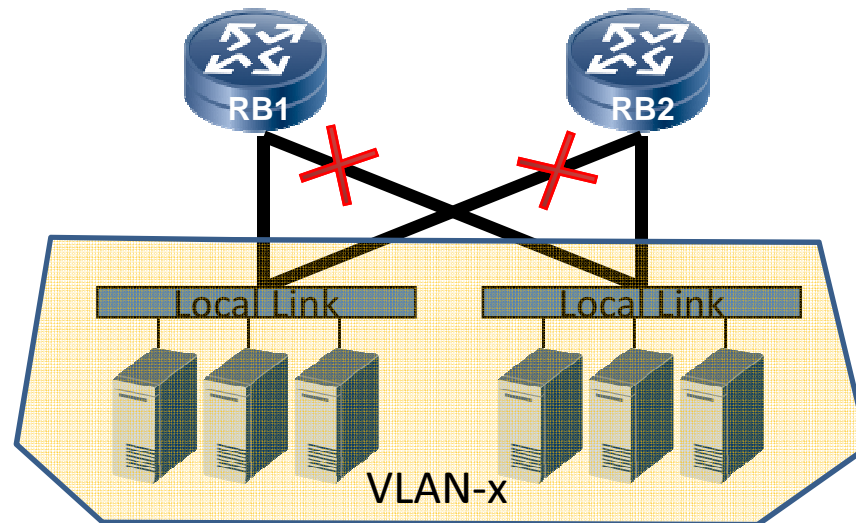


# RBridge Aggregation

Mingui Zhang, Donald Eastlake  
zhangmingui@huawei.com

# Single AF: Loop Avoidance

- To avoid loops involving native frames, TRILL allows only a single Appointed Forwarder for one VLAN on a local link.
- Single AF means single ingress & single egress for a specific VLAN.

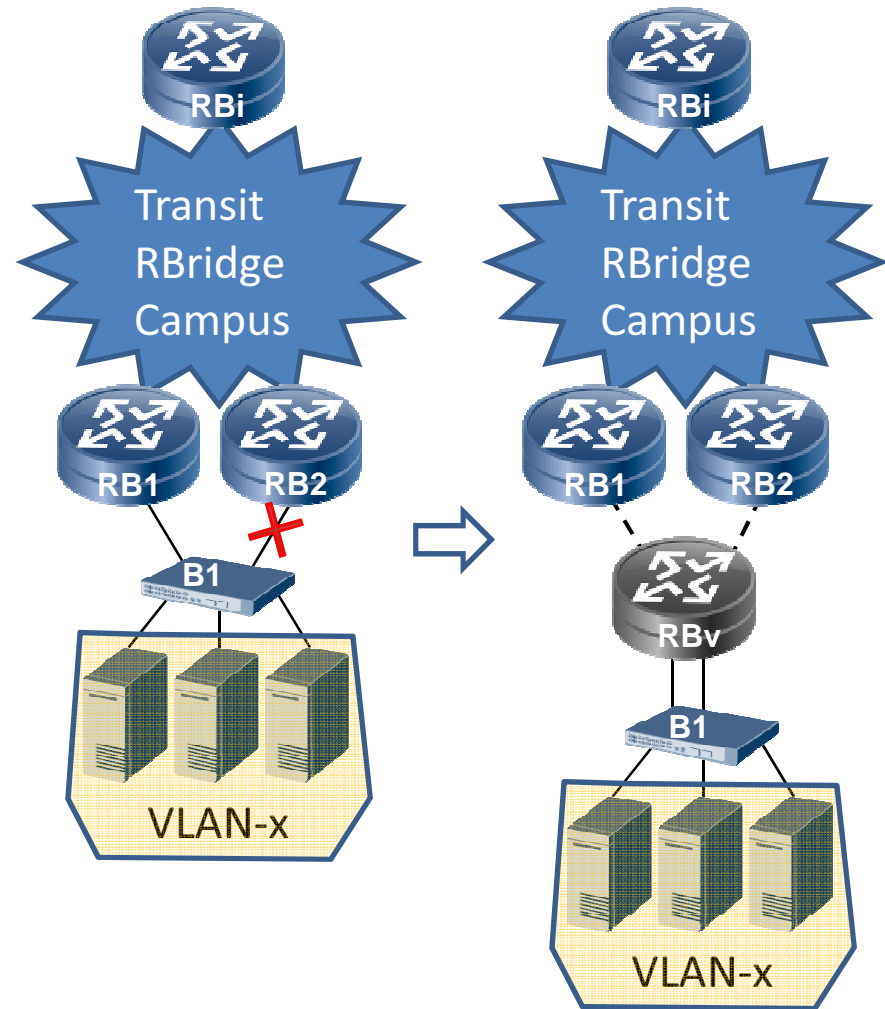


# Purpose of Aggregation

- Provide active/active multi-homing to a specific VLAN on a local link
- Increase the reliability of TRILL edge
- Increase the access bandwidth of RBridge campus

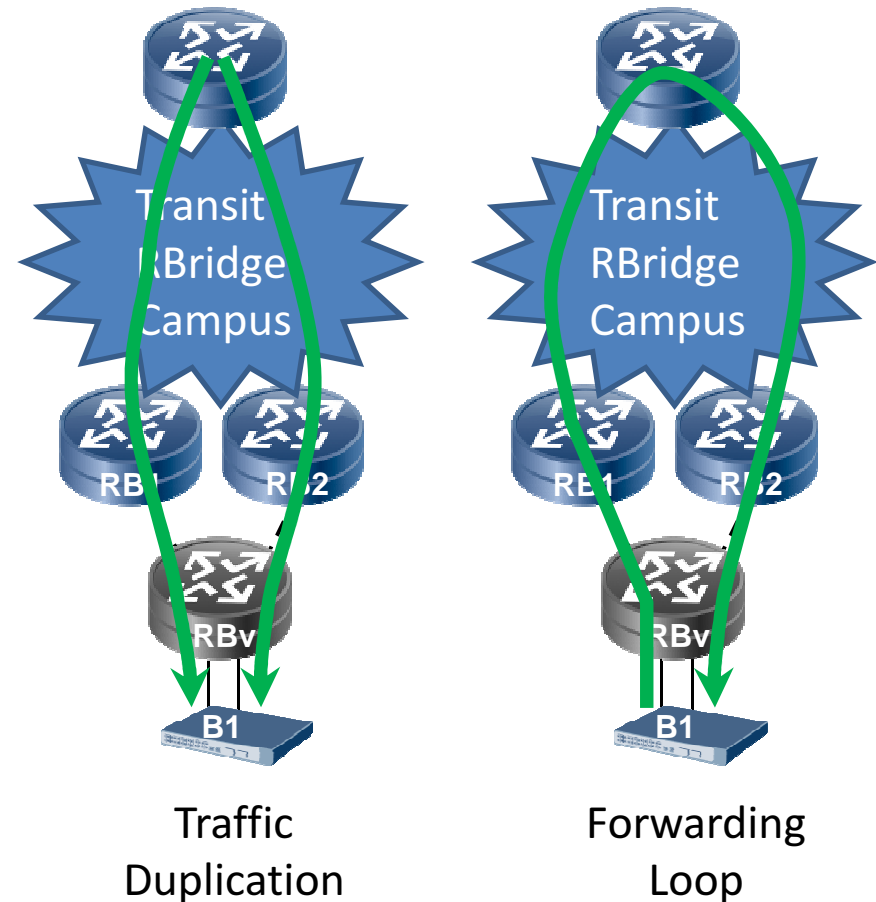
# RBridge Aggregation

- Two RBridges use one nickname, pseudonode nickname, to ingress frames.
- Advertise virtual links (e.g., RB1-RBv, RB2-RBv)
- They can forward data frame for VLAN-x at the same time (active-active).



# Two Possible Issues

- Traffic Duplication
  - Multicast frames egressed by both aggregated members
- Forwarding Loop
  - Multicast frames from the local link are egressed back to the local link by another aggregated member




# Link Aggregation [802.1AX] or Hashing

- The access links of the bridge are configured as link aggregation [802.1AX].
  - A frame is transmitted by one link one time
- Choose a single member using local hashing
  - A frame goes through one RBridge one time
  - Current version of our draft adopt this solution

# Frame Processing: Known Unicast

- Ingress
  - Set the nickname to the pseudonoe nickname “RBv”
  - Receiver ingresses the frame to the TRILL campus
- Egress
  - Receiver egresses the frame to the local link



# Frame Processing: Multicast

- Ingress

- If hashing matches

- Set the nickname to “RBv” and send it



- If not, drop it

- Egress

- If hashing matches

- send the frame



- If not, drop it





# Frame Processing: Unknown Unicast

- Ingress
  - The access bridge sends one copy of the frame with *unknown unicast blocking* technique
  - Aggregated RBridges treat the frame as unicast
- Egress
  - Same as multicast egress




# MAC Learning

- MAC addresses of the aggregated RBridges SHOULD be synchronized using ESADI.
- Before a MAC address is got from another aggregation member's ESADI, frames destined to this MAC address will be sent as unknown unicast.

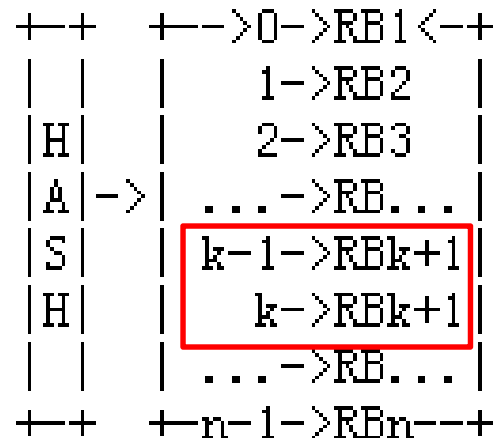
# Hashing Function Configuration

- An aggregation member can only send multicast frames with a specific hashing value.
  - E.g., RB1 only sends multicast frames that the last bit of their source MAC is “0”.



# Hashing under Link Failures

- When a connection to one aggregation member failed, the next member on the list takes the responsibility to send multicast frames for the aggregation.



Thanks!