

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

draft-dunbar-trill-directory-assisted-encap-08


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”
- ” July 2013:
 - . Emphasized the benefits of pre-encapsulation
 - . Described why/how pre-encapsulation can limit MAC Address table size as VLANs getting more fragmented.
- ” March 2014:
 - . Wording refinement to reflect the changes of RFC7067 and the directory scheme draft
- ” Oct 2014:
 - . Revised the draft to reflect changes made in draft-ietf-trill-directory-assist-mechanisms-01

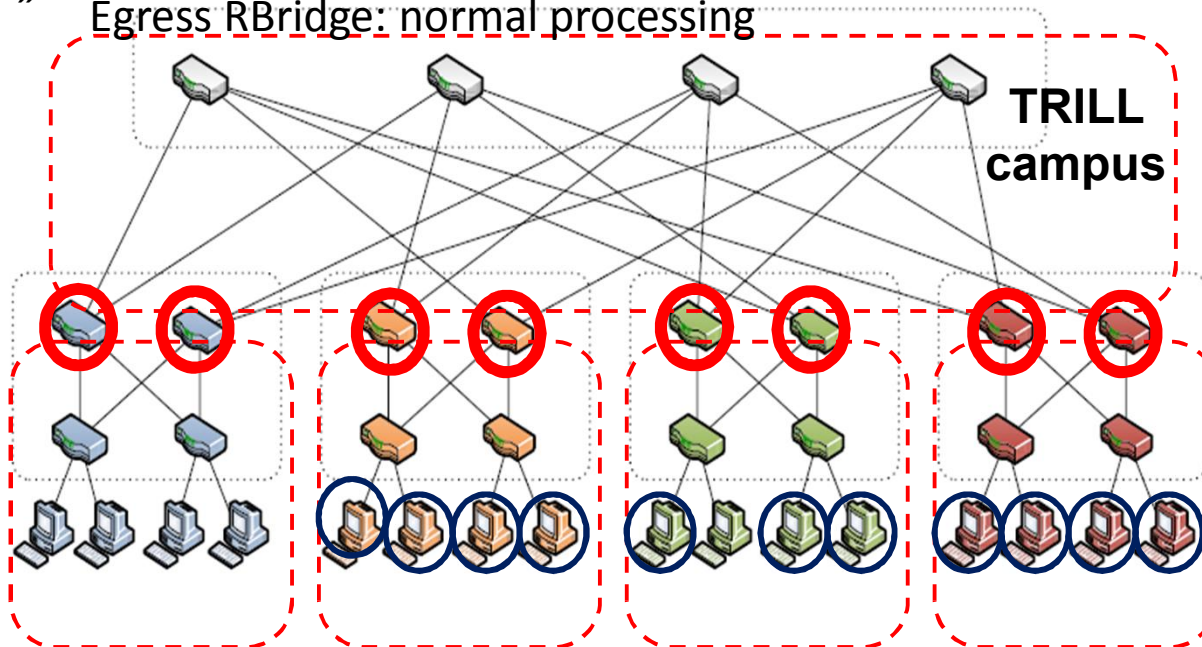
Key Benefits

- “ **Reduce MAC learning table size for Rbridge Edges and switches on Bridged LANs**
 - “ Local switches don't need to learn MAC addresses of hosts in remote domains.
- “ **Better load spreading (Ingress ports) because the VM/hypervisor can send pre-encapsulated data frames to any edge Rbridges (avoid Appointed Forwarder)**
- “ **Avoid Nickname Exhaustion issue in Data Centers with very large number of VMs and servers.**



Detailed Mechanism

- “ End node (e.g. Virtual switch on servers) uses Pull Directory to determine if a packet should be sent across TRILL domain to another Rbridge egress node.
 - “ If yes, pre-encapsulate the frame with TRILL header and send to the local ingress Rbridge.
 - “ All nodes under RBridge use a common local ingress 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



vSwitch gets mapping from Directory Service
DA = Egress RBridge
SA = Ingress RBridge

Local Decision



“ Directory:

“ End node attached to {R1, R2, R3, R4}

Benefits

- “ Local switches don't need to keep remote hosts MAC addresses in their FDB.
- “ Allow all ports to forward TRILL frames (i.e. no need to designate one port as AF port)
- “ Simplified processing on Ingress RBridge
 - . Utilize directory service on virtual switches
- “ No change to edge RBridge processing