

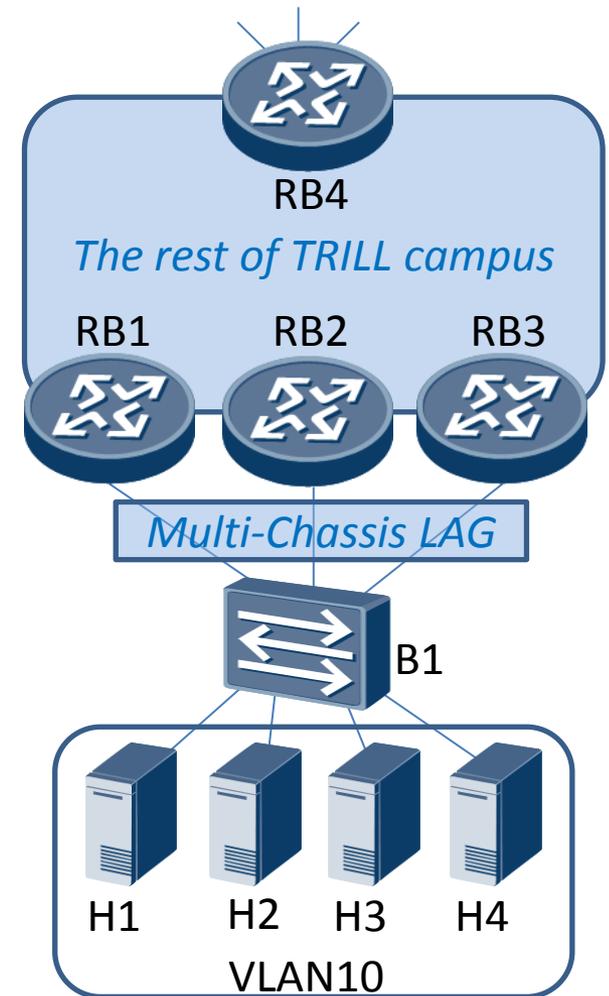
# TRILL Active-Active Edge (AAE) Using Multiple MAC Attachments

draft-zhang-trill-aa-multi-attach-04

Mingui, Radia, Hongjun, Muhammad, Mukhtiar, Sujay

# Overview

- AAE can balance the load of a VLAN among members.
- AAE can locally recover a link failure without triggering routing convergence.
- Issue: if RB1,2,3 ingress a packet using their own nicknames, RB4 will observe MAC flip-flop.



# To eliminate the MAC flip-flopping

- The remote RBridge keeps all MAC locations.
- It keeps using the closest one, so no flip-flop.
  - E.g., RB4 keeps using RB1 as the egress RB for H1.

OLD: RB4's MAC table

MAC	VLAN	Egress RB
H1	10	RB1/RB2/RB3

Egress RB of H1 keeps changing

NEW: RB4's MAC-table/LSDB

#	MAC	VLAN	Egress RB
1	H1	10	RB1
2	H1	10	RB2
3	H1	10	RB3

All egress RBs are remembered, while RB4 keeps using a single one.

# Discovery

- Member RBridges in the same group discover each other using the following TLV.

```
+++++
| Type = MC-LAG-GROUP-RBRIDGES | (2 bytes)
+++++
| Length | (2 bytes)
+++++
| Sender Nickname | (2 bytes)
+++++...+++++
| MC-LAG System ID (8 bytes) |
+++++...+++++
```

# Using ESADI to distribute MACs

- Utilize ESADI to distribute attachments of MACs.
- MAC Reach-ability TLV [RFC6165] is included as a sub TLV of the MC-LAG Group MAC TLV.

```
+++++
| Type = MC-LAG-GROUP-MAC          | (2 bytes)
+++++
| Length                            | (2 bytes)
+++++...+++++
| MC-LAG System ID                 | (8 bytes) |
+++++...+++++
| MAC-Reachability TLV             | (7 + 6*n bytes) |
+++++...+++++
```

# Remember multiple attachments

- The remote RBridge remembers multiple attachments in the LSDB.
- But it always installs a single one into the FIB. E.g., the closest one, breaks the tie according to a pseudo-random selection method

LSDB

#	MAC	VLAN	Egress RB
1	H1	10	RB1
2	H1	10	RB2
3	H1	10	RB3



FIB

#	MAC	VLAN	Egress RB
1	H1	10	RB1

# Capability

- The remote RBridge advertises the capability to remember multiple locations of a MAC in its LSP.

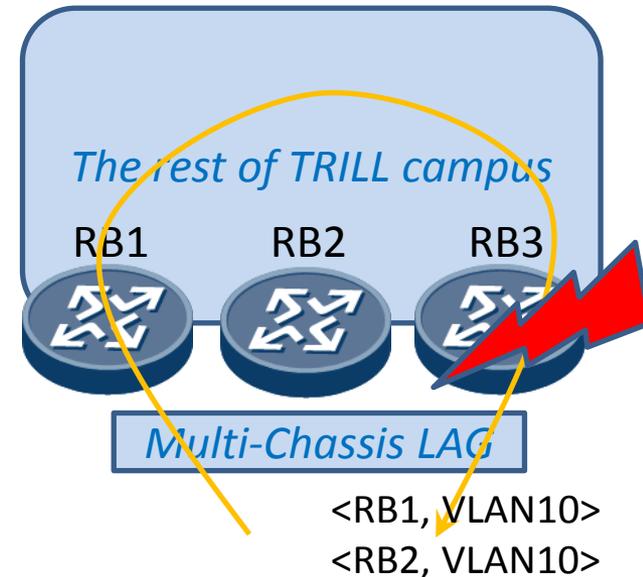
```
+-----+
| Type = MULTI-MAC-ATTACH-CAP | (2 bytes)
+-----+
| Length | (2 bytes)
+-----+
|E|H| Reserved | (1 byte)
+-----+
```

# No duplication (single exit point)

- When a multicast is going to be egressed onto the MC-LAG, only one of the AAE members can act as the exit point (otherwise traffic duplication will happen).
- It utilizes the hashing function of MC-LAG to determine the single exit point.
  - If the output of the hashing function points to the port attached to the receiver RBridge itself (i.e., the packet should be egressed out of this node), it egresses this packet. Otherwise, the packet **MUST** be dropped.

# No Echo (Split Horizon)

- Echo of a multicast frame will cause a forwarding loop.
- Split horizon is the answer.
- The key for split horizon
  - ingress nickname+VLAN
- It's set up through the discovery.
  - The “MC-LAG Group RBridges TLV” includes the sender nickname for the group identified by the MC-LAG System ID.



RB3 splits horizon to prevent the forwarding loop.

# No black-hole

- If a link in the MC-LAG fails, the corresponding RBridge sends out MAC withdraw message to flush the affected MAC addresses remembered by remote RBridges. This prevents the black-hole issue.

# Load Balance Towards the AAE

- The remote ingress RBridge may install any one of the attachments from the LSDB to the FIB according to the pseudo-random selection method. This provides the load-balancing.
- It achieves a south-bound flow-level traffic spreading.

Thanks!