Directory Assisted TRILL Encapsulation on non-TRILL nodes (Directory Reliant Smart End Node)

Linda Dunbar
Donald Eastlake
Radia Perlman
Igor Gashinsky

History

- 2011: stripped from “draft-dunbar-trill-directory-assisted-edge”
- March 2013:
  - revised to reflect progress made by “draft-ietf-trill-directory-framework” and “draft-dunbar-trill-scheme-for-directory-assist”

Motivation

- Reduce the amount of MAC&VLAN <-> RBridgeEdge mapping maintained by RBridge edge.
  - ToR with 40 servers: 40*20VMs=800 hosts, Worst case: 800 VLANs →800 *200 (hosts/VLAN) = 160000 entries
- Allow all ports connected to a Bridged LAN to forward TRILL traffic
  - Today: Unknown DA frames are flooded to both TRILL domain and Ethernet domain

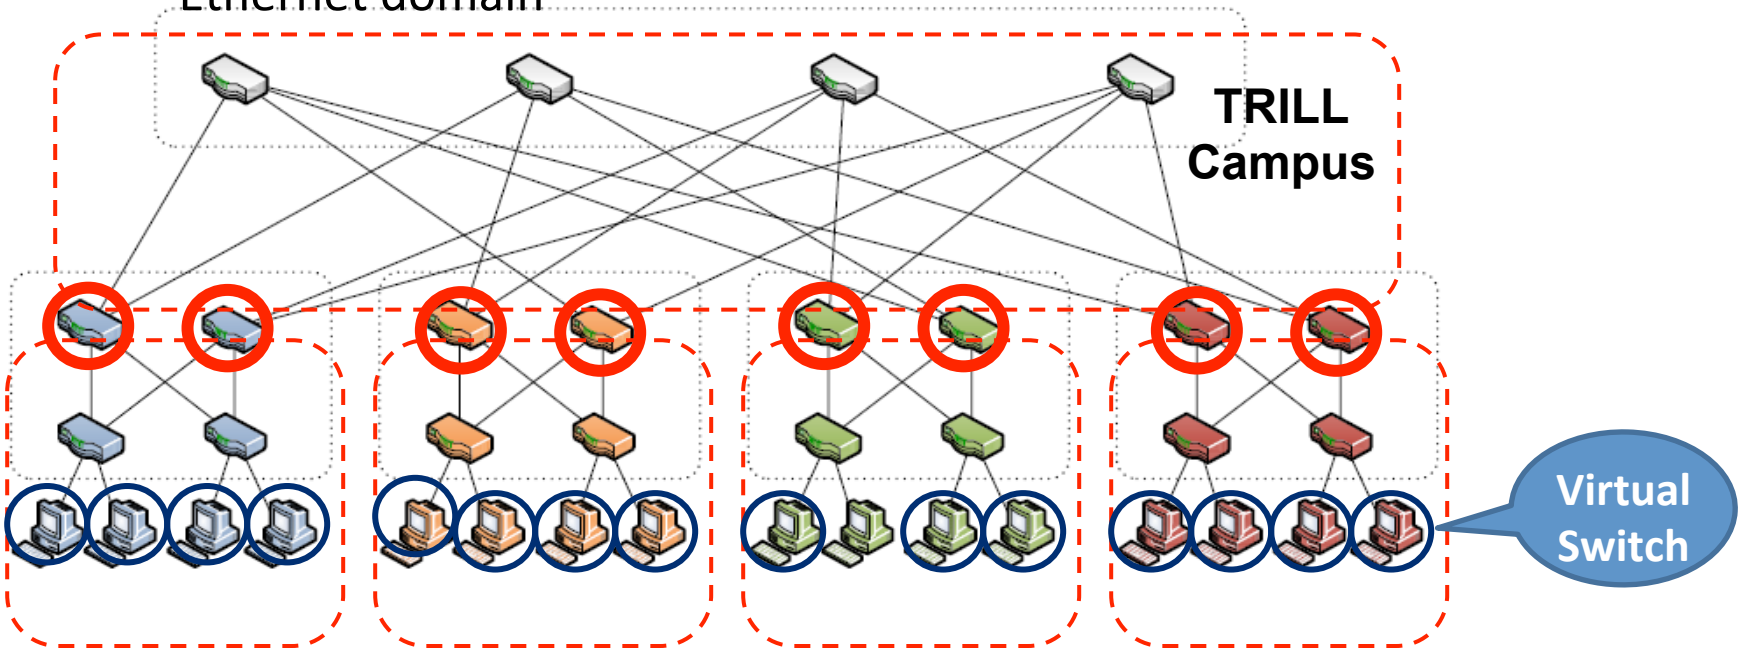

Detailed Mechanism

- All hosts under RBridge use RBridge's nickname
- Ingress RBridge: forward TRILL encapsulated frames to TRILL campus. Native Ethernet frames stay locally (or pull directory for unknown DA)
- Egress RBridge: normal processing

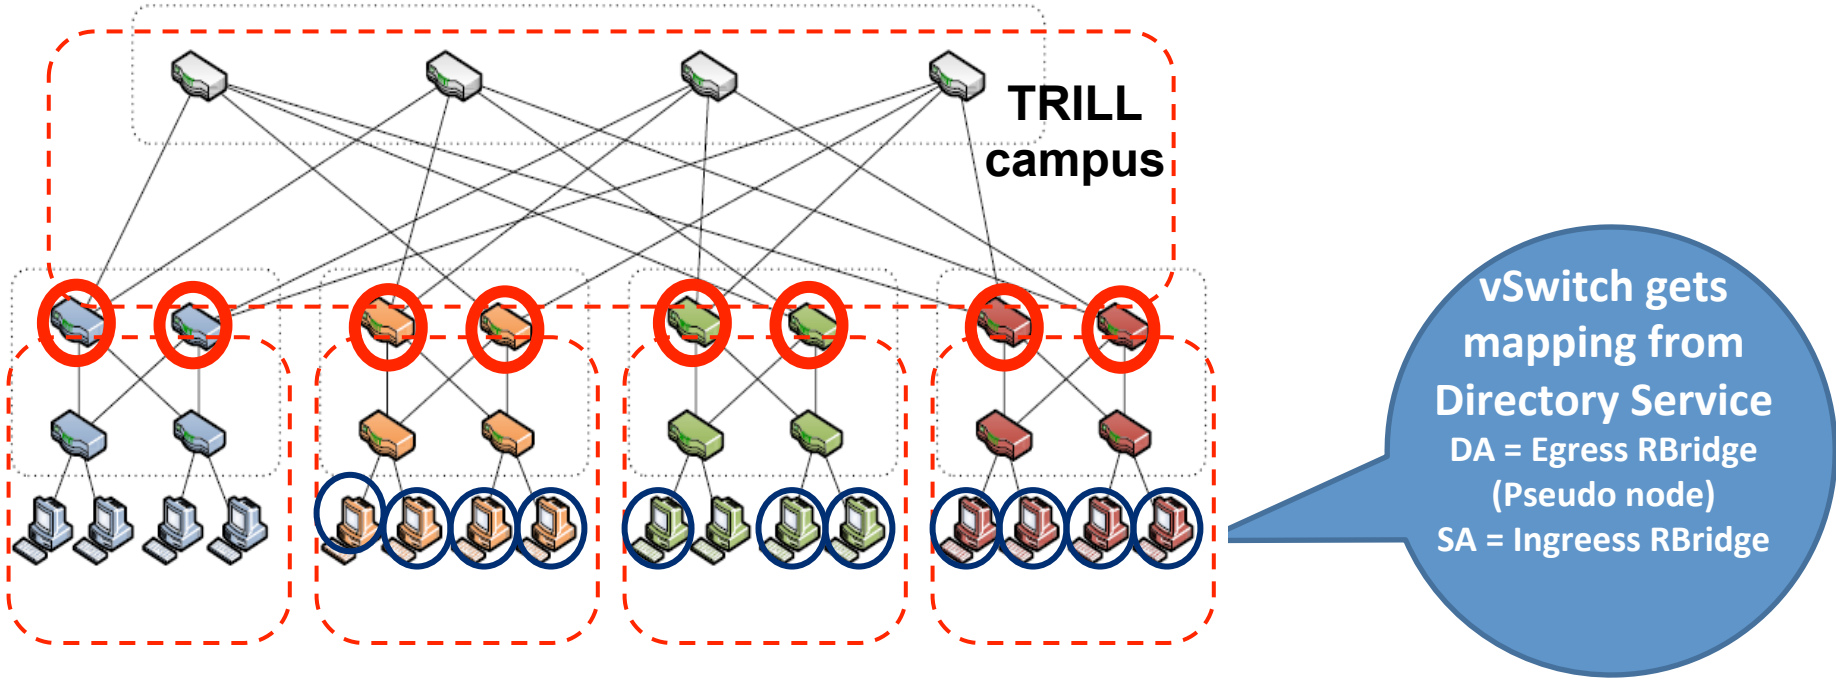

Local Decision

Choose any of directly attached Edge RBridge as Ingress SA (local decision, same as ECMP)

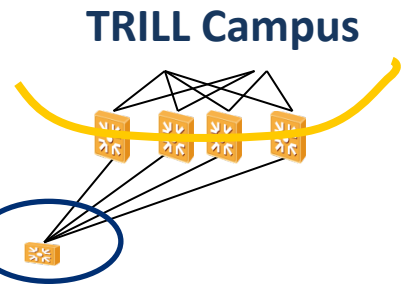

- Directory:
  - End node attached to {R1, R2, R3, R4}

Benefits

- Simplified processing on Ingress RBridge
  - Utilize directory service on virtual switches
- No change to edge RBridge processing
- Allow all ports to forward TRILL frames (i.e. no need to designate one port as AF port)

Next Step

- Need more comments from WG
- WG adoption?