

Resolving ETS Configuration Inconsistencies

Joe Pelissier

Pat Thaler

January, 2010

az-pelissier-ETSConfigInconsistencies-0310-v02

What we have today:

- Three tables are involved in the configuration of ETS:
 - A Transmission Algorithm Selection (TSA) table that specifies the Transmission Selection Algorithm per Traffic Class
 - A Traffic Class Grouping (TCG) table that:
 - Specifies the Transmission Selection Algorithm of being ETS or not ETS
 - If ETS, specifies a Traffic Class to Traffic Class Group Mapping
 - A table that specifies the ETS Bandwidth per TCG
 - Every TCG gets an allocation even if no traffic classes are currently assigned to it
- The problem
 - The first two tables are interdependent. The specification today inadequately specifies how setting a parameter in one table affects the other
 - With DCBX, if one side is willing and observes that the other end is not ETS (i.e. its TG15), there is currently no way to indicate what the local side should do
 - There is currently no DCBX mechanism to detect that one side is configured AVB while the other is configured strict priority
- Three possible solutions:
 - Specify the effects that programming one has on the other
 - Specify that one table overrides the other (or vice versa)

Specifying effects (option 1)

- Here is one possible set of rules:

In the TSA table:

If a TC is changed from Strict to AVB or from AVB to Strict, then no change to the TCG table

If a TC is changed from Strict to ETS, the TCG table is modified to move that TC from TCG15 to TCG0

Implies the admin would in a second step reassign the TC from TCG0 to the desired TC if it was not TCG0

If a TC is changed from ETS to Strict or AVB, the TCG is modified to move that TC from TCG0..7 to TCG15

In the TCG table:

If a TC is changed from TCG0..7 to TCG15, the TSA table is modified to change the TC's algorithm from ETS to Strict

If a TC is changed from TCG15 to TCG0..7, the TSA table is modified to change the TC's algorithm from Strict or AVB to ETS

In DCBX, if the local side is willing, the remote side is not willing, and the local side has some TCs assigned to AVB, those TCs will not change

(In other words, even if you are willing in general, you are never willing to change a priority assigned to a TC that is using AVB)

Specifying TCG table dominance (option 2)

- Remove the option in the Transmission Algorithm Selection (TSA) table to select ETS
- If the TCG table indicates that a TC is in TCG15, then the TSA table specifies the transmission selection algorithm
- If the TCG table indicates that a TC is in TCG0-7, then the TSA table is ignored and the ETS algorithm is used

The TSA table remains unchanged, i.e., it would still indicate that the Strict, AVB, or vendor specific algorithm is being used

- A change in either table has no effect on the other table

However, a change in TSA table may or may not change the actual algorithm used

Depends on status of the TCG table

Specifying TSA table dominance (option 3)

- Leave the TSA table and the ETS bandwidth table as is
- Eliminate TCG15. In the TCG table, TCs may be assigned to TCG0..7
 - If the TC is not currently utilizing ETS, then the assignment to the TCG has no effect
 - Similar to the ETS bandwidth table for a TCG that has no TCs assigned to it
 - Similar to the TCG table for a traffic class that has no priorities assigned to it
- Changing a value in the TSA table does not cause a change in the TCG table, and vice-versa
- This would require the TSA table to be part of DCBX asymmetric parameter passing

Detecting Mis-configuration (options 1 & 2)

- In options 1 & 2 it may be desirable to observe the TSA table to detect a mis-configuration
- DCBX can easily provide this capability