Revision of the Appointed Forwarder RFC

draft-eastlake-trill-rfc6439-00.txt

Donald E. Eastlake, 3rd
d3e3e3@gmail.com
+1-508-333-2270

Background

- TRILL base protocol supports multiple TRILL switches (RBridges) on a link providing end station service where traffic is allocated among the RBridges by VLAN.
- The RBridge handling native end station traffic to/from the link in VLAN-x is called the Appointed Forwarder for VLAN-x
- The current details are in RFC 6439 which was updated by RFC 7180.

Background

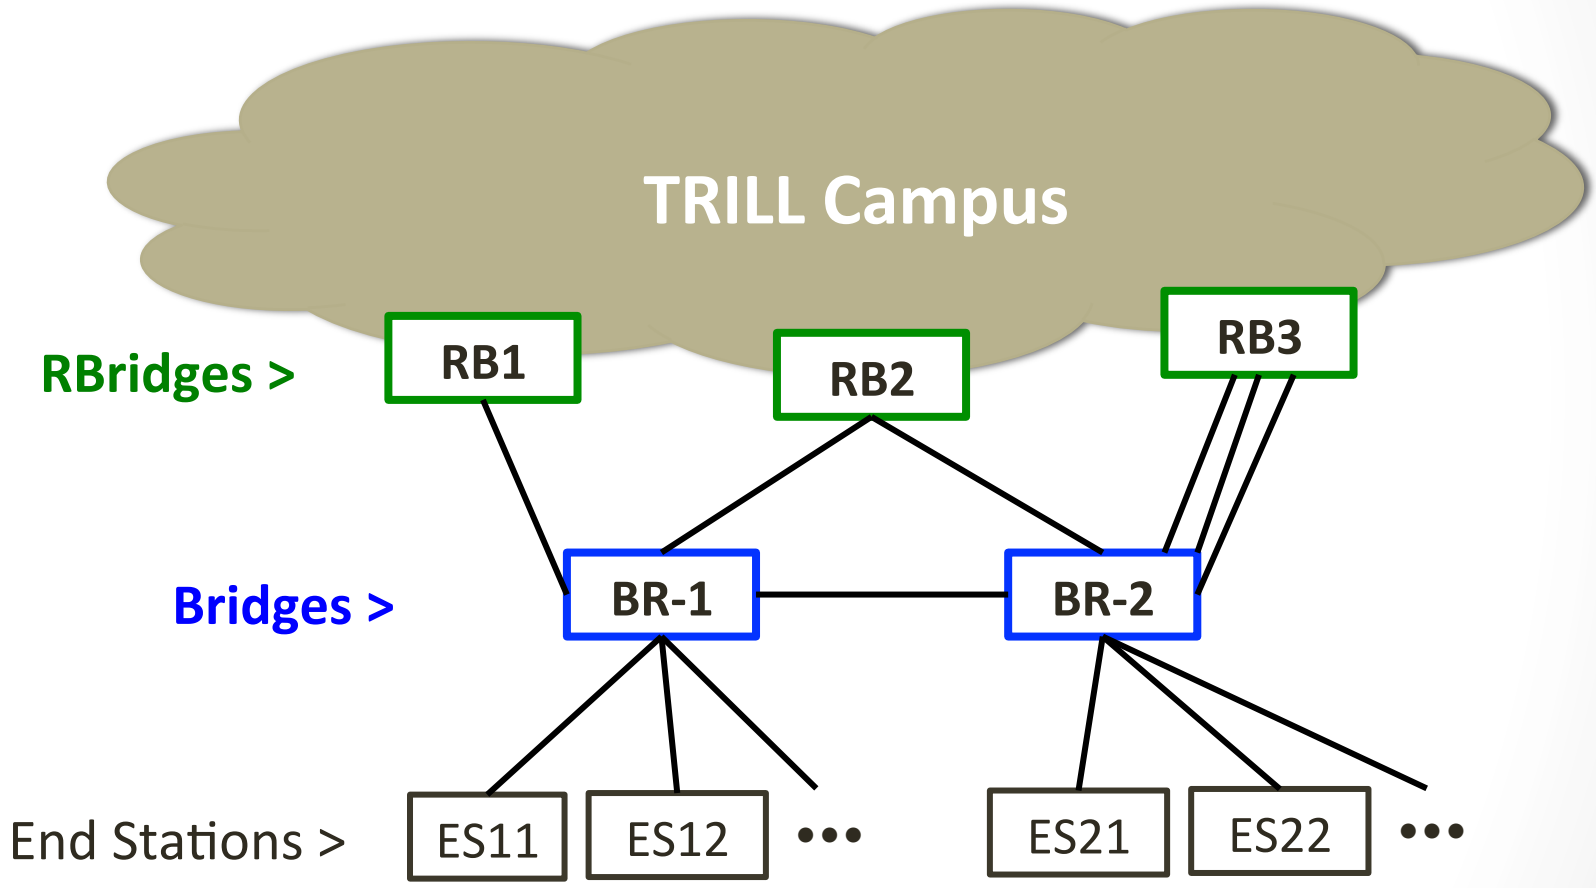
TRILL Campus
RBridges >
RB1
RB2
RB3
Bridges >
BR-1
BR-2
End Stations >
ES11
ES12
ES21
ES22

Current rfc6439bis Draft

- Add mandatory E-L1CS FS-LSP support. The current rfc6439bis draft requires support of Extended Level 1 Circuit Scoped (E-L1CS) Flooding Scope LSPs as specified in RFC 7356.
  - Expands space for link local state and configuration advertisements from the very limited space in the IS-IS Hello up to 2**16 link local LSP fragments.
  - This eliminates the current restrictions on appointing forwarders on a link due to lack of space in a Hello for appointment information.

Current rfc6439bis Draft (cont.)

- Specifies APPsub-TLVs for sending efficiently encoded forwarder appointments in E-L1CS in cases other than contiguous blocks of VLANs.
  - Blocks of contiguous VLANs efficiently encoded by Appointed Forwarders sub-TLV [RFC7176].
  - Scattered VLANs over a narrow range efficiently encoded by Appointment Bitmap APPsub-TLV in rfc6439bis.
  - Scattered VLANs over a large range efficiently encoded by Appointment List APPsub-TLV in rfc6439bis.
- Incorporates Appointed Forwarder material from RFC 7180 that optionally reduced the number of VLANs in which Hellos must be sent.
- Adds optional FGL mapping consistency check on a link.

Further Possible Improvements

- Two further optional improvements that could be incorporated into this draft:
  1. Port failure/shutdown notification
  2. Root bridge change optimizations

- There is IPR on these improvements that will be disclosed.

1. Port failure/shutdown notification

- An RBridge may have a planned shut down of a port on a link or might detect port failure at the physical or other level.
- Port failure is not currently noticed by other RBridges until multiple Hellos from that port are missed, which can take many seconds.

1. Port failure/shutdown notification

- The proposal is to send an RBridge Channel message indicating port failure/shutdown.
  - Message is serially unicast through the TRILL campus to other RBridges on the link and/or (in the case of shutdown) broadcast on the link.
  - The port can be unambiguously identified by the RBridge's IS-IS System ID and the port's Port ID.
  - This message can be secured using the RBridge Channel Tunnel protocol.

1. Port failure/shutdown notification

- Other RBridges on the link, on receipt of this message, drop adjacency to the port that failed or is being shutdown.
  - If that RBridge/port had been elected Designated RBridge (DRB) on the link, a new DRB is elected.
  - If that RBridge/port was Appointed Forwarder, a new forwarder can be appointed immediately .
- [Similar steps can be taken if an RBridge on the link crashes but no new message is needed in that case as this is visible via LSPs.]

2. Root bridge change optimizations

- Many links are point-to-point but, in the case of multiple RBridges and end stations on a link, are likely to be bridge LANs.
- Bridges and links can go down or come up. As a result a bridged LAN could partition or bridged LANs could merge.
- While bridging protocols are settling for the new bridge LAN topology, there could be multiple VLAN-x Appointed Forwarders, TRILL Hellos can be blocked, etc..

2. Root bridge change optimizations

- As a result, when a CIST root bridge change is detected by an RBridge port, current TRILL protocol requires that the port be “inhibited” for a configurable fixed length of time.
  - Note: “Inhibition” has no effect on TRILL Data or IS-IS packets. It only stops ingress and egress of native frames from/to end stations.
- There are at least three cases where this is **not** necessary.

2. Root bridge change optimizations

2. Three cases where the CIST root Bridge Identifier changes but inhibition is unnecessary or can be terminated early:
   a) If the Bridge Identifier changes to a lower priority value.
   b) If the the root Bridge Address stays the same.
   c) If, since the Identifier change, Hellos have been received from all known RBridges on the bridged LAN.

2. Root bridge change optimizations

a) If the Bridge Identifier changes to a lower priority value.
   - This implies that a bridge or link crashed, possibly partitioning the bridge LAN. In this case, no inhibition is required.
   - A change is unsafe only for an RBridge port that sees a new higher priority CIST root Bridge Identifier.

2. Root bridge change optimizations

b) If, when the CIST Bridge Identifier changes, the Bridge Address part stays the same.
   - This implies that the only thing that happened was that the other bits in the Bridge Identifier have been configured to a different value but he bridge that was CIST root is unchanged.
   - In this case, no inhibition is required.

2. Root bridge change optimizations

c) If, since the Identifier change, Hellos have been received from all known RBridges on the bridged LAN.

- RBridges are required to report the CIST root Bridge Address that they see in their LSPs. All RBridges attached to the same bridged LAN will see the same root.
- Either there is connectivity outside the bridge LAN to another RBridges on the bridge LAN or not.
  - In the first case, you will know about the other RBridge and can wait for a new Hello or the inhibition timeout.
  - The second case, where connectivity is only through the bridged LAN, is loop safe.
- So, inhibition can be terminated as soon as new Hellos have been received from all other RBridges on the bridged LAN as known through LSP root Bridge Address reports.

END

Donald E. Eastlake, 3rd
d3e3e3@gmail.com
+1-508-333-2270