Transmit selection and traffic shaping for Privacy Channels

(Work in progress)

Mick Seaman
mickseaman@gmail.com

Privacy Channels (background)

- Adversaries can use frame sizes and transmission timing to detect changes in activity, fingerprint applications, and guess content
- A Privacy Channel comprises fixed size ‘MPPDUs’ (MAC Privacy protection PDUs) transmitted [scheduled for transmission] at regular intervals
- Each MPPDU encapsulates zero or more user data frames (or fragments of user data frames) and padding
- MPPDUs are confidentiality protected by MACsec
- Reduces activity level, frame size, timing exposure
  - — Provided Privacy Channel parameters (MPPDU sizes, transmission interval) are not tuned/retuned to current application requirements

Privacy Channel properties

- Recommended MPPDU size (prior to MACsec protection)
  - — 2 [MAC Privacy protection EtherType] + 2 * 6 [overhead for encapsulating two Frame Fragments] + 1518 [encapsulated user data frame] - 4 [FCS not encapsulated] + 4 [VLAN Tag] = 1532 [for IP MTU Size 1500]
  - — ‘On the wire, with MACsec protection’ add 12 [DA, SA] + 16 [SecTAG] + 16 [ICV] + 4 [FCS] = 1580
  - — ‘On the wire, to the service provider’ add 4 octets [outer VLAN TAG for priority and service selection] = 1584
  - — Needs discussion. Rationale for the above is to support 100% load with ‘1518’ user data frames initially fragmented without spill to later MPDUs. Could reduce to 1522/1570/1574 if objective merely to encapsulate original ‘1518’ frame without fragmentation. Maximum possible throughout size depends on large/small user data frame mix.
  - — MPPDU size at least capable of encapsulating single IP MTU without fragmentation (no point in jumbo frames, if fragmentation increases loss probability).
  - — Given the MPPDU size desired bandwidth implies transmission interval (and v.v.).

Priority-based selection and shaping

Transmission selection and traffic shaping algorithms can use the ***priority*** of user data frames as an indication of both:

- ***urgency*** [how quickly should a given frame be forwarded]
- ***importance*** [what should be done to minimize its loss]

Priority attributes can be refined by identifying frames as part of a flow or ***stream*** with resource reservation (and limitation).

Privacy Channel encapsulation impacts transmission selection and transit delay for resource and time critical flows. Priority support mandated to allow TSN:

- ***coexistence*** [can send individual frames with/without privacy]
- ***support*** [operational specification might be sufficient in some cases]

Support for stream identification and parameters **may** (explicit option) be implemented in addition to priority support.

Selection and shaping algorithms

In a bridge, the Forwarding Process (.1Q-2018 8.6) does transmit selection and shaping. Privacy protection has minimal buffering for transmit Privacy Channel encapsulation and received fragment reassembly.

Need to understand how Privacy Channels work with existing transmission selection and shaping algorithms:

- Strict priority (default,.1Q mandatory to implement;.1Q-2018 8.6.8 without PFC/Priority_Paused[n], 8.6.8.1)
- Simple priority (as Strict priority but allowing local functionality equivalent to Priority_Paused[n])
- Credit-based shaper (.1Q 8.6.8.2, Annex L)
- ETS (Enhanced Transmission Selection,.1Q-2018 8.6.8.3, Clause 37)
- Enhancements for scheduled traffic (transmission gates,.1Q-2018 8.6.8.4)
- ATS (Asynchronous Traffic Shaping,.1Qcr-2020 8.6.8.5)

Traffic shaping objectives

Strict priority transmission selection does not distinguish ***urgency*** and ***importance***. Every node (bridge) assumes that traffic sources are naturally burst limited, the next node can always buffer *important* frames (possibly discarding lesser priority frames) if the available transmission bandwidth is less than the receive link rate.

Traffic shaping attempts (in varying degrees by algorithm) the following, for frames of an identified class or flow:

- Allocation of sufficient link bandwidth
- Availability of sufficient buffering

to avoid (when given sufficient control over the aggregate load through some form of reservation):

- Loss due to buffer exhaustion or (rarely) to node transit delay
- Excessive end-to-end transit delay
- Excessive delay variance (for frames of a given priority or flow)

Traffic shaping assumptions

Performance analysis usually makes (and states) some or all of:

- Frame transmission begins immediately after selection
- The highest priority/most eligible frame will have been selected for that transmission
- ‘Time on the wire’ (i.e. transmit opportunity denied to competing flows) does not depend on when the frame is transmitted or on the priority or other characteristics of the frame
- Frame size does not change between reception and transmission at a given node.

Privacy Channel encapsulation challenges these assumptions:

- The next node’s Forwarding Process does not ‘receive’ a frame until the MPPDU’s trailing ICV has been validated
- Packing the MPPDU can take time
- Can model these delays in various (shaper dependent) ways

Privacy Channel MPPDU transmission

- Timing cannot depend on content (frames, fragments, or pads)
  - — Care required to not accidentally reveal as a consequence of time needed for internal operations
- Timing can be subject to subsequent traffic shaping of MPPDUs (e.g. by an EDE’s network component) if completely ignorant of content
- Timing can be affected by other Privacy Channel frames or by individual Privacy Frames
  - — To the extent that externally visible priority already discloses the priority relationship

Strict priority transmission selection

Cases:

- Single Priority Channel
  - Viable
  - Questions around when eligibility decisions for the next MPPDU are made, how should these parameterized, should they be parameterized in the standard
- Two Priority Channels
  - ‘Wrong priority’ availability can reduce lower priority channel use
- Single Priority Channel plus individual higher

Simple priority transmission selection

Transmission gates

ATS