

# 802.3 DMLT BELOW THE MAC

## A Technical Feasibility Presentation


2013-03-18

Pat Thaler



# MAC Merge layer encapsulation goals

- Preserve frame integrity
  - No increase in undetected errors
- Indicate which MAC receive frame belongs to
- Minimize impact on throughput
- Transparent to existing non-deprecated PHYs above 10 Mb/s



# Terminology



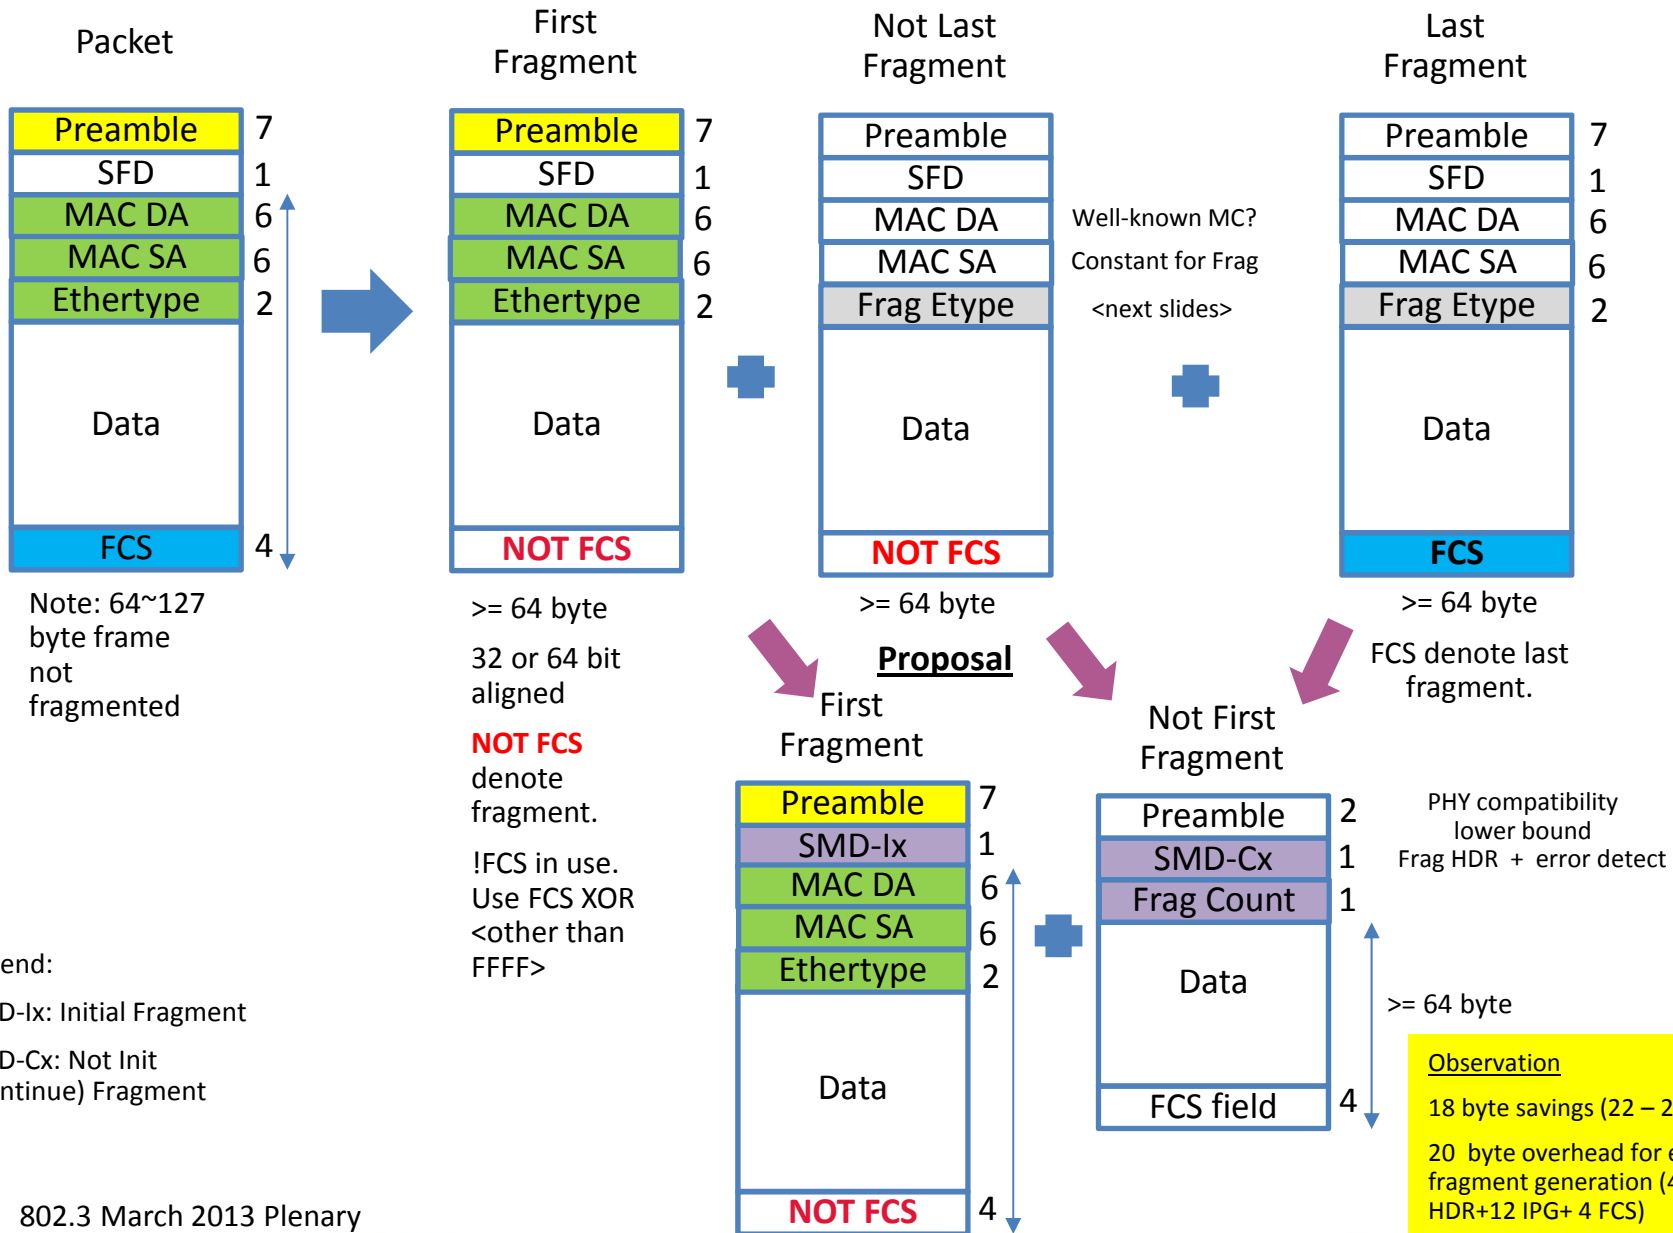
- DMLT frame – frames with the lowest latency
- Preemptable frame – frames that are not distinguished to be low latency.
- Mframe -- A transmitted unit from MAC Merge that includes both whole frames and fragments of preemptable frames – Mframe
  - Which stands for MAC Merge frame – a unit that looks like a frame at the PHY layer but may contain a whole frame or a fragment of a MAC layer preemptable frame. I'm leaning toward this one.

# Assumptions



- MAC Merge does not have to be supported over deprecated PHY Clauses

# Conceptual MAC Merge Frame Consideration



# MAC encapsulation elements



- Preamble
  - Minimize use of preamble bytes for frame start
  - Provide at least 2 bytes of preamble for fragments
    - The 1000BASE-X PHY can drop up to 2 bytes of preamble and insert the SFD over another byte. In practice, implementations only drop 1 and many delay the start of preamble rather than drop and only overwrite 1 byte (similar to 10GBASE-X alignment)
- Identify Mframe as DMLT or start of a preemptable frame or a later fragment of a preemptable frame
- Protection for reassembly errors when an Mframe is lost
  - Frame number – circular count from 0 to 3
  - Fragment number – circular count from 0 to 3
- Identify last Mframe of a Frame
  - Mark end of preemptable Mframe

# Mframe start

- For start of non-initial fragments
  - Insert 2 bytes of preamble followed by
  - SMD byte (Start Mframe Delimiter)
  - Frag byte (Fragment count)
- Preemptable Frame start and DMLT frame
  - Replace SFD with SMD
- SMD values have Hamming distance 4 from each other
- Frag values have Hamming distance 4 from each other

# SMD and Count byte encodings



Mframe type	Frame #	SMD
DMLT	NA	0x33
Preemptable frame start	0	0x66
	1	0xCC
	2	0xFF
	3	0xAA
Non-initial fragment	0	0xE1
	1	0xD2
	2	0x1E
	3	0x2D

Fragment #	Frag
0	0x66
1	0xCC
2	0xFF
3	0xAA





# Frame CRC



- Frame CRC is generated by the MAC over the same frame bits as always. It is not altered by the MAC Merge sublayer

# Indicating end of fragment



- Since one doesn't know that a frame will be preempted until transmission of the frame is in progress, a marking at the end transmission indicating that this is not the final fragment is needed.
- Two alternatives
  - All preemptable frames have a fragment trailer that indicates whether this is the end of the frame.
  - Keep a running CRC calculation in the MAC Merge sublayer (identical to the CRC calculation done in the MAC sublayer) and append the intermediate **result altered in some way** to the frame.
- Note: Preempted fragment size should be no smaller than 64 bytes, and each fragment should be 32 or 64 bit aligned (except for the last one), and packet less than 127 bytes should not be preempted.

# Fun with CRCs



- IEEE 802.3 inverts the calculated CRC and appends the result to the frame.
- Transmitting MAC Merge sublayer could invert the second two bytes of the intermediate CRC result (the CRC computed over the bytes of the MAC frame that have been transmitted so far) at the end of non-final fragments.
- Receiving MAC Merge sublayer runs a CRC calculation as the frame is received. When an Mframe ends, it compares the calculated value with the
  - If the difference between that and the last 4 bytes of the Mframe is 0xFFFF, it's the end of a MAC frame
  - If the difference between that and the last 4 bytes of the Mframe is 0x00FF, it's a non-final fragment.

Thank You.

Questions or Comments?

# Fragment trailer



- 1 byte – all zeros indicates frame incomplete, all ones indicates frame complete. (Or use 4 bits and leave 4 bits reserved).
  - Another alternative if we use shortened header for frame start would be to encode frame number for incomplete frames by use of 5 values with Hamming distance 4:
    - Frame complete
    - Frame incomplete, Frame #0
    - Frame incomplete, Frame #1
    - Frame incomplete, Frame #2
    - Frame incomplete, Frame #3