



Configuring ETS Bandwidth with DCBX

Joe Pelissier

az-pelissier-DCBX-Bandwidth-0309

Background

- **ETS Bandwidth configuration is asymmetric**

Each port on a link configures its *transmit* bandwidth allocations per Priority Group

There is no need for these to be configured with the same values on each end of the link

In fact, it will typically be the case that different values on each end of the link will be optimal

In many cases, the port on each end of the link best knows how the far end should be configured

Implies that in some cases it is best for each end to suggest the configuration to the other

- **Priority to Traffic Class configuration is also a necessary component of this configuration**

- **Current DCBX proposals support configuration in one direction only**

And do not provide the Traffic Class mapping configuration

Overview of Proposal

- **Provide two TLVs in DCBX:**

ETS Configuration TLV provides the transmitting port's:

Priority to Traffic Class Mapping

Traffic Class to Priority Group Mapping

Currently configured Priority Group Bandwidths

ETS Recommendation TLV provides a recommendation for the programming of the far end port:

Recommendation valid (I'll explain why we need this in a bit), indicates whether the TLV is providing a recommendation (otherwise the TLV is ignored).

Recommended Priority Group Bandwidth Assignments

Why two TLVs?

- **It is not always desirable for this operation to work asymmetrically**

For example, a bridge may be configured such that it updates its configuration based on an end station's recommendation, but it provides no recommendation to the end station.

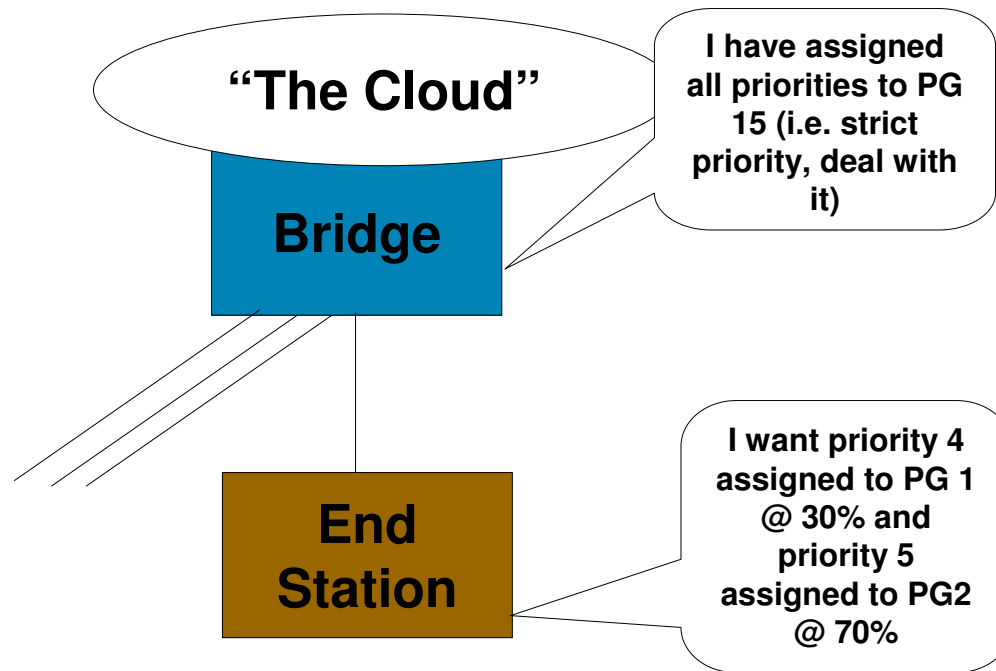
In this case, the bridge would transmit only the configuration TLV (so the end station knows the Priority to TC mapping) and the end station would transmit only the recommendation TLV.

Why the Recommendation Valid bit?

- **LLDP requires that a TLV always be transmitted or never be transmitted (based on administrative setting)**
- **Before providing a recommendation, a device may need to know the Priority to Traffic Class setting of the far end port (along with number of TCs supported and the Application TLVs)**
- **Therefore, there is a period during which the TLV needs to be transmitted before sufficient knowledge is known to populate the TLV**

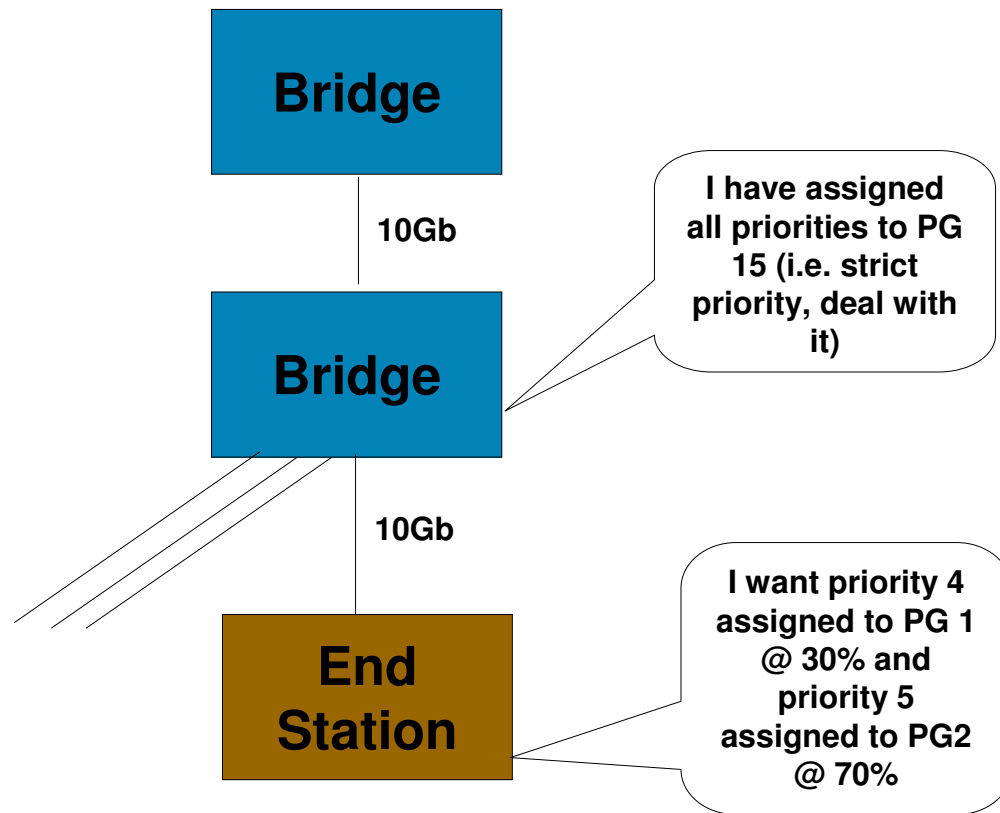
The valid bit enables this

An observation on configuration checking



Should the endstation do anything about this?

An observation on configuration checking



Should the endstation do anything about this?

The lower bridge is stuck in the back of a blade rack.

It has a single 10G uplink

The network is such that the majority of traffic goes over the uplink

Therefore, there is no point in configuring the downlinks to any specific bandwidth allocation (it is controlled completely by the northern bridge)

Therefore, this is a perfectly reasonable configuration

Moral: The end station has insufficient knowledge to make any judgment as to whether the BW allocations being provided to it are appropriate

Thank You!