


Multi-Level TRILL

draft-perlman-trill-rbridge-multilevel-10.txt

Radia Perlman, Donald Eastlake, Mingui Zhang,
Anoop Ghanwani, Hingjun Zhai

Multi Level

- TRILL is specified as a single level IS-IS routing domain.
- draft-perlman-trill-rbridge-multilevel-10 is an Informational draft surveying alternatives for multi-level TRILL. This presentation just covers some highlights from the draft.
- Multilevel TRILL is built on multi-level IS-IS.



Advantages of Multi-Level

- Reduced routing computation load
 - With optimizations, from $N \cdot \log(N)$ to $\sqrt{N} \cdot \log(N)$, for N routers
- Reduced link state volatility at each router
 - Less control traffic
 - More often in a converged state
- Reduce the size of the link state database at each router
- Eliminate the limit on the number of TRILL switches due to 16-bit TRILL nicknames
- Can be used to build facilities to reduce some of the load from multi-destination traffic

One Implementation Choice

- An important choice is whether all TRILL switches in a multi-level TRILL campus have unique nicknames or not.
 - **Unique nicknames:** Each TRILL switch has a nickname unique across the campus. Easier to implement, no special modifications to TRILL data packets at border RBridges. But, you cannot have more TRILL switches than will fit in the 16-bit nickname space.
 - **Aggregated nicknames:** Nicknames can be re-used in different Level 1 areas. Level 1 areas are represented by a nickname or small set of nicknames in Level 2. Requires modification to the TRILL header in TRILL data packets as they pass through border RBridges. Virtually eliminates the nickname space constraint on the number of TRILL switches in the campus.

One Implementation Choice

- This draft suggests a hybrid scheme where some Level 1 areas can use unique globally visible nicknames while other Level 1 areas can be aggregated.
- For example, as a TRILL campus grew and became multi-level, it could initially use the simpler unique nicknames. Then, as it continues to grow, some Level 1 areas could be aggregated and eventually most/all Level 1 areas could be aggregated.

A Few Other Issues / Features

- Area Addresses: TRILL currently specifies a fixed area address of zero . The draft surveys alternatives for area addresses in a multi-area TRILL campus.
- Distribution Trees: TRILL uses trees for the distribution of multi-destination TRILL data packets. There are a variety of ways to handle the building of trees and any transition between trees at the Level 1 / Level 2 boundary.
- The partition of a TRILL campus into multiple areas can be the used to limit the broadcast domain for some protocols to an area by blocking the protocol at the area boundary and including a proxy server in the area.

A Few More Issues

- Co-existence with legacy TRILL switches: This is somewhat difficult, partly due to nickname reservation problems, but alternatives are discussed in the draft.
- Multi-access links with end stations: You could have end stations and multiple TRILL switches from different Level 1 areas on the same link. A simple rule is needed to determine which area's TRILL switches will handle traffic to/from such end stations. The draft suggests that the lowest area address wins.

End

- This draft is currently in call for WG adoption.
- Questions?