IEEE 802.1 Security
MACsec Privacy Rate Specification
Deep Dive Version 01

Don Fedyk – don.Fedyk@labn.net

Disclaimer

- This is a work in progress. The material here is for discussion purposes and may contain errors.

Configuration parameters for MAC Privacy 802.1AEdk

- How big should an MPPDU be?
  - Encapsulation Arithmetic
- Bandwidth, Interval, Delay, Burst
- Based on some ideas in Mick's Updated Clause 20 text

Determining Frames Size and Rates for MAC Privacy Channels and Frames

- PrY Channel Frames:
  - Fragmentation
  - Fragmentation is default enabled. (No config option in standard but users can add one)
    - Allows Higher efficiency, (allows late addition) – not in the standard.
  - Setting an MPPDU too small can force fragmentation when Max size user frames are encountered.
  - Determine the maximum user frame size "User Data Frame size".
- PRY Frames
  - No Fragmentation
  - Determine the maximum user frame size "User Data Frame size".
  - This must be greater that or equal to the User Data Frame size or larger frames could be dropped.

MAC Privacy – Which Length?

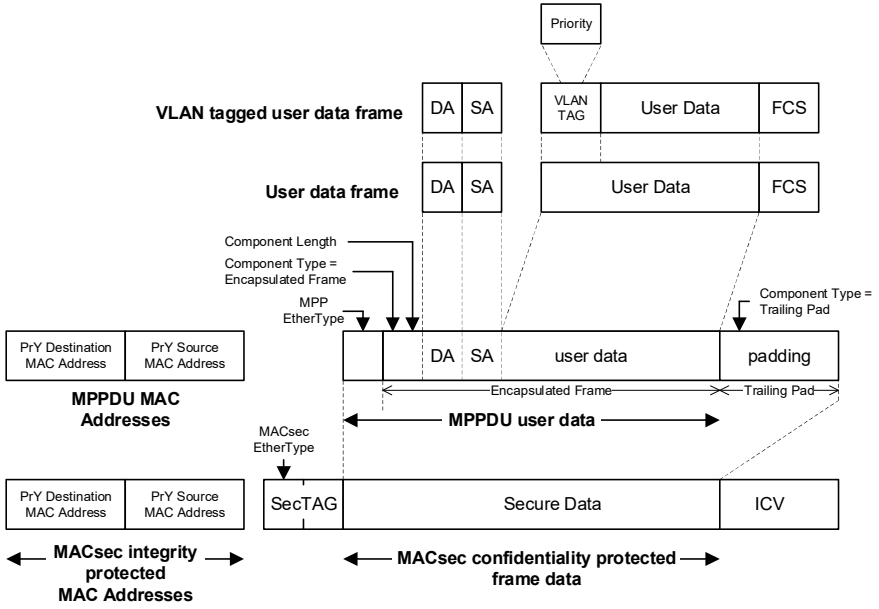

This is the source of the defining length

MAC Privacy

MAC Privacy & MACsec

Clause 17 Draft figure

Standard Ethernet Frames

Standard Ethernet encapsulation:

- Frame sizes are dependent on media
- Ethernet Standard are 1500 octets of user data
- Ethernet Jumbo is 9000 octets of user data
- Uses the Media overhead bytes

Goal is to determine the MTU for the situation

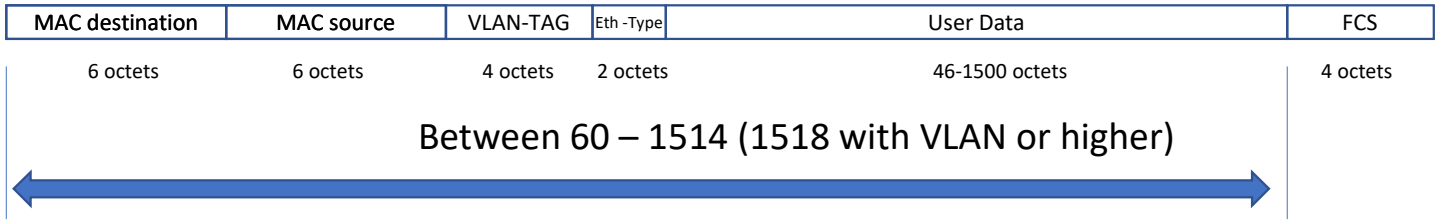

User data no FCS
a.k.a. L2 MTU

MAC Privacy allows for an unfragmented Max User data frame
This means encapsulating nominally up to 1518 octets but possibly higher.
Other formats (e.g. LLC, SNAP) are supported as well but are less than or equal to 1500 octets. IEEE 802.3 allows up to 2000 for envelope frames 802.3as-2006

MAC Privacy MPPDU

- MAC Privacy Encapsulation adds
- 6 octets for Fixed Full frames
- +2 octets for a fragment – but that can be absorbed since fragments are variable
- Therefore 6 octets plus user data frame 1518 + 6.

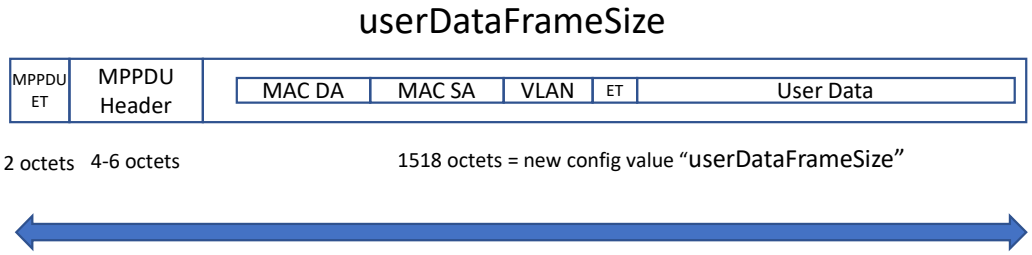

This is not a frame. A PrY Frame in non collocated PrY adds 16 octets more.

Notes

- I argue that userDataFrameSize is what a user should configure. This value is what an MPPDU should carry.
- Logically this is 1500 (9000) max plus the Ethernet L2 overhead but since we can have various formats and various MTUs I think configuring the actual payload of the MPPDU makes sense. ~~**It is a discussion point though.**~~ **Agreed**
- However, there are other values which relate to the total bandwidth on the wire coming next.
- We need to keep configuration to a minimum – yet accurately represent bits if we are to follow the algorithms currently suggest by the text.

Summary

- Only need to configure the PrY L2 MTU e.g. 1518 or similar (e.g. 9018)
  - This number is whatever the base traffic user traffic format is for example PBB frames would use a larger number.
  - MPPDU length will be fixed at 1518 + 6 octets. (1524)
- As far as MAC privacy 1518 is the configurable encapsulation payload.
  - userDataFrameSize, user-data-frame-size
- If this number is configured smaller than source user traffic, some large frames may be fragmented – for channels.
- This number must be supported for privacy frames or the 1500 (9000) MTU of the user frame is impacted.
- If a smaller MTU is required for other reasons for PrY channels this number can be adjusted downwards this guideline is merely to prevent fragmentation of whole frames, but implementations may fragment anyway in the interest of reducing delay or increasing efficiency.

Part 2 Channel Frame Timing

- Whatever the Channel MPPDU length is, the MAC Pry Frames size determines the interval for a given data rate.
- We need to determine the size of the MAC PrY channel frame on the wire.

MAC Privacy Frame (using MACsec)

- MACsec has 8-16 octets of Sec TAG header
- Additional addresses (MAC Privacy) 12 octets
- Plus ICV 8-16 Octets

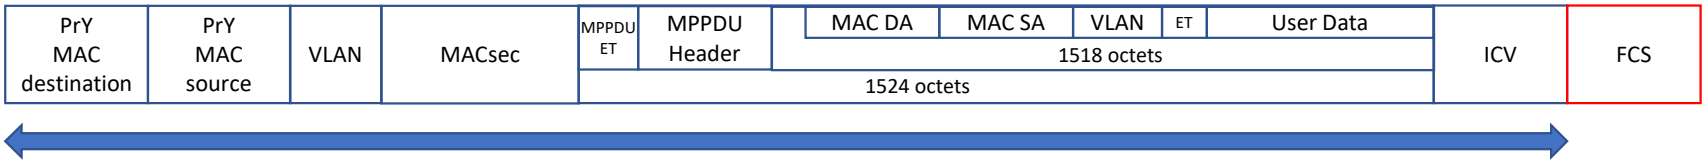

1552-1568 octets (16-32 octets MACsec encapsulation + 12 PrY MAC) (32 or 40 or 48)
4 octets VLAN likely
Maximum as high as ~1596 (all well within Frame overhead maximum)

How to compute a Frame size Base on Pry MTU

- Frame size is the PrY User data frame with no FCS
  - Ethernet with 1 VLAN Tag = 1518 (from previous slides)
  - For Rate computation the system adds:
    - 6 octets MAC PrY + 4 octets VLAN 44 octets MAC Sec+ ⟸ Constant for a configuration
    - FCS (4 octets) + Media overhead (20 octets). ⟸ Media dependent
    - 1596 MAC PrY

Translation to a Media Rate/Interval

- 802.3 Media example
  - 8 Octets of preamble
  - 12 Octets of Inter Frame Gap
- These values plus the frame size determine the frame rate.
- I.E. 10Bbit/s = 10,000,000,000/(1596))*8))

  = 783208 Frames per second.

  (original was 10,000,000,000/((20+1518)*8) = 810635

  ~3.3 percent overhead.

Arithmetic for Interval

Description	Value	Octets Tally	Bits	Bandwidth bps	Interval ns	Actual BW
L2 MTU User		1500	12000	100,000	127,680,000	100000
L2 MTU Size Includes Address and Ethertype		1514	12112	500,000	25,536,000	500000
Number of User VLAN tags	1	1518	12144	1,000,000	12,768,000	1000000
Other Header Overhead	0	1518	12144	10,000,000	1,276,800	10000000
MAC Pry MTU max-mppdu-payload-size		1518	12144	100,000,000	127,680	100000000
MAC PrY MPPDU Header (constant)	6	1524	12192	1,000,000,000	12,768	1000000000
MAC DA+SA (constant)	12	1536	12288	10,000,000,000	1,277	9998433829
MACsec VLAN	1	1540	12320	40,000,000,000	320	39900000000
MACsec Header (constant)	8	1548	12384	100,000,000,000	128	99750000000
SCI Additional (Configured by Environment)	8	1556	12448			
ICV Size (Configured by Cipher Suite)	16	1572	12576			
FCS Frame on the Wire	4	1576	12608			
Media Preamble	8	1584	12672			
Media Interframe Gap	12	1596	12768			
Total Bits per Frame on the wire			12768			

User Configures:

- Payload for MAC Pry to enable full frame.

Clarified we do not need interval

Note that total bits per frame maps to interval but a User does necessarily know this immediately . This is one reason bandwidth was in the original proposal. Interval can be computed but it is not solely determined by MAC PrY parameters alone.

Notes: For User Configuration

- Payload for MAC PrY
- :
  - A bandwidth that allows us to compute the size and interval to suite.
    - An interval is calculated and rounded to the nearest nanosecond or whatever timescale is supported. A corresponding actual bandwidth is calculated from the interval and bits on the wire.
    - If the bandwidth is too large it – exceeds capability it can be reduced
- The Shaping algorithm
  - Next,

Burst Size

- We want a burst-size that reflects a frame sent on the wire.
- "The notional bucket contains a number of bits. When a Privacy Channel MPPDU is generated, 8 bits are subtracted from the bucket for each of the userDataFrameSize plus overhead octets. Bits are added to the bucket at the channel requestedBitRate, expressed in bits per second, up to a maximum of channel BurstSize octets. "
- If the burst size was based on frames, we would not need to know the actual bits on the wire –it can be computed
- Current text suggests we know that 1518 octets MMPDU results in ~1592 octets on the wire (give or take 16 octets).
- Does the burst size need to be exact?
  - Trying to avoid configuring multiple packet sizes.
- Not that if multiple frames sizes are being encapsulated this is the largest Frame size bits on the wires time burst frames.

Burst Size

Burst Size Should be based on this if it is in bits

Bits on the Wire – 1596 octets

Bandwidth is based on this too

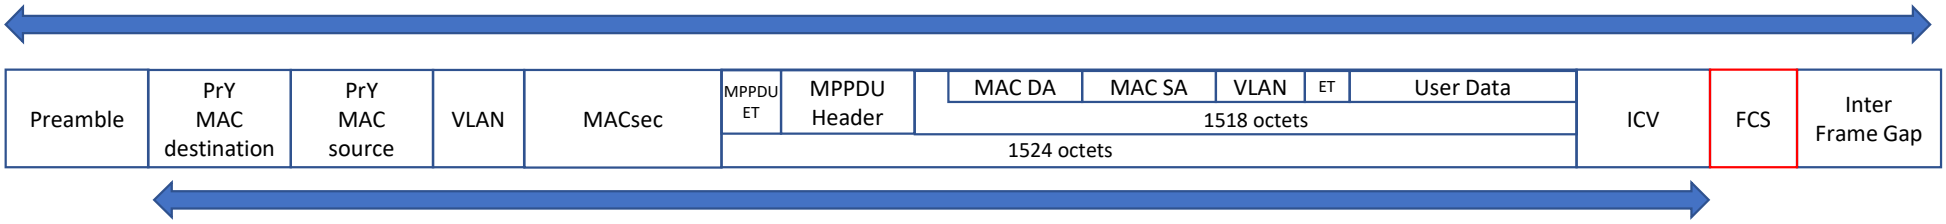

MACsec Size 1572 octets fixed includes padding

Burst Size could be in integral frames and then it does not matter?

User Bandwidth is 1500/1518 vs 1500/1596

Transmission Cases

- MAC PrY Channel is the whole link
  - Rate is line rate
  - Interval can be computed
  - Burst rate – Not really required ?
- MAC PrY Channel is a portion of the link
  - Interval is configured and rate can be computed.
  - Rate is some fraction of the line rate (interval could be computed)
  - Burst rate must be specified to average the interval/rate over a time period
- Burst size is limited by delay at lower speeds and memory buffers at high speeds.

Transmission Timing
(Time Sensitive Networking for example)

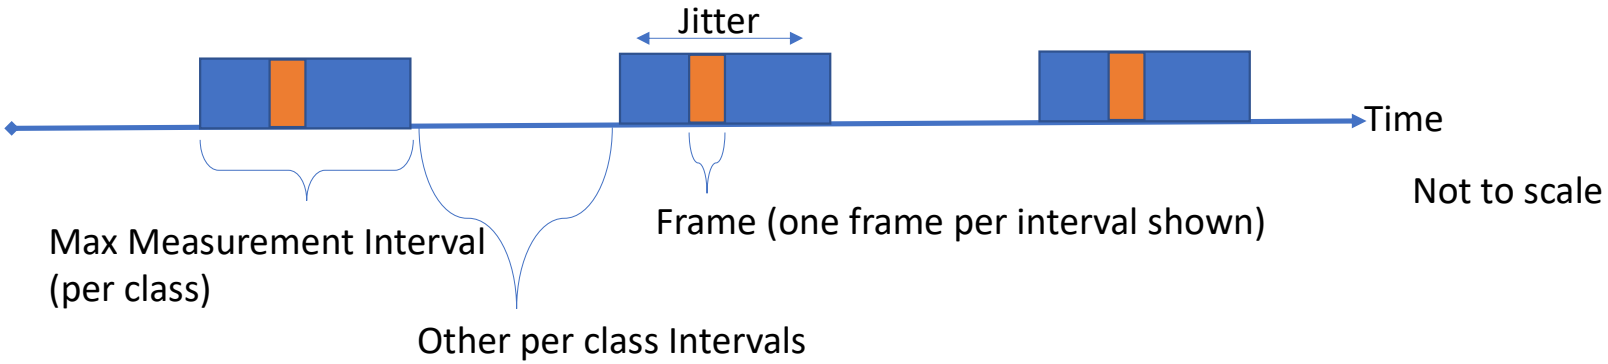

Controls / Per Class

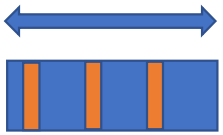

- Measurement interval 500us – 250us Typical
- Class Bandwidth - Percent of port bandwidth (locked or shared with other classes)
- Integral number of Frames per interval

Note this is my interpretation

One Example Delays that cause Burst

MAC PrY Interval may be different than TSN intervals and scheduling may overlap:

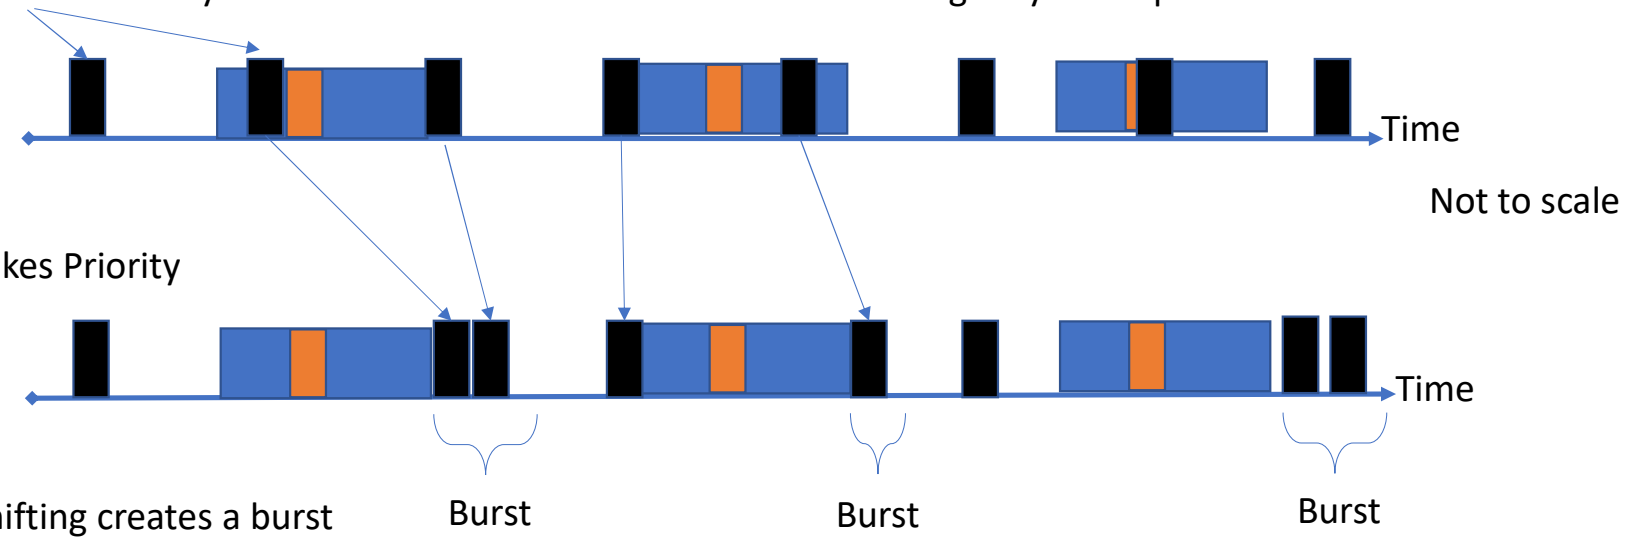

Average transmission rate is the same over a time span need to configure a burst limit.

One Example Delays that cause Bursts

MAC PrY Interval may be different than TSN intervals

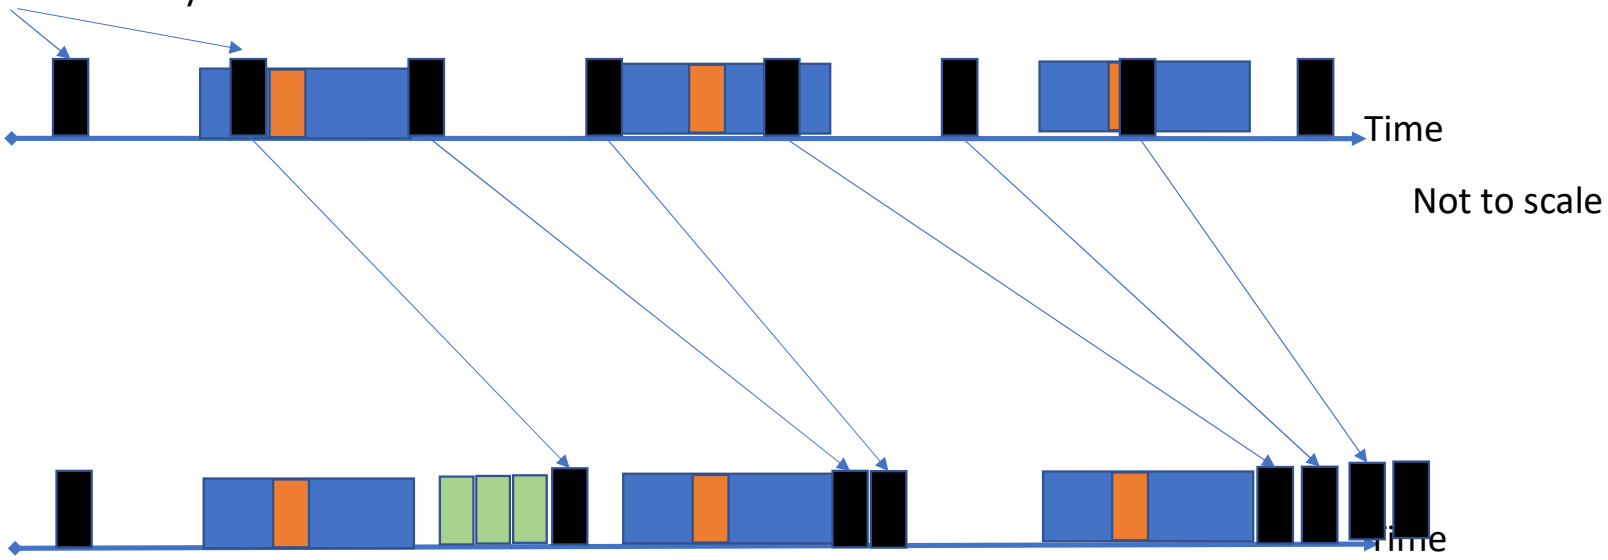

Other Traffic May cause the MAC Pry Frames to be pushed even further out. A maximum Burst size limits the number of frames that can be delayed from the original transmission time. Burst is based on effective bandwidth to the MAC PrY class.

Arithmetic for Channel Burst

Frame Size Bits on the Wire	1592			
Link Rate 10 gbit/s	10000000000			
Transmission Time/Frame	1.2736E-06			
Effective Bandwidth available at or equal to MAC Pry Transmission Priority		Elapsed Time/ Frame	Frames/Sec	Memory octets
10 kbit/s	10000	1.2736	0	0
100 kbit/s	100000	0.12736	7	11144
1 mbit/s	1000000	0.012736	78	124176
10 mbit/s	10000000	0.0012736	785	1249720
100 mbit/s	100000000	0.00012736	7851	12498792
1 gbit/s	1000000000	1.2736E-05	78517	124999064
10 gbit/s	10000000000	1.2736E-06	785175	1249998600
100 gbit/s	1E+11	1.2736E-07	7851758	12499998736

Relative Delay
Aggregation Delay/Elapsed Time + Transmission Delay

Frames or Octets?

What determines Burst octets or Frames?

- It is effective BW not Actual link rate when effective bandwidth is less than the link rate.

What else ?

- Burst Limit Memory?
- Burst limit aggregation delay?

Offer configuration of burst frames
0 to Max 100 Thousand frames?
~20 MB of memory

Channel Configuration Parameters

- user-data-frame-size
- Bitrate – available
- Channel Burst-size
  - Uint32 octets ~~(or Frames)~~ Not using range in YANG.

Use a choice statement.
Configure one Compute the other

- Actual bitrate and ~~actual Interval~~ can be computed by the system and reported. The actual frame size on the wire is used to compute the rate = interval x effective frame size on the wire. System has to add the MACsec frame sizes and media overhead.
- One problem – Non–collocated MAC PrY does not know MACsec/Media Overhead
  - If entering ~~Interval,~~ a slightly higher value (~2%) can be used. If entering bandwidth, a slight smaller value (~2%) can be used. Only a factor when using full link bandwidth.

Frame Configuration Parameters

- user-data-frame-size
- nearest-multiple-pad
  - range "0|16|32|64";
  - Pad to “32” means anything <= 64 = 64, anything 65 to 96= 96
- A 64 octet user frame without FCS could be as small as 60 octets
- The values 16, 32, 64 would result in 4 octets of padding to 64.
- A 68 octet user frame without FCS would be 64 octets.
- Any larger frame of 65 or more octets(after FCS removal) would result in 80 octets (16), 96 octets (32) and 128 octets (64).

Revised Prototype YANG Model (snippet)

```
   module: ieee802-dot1ae-pry

  augment /if:interfaces/if:interface:
    +--rw pry {macsec-priv}?
       +--rw mac-privacy-enabled?          boolean
       +--rw pry-source-address?           ieee:mac-address
       +--rw pry-destination-address?      ieee:mac-address
       +--rw user-priority-to-pry* [user-priority]
       |  +--rw user-priority    uint8
       |  +--rw privacy-type?    union
       +--rw privacy-channel* [channel-id]
       |  +--rw channel-id                 identityref
       |  +--rw user-data-frame-size?      uint16
       |  +--rw mppdu-priority?            dot1q-types:priority-type
       |  +--rw requested-bit-rate?        uint64
       |  +--ro actual-bit-rate?           uint64
       |  +--rw burst-size?                uint32
       |  +--ro total-size-on-wire?        uint16
       +--rw privacy-frame* [frame-id]
       |  +--rw frame-id                   identityref
       |  +--rw user-data-frame-size?      uint16
       |  +--rw mppdu-priority?            dot1q-types:priority-type
       |  +--rw nearest-multiple-pad?      uint16
augment /if:interfaces/if:interface:
```

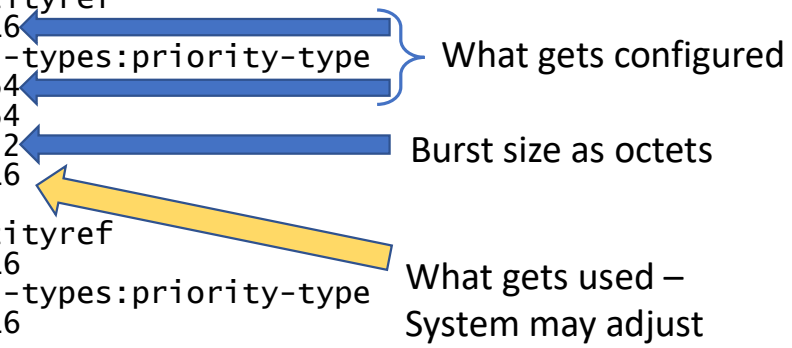

UML Config

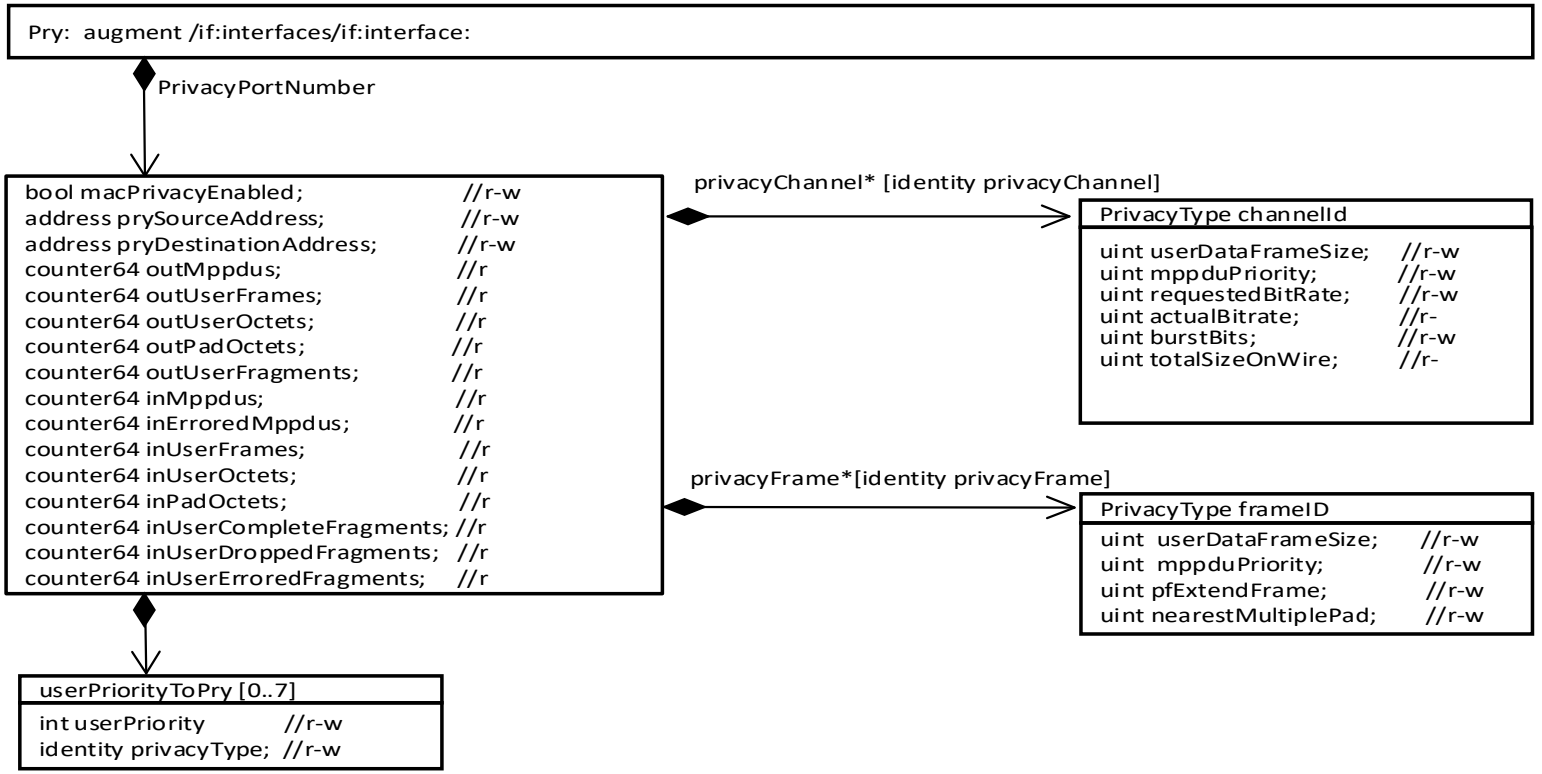
Pry: augment /if:interfaces/if:interface:
PrivacyPortNumber
bool macPrivacyEnabled; //r-w
address prySourceAddress; //r-w
address pryDestinationAddress; //r-w
counter64 outMppdus; //r
counter64 outUserFrames; //r
counter64 outUserOctets; //r
counter64 outPadOctets; //r
counter64 outUserFragments; //r
counter64 inMppdus; //r
counter64 inErroredMppdus; //r
counter64 inUserFrames; //r
counter64 inUserOctets; //r
counter64 inPadOctets; //r
counter64 inUserCompleteFragments; //r
counter64 inUserDroppedFragments; //r
counter64 inUserErroredFragments; //r
privacyChannel* [identity privacyChannel]
PrivacyType channelId
uint userDataFrameSize; //r-w
uint mppduPriority; //r-w
uint requestedBitRate; //r-w
uint actualBitrate; //r-
uint burstBits; //r-w
uint totalSizeOnWire; //r-
privacyFrame*[identity privacyFrame]
PrivacyType frameID
uint userDataFrameSize; //r-w
uint mppduPriority; //r-w
uint pfExtendFrame; //r-w
uint nearestMultiplePad; //r-w
userPriorityToPry [0..7]
int userPriority //r-w
identity privacyType; //r-w

What if MAC PrY traffic is TSN?

- Should use a MAC Privacy Frame and timing has little impact (but frame sizes are slightly larger).
  - ieee8021SrpStreamTspecMaxFrameSize would be adjusted to account for MAC PrY and MACsec overhead.
  - No interval specified
- What if the TSN traffic is in a channel ?
  - Usually Not recommended. This would add delay to TSN frames. Any frame interval would be mapped to TSN parameters.
  - If the Channel used a small frame size (~256-512) octets and fragmentation was enabled? Not intended operation but possible.

Comments?
Thank You