Starvation-free priority precedence

Mick Seaman
mickseaman@gmail.com

802.1Q provides options for transmission selection. All except for the ‘strict priority’ default require management control. This presentation describes a simple approach that does not require management and imposes a worse case delay on latency sensitive frames that is no worse than for ‘strict priority’ but removes the ‘starvation’ risk for lower priority frames.

Priority precedence – Algorithm

As for simple strict priority:

- Transmission queues per output port per traffic class
- Received frames queued in reception order
- FIFO queuing, prior to transmission

But:

- The frame at the head of each output queue other than that of the highest priority is added to the tail of the next highest queue if (and once) higher priority queues currently contains no such promoted frame
- Transmission only occurs from the highest priority queue

Priority precedence – Result

As for simple strict priority:

- Transmission of any frame of priority P can be delayed by all previously queued frames of higher priority P^+, by previously queued frame of priority P, and by one previously queued frame of lower priority P^-.

Gain:

- Worst-case delay is the same as for simple strict priority.
- Avoids the well-known strict priority starvation issue.
- Practical replacement for ETS.

Cost:

- The expected/mean delay experienced by frames of priority P will be higher.