Tspec assumptions

Michael Johas Teener
mikejt@broadcom.com

Original Tspec assumptions

- Only two parameters matter to the endpoint applications:
  - Bandwidth and latency
- Latency was OK to lump into two classes: A and B
  - Class A for < 2ms "through worst case Ethernet home net"
  - Class B for < 20ms "through *typical* worst case home network"
    - say, two WiFi hops and two Ethernet hops
    - "fuzzy" upper bound
- Bandwidth needs to be measured over a period, and the period depends on latency
  - longer period > longer bunches > longer latency
  - so, Class A used 125us, Class B used 1ms

Background

- For some time, low latency traffic (class A) has had a worst case latency of 2ms through 7 hops *on Fast Ethernet*
  - average worst case latency of a bit more than 250us per link
  - assumed some kind of traffic shaping would limit stream traffic bursts on all ingress ports to less than 125us (actually, less than 100us to allow for a guaranteed window for best effort traffic)
  - works fine since 100us + worst case best effort packet is substantially less than 250us
- So, class A shaping requires some kind of credit building based on 125us assumption for "bandwidth measurement window"
- Similar thinking gave us something like 1ms for Class B

Simple Tspec!

- Let's just use traffic class and bandwidth ...
  - bandwidth would be expressed as bytes/measurement period
- Ah, but there is packet overhead ....
  - ... and packet overhead is different for each layer 2
- So let's use traffic class, max bytes/class measurement period, max packets/class measurement period
- Bridges could use link speed and link technology to figure out the effect on link capacity
  - simple!

So, how have things changed?

????