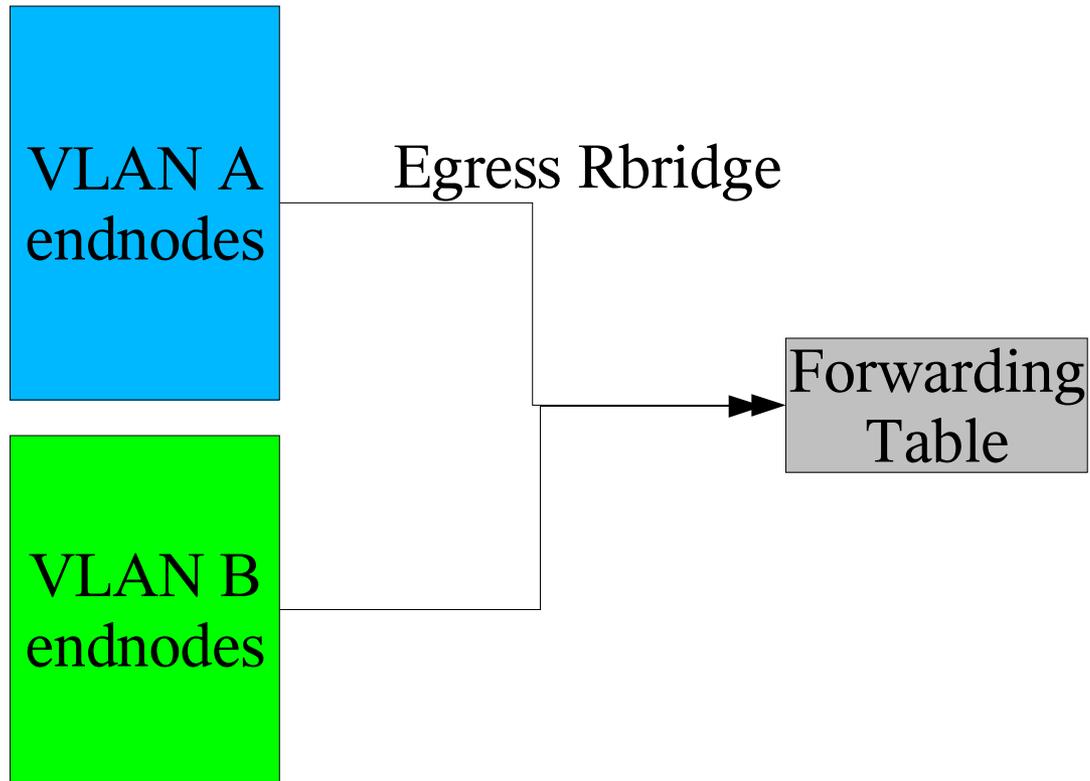
# Avoiding keeping endnode info

- Suppose we want Rbridge R to only know about endnodes in VLAN(s) R supports
- One proposal:
  - Separately flood VLAN endnode info, and only within that VLAN
  - Add “destination Rbridge” to encapsulation header

# Packet forwarded on Ethernet

- Outer hdr (must look kosher to bridges)
  - source=transmitting Rbridge
  - destination=next hop Rbridge
  - Protocol type=new, to be defined
- Encapsulation header
  - TTL
  - Egress Rbridge
- Original frame

# Rbridge Tables



# Conclusion

- Forwarding table only contains Rbridges
- First hop Rbridge determines egress Rbridge
- Only need to keep link state info about endnodes in VLAN(s) you support