One person's comparison of TRILL and MAC-in-MAC

Rev. 1

Norman Finn

nfinn@cisco.com

References

- This document: http://www.ieee802.org/1/files/public/docs2011/new-nfinn-TRILL-MinM-0611-v1.pdf.
- TRILL: http://www.ietf.org/rfc/rfc5556.txt http://datatracker.ietf.org/doc/draft-ietf-trill-rbridge-protocol/?include_text=1
- MAC-in-MAC: http://standards.ieee.org/getieee802/download/802.1ah-2008.pdf
- This work is the opinion of the author. It does not represent the opinion of this station, its owners, network, or affiliates.

Control plane

- TRILL control plane: IS-IS
- M-in-M control plane: IS-IS

Network diagram

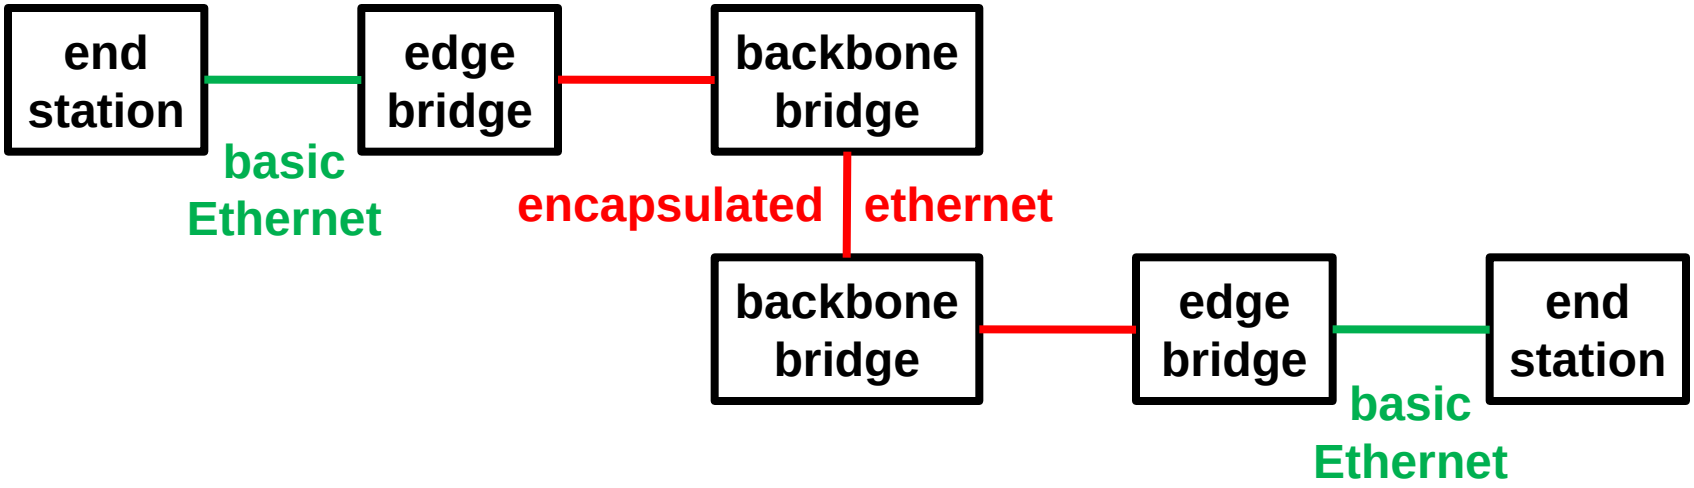

- TRILL: station – encapsulation – backbone – decapsulation – station
- M-in-M: station – encapsulation – backbone – decapsulation – station

Learning

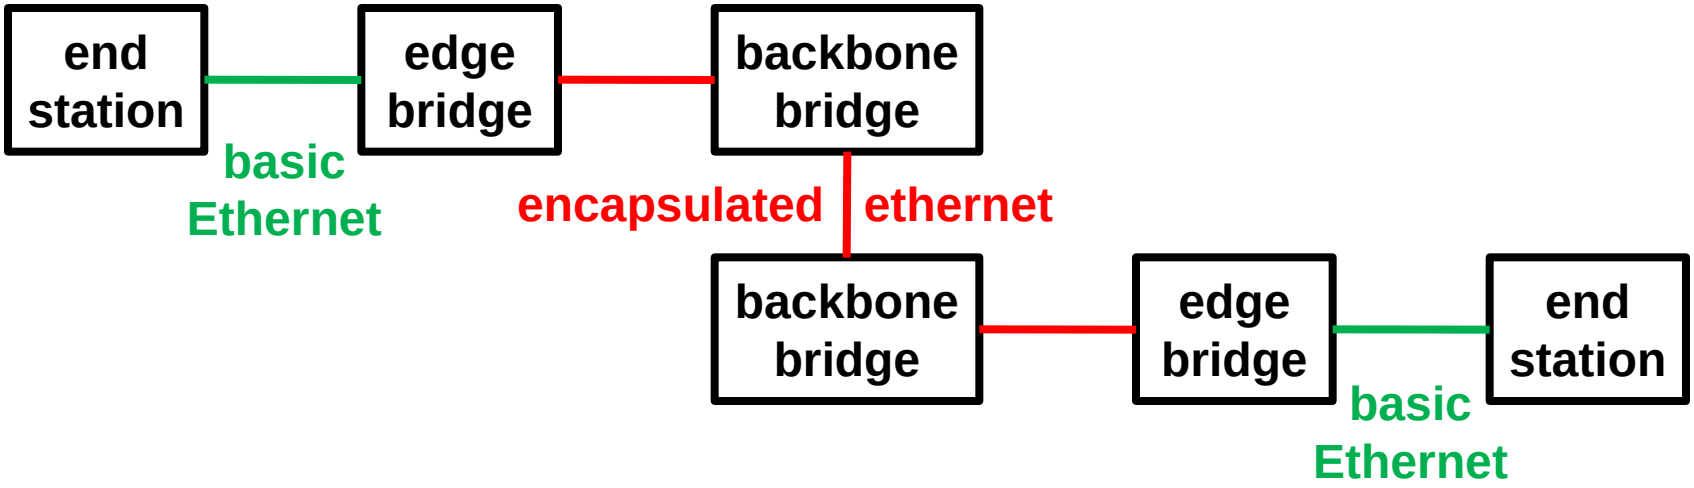

- TRILL: Destination edge bridge learns association between customer source address and source edge bridge identity.
- M-in-M: Destination edge bridge learns association between customer source address and source edge bridge identity.

Learning

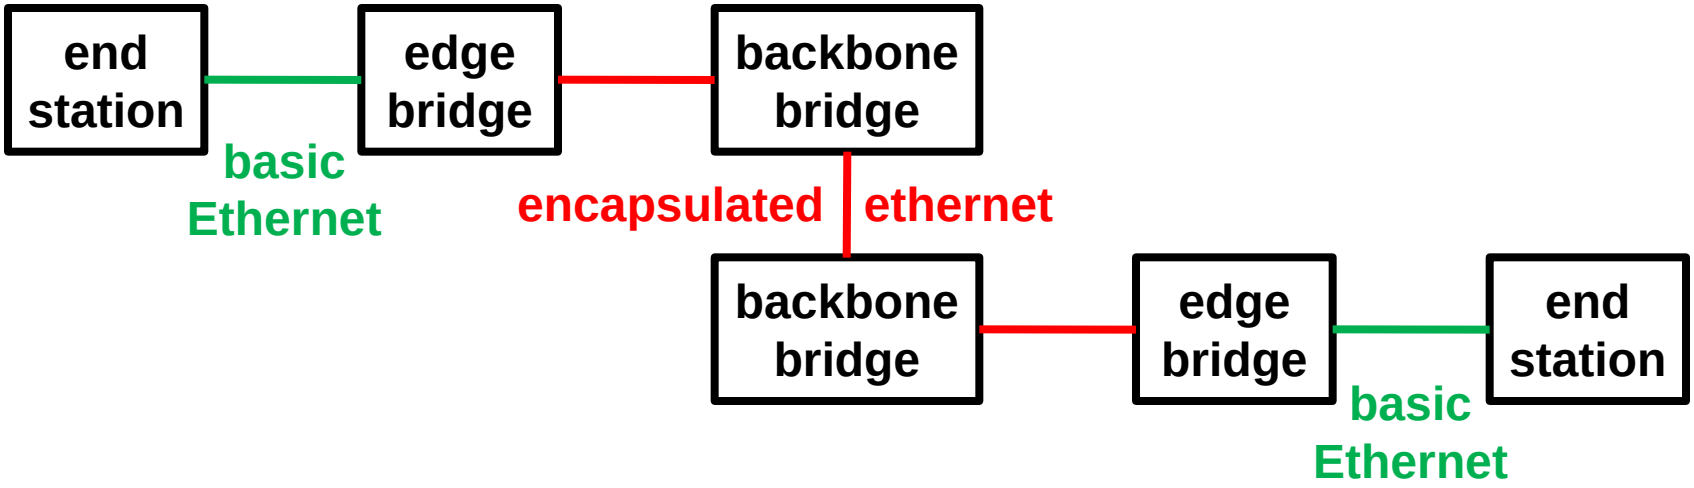

- TRILL: No learning of backbone (edge bridge) addresses.
- M-in-M: No learning of backbone (edge bridge) addresses.

High-level frame formats

M-in-M

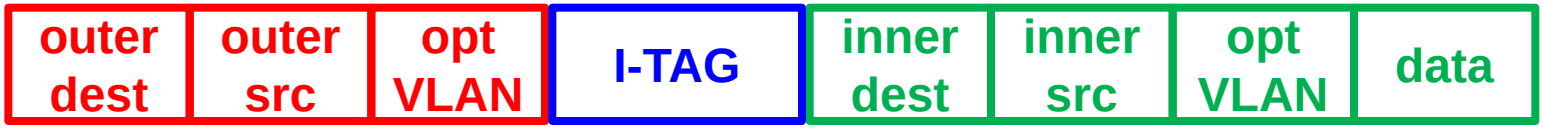

TRILL

- They both have an outer and an inner MAC address, they both encapsulate a customer frame, they both use an IS-IS control plane, so they are **basically the same**.
- **Well ... not exactly**

TRILL format

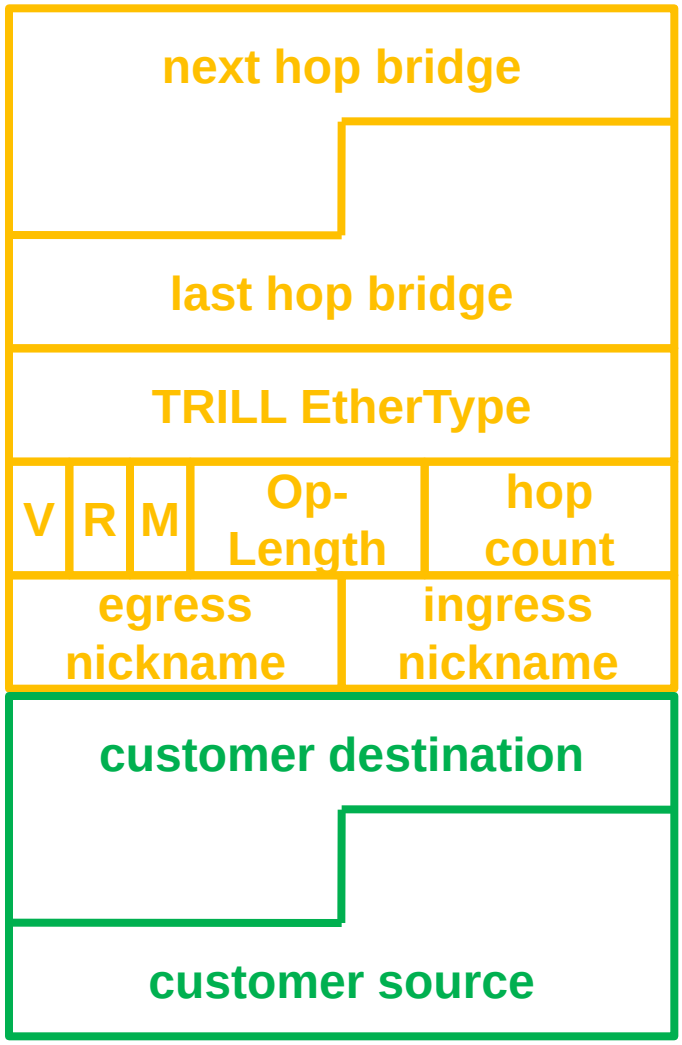

M-in-M format

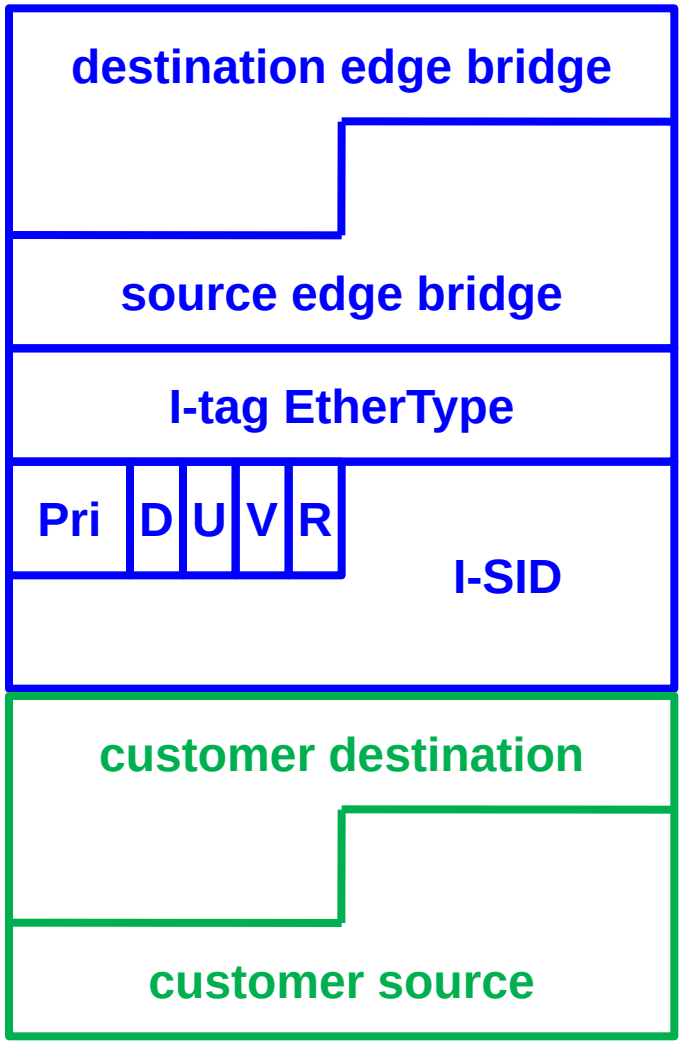

(omitting the VLANs)

What are the essential differences?

- TRILL is **routing**, and M-in-M is **bridging**, in the sense that:
  - A TRILL packet carries an explicit next hop identifier to prevent multiple deliveries, and a hop count field to mitigate forwarding loops.
  - M-in-M utilizes interlocked routing protocol states to make use of shared media links and to prevent multiple deliveries or forwarding loops.
- The other differences (or lack thereof) depend on whether you are asking about **unicast** forwarding or **multicast** forwarding.

What's the difference? Unicast

- What does a backbone forwarding function use to select the output port?

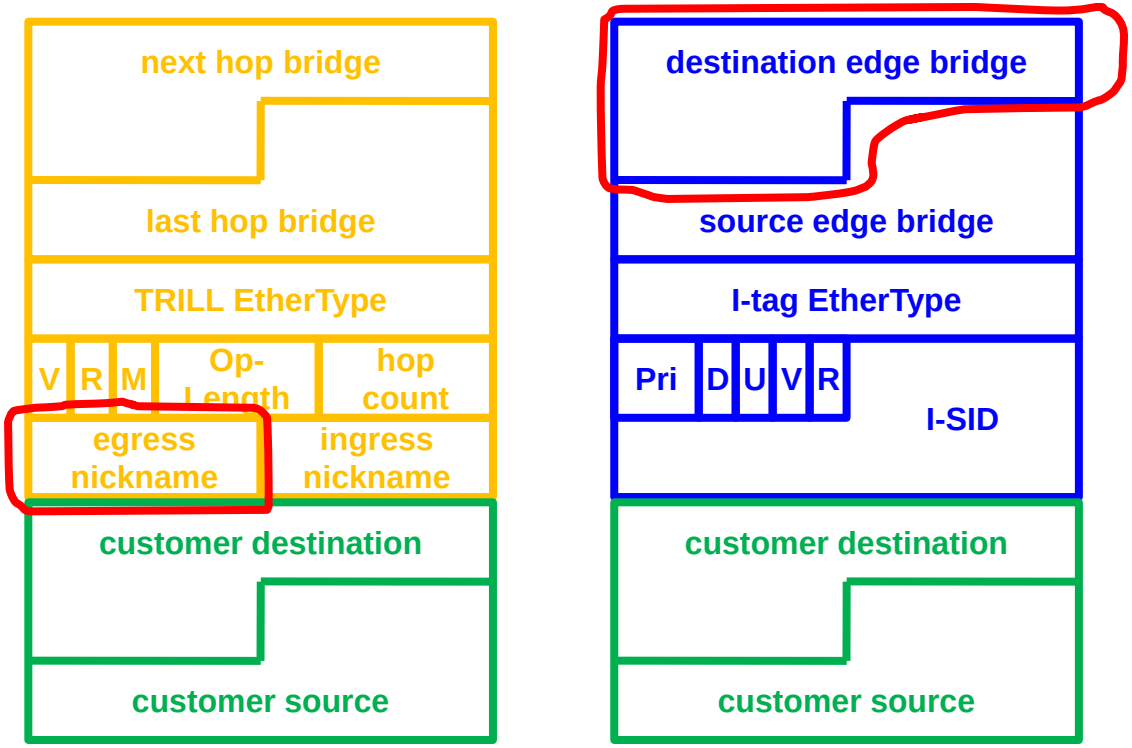

- **No fundamental difference**

What's the difference? Multicast

- What does a backbone forwarding function use to select the output port?

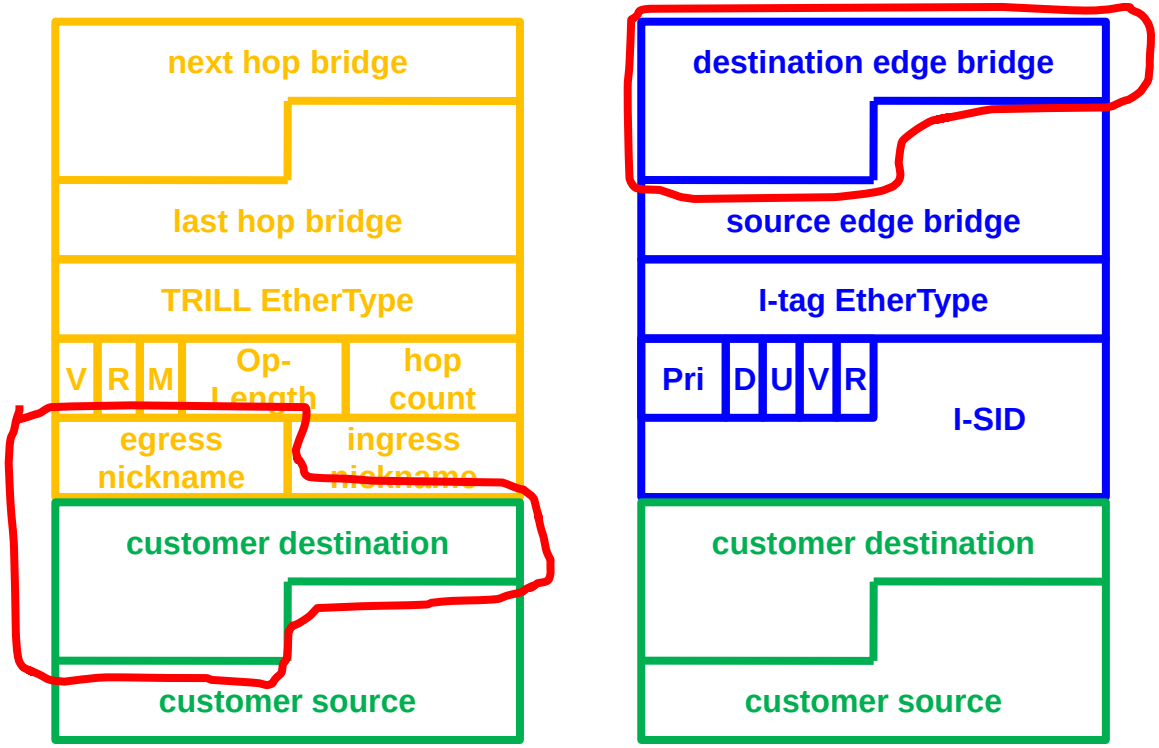

- **Fundamental difference**

What's the difference?

- For **unicast** frames, the fact that the customer destination address is not used by TRILL forwarding makes it reasonable to say that the nickname is equivalent to the M-in-M outer address, so the two really are the same.
- For **multicast** frames, TRILL uses the customer destination MAC address across the backbone, not just at the edges. M-in-M uses only the outer MAC address. The two are really different.

Multicast difference

- TRILL uses the egress nickname to carry an explicit tree ID. The TRILL backbone switches must have multicast state based on (at least) customer destination address, VLAN, and tree ID in order to prune.
- In M-in-M, backbone switches do not use the customer address; they prune based on the outer multicast destination address (and B-VLAN), which can be unicast or multicast.
- The M-in-M backbone is **isolated** from the customer addresses and VLAN; the TRILL backbone is not.
- The relationship between customer VLAN/multicast MAC address and backbone MAC address must be **managed** in M-in-M, but not in TRILL.

Summary

- There are many minor differences.
- Routing vs. bridging (TRILL next hop + hop count vs. M-in-M interlocked protocol state) is a significant difference.
- Multicast delivery (TRILL customer address vs. M-in-M encapsulation) is a significant difference.