

802.1Q Queues for non-bridges

Norman Finn
Huawei Technologies Co. Ltd

Introduction

- This presentation is offered as food for thought.
- It anticipates a new PAR for adoption in March, 2018, but does not go so far as to offer a scope.
- Further off-line discussion is encouraged.

Zero congestion loss in a **router**

- IETF DetNet is defining deterministic networking for at least:
 - Routers
 - Label switches
 - Hosts
- Traditionally, IETF does not specify low-level queuing, such as is implemented in port ASICs; it leaves that (for Ethernet) to IEEE.
- But, for TSN/DetNet, the **queuing algorithms are critical**.

Zero congestion loss in a **bridge**

- IEEE 802.1Q-2014, plus past and ongoing amendments, define the queuing techniques required to achieve zero congestion loss.
- But, these are fully defined only for bridges.
 - L2 priority is the only criterion for selecting a class of service.
 - Except for 802.1Qci, which is still tied to a bridge.
 - Transmission selection is part of the bridge forwarding process (8.6).

Zero congestion loss in an **end system**

- Makers of end systems also want to implement many of the 802.1Q queuing techniques.
- Such implementers are often much less sophisticated than those building routers or bridges.
- They also need explanations on how to apply the descriptions in IEEE 802.1Q.

Other TSN functions

- 802.1Qci input gates + 802.1Qch Cyclic Queuing and Forwarding are an input-output pair that are most useful in a “relay system” (to use 802.1CB terminology) such as a bridge, a router, or a label switch — but are defined only for a bridge.
- 802.1CB Frame replication and elimination is **not** tied to bridges.
 - It also provides features required for 802.1Qci input gates.
 - It can easily and usefully be referenced by RFCs for application to L3 devices.
 - It illustrates one path for moving forward.

DetNet requires 802.1Q

- DetNet routers and label switches need normative references to 802.1Q in order to have standards for L3 deterministic networking.
 - Or, they need to re-do everything we've been doing in TSN for the last few years.
 - Nobody wants that.
- **Creating normative references from an RFC to 802.1Q to describe a “deterministic router” is impossible.**
 - The queue definitions and usage in 802.1Q are all tied to the forwarding function, which a router obviously does not have.
 - Any references would require “Do clause 8.x.y.z, except where the text says, “blah blah blah ...”, do “mumble, mumble, mumble ...”, instead.

P802.1XY: Use of 802.1Q Quality of Service by non-802.1Q Systems

(Consider that title a trial balloon – does it float?)

- This author considered restructuring 802.1Q so that the definitions of queuing are applicable to any device, and restructuring the rest of 802.1Q to use that feature.
 - That looks like too much work for too little benefit.
- An alternative is to write a new standard that tells an implementer of queuing in a non-bridge (either an “end system” or a “relay system” [to use 802.1CB terminology]) how to make sense of 802.1Q.

Standard vs. Profile

Standard, not profile, because:

- The "Do clause 8.x.y.z, except where the text says, "blah blah blah ...", do "mumble, mumble, mumble ..." kinds of clauses will contain normative text.
- A conformance clause and a PICS proforma, with options, are appropriate.

DetNet then writes ...

- Give 802.1XY, the DetNet WG can reference specific capabilities by name and clause number, along with a bi-lingual terminology dictionary, and get a clear standards trail for a deterministic router or label switch.

Requirements and enhancements

- Clearly, the packet characteristics that select the class of service are not limited to L2 priority – that becomes a question for the standard that references 802.1XY.
- It seems that 802.1XY need not be limited to 8 classes of service.
- One of the greatest benefits of the 802.1Q queuing structure is that the interactions among all queuing techniques are well-defined. To maintain this, there must be handles to allow connecting queuing techniques defined outside 802.1Q (e.g. non-port queues for L3 data streams) to 802.1Q techniques.

What about new queuing techniques?

- For a while, any new queuing technique would require modifying both 802.1Q and 802.1XY.
- This is why refactoring 802.1Q, similarly to the manner in which 802.1AC was split away from the old 802.1Q Clause 6.4, is the better answer.
- But, the need for 802.1XY is immediate, and in writing it, we will learn a lot about how 802.1Q might be refactored in the future.

Summary

- Further discussion may open up better ideas.
- We should act with reasoned alacrity.
- This seems a good solution for the next few years. Then, we may want to bust up 802.1Q and put our standards together differently.

Thank you