PHY latency and AUI BER recap

Matt Brown, Huawei, P802.3dj Editor-In-Chief
Kent Lusted, Intel

IEEE P802.3dj Task Force
April 20, 2023 electrical ad hoc meeting

Introduction

- This presentation is a condensed recap of the following three presentations:
  - #1 PHY/FEC architecture considerations V2
    https://www.ieee802.org/3/dj/public/23_03/brown_3dj_01a_2303.pdf
  - #2 BER considerations for 200 Gb/s per lane AUIs
    https://www.ieee802.org/3/dj/public/adhoc/electrical/23_0406/brown_3dj_elec_01a_230406.pdf
  - #3 MAC link latency considerations
    https://www.ieee802.org/3/dj/public/adhoc/optics/0423_OPTX/brown_3dj_optx_01b_230413.pdf
- Provides a summary of the trade-off between the AUI BER choice and the MAC-link latency.

#1 PHY/FEC architecture considerations V2

PHY/FEC architecture considerations V2

Matt Brown, Huawei
Gary Nicholl, Cisco
John D'Ambrosia, Futurewei, US Subsidiary of Huawei

#1 PHY/FEC architecture considerations V2

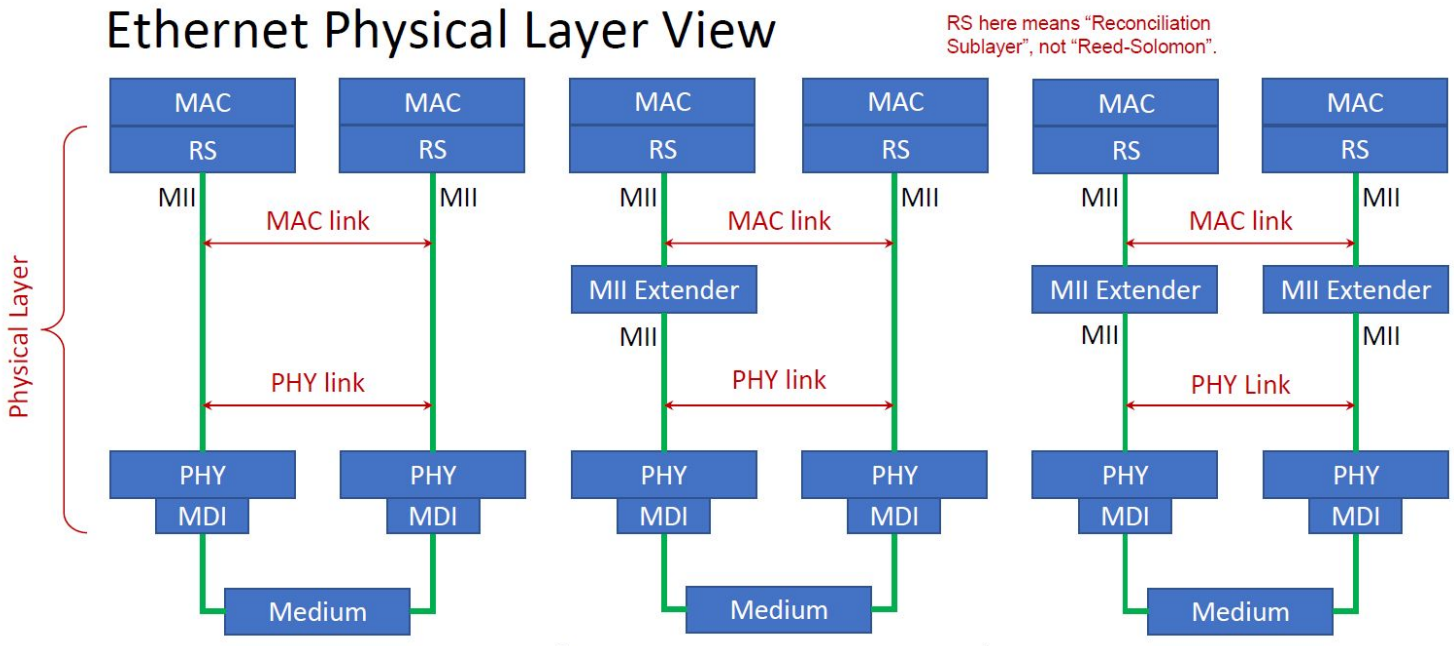

#1 PHY/FEC architecture considerations V2

Type 1 PHY/FEC

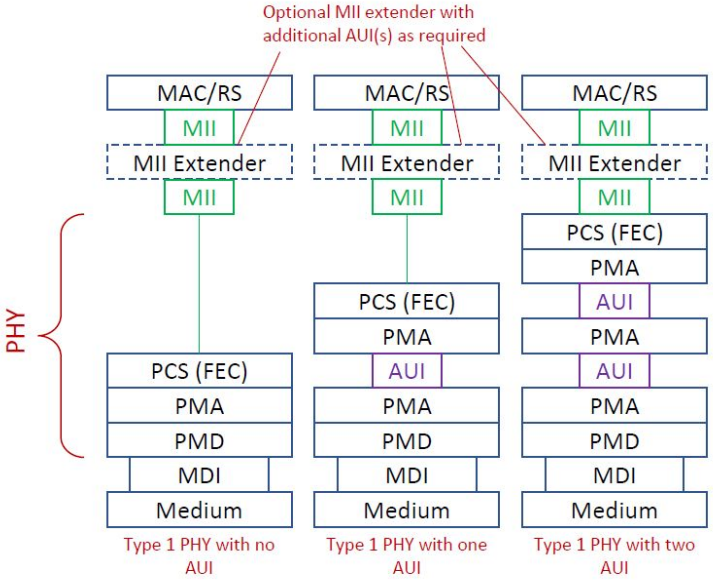

- A single FEC spans the PHY link (PCS to PCS) which may include up to four AUIs.
- FEC corrects errors that are contributed by the PMD link and the AUIs.
- PMD and Medium characteristics are defined with AUI errors in mind.
- BER trade off between the AUIs and the PMD link.
- More AUIs may be added above the PHY using the optional MII Extender without affecting PHY performance.
- The following PHYs are Type 1 PHY/FEC:
  all 200GBASE-R in 802.3, 802.3ck, 802.3db
  all 400GBASE-R in 802.3, 802.3ck, 802.3db
  all 800GBASE-R in 802.3df

2023/3/13 IEEE 802.3dj Task Force 8

#1 PHY/FEC architecture considerations V2

Type 2 PHY/FEC

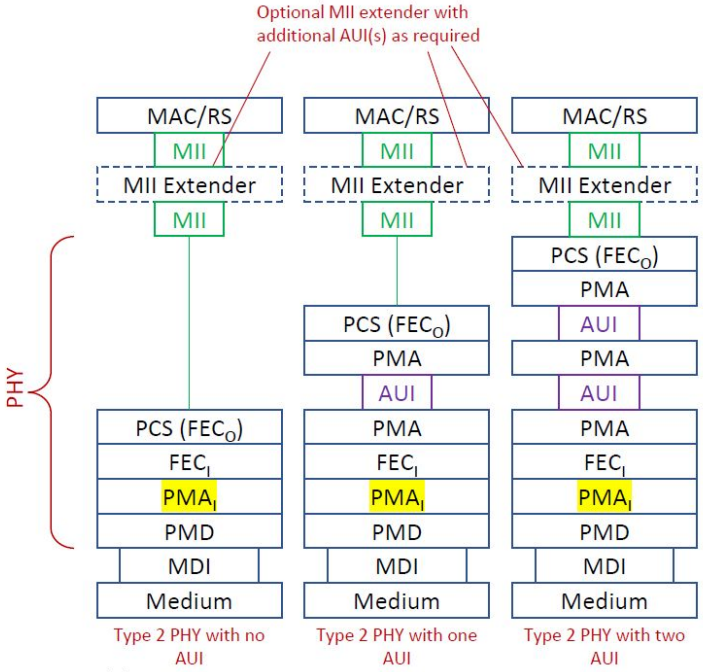

- An outer FEC (FEC$_O$) spans the PHY link (PCS to PCS) including up to four optional AUIs (like Type 1)
- An inner FEC (FEC$_I$) spans only the PMD link (PMD to PMD)
- The PMA below FEC$_I$ (PAM$_I$) is different from the PMAs above FECI.
- FEC$_I$ corrects "most" errors contributed by the PMD link
- FEC$_O$ corrects errors not corrected by FEC$_I$ and error contributed by the AUIs
- The combined effect of FEC$_I$ and FEC$_O$ results in the target frame loss ratio (FLR) for the PHY.
- FEC$_I$ and FEC$_O$ defined in conjunction with each other.
- PMD and Medium characteristics defined with AUI errors in mind.
- BER trade off between the AUIs and the PMD link.
- More AUIs may be added above the PHY using the MII Extender without affecting PHY performance.
- This PHY/FEC type is new for 802.3.

2023/3/13 IEEE 802.3dj Task Force 11

#1 PHY/FEC architecture considerations V2

Type 3 PHY/FEC

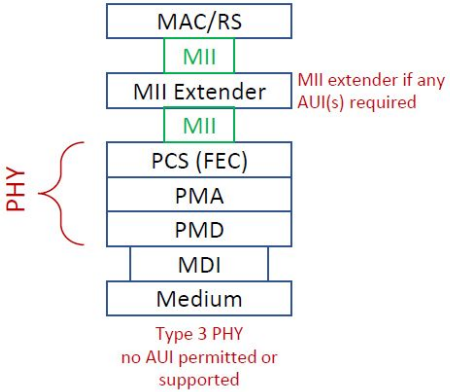

- An FEC spans the PHY link (PCS to PCS) with no AUIs in either PHY.
- If one or more AUIs are required at either end, then an MII Extender is always required.
- The FEC corrects errors contributed ONLY by the PMD link.
- The FEC may take many forms, e.g., RS only, RS + Hamming/BCH (like Type 2), oFEC, etc.
- FEC may be defined independently of other encoding sublayers.
- PMD and Medium characteristics defined independent of AUI characteristics.
- No trade off between the AUIs and the PMD link is required.
- The following PHY is a Type 3 PHY/FEC:
  400GBASE-ZR in 802.3cw

2023/3/13 IEEE 802.3dj Task Force 13

#2 BER considerations for 200 Gb/s per lane AUIs

BER considerations for 200 Gb/s per lane AUIs

Matt Brown, Huawei
Kent Lusted, Intel
Mike Dudek, Marvell
Gary Nicholl, Cisco
Adam Healey, Broadcom

802.3dj electrical ad hoc meeting

#2 BER considerations for 200 Gb/s per lane AUIs

Straw poll

Straw polls #9 and #10 taken at the March 2023 Plenary meeting are shown to the right.

https://www.ieee802.org/3/dj/public/23_03/motions_3dfdj_2303.pdf

Straw poll #9 demonstrates strong desire for a AUI BER limit of 1E-5 and less so 5E-5.

Straw poll #10 demonstrates a split opinion on what to allow for the "high-BER" BER case:

- BER target to be 1E-5
- BER target to be 5E-5
- BER target to be 1E-4

Straw Poll #9

I believe 200G Medium BER C2M AUI specifications will require support for:

A. BER ≤ 1e-5 (per segment)
B. BER ≤ 5e-5 (per segment)
C. BER ≤ 1e-4 (per segment)
D. BER ≥ 1e-4 (per segment)
E. Need more information

(pick one)

Results (all): A: 49 , B: 30 , C: 0 , D: 0 , E: 25

Straw Poll #10

I believe 200G High BER C2M AUI specifications will require support for:

A. BER ≤ 1e-5 (per segment)
B. BER ≤ 5e-5 (per segment)
C. BER ≤ 1e-4 (per segment)
D. BER ≥ 1e-4 (per segment)
E. Need more information

(pick one)

Results (all): A: 16 , B: 47 , C: 17 , D: 1 , E:23

22

2023/4/6 IEEE 802.3dj Task Force 5

#2 BER considerations for 200 Gb/s per lane AUIs

AUI BER Target Options

(A) C2M and C2C AUI BER 1E-5
- Up to 2 AUI per Type 1 or Type 2 PHY
- Minimum channel reach/tolerance per AUI

(B) C2M and C2C AUI BER 2E-5
- Up to 1 AUI per Type 1 or Type 2 PHY
- Improved channel reach/tolerance per AUI

(C) C2M and C2C AUI BER 5E-5
- Up to 2 AUI per Type 2 PHY
- Extender always required for Type 1 PHY
- More improved channel reach/tolerance per AUI

(D) C2M and C2C AUI BER 1E-4
- Up to 1 AUI per Type 2 PHY
- Extender always required for Type 1 PHY
- Most improved channel reach/tolerance per AUI

(E) C2M AUI BER 8E-5 and C2C AUI BER 2E-5
- Up to 1 C2C AUI per Type 1 PHY
- Up to 1 C2C AUI + 1 C2M AUI per Type 2 PHY
- Extender required for C2M AUI for Type 1 PHY
- Best compromise channel reach/tolerance

2023/4/6 IEEE 802.3dj Task Force 11

#2 BER considerations for 200 Gb/s per lane AUIs

AUIs and Type 1 modular PHYs

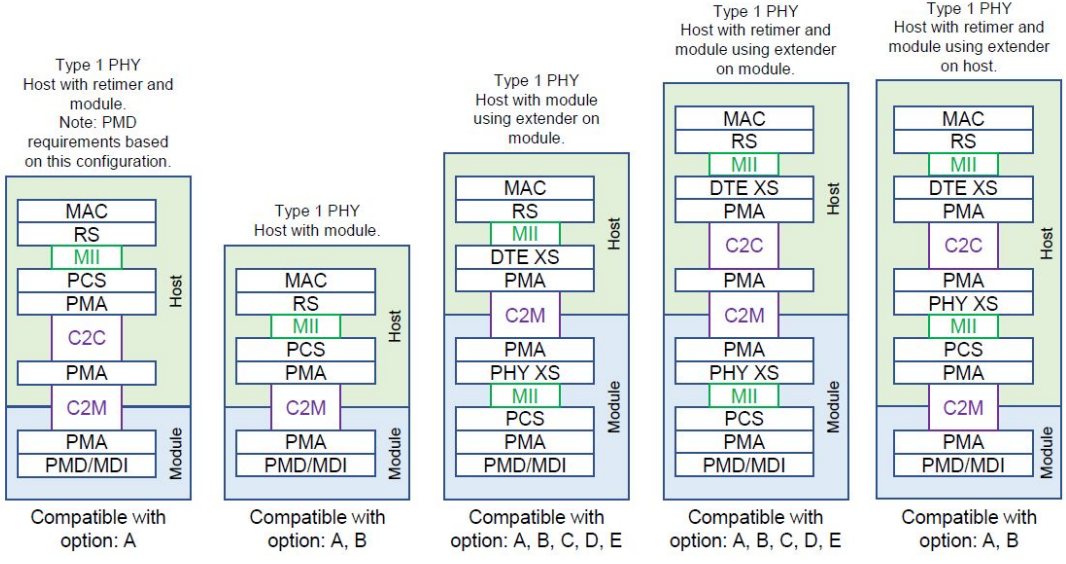

Option	C2M BER	C2C BER
A	1E-5	1E-5
B	2E-5	2E-5
C	5E-5	5E-5
D	1E-4	1E-4
E	8E-5	2E-5

Type 1 PHY (e.g., 800GBASE-DR8)
Total AUI BER per PHY must be less than 2E-5.
Higher BER AUIs force use of an extender with PHY XS and PCS on the module.
Applies to any 100 Gb/s per lane PMD specified thus far.
Applies to any 200 Gb/s per lane Type 1 PMD that might be specified.

2023/4/6 IEEE 802.3dj Task Force 12

#2 BER considerations for 200 Gb/s per lane AUIs

AUIs and Type 1 on-board PHYs

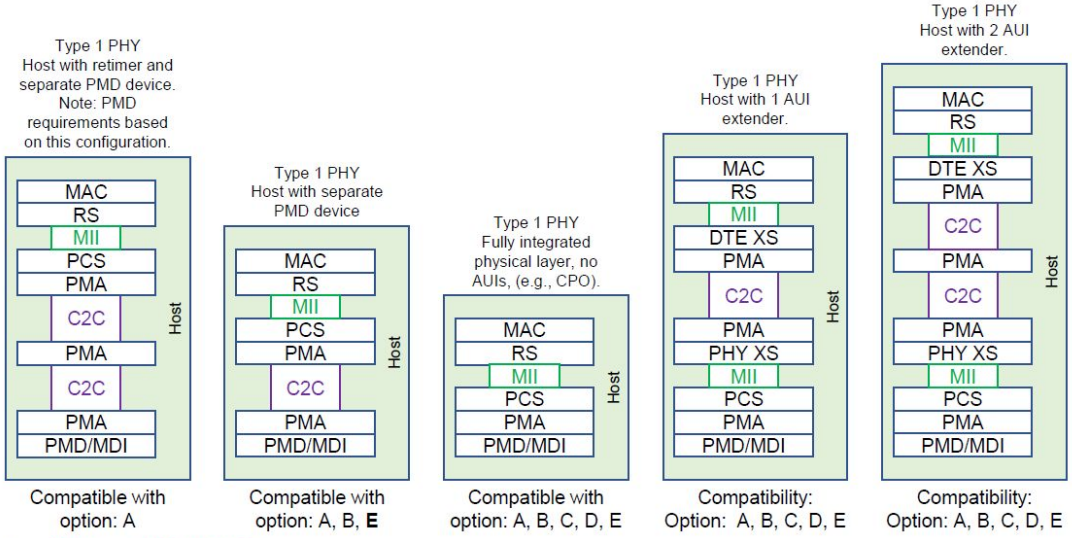

Option	C2M BER	C2C BER
A	1E-5	1E-5
B	2E-5	2E-5
C	5E-5	5E-5
D	1E-4	1E-4
E	8E-5	2E-5

Type 1 PHY (e.g., 800GBASE-CR8)
Total AUI BER per PHY must be less than 2E-5.
Higher BER AUIs force use of an extender with PHY XS and PCS near the PMD.
Applies to any 100 Gb/s per lane PMD specified thus far.
Applies to any 200 Gb/s per lane Type 1 PMD that might be specified.

2023/4/6 IEEE 802.3dj Task Force 13

#2 BER considerations for 200 Gb/s per lane AUIs

AUIs and Type 2 modular PHYs

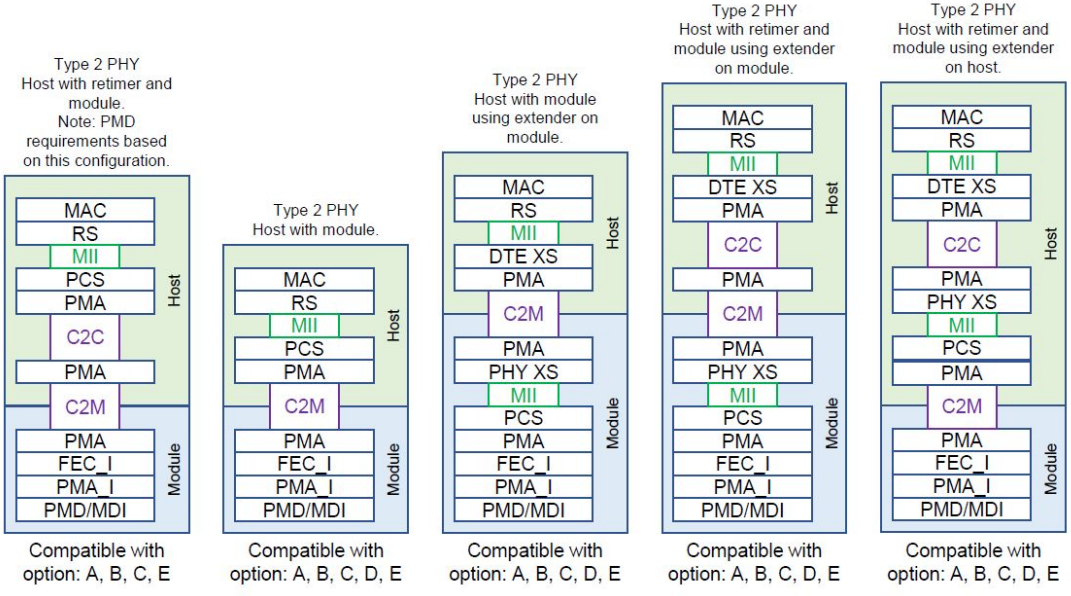

Option	C2M BER	C2C BER
A	1E-5	1E-5
B	2E-5	2E-5
C	5E-5	5E-5
D	1E-4	1E-4
E	8E-5	2E-5

Type 2 PHY (e.g., 800GBASE-FR4)
Total AUI BER per PHY must be less than 1E-4.
Higher BER AUIs force use of an extender with PHY XS and PCS on the module.
Applies to new PHYs with 200 Gb/s per lane PMD and concatenated inner FEC.

2023/4/6 IEEE 802.3dj Task Force 14

#3 MAC link latency considerations

MAC link latency considerations

Matt Brown, Huawei

802.3dj joint optics/logic ad hoc meeting

2023/4/13 IEEE 802.3dj Task Force 1

#3 MAC link latency considerations

MAC-link Latency, no extenders

Case #1: Type 1, no extenders	1.6T	800G	400G	200G
PCS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
Total (ns)	49.8	62.6	62.6	88.2
Case #2: Type 2, 4 CW interleaving, no extenders	**1.6T**	**800G**	**400G**	**200G**
PCS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
FEC_I: interleaver/deinterleaver	0.0	0.0	25.6	51.2
FEC_I: encoder/decoder	23.5	23.5	23.5	23.5
Total (ns)	73.3	86.1	111.7	162.9
Case #3: Type 2, 12 CW interleaving, no extenders	**1.6T**	**800G**	**400G**	**200G**
PCS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
FEC_I: interleaver/deinterleaver	25.6	51.2	128	256
FEC_I: encoder/decoder	23.5	23.5	23.5	23.5
Total (ns)	98.9	137.3	214.1	367.7

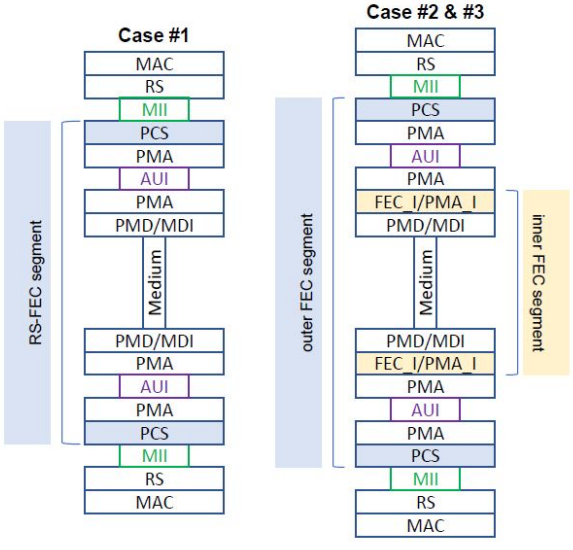

#3 MAC link latency considerations

MAC-link latency, with extenders

Case #4: Type 1, extender at each end	1.6T	800G	400G	200G
XS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
PCS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
XS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
Total	149.4	187.8	187.7	264.6
Case #5: Type 2, 4 CW interleaving, extender at each end	**1.6T**	**800G**	**400G**	**200G**
XS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
PCS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
FEC_I: interleaver/deinterleaver	0.0	0.0	25.6	51.2
FEC_I: encoder/decoder	23.5	23.5	23.5	23.5
XS: RS FEC encoder/decoder	49.8	62.6	88.2	139.4
Total	172.9	211.3	262.5	390.5
Case #6: Type 2, 12 CW interleaving, extender at each end	**1.6T**	**800G**	**400G**	**200G**
XS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
PCS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
FEC_I: interleaver/deinterleaver	25.6	51.2	128	256
FEC_I: encoder/decoder	23.5	23.5	23.5	23.5
XS: RS FEC encoder/decoder	49.8	62.6	62.6	88.2
Total	198.5	262.5	339.3	544.1

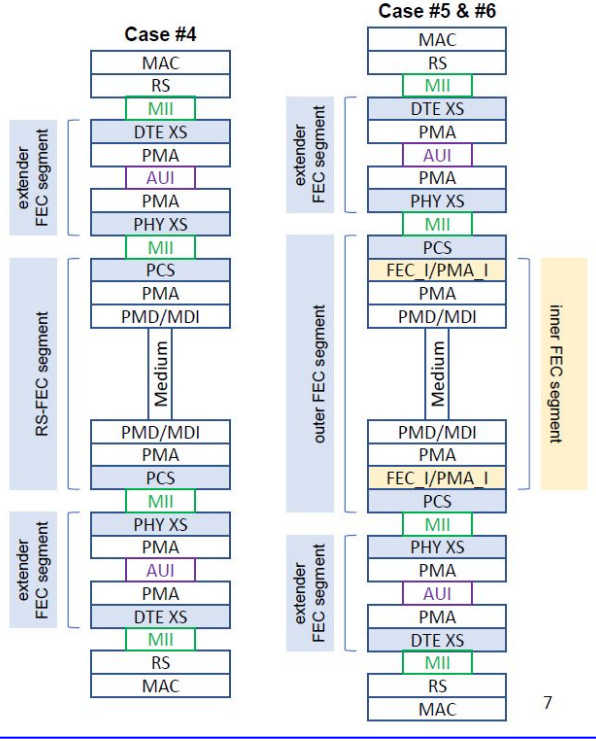

2023/4/13 IEEE 802.3dj Task Force 7

#3 MAC link latency considerations

MAC-link latency, summary and observations

Case	Latency (ns)			
	1.6T	800G	400G	200G
Case #1: Type 1, 4 CW interleaving, no extenders	49.8	62.6	62.6	88.2
Case #2: Type 2, 4 CW interleaving, no extenders	73.3	86.1	111.7	162.9
Case #3: Type 2, 12CW interleaving, no extenders	98.9	137.3	214.1	367.7
Case #4: Type 1, 4 CW interleaving, extender at each end	149.4	187.8	187.8	264.6
Case #5: Type 2, 4 CW interleaving, extender at each end	172.9	211.3	262.5	390.5
Case #6: Type 2, 12 CW interleaving, extender at each end	198.5	262.5	339.3	544.1
Legend | Blue: < 100 ns | Green: 100 ns to 200 ns | Yellow: 200 ns to 300 ns | Red: > 300 ns |				

1. Case #1 provides a minimum latency baseline for comparison.
2. Small relative latency increment from Type 1 (#1) to Type 2 (#2) if interleaving limited to 4 RS CWs.
3. Latency, esp. for 200GE and 400GE, is getting out of hand for Type 2 with 12 CW interleaving (#3) or when using extenders (#4, #5, #6).

2023/4/13 IEEE 802.3dj Task Force 8

should be green

Note that the latency numbers on this slide are only for the physical layer between the MAC/RS and the MDI. It does not include the medium.

Trade-offs

(A) C2M and C2C AUI BER 1E-5

- Up to 2 AUI per Type 1 or Type 2 PHY
- Minimum channel reach/tolerance per AUI

(B) C2M and C2C AUI BER 2E-5

- Up to 1 AUI per Type 1 or Type 2 PHY
- Improved channel reach/tolerance per AUI

(C) C2M and C2C AUI BER 5E-5

- Up to 2 AUI per Type 2 PHY
- Extender always required for Type 1 PHY
- More improved channel reach/tolerance per AUI

(D) C2M and C2C AUI BER 1E-4

- Up to 1 AUI per Type 2 PHY
- Extender always required for Type 1 PHY
- (new) Extender required for Type 2 PHY with 2 AUIs
- Most improved channel reach/tolerance per AUI

(E) C2M AUI BER 8E-5 and C2C AUI BER 2E-5

- Up to 1 C2C AUI per Type 1 PHY
- Up to 1 C2C AUI + 1 C2M AUI per Type 2 PHY
- Extender required for C2M AUI for Type 1 PHY
- Best compromise channel reach/tolerance

Case	Latency (ns)			
	1.6T	800G	400G	200G
Case #1: Type 1, 4 CW interleaving, no extenders	49.8	62.6	62.6	88.2
Case #2: Type 2, 4 CW interleaving, no extenders	73.3	86.1	111.7	162.9
Case #3: Type 2, 12CW interleaving, no extenders	98.9	137.3	214.1	367.7
Case #1: Type 1, 4 CW interleaving, no extenders	49.8	62.6	62.6	88.2
Case #2: Type 2, 4 CW interleaving, no extenders	73.3	86.1	111.7	162.9
Case #3: Type 2, 12CW interleaving, no extenders	98.9	137.3	214.1	367.7
Case #2: Type 2, 4 CW interleaving, no extenders	73.3	86.1	111.7	162.9
Case #3: Type 2, 12CW interleaving, no extenders	98.9	137.3	214.1	367.7
Case #4: Type 1, 4 CW interleaving, extender at each end	149.4	187.8	187.8	264.6
Case #2: Type 2, 4 CW interleaving, no extenders	73.3	86.1	111.7	162.9
Case #3: Type 2, 12CW interleaving, no extenders	98.9	137.3	214.1	367.7
Case #4: Type 1, 4 CW interleaving, extender at each end	149.4	187.8	187.8	264.6
Case #5: Type 2, 4 CW interleaving, extender at each end	172.9	211.3	262.5	390.5
Case #6: Type 2, 12 CW interleaving, extender at each end	198.5	262.5	339.3	544.1
Case #1: Type 1, 4 CW interleaving, no extenders	49.8	62.6	62.6	88.2
Case #2: Type 2, 4 CW interleaving, no extenders	73.3	86.1	111.7	162.9
Case #3: Type 2, 12CW interleaving, no extenders	98.9	137.3	214.1	367.7
Case #4: Type 1, 4 CW interleaving, extender at each end	149.4	187.8	187.8	264.6

Next steps?

- Illustrate trade-off between AUI BER limit and PMD BER limit.
- Regarding the proposed AUI BER cases, determine if the following are acceptable:
  - One AUI per PHY instead of two (see cases B and D)
  - Asymmetric AUI BER limit for C2M vs. C2C (see case E)
  - BER targets realistics (e.g., case B/E 2E-5, case A 1E-5)
  - Others
- Regarding latency:
  - What is a reasonable value?
  - Is it different depending PMD type?

Thanks!