Problem Statement: TRILL Active/Active Edge

draft-zhang-trill-aggregation-03.txt

Mingui Zhang

Donald Eastlake

Overview

- An MCLAG link is used to connect an edge device to multiple RBridges.
- The edge RBridge group is represented by a virtual RBridge RBv
- All member RBridges use RBv as the ingress nickname in TRILL data encapsulation to avoid the MAC move issue at the remote RBridges.

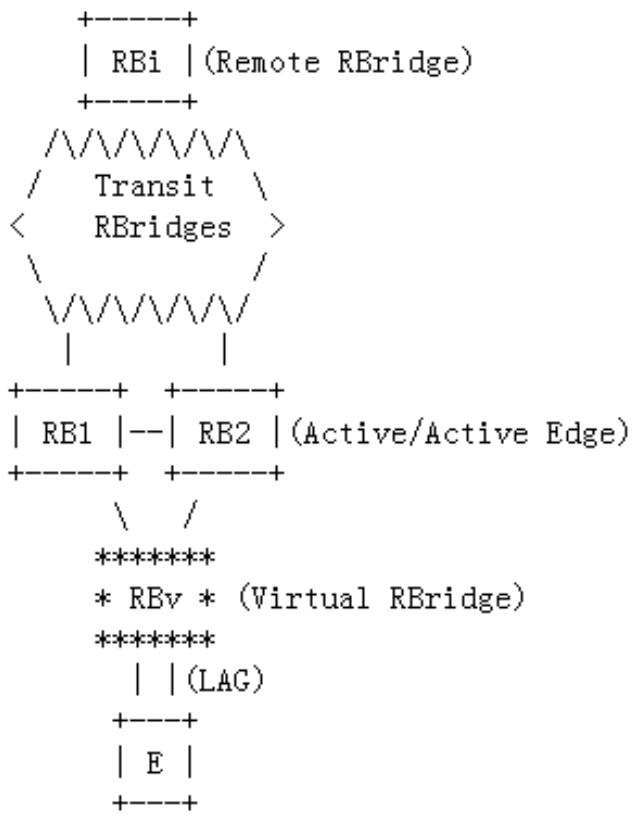

Purpose

- Provide edge devices with active/active connection to multiple RBridges
  - Increase the reliability of TRILL edge
  - Increase the access bandwidth of RBridge campus
- It's different from the Active/Standby connection for LAN links

Frame Processing: Unicast Ingressing

- Encapsulate native frames using RBv as their ingress nickname.

- Remote RBridges will regard the RBv as the egress RBridge for the edge node.

Frame Processing: Unicast Egressing

- Member RBridges should egress TRILL data frames whose egress nickname is RBv.

- Discard frames whose ingress nickname is RBv to avoid loops.

Frame Processing: Multicast Ingressing

- Member RBridges encapsulate the native frames using RBv as their ingress nickname.

- Member RBridges must not share a distribution tree to avoid the Reverse Path Forwarding Check issue.

Frame Processing: Multicast Egressing

- Only one member RBridge can egress a data frame. This avoids duplication.

- Discard frames whose ingress nickname is RBv to avoid loops.

DRB and Pseudonode

- Since there is no HELLO exchanging on the MCLAG link, member RBridges SHOULD have other signaling method to discover each other and elect the DRB.
- Each MCLAG link should be allocated with a pseudonode nickname, otherwise, Component Links from Different MCLAG Links Cannot be Distinguished by the same RBridge.

MAC Addresses Sharing

- Northbound Sharing
  - MAC address learnt from the local end node
- Southbound Sharing
  - MAC addresses learnt from remote RBridges
- MAC Addresses Sharing helps to reduce the multicast frames.

Failures and Self-healing

- The failure of a component link/node of the active/active edge group will be handled by ISIS.
- This provides the self-healing characteristic of the edge group.

Reverse Path Forwarding Check

- According to RFC 6325, RBridges MUST drop multicast frames that fail the RPFC.
- For the distribution tree on the left, multicast frames from RBv should only come from the port of RBi connecting to RB1.
- If RB2 uses this distribution tree to ingress multicast frame, these frames will be discarded by RBi.

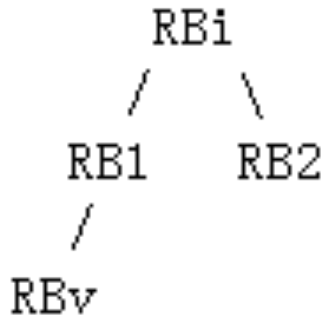

Reverse Path Forwarding Check (cont.)

- Leverage the feature that one RBridge can compute multiple Distribution Trees.
- Each member RBridge gets an unique distribution tree [CMT].
- This avoids the RPFC issue.

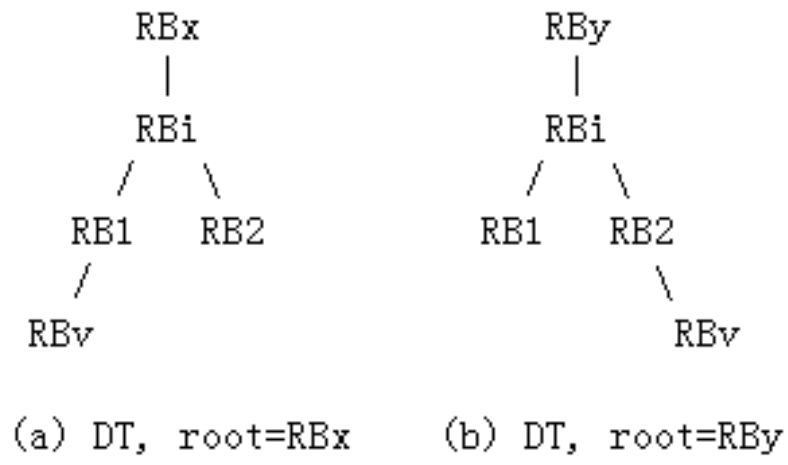

Next Step

- WG adoption.

Thanks!