

TRILL Resilient Distribution Trees

draft-zhang-trill-resilient-trees-01.txt

Mingui Zhang

Tissa Senevirathne

Janardhanan Pathangi

Ayan Banerjee

Anoop Ghanwani

Backup Tree Roots Selection

- Before a backup DT is calculated, its tree root should be selected, there are two options
 - Option 1: use the same root as the primary DT
 - Root RBridge owns multiple nicknames
 - Option 2: RB1 which has the highest root priority nickname explicitly advertise the backup DT root.
 - Tree Identifiers Sub-TLV

Backup DT Calculation using Distributed Algorithm

- When backup DT is calculated, the weights of the links on the primary DT should be increased by a multiplier of 10x or 100x.
- This will ensure that the backup DT does not use any link on the primary DT unless it is the cut link.
- Each RBridge can independently run this algorithm. No configuration is necessary.

Backup DT Installation

- TRILL switches can intrinsically calculate multiple trees, but they do not surely install all trees.
 - As specified in RFC 6325 Section 4.5.2, an ingress RBridge MUST announce the distribution trees it may choose to ingress multicast frames. Thus other RBridges in the campus can limit the amount of states which are necessary for RPF check.
- For RDT, the announced distribution trees MUST include backup DTs, therefore they can be installed in advance by other RBridges.

Switching Back to the Primary Distribution Tree

- When the backup DT is used, the convergence of the primary DT is going on in parallel. Each RBridges should continue to use the backup DT until the convergence finishes.
- For safety, a timer is set upon the arrival of LSP which indicate a link failure. When the timer expires, the convergence should have finished and RBridges might switch back to the new primary DT.
- A default value (T_s) is offered for this timer and operators may change it.

Next Step

- Will incorporate further comments.
- Call for WG adoption?

Thanks!