Cl **FM** *SC* **FM** *P* **1** *L* **31** # 47

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

"adds MAC parameters, Physical Layers, and management parameters" but we talk about "the Physical Layer" like "the sky", although we have many "Physical Layer types" (and Physical Layer device types). This should be more like the text in the PAR 5.2.b.
Compare other projects' self descriptions:
adds Physical Layer specifications and management parameters;
includes Physical Layer specifications and management parameters;
adds 2.5 Gb/s, 5 Gb/s, 10 Gb/s, 25 Gb/s and 50 Gb/s Physical Layer specifications and management parameters;
adds 400 Gb/s Physical Layer specifications and management parameters;
adds physical layer specifications and management parameters;
includes Physical Layer specifications and management parameters.
As the PAR says, a feature of this project is "based on 100 Gb/s per lane signaling technology".
I don't see that we are adding any MAC parameters (the PAR says "Define Ethernet MAC parameters" and it looks like we are re-using what we have).

SuggestedRemedy

Change these three texts:
Page 1 line 30:
This amendment includes Media Access Control parameters for 800 Gb/s and Physical Layers and management parameters for 400 Gb/s and 800 Gb/s operation.
Page 3, Abstract:
The amendment adds MAC parameters, Physical Layers, and management parameters for the transfer of IEEE 802.3 format frames at 400 Gb/s and 800 Gb/s.
Page 13, self description:
This amendment includes Physical Layer specifications and management parameters for 400 Gb/s and 800 Gb/s operation.
All to:
This amendment adds Physical Layer specifications and management parameters for 400 Gb/s and 800 Gb/s based on based on 100 Gb/s per lane signaling.

Proposed Response *Response Status* **O**

Cl **FM** *SC* **FM** *P* **8** *L* **15** # 40

Nicholl, Shawn AMD

Comment Type **ER** *Comment Status* **X**

There is a typo in "Gary Nichol".

SuggestedRemedy

It should be "Gary Nicholl".

Proposed Response *Response Status* **O**

Cl **1** *SC* **1.4.145a** *P* **31** *L* **1** # 48

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Missing definitions for 800GAUI-n C2C and 800GAUI-n C2M

SuggestedRemedy

Add 1.4.145a 800 Gb/s Attachment Unit Interface (800GAUI-n): Two kinds of physical instantiation of the PMA service interface to extend the connection between 800 Gb/s capable PMAs over n lanes, used for chip-to-chip (C2C) or chip-to-module (C2M) interconnections. One width of 800GAUI-n is defined: the eight-lane 800GAUI-8 C2C and 800GAUI-8 C2M. (See IEEE Std 802.3, Annex 120E.)

Proposed Response *Response Status* **O**

Cl **45** *SC* **45.2.1.7.5** *P* **40** *L* **3** # 49

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

D1.0 comment 118: Missing entries in transmit fault, *receive fault and transmit disable tables*

SuggestedRemedy

In the tables for receive fault and transmit disable, include rows for 100GBASE-VR1, 100GBASE-SR1, 200GBASE-VR2, 200GBASE-SR2, 400GBASE-VR4, 400GBASE-SR4, 800GBASE-VR8, 800GBASE-SR8 and 400GBASE-DR4, 400GBASE-DR4-2, 800GBASE-DR8, 800GBASE-DR8-2 Revise the rubrics.

Proposed Response *Response Status* **O**

Cl **45** *SC* **45.2.1.138** *P* **44** *L* **25** # 50

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

It's not clear if Table 45-107 - 50GAUI-n, 100GAUI-2, 200GAUI-n, and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 register bit definitions - applies for 100G/lane AUIs or not. Most of 120F implies it doesn't except 120F.3.2.4 Receiver interference tolerance "Receiver interference tolerance is defined by the procedure in Annex 93C with the exception that transmitter equalization is configured by management (see 120D.3.2.3)".

SuggestedRemedy

If it applies, update 45.2.1.135, 45.2.1.136, 45.2.1.137, 45.2.1.138 to include 800GAUI-n.
If it doesn't, say so in these sections because the terms "100GAUI-2, 200GAUI-n, and 400GAUI-n" with unqualified n are too wide now, and address their use (or not) in 120F.3.2.4.
It would help to add these registers to MDIO/PMA variable mapping tables, either in the PMA clauses where there are such tables already, or the AUI annexes.

Proposed Response *Response Status* **O**

Cl **45** *SC* **45.2.3** *P* **46** *L* **26** # 45

Huber, Tom Nokia

Comment Type **E** *Comment Status* **X**

There is some ambiguity in the use of green vs black coloring for the clause references in Table 45-233. In my understanding, green text is used to indicate a reference to a clause (or a table or figure) that is not itself present in this amendment

SuggestedRemedy

Assuming my understanding of the convention is correct, since 45.2.3.25, 45.2.3.49, and 45.2.3.58 are all present in 802.3df (because they are being modified), they should be in black text rather than green text.

Proposed Response *Response Status* **O**

Cl **45** *SC* **45.2.3.19** *P* **47** *L* **28** # 51

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

BASE-R PCS test-pattern control register (Register 3.42)
... Scrambled idle test patterns are defined for 25/40/50/100/200/400GBASE-R PCS only.

SuggestedRemedy

Add 800G

Proposed Response *Response Status* **O**

Cl **45** *SC* **45.2.3.25** *P* **47** *L* **31** # 33

Ran, Adee Cisco

Comment Type **E** *Comment Status* **X**

45.2.3.25 describes the lane alignment register, with one subclause per bit; this continues in 45.2.3.26 and in the new 45.2.3.26a. With 32 lanes, we have 32 subclauses that are essentially the same.

This is repetitive, not helpful for readers, and will require further editorial work when future PCSs are defined (for example 1.6TBASE-R).

It may be better to have one subclause, 45.2.3.25.1, with a full definition of "lane 7 aligned", and have all the remaining bits defined together using something like "defined similarly to 45.2.3.25.1" - as done for example in 45.2.3.49 and 45.2.3.50.

This can remove most of the text in 45.2.3.25 (for register 3.52), 45.2.3.26 (for register 3.53), and 45.2.3.26a (for register 3.54). It may also be possible to merge these three subclauses into one (similar to 45.2.3.50).

The new text should address the number of lanes that exist in every PCS when referring to clause 82, clause 119, and clause 172.

Similar changes can be applied in 45.2.4.16 and 45.2.4.16a for PHY XS, and in 45.2.5.16 and 45.2.5.16a for DTE XS.

SuggestedRemedy

Change the structure as suggested in the comment, with editorial license.

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.3.26.11** P **51** L **34** # 35

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

Stray "1" in "(see 1119.2.6.2.2 and 172.2.6.2.2)."

SuggestedRemedy

Change "1119" to "119".

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.3.26.a** P **49** L **39** # 34

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **X**

The new subclauses 45.2.3.2.26.a through 45.2.3.2.26.d refer to lanes 23 through 20, which exist only in the 800G PCS (clause 172). References to 82.2.19.2.2 are not required in these subclauses.

Similarly in 45.2.3.26a.1 through 45.2.3.26a.8 for lanes 31 through 24.

SuggestedRemedy

In 45.2.3.26.a, change "This bit reflects the state of am_lock[19] (see 82.2.19.2.2) or amps_lock[19] (see 172.2.6.2.2)" to "This bit reflects the state of amps_lock[19] (see 172.2.6.2.2)".

Apply similar changes in 45.2.3.26.b through 45.2.3.26.d and in 45.2.3.26a.1 through 45.2.3.26a.8.

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.3.26a** P **49** L **39** # 123

Slavick, Jeff Broadcom

Comment Type **T** *Comment Status* **X**

df added PCS lanes 20-31, they do not exist in clause 82.

SuggestedRemedy

Remove "am_lock[##] (see 82..2.19.2.2) or" from PCS lanes 20-31

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.3.48a** P **53** L **46** # 52

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

The text should mention that this is an optional feature. Also, 172.3.5 doesn't define the register (Clause 45 does that), it defines the counter.

SuggestedRemedy

For example, change
See 172.3.5 for a definition of this register.
to
See 172.3.5 for a definition of this optional counter.

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.3.48b** P **54** L **20** # 53

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

assignment of bits ... is identical to that of bin 1

SuggestedRemedy

for bin 1 ?

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.3.48b** P **54** L **23** # 54

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

The text should mention that this is an optional feature.

SuggestedRemedy

Add: these counters are optional.

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.3.49** P **54** L **51** # 97

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Subject and verbs number don't match (editorial bug in base document)

SuggestedRemedy

Consider changing
The contents of the Lane 0 mapping register is valid when Lane 0 aligned bit (3.52.0) is set to one and is invalid otherwise.
to content ... is ... is or contents ... are ... are
At some stage, a wider clean-up and harmonisation (contents vs. values) would be helpful.

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.3.63** P **57** L **8** # 98

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

See 119.3.3 and 172.3.3 for a definition of this counter.

SuggestedRemedy

See 119.3.3 or 172.3.3 for a definition of this counter.

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.4.4.a** P **59** L **3** # 46

Huber, Tom Nokia

Comment Type **E** *Comment Status* **X**

The title of the new clause should be 800G capable rather than 400G capable

SuggestedRemedy

Change 400G to 800G.

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.4.4.a** P **59** L **59** # 44

Dudek, Mike Marvell

Comment Type **T** *Comment Status* **X**

The sub-clause title is wrong

SuggestedRemedy

Change "400G capable" to "800G capable"

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.4.16a** P **63** L **25** # 99

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

5register

SuggestedRemedy

insert space. Also in 45.2.5.16a.

Proposed Response *Response Status* **O**

Cl **45** SC **45.2.4.16a.1** P **64** L **18** # 36

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **X**

The new subclauses 45.2.4.16a.1 through 45.2.4.16a.8 refer to lanes 31 through 24, which exist only in the 800GXS (clause 171, based on clause 172 PCS). References to 119.2.6.2.2 are not required in these subclauses.

Also in 45.2.5.16a subclauses for the DTE XS.

SuggestedRemedy

In 45.2.4.16a.1, change "This bit reflects the state of amps_lock[31] (see 119.2.6.2.2 and 172.2.6.2.2)." to "This bit reflects the state of amps_lock[31] (see 172.2.6.2.2).".

Apply similar changes in 45.2.4.16a.2 through 45.2.4.16a.8 and in 45.2.5.16a.1 through 45.2.5.16a.8.

Proposed Response *Response Status* **O**

Cl **45** *SC* **45.2.4.17** *P* **65** *L* **25** # 100

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

"XS described in Clause 118 and Clause 171"
But a product complies to applies to one or the other, at any time.

SuggestedRemedy

XS described in Clause 118 or Clause 171
Also in 45.2.5.17, 45.2.5.22.2, 45.2.5.22.3 and so on

Proposed Response *Response Status* **O**

Cl **45** *SC* **45.2.7.12.3** *P* **78** *L* **10** # 101

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

Base text says "these bits in register 7.48 and register 7.49 indicate the negotiated port type. Only one of these bits is set depending on the priority resolution function" but is this correct? There are FEC options in these registers as well as port types.

SuggestedRemedy

Revise text if appropriate

Proposed Response *Response Status* **O**

Cl **90** *SC* **90** *P* **86** *L* **8** # 5

Brown, Matt Huawei

Comment Type **T** *Comment Status* **X**

IEEE 802.3cx has introduced two new optional signals (RX_NUM_BIT_CHANGE and TX_NUM_BIT_CHANGE) at the PCS service interface (xMII) used for time synchronization that are not defined in the 800GBASE-R PCS or the DTE/PHY 800GXS.

SuggestedRemedy

Define these optional signals in the 800GBASE-R PCS and DTE 800GXS service interfaces (800GMII) and as inputs to the PHY 800GXS (service interface below). For a definition of these signals refer to Clause 90 as appropriate. A presentation will be provided.

Proposed Response *Response Status* **O**

Cl **116** *SC* **116.1.4** *P* **89** *L* **9** # 37

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

Table 116-5 column order is different from the order in the published Std 802.3db-2022 and Std 802.3ck-2022.

SuggestedRemedy

Reorder the columns to align with the published standard.

Proposed Response *Response Status* **O**

Cl **120** *SC* **120.5.6** *P* **90** *L* **6** # 102

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Annex 120F, which specifies the 200GAUI-2 and 400GAUI-4 interfaces for chip-to-chip applications.
Annex 120G, which specifies the 200GAUI-2 and 400GAUI-4 interfaces for chip-to-module applications.

SuggestedRemedy

Add 800GAUI-8

Proposed Response *Response Status* **O**

Cl **120G** *SC* **120G.3.1.5** *P* **246** *L* **26** # 32

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

120.5.11.2.2 is now included in this draft.
Also in 120G.3.2.2, 120G.3.3.5.2, 120G.3.3.5.3, 120G.3.4.3.2, and 120G.3.4.3.3.

SuggestedRemedy

Make all instances of 120.5.11.2.2 active cross references.

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.1** *P* **91** *L* **21** # 103

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

Need a section to explain interoperability of DRn and DRn-2. Compare 140.11 and 151.12 but this is simpler.

SuggestedRemedy

Add a new sentence "The 400GBASE-DR4 and 400GBASE-DR4-2 PMDs can interoperate with each other provided that the fiber optic cabling (channel) characteristics for 400GBASE-DR4 are met, and similarly for 800GBASE-DR8 and 800GBASE-DR8-2". This could be a new subclause 124.11a but because it's so simple this time and it helps the reader understand what these PMDs can be used for, it could be added to 124.1 before 124.1.1 Bit error ratio.

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.1.1** *P* **94** *L* **3** # 130

Opsasnick, Eugene Broadcom

Comment Type **TR** *Comment Status* **X**

Same as previous comment

SuggestedRemedy

Change 1.7E-12 to 3.4E-12

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.2** *P* **94** *L* **39** # 104

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

If as we hope and expect, we set the bit multiplexing rules so that the transition density problem won't happen on 8-lane 800GBASE-R, this sentence and similar ones will need modification. But it remains for 200GBASE-R and 400GBASE-R, so the same point should be made in Clause 167.

SuggestedRemedy

Change: See NOTE
to: For 400GBASE-DR4 and 400GBASE-DR4-2, NOTE
Similarly in 124.7.2
Add equivalent texts in Clause 167

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.7.1** *P* **101** *L* **27** # 105

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

The OMAouter (max) limits are all the same (deliberately, for interoperability)

SuggestedRemedy

Change "values" to "value"

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.7.2** *P* **104** *L* **27** # 106

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

800GBASE-
DR8

SuggestedRemedy

Use non-breaking hyphen?

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.8.1** *P* **107** *L* **9** # 107

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

This has e.g. "3, 5, 6, valid 400GBASE-R signal, or 800GBASE-R signal". 138 has "3, 4, 5, 6, or valid 50GBASE-SR, 100GBASE-SR2, 200GBASE-SR4, or 400GBASE-SR8 signal". 167 has "3, 4, 5, 6, or valid 100GBASE-VR1, 200GBASE-VR2, 400GBASE-VR4, 800GBASE-VR8, 100GBASE-SR1, 200GBASE-SR2, 400GBASE-SR4, or 800GBASE-SR8 signal". Is a non-valid 800GBASE-R signal allowed?

SuggestedRemedy

Change "valid 400GBASE-R signal, or 800GBASE-R signal" to "or valid 400GBASE-R or 800GBASE-R signal" three times.
Maybe in maintenance we should delete "valid" in multiple clauses.

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.8.5** *P* **107** *L* **1** # 1

Stassar, Peter Huawei

Comment Type **TR** *Comment Status* **X**

The text in the last bullet under 124.8.5 "The 400GBASE-DR4-2 or 800GBASE-DR8-2 transmitter is tested using an optical channel with dispersion and insertion loss as specified for 100GBASE-FR1 in 140.7.5.2, and optical return loss at the maximum for optical return loss tolerance specified in Table124–6." was agreed as a resolution to comment #130 to D1.0. The embedded compliance channel requirements are somewhat indirect and it would be much clearer if a special section be created with details and especially a Table with channel requirements, following the style of 151.8.5.1, especially because there is no precedence for channel requirements for DR type PMDs over 2 km.

SuggestedRemedy

Create a new subclause 124.8.5.1 with channel requirements for 400GBASE-DR4, 400GBASE-DR4-2, 800GBASE-DR8, and 800GBASE-DR8-2, following the specific proposal in a presentation

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.8.9** *P* **109** *L* **1** # 108

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Missing tab or format issue

SuggestedRemedy

fix

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.8.9.1** *P* **109** *L* **11** # 38

Ran, Adee Cisco

Comment Type **E** *Comment Status* **X**

The parameter in this subclause is called "receiver sensitivity (OMA_outer)" in Table 124-7 and in 124.8.9.2. For 400GBASE-DR4 it is optional, but I assume the name should be the same.

SuggestedRemedy

Insert "(OMA_outer)" after "receiver sensitivity", 3 instances in this subclause.

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.11.3.3** *P* **113** *L* **33** # 39

Ran, Adee Cisco

Comment Type **E** *Comment Status* **X**

IEC 61754-7-4 does not appear in the normative references list (1.3); only 7-1 and 7-2 are listed.

SuggestedRemedy

Add a reference to the appropriate document in 1.3

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.11.3.3** *P* **113** *L* **35** # 6

Ran, Adee Cisco

Comment Type **E** *Comment Status* **X**

"interface 7-4-1: <…>" - where is that one defined? Is it also IEC 61754-7-4?

SuggestedRemedy

Add "as defined in IEC 61754-7-4" after the interface name.

(If it's another document, add that instead, and make sure the document is listed in 1.3).

Proposed Response *Response Status* **O**

Cl **124** *SC* **124.12.4.4** *P* **115** *L* **24** # 109

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Items to OM12 depend on PMD type

SuggestedRemedy

Add major options for PMD types. These items will be conditionally mandatory.
Also, adjust:
124.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-DR4
F1 Compatible with 400GBASE-R PCS and PMA

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.1** *P* **116** *L* **39** # 110

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

The document uses a mixture of 800GMII extender and 800GMII Extender (aside from "800GMII Extender Sublayer"

SuggestedRemedy

Make consistent

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.1** *P* **117** *L* **4** # 7

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

In the published 802.3ck-2022, the definition of frame loss ratio is in 1.4.344. Also in 163.1.

SuggestedRemedy

Change "1.4.275" to "1.4.344", in both clauses.

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.1** *P* **117** *L* **7** # 131

Opsasnick, Eugene Broadcom

Comment Type **TR** *Comment Status* **X**

The FLR value that results from 2.4E-4 BER is referred to in two places, in lines 7 and 10:

"This BER allocation enables a frame loss ratio lower tha 9.2 x 10^-13 after processing by the PCS …".

And on line #10. "… to maintain a frame loss ratio lower than 9.2 x 10^-13."

This FLR value, 9.2E-13, corresponds to a "non-interleaved" RS(544,514) FEC as used in the 50G & 100G PCS. The value should be changed to 1.7E-12 for 200G and 400G PCS which have 2-way interleaved FEC, and should be changed to 3.4E-12 for 800G PCS with 4-way interleave FEC.

This same issue was addressed in comment #62 of 802.3bs D1.3:
https://www.ieee802.org/3/bs/comments/P802d3bs_D1p3_comments_final_ID.pdf#page=13

The FLR scaling factor of (1 +MFC)/MFC should be modified to be (1 + 2*MFC)/MFC for the 2-way interleave PCS and to (1 + 4*MFC)/MFC for the 4-way interleaved PCS.

SuggestedRemedy

Remove 800G from this paragraph. Keep origin paragraph referring to 200G/400G, but change the FLR value to 1.7E-12.

Add a similar pargraph after this one with references changed from 200G/400G to 800G and FLR value to 3.4E-12.

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.7** *P* **122** *L* **47** # 111

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Register for lanes 1 to 3 7 are located at an offset from the lane 0 register.

SuggestedRemedy

Suggest: Registers for lanes 1 to 3 7 are located at offsets from the lane 0 register.

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.8.1** *P* **123** *L* **37** # 8

Ran, Adee Cisco

Comment Type **E** *Comment Status* **X**

The location of the "NOTE" in Figure 162-2 is unusual.

SuggestedRemedy

Move the NOTE label to the lower left of the figure.

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.8.11.1** *P* **130** *L* **11** # 112

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **X**

These default seeds are different to the ETC defaults. Also, as the Training state machines on each lane are independent, there is no guarantee that setting the seed will have the desired effect of de-correlating the signals of lanes that share a polynomial. It would be better to give the implementer the freedom to make a good choice for his implementation. 45.2.1.168 already says "should".

SuggestedRemedy

Change "the default value of seed_i" to "the recommended default value of seed_i"

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.9.4** *P* **125** *L* **15** # 9

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

In the published 802.3ck-2022, the subclause reference for "Signaling rate" in Table 162-11 has been deleted. The change in the first row is not required anymore.

SuggestedRemedy

Delete the struck-out subclause reference, and delete "the first row and" in the editorial instruction.

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.14.3** *P* **129** *L* **27** # 113

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

!CR4:O.2 looks like a copy and paste from 802.3cd

SuggestedRemedy

I think it should be CR1:O.2. Also for KR in 163.13.3

Proposed Response *Response Status* **O**

Cl **162** *SC* **162.14.3** *P* **129** *L* **35** # 10

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

In the published 802.3ck-2022, the reference for item PCS400 is 162.1

SuggestedRemedy

Change 162.9.4.8 to 162.1

Proposed Response *Response Status* **O**

Cl **163** *SC* **163.1** *P* **131** *L* **7** # 132

Opsasnick, Eugene Broadcom

Comment Type **TR** *Comment Status* **X**

Same as previous comment.

SuggestedRemedy

FLR for 200G/400G should be changed to 1.7E-12. For 800G, FLR should be changed to 3.4E-12.

Proposed Response *Response Status* **O**

Cl **167** SC **167.1.1** P **141** L **46** # 133

Opsasnick, Eugene Broadcom

Comment Type **TR** *Comment Status* **X**

Same as previous comment, except the value is already updated to 1.7E-12 in part that instructs to "Insert a new third paragraph in 167.1.1"

SuggestedRemedy

Change 1.7E-12 to 3.4E-12 in two places

Proposed Response *Response Status* **O**

Cl **167** SC **167.8.1** P **148** L **41** # 11

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

120.5.11.2.2 is now included in this draft.

SuggestedRemedy

Make 120.5.11.2.2 an active cross reference.

Proposed Response *Response Status* **O**

Cl **167** SC **167.9.2** P **150** L **41** # 114

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

800GBASR-VR8

SuggestedRemedy

800GBASE-VR8

Proposed Response *Response Status* **O**

Cl **167** SC **167.10.3.1a** P L # 115

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

Discussions at the last round indicated that "Option A (24 fibers in two rows in one connector shell) is the least used of three connector formats for 8-lane multimode. It should not be the first option.

SuggestedRemedy

Take whatever polls are necessary to establish consensus and delete Option A.

Proposed Response *Response Status* **O**

Cl **167** SC **167.10.3.4** P **155** L **12** # 12

Ran, Adee Cisco

Comment Type **E** *Comment Status* **X**

"interface 7-4-1: <…>" - where is that one defined? Is it also IEC 61754-7-4?

SuggestedRemedy

Add "as defined in IEC 61754-7-4" after the interface name.

(If it's another document, add that instead, and make sure the document is listed in 1.3).

Proposed Response *Response Status* **O**

Cl **167** SC **167.11.4.6** P **158** L **13** # 116

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

These PICS need work to align them to the clause

SuggestedRemedy

Removing Option A will make this task simpler

Proposed Response *Response Status* **O**

Cl **167** *SC* **167.11.4.6** *P* **158** *L* **31** # 13

Ran, Adee Cisco

Comment Type **T** *Comment Status* **X**

The status of items OC15 through OC20 includes "AFI:", which makes them conditional on an angled fiber interface. However, the reference 167.10.3.4 also specifies flat fiber interfaces.

The value/comment needs to be different for angled and flat.

SuggestedRemedy

Add or change PICS items for 167.10.3.4 as appropriate.

Proposed Response *Response Status* **O**

Cl **167** *SC* **167.11.4.6** *P* **158** *L* **37** # 14

Ran, Adee Cisco

Comment Type **T** *Comment Status* **X**

The value/comment for OC18 includes "or per ANSI/TIA-604-18-A designation FOCIS 18 A-1-0 or FOCIS 18 R-1x16-1-0-1-2-0".

These do not appear in the referenced subclause 167.10.3.4.

Also in OC19.

SuggestedRemedy

Align the value/comment and the subclause text.

Proposed Response *Response Status* **O**

Cl **169** *SC* **169.5** *P* **167** *L* **14** # 117

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

"as illustrated in Figure 169–7 (single 800GAUI-n interface) and Figure 169–8 (multiple 800GAUI-n interfaces)": tautology, ambiguous as one could say that a physically instantiated AUI has an interface at each end, and the figure titles do this differently.

SuggestedRemedy

Change to "as illustrated in Figure 169-7 for a PHY with a single 800GAUI-n and in Figure 169-8 for a PHY with multiple 800GAUI-n"
In Annex 173A, adjust figure titles to be consistent with the way Figure 169-7 and Figure 169-8 are done.

Proposed Response *Response Status* **O**

Cl **169** *SC* **169.5** *P* **169** *L* **8** # 118

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **X**

These Skew limits were created 14 years ago assuming FPGAs clocked at 160 Mb/s (see e.g. https://ieee802.org/3/ba/public/may08/giannakopoulos_01_0508.pdf). As the number of bits to buffer goes up with the width, we should revisit this and take out the padding that modern FPGAs don't need. For example, if we assume 644 Mb/s clocking, we might save 38 ns out of a total of 180 ns, which is enough to be interesting.
With the current limits, the Skew can be significantly more than the FEC block time (25.6 ns), which is unfortunate; we would get better protection against error bursts on the line if the four FEC streams overlapped in time.

SuggestedRemedy

Take out the allocation for slow wide FPGA internal interfaces, that are no longer necessary, from the allocations for PMA Skew. This could be 3/4 * 12.8 ns for each PMA. Make coordinated changes in the subclauses that repeat the Skew limits (e.g. 120.5.3, 124.3.2, 162.6.2, 163.6.2, 167.3.2, 171.8.4.2).

Proposed Response *Response Status* **O**

Cl **169** *SC* **169.5** *P* **169** *L* **9** # 15

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **X**

The skew constraints for 800 Gb/s in ns are the same as those for earlier generations, as early as 40 Gb/s, Table 80-8.

The size of PCS buffers required for deskewing grows linearly with the data rate; the size is quite large even at 400G, and would be doubled at 800G, due to the doubling of the number of PCS lanes. The current skew limit of 160 ns at the PCS receive requires about 150 kilobits per port just for deskewing. This affects both latency and power consumption across the industry.

The original skew limits were probably exaggerated even for 40G, and there is no need to carry them on for new technologies and new PCS designs. The numbers we set in 802.3df will also affect hosts and modules (with XS) in 802.3dj, so are worth considering carefully now.

The numbers below are in "UI" of a PCS lane equal to 37.64706, although most skews are created on physical interfaces where the real UI is 18.82 ps.

- Limit of Skew generated at SP1 is currently 770 "UI", it can safely be reduced to 256 "UI" (512 UI of a PMD, or 8 clock cycles in a typical SerDes).
- Limit of Skew generated at SP2 is currently 1142 "UI", allowing additional skew of ~350 "UI" by the PMA in the module; this can safely be reduced to 128 "UI" (4 clock cycles of a typical SerDes; 384 "UI" including the reduced SP1)
- Limit of Skew generated at SP3 is currently 1434 UI, allowing additional skew of ~290 "UI" by the PMD; this can safely be reduced to 128 "UI" (4 clock cycles of a typical SerDes; 512 "UI" including the reduced SP2)
- Limit of Skew generated at SP4 is currently 3559 UI, allowing additional skew of 2125 "UI" (80 ns, ~16 m of fiber) by the media; this can safely be reduced to ~4 m of fiber or 512 "UI" (1024 "UI" including the reduced SP3)
- Limit of Skew generated at SP5 is currently 3852 UI, allowing additional skew of ~300 "UI" by the PMD; this can safely be reduced to 128 "UI" (4 clock cycles of a typical SerDes; 1152 "UI" including the reduced SP4)
- Limit of Skew generated at SP6 is currently 4250 UI, allowing additional skew of ~400 "UI" by the PMA; this can safely be reduced to 128 "UI" (4 clock cycles of a typical SerDes; 1280 "UI" including the reduced SP5)
- Limit of Skew generated at the PCS receive is currently 4781 UI, allowing additional skew of ~530 "UI" by the PMA collocated with the PCS; this can safely be reduced to 128 "UI" (4 clock cycles of a typical SerDes; 1408 "UI" including the reduced SP6)

The result could be a reduction of the allowed skew by 70%, which allows a significant saving in PCS buffer size.

The suggested remedy lists skew as an exact number of "UI" and an approximate number in ns (unlike the current table). It can also be the other way around.

SuggestedRemedy

Change the skew table to
Skew point | max skew ns (approx.) | max skew UI
==========
SP1 | 6.8 | 256
SP2 | 10.2 | 384
SP3 | 13.6 | 512
SP4 | 27.2 | 1024
SP5 | 30.6 | 1152
SP6 | 34 | 1280
PCS input | 37.4 | 1408

Change skew limits in the PCS, PMA, and PMD clauses accordingly.

Proposed Response *Response Status* **O**

Cl **169** *SC* **169.5** *P* **169** *L* **18** # 3

de Koos, Andras Microchip Technology

Comment Type **T** *Comment Status* **X**

As explained in 802.3cx (D3.3) Clause 90.7.3, transmitter skew can be problematic for timestamping. This should be flagged when discussing the skew limits for SP1, SP2, SP3.

"Lane skew is possible on a transmitter with multiple PCS and PMA/PMD lanes when these lanes have different static latencies such that their alignment markers appear staggered as they depart the device at the MDI output. Since transmit skew in series with medium skew is not strictly additive, transmit skew can contribute to time synchronization error by obscuring the actual latency of the medium. Transmit skew is expected to be minimized, ideally to zero, representing an ideal case for the accuracy of a TimeSync Client."

SuggestedRemedy

After Table 169-5, add a note that for 800GEGb/s devices that implement timestamping, transmitter skew (skew points SP1, SP2 and SP3) should be minimized, ideally to zero. Can point to Clause 90.7.3.

Proposed Response *Response Status* **O**

Cl **169** *SC* **169.5** *P* **169** *L* **38** # 16

Ran, Adee Cisco

Comment Type **T** *Comment Status* **X**

Skew variation is dominated by SP4 minus SP3 - the media contribution - which is currently 3.4-0.6=2.8 ns, corresponding to more than 0.5 m of fiber.

It seems unlikely that fibers dynamically "shrink" or "expand" (effectively) that much.

It is suggested to reduce this contribution by a factor of 4, to 0.7 ns (about 14 cm of fiber). This will affect the maximum skew variation at points below SP4 too.

SuggestedRemedy

Change the values in the SP4 row and below:
SP4 | 1.3 | 69
SP5 | 1.5 | 80
SP6 | 1.7 | N/A
At PCS receive | 1.9 | N/A

Change skew variation limits in the PCS, PMA, and PMD clauses accordingly.

Proposed Response *Response Status* **O**

Cl **169** *SC* **169.6** *P* **169** *L* **48** # 17

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **X**

The FEC degrade functionality in clause 119 is not useful, especially at 100 Gb/s per lane signaling. It is now common knowledge that correlated errors (which can occur due to DFEs and other reasons) can cause FEC failure even when the average SER is "good", so the average SER that this feature measures is not enough to predict when errors are going to occur.

We now have a better way to predict FEC performance through the codeword bin counters, which can be accessed through management; the FEC degrade "feature" should not be carried over to 800G Ethernet.

SuggestedRemedy

Delete 169.6 and 171.5, and edit other places where FEC degrade is mentioned in this draft to remove this feature.

Replace all references to the FEC degrade in clause 119 with text stating that FEC degrade is not defined for the 800GBASE-R PCS and XS.

Proposed Response *Response Status* **O**

Cl **169** *SC* **169.8** *P* **171** *L* **9** # 119

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Same as what?

SuggestedRemedy

Change "conforms to the same notation and conventions used in 21.6" to "conforms to the notation and conventions used in 21.6" or "conforms to the same notation and conventions as used in 21.6".

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.1** *P* **179** *L* **26** # 124

Slavick, Jeff Broadcom

Comment Type **T** *Comment Status* **X**

Table 171-1 lists the AUI as Optional but at least one of them must exist.

SuggestedRemedy

Attach a footnote to each Optional that specifies that at least one is required.

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.1.1** *P* **180** *L* **39** # 120

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Some more basic, strategic concepts are missing from this list

SuggestedRemedy

Say that the 800GMII Extender uses two PCS-like entities, DTE 800GXS and PHY 800GXS, that communicate to each other over an 800GAUI-n. Say that the DTE 800GXS is similar to the Clause 72 PCS, and the PHY 800GXS is similar but used "upside down".

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.1.1** *P* **180** *L* **40** # 121

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

The 800GXS doesn't support physical instantiations of the 800GAUI-n. The 800GMII Extender uses them, or it. The XGSs connect to them or it. There are two 800GXS, not the same as each other. A 800GAUI-n has to be physical.

SuggestedRemedy

Change "The 800GXS leverages all functions in the Clause 172 PCS and supports physical instantiations of the 800GAUI-n" to "Each 800GXS leverages all functions in the Clause 172 PCS and connects to a 800GAUI-n, as shown in Figure 171-1"

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.2** *P* **180** *L* **45** # 55

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

FEC degrade is an optional feature of the PCS. As the AUI inside the 800GMII Extender shouldn't be making many errors, the main interest for the DTE 800GXS is in receiving any FEC degrade from the line PCS in the module. The host could have got similar information from the module's management interface. So if it's optional for the PCS it should be optional for the DTE 800GXS, although one could split receiving a FEC degrade signal, and generating FEC degrade from a bad BER, into two separate options.

SuggestedRemedy

Delete "with the additional FEC degrade signaling defined in 171.5"

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.3** *P* **181** *L* **3** # 2

de Koos, Andras Microchip Technology

Comment Type **T** *Comment Status* **X**

From 802.3cx (D3.3) Clause 90.7.2, an MII extender device should avoid insertion/deletion of alignment markers and idles. But as described in Clause 171, there is no provision to do this in the 800GXS Sublayer.

I can make a presentation to explain this further, if needed.

"NOTE 5—When TX_NUM_BIT_CHANGE and RX_NUM_BIT_CHANGE are not available (e.g., over physical interfaces such as instantiated xMII or AUI), it is recommended to avoid insertion and removal of Idles, alignment markers, and codeword markers in the sublayers below the xMII/AUI, when possible, to reduce timestamping accuracy impairments (see Annex 90A)."

SuggestedRemedy

There should be a provision that an MII Extender device (PHY 800GXS + standard 800G PHY) can optionally avoid any modification to the MII stream, and any modification of the position of alignment markers or codeword markers with respect to the MII, between the input and output.

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.3** *P* **181** *L* **8** # 56

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

The FEC degrade feature is not very interesting for the errors on the AUI inside the 800GMII Extender, and if it is optional for the PCS, it should be optional for the PHY 800GXS in the same module.

SuggestedRemedy

Delete "Additional FEC degrade signaling defined in 171.5 is included."

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.3** *P* **182** *L* **9** # 57

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Figure 171-2 contains the rogue capitals that have just been removed from Figure 172-2. Also, "66B" should be "66-bit", twice

SuggestedRemedy

Fix

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.3** *P* **182** *L* **45** # 58

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

As in Figure 172-2, functional block diagram for the PCS

SuggestedRemedy

Please indicate the position of the 800GMII

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.3.1** *P* **183** *L* **3** # 126

Slavick, Jeff Broadcom

Comment Type **T** *Comment Status* **X**

Isn't Figure 169-3 a better reference?

SuggestedRemedy

Change the Figure referecne to 169-3

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.3.2** *P* **183** *L* **23** # 18

Ran, Adee Cisco

Comment Type **E** *Comment Status* **X**

"defined for the 32:8 PMA defined in 173.3"
The first "defined" is superfluous. Compare to the previous paragraphs, which do not have it.

SuggestedRemedy

Delete the first instance of "defined".

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.5** *P* **183** *L* **46** # 4

Brown, Matt Huawei

Comment Type **T** *Comment Status* **X**

Support of FEC degrade in the 800GMII extender sublayers requires that the 800GXS uses monitor states in the PCS below, but the base standard (Clause 117, 118, and 119) do not define a signal across the 800GBASE-R PCS and DTE 800GXS service interfaces (800GMII). Instead Clause 118 makes reference to status bits in the PCS (Clause) 119. Keeping with common conventions, signals across the PCS service interface should be defined to convey the degrade state and the signal referenced in each sublayer.

SuggestedRemedy

Update 800GMII to include FEC degrade signaling across the 800GMII. Update the 800GBASE-R PCS to include the generation of the FEC degrade signal. Update the PHY 800GXS to use the new FEC degrade signal rather than the status bit(s) in the PCS. A presentation will be provided.

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.5** *P* **183** *L* **49** # 59

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

According to 171.8.3, FEC degrade for 800GXS. According to 116.6 and 118.5.3, it's optional for 200GXS and 400GXS. It's optional for the 800GBASE-R PCS too.

SuggestedRemedy

Add a sentence: FEC degrade signaling is optional.

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.7** *P* **185** *L* **46** # 60

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Broken variable name but it looks like there is space in this table to avoid it

SuggestedRemedy

Make the right column two characters wider, making the third column narrower.

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.7** *P* **186** *L* **6** # 125

Slavick, Jeff Broadcom

Comment Type **T** *Comment Status* **X**

Table 171-3 and 171-5 map the FEC_cw_counter and FEC_codeword_error_bin counters to PCS space.

SuggestedRemedy

Create new registers in the PHY XS and DTE XS MDIO space for these counters and map them to the new registers appropriately.

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.8.3** *P* **189** *L* **12** # 41

Nicholl, Shawn AMD

Comment Type **E** *Comment Status* **X**

Fourth row of table has text wrapped in first column.

SuggestedRemedy

Propose to widen the first column slightly to prevent wrap of *800GXS text.

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.8.4.3** *P* **190** *L* **50** # 61

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

According to 82.2.3.6, "deletion" doesn't get a special capital letter

SuggestedRemedy

Change Deletion to deletion

Proposed Response *Response Status* **O**

Cl **171** *SC* **171.8.4.4** *P* **191** *L* **5** # 62

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

The two scramblers must be desynchronised to it's not exactly as in Clause 49 without qualification

SuggestedRemedy

Point to 172 instead of 49

Proposed Response *Response Status* **O**

Cl **172** *SC* **172** *P* **194** *L* **1** # 63

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

This style of title follows 49. Physical Coding Sublayer (PCS) for 64B/66B, type 10GBASE-R. "for" isn't great but I see why it was there in 49. Back then, 64B/66B was new and a big thing, to be contrasted with 8B/10B. Here, it's only an internal step on the way to 256B/257B with RS-FEC. Type R is very familiar now.
By the way, the copy in 172.7.2.2 differs.

SuggestedRemedy

Change the title of 172 from "172. Physical Coding Sublayer (PCS) for 64B/66B, type 800GBASE-R" to 172. Physical Coding Sublayer (PCS), type 800GBASE-R"
Here and in the PICS.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.1.3** *P* **194** *L* **47** # 64

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

There are three things with essentially the same title:
172. Physical Coding Sublayer (PCS) for 64B/66B, type 800GBASE-R
172.1.3 Physical Coding Sublayer (PCS)
172.2 Physical Coding Sublayer (PCS)
A new reader does not see something that indicates it's an introduction.
Compare e.g. 171:
171. 800GMII Extender and 800GMII Extender Sublayer (800GXS)
171.1.1 Summary of major concepts
(and then the various hard specification subclauses are one level higher)
Also note
173.1.3 Summary of functions
173.4 Functions within the PMA

SuggestedRemedy

Change the title of 172.1.3 to "Summary of major concepts", "Principal features of the 800GBASE-R PCS" or equivalent
Change the title of 172.2 to "Detailed specifications of the 800GBASE-R PCS" or equivalent
For consistency, 137.4 Functions within the PMA could be something like Detailed specifications of functions within the PMA

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.1.3** *P* **194** *L* **53** # 128

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

In Section 8, "based on" appears 75 times, "based upon" 9 times. In this document, "based on" appears 11 times, "based upon" 5 times

SuggestedRemedy

Maybe we should change all the new "based upon" to "based on"

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.1.3** *P* **195** *L* **5** # 66

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Scrambling, lane synchronisation and lane re-ordering (or identification) are important enough that they should appear in this list, particularly as alignment markers appear without explanation at item e.

SuggestedRemedy

Please add them

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.1.3** *P* **195** *L* **5** # 65

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Reed-Solomon encoding (decoding) the 257-bit blocks. As this code is "systematic", it can be decoded by throwing away the parity block, but that's not the point. Also, it would be good to mention FEC.

SuggestedRemedy

Change to "Encoding (decoding with correction) the 257-bit blocks with Reed-Solomon FEC

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.1.4** *P* **195** *L* **21** # 67

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

"It is important to note that": pompous fluff, and singling out a point that isn't so special. Section 8, for example, uses "while this specification defines" three times with "It is important to note that" and three times without.

SuggestedRemedy

Delete. This is the only one in this draft.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.1** *P* **197** *L* **31** # 68

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Change of subject without indication. According to line 5, there are only two processes, Tx and Rx.

SuggestedRemedy

Insert "In | for the receive direction | Receive process". Reconcile whether PCS Synchronization process is a component of the Receive process or not.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.1** *P* **197** *L* **36** # 69

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

and then reordered, deskewed, and the align_status flag is set.

SuggestedRemedy

and then reordered and deskewed, and the align_status flag is set.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.1.1** *P* **198** *L* **28** # 42

Nicholl, Shawn AMD

Comment Type **TR** *Comment Status* **X**

To allow use of the PCS stateless encoder at both 400 Gb/s and 800 Gb/s data rates, place the new sub-clause 172.2.4.1.1 (PCS stateless encoder) into Clause 119 directly.

SuggestedRemedy

Propose to create a new sub-clause 119.2.4.1.1 containing the current text of 172.2.4.1.1 (PCS stateless encoder), except replace (twice) "800GMII vector(s)" with "MII vector(s)". Or replace with "tx_raw vector(s)" instead.

In sub-clause 119.2.4.1 (Encode and rate matching), change "... state diagram as shown in Figure 119-14." to "... state diagram as shown in Figure 119-14 or (for 400GBASE-R PCS or 800GBASE-R PCS) by the stateless encoder specified in 119.2.4.1.1."

In sub-clause 172.2.4.1 (Encode, rate matching, and block distribution), change "stateless encoder specified in 172.2.4.1.1." to "stateless encoder specified in 119.2.4.1.1."

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.1.1** *P* **198** *L* **32** # 70

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

alternate ... alternative: shouldn't it be the same word each time? But the second one is unnecessary and there is no other stateless encoder.

SuggestedRemedy

Delete "alternative". Also in 172.2.5.8.1.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.1.1** *P* **198** *L* **37** # 19

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **X**

Table 172–1 has "reset" as the first column, but reset is not defined in clause 172.

Similarly, LBLOCK_T, EBLOCK_T, T_TYPE and the block types C, T, S, D, ENCODE, and tx_raw are not defined anywhere in this draft.

SuggestedRemedy

Add text pointing to the definitions of LBLOCK_T and EBLOCK_T in 119.2.6.2.1, reset and tx_raw in 119.2.6.2.2, and T_TYPE and ENCODE in 119.2.6.2.3.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.1.1** *P* **198** *L* **37** # 71

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Usually we write function(something) with no space

SuggestedRemedy

Delete "alternative". Also in Table 172-4.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.1.1** *P* **198** *L* **39** # 72

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

Because Figure 119-14 specifically doesn't apply, we need cross-references to define LBLOCK_T, C, T, S, ENCODE and so on

SuggestedRemedy

Provide the cross-references. Also for the stateless decoder in 172.2.5.8.1.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.1.1** *P* **198** *L* **40** # 73

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

No indication as to how to add block types

SuggestedRemedy

If you mean "or" as in Table 172-4, change + to or, 4 times.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.1.1** *P* **198** *L* **40** # 20

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **X**

Table 172-1 column "T_TYPE (tx_raw_i-1)" has cells with the strings "C + T" and "S + D". These seem to be based on the state diagram convention that "+" is a logical-OR, but this is not a state diagram, and the letters are not conditions, so it isn't very clear. Using "or" would be preferable (as in the similar Table 172–4).

In addition, for each of these two strings there are two rows with two values in "T_TYPE (tx_raw_i)" column; these can be merged with the word "or" as well.

SuggestedRemedy

Merge rows 2 and 5 to a single row with columns:
"0 | C or T | C or S | ENCODE (tx_raw_i)".
Merge rows 3 and 4 to a single row with columns:
"0 | S or D | D or T | ENCODE (tx_raw_i)".

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.3** *P* **199** *L* **10** # 21

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **X**

If the two scramblers are initialized to the same value and have the same input, their outputs will be equal. This may cause various problems when PCSLs from the two flows are muxed together into the same physical lane, such as pairs of identical PAM4 symbols.

The scrambler specification goes back to 49.2.6 which says "there is no requirement on the initial value for the scrambler". But implementations may force some initial value, e.g. during reset, and with the new concern, some guidance should be given.

A presentation with more details will be supplied.

SuggestedRemedy

Add the following paragraph in 172.2.4.3:
"Although there is no requirement on the initial value of each scrambler, if an implementation sets the scrambler state at any time (e.g., when reset is asserted), the two scramblers should be set to different states."

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.3** *P* **199** *L* **10** # 74

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **X**

The two scramblers must be desynchronised to avoid a gross failure of signal statistics after restricted bit multiplexing the two flows. It is hard to say whether they need to be offset by more than the Skew limit at SP1 or whether any offset is enough. However, it's very easy to provide a big offset by choosing the scramblers' initial conditions appropriately.

SuggestedRemedy

Say that the two scramblers should be started so that their outputs are offset by at least enough so that they will not be aligned when Skewed as allowed when forming the 8-lane PMA/PMD signals.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.4** *P* **199** *L* **23** # 75

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

"n"

SuggestedRemedy

Usually n is a number of things (cardinal number) and i is an index (ordinal) number. Wouldn't i (italic) be more usual?

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.4** *P* **199** *L* **25** # 76

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

It would help the reader understand tables 172-2 and 3 to provide some of the information from 119.2.4.4. Also to save reverse engineering the tables, we can say what the difference between the tables is.

SuggestedRemedy

Add: In Table 172-2 and Table 172-3, CM0 to CM5 are the same for all PCS lanes, UM0 to UM5 are unique per lane, and UP0 to UP2 are a pad per lane. UP0 to UP2 for lanes 16 to 31 are the same as those for lanes 0 to 15, respectively.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.4** *P* **200** *L* **4** # 22

Ran, Adee Cisco

Comment Type **E** *Comment Status* **X**

The PCS AM tables do not convey to the reader the structure of the AMs (common and unique contents).

This can be improved by splitting the "Encoding" column into 4 columns:
- CM0, CM1, CM2 (straddled, the same values for all lanes)
- UP0 (unique per lane)
- CM3, CM4, CM5 (straddled, the same values for all lanes)
- The rest (unique per lane)

The two tables can also be joined to one table with 32 rows.

SuggestedRemedy

Change tables 172-2 and 172-3 as described.
Consider merging the two tables.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.4** *P* **200** *L* **5** # 77

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

These tables are still very hard to use because the ~headers don't line up with the ~columns

SuggestedRemedy

For the header row, insert a space after each comma

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.4** *P* **201** *L* **39** # 78

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

x

SuggestedRemedy

Use multiplication symbol, twice

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.9** *P* **202** *L* **48** # 122

Slavick, Jeff — Broadcom

Comment Type **T** *Comment Status* **X**

To make this section agnostic to the MII rate for referencing in the future. We could refer to the service interface insteead.

SuggestedRemedy

Change "PCS at the 800GMII" to "PCS, at the PCS service interface,"

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.4.9** *P* **202** *L* **52** # 79

Dawe, Piers — Nvidia

Comment Type **T** *Comment Status* **X**

This mentions the test-pattern control register (bit 3.42.3). But does 3.42.7 Scrambled idle test-pattern apply also?

SuggestedRemedy

Please clarify, and please refer to 172.3.1 PCS MDIO function mapping

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.5.2** *P* **203** *L* **12** # 80

Dawe, Piers — Nvidia

Comment Type **E** *Comment Status* **X**

PCS lanes can be received on different lanes of the service interface from which they were originally transmitted - needs rewording?

SuggestedRemedy

Suggest:
The signals received by a PCS can contain PCSLs in a different arrangement to the lane ordering at the transmitting PCS. The PCS receiver is capable of receiving PCSLs in any arrangement.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.5.8.1** *P* **204** *L* **10** # 43

Nicholl, Shawn — AMD

Comment Type **TR** *Comment Status* **X**

To allow use of the PCS stateless decoder at both 400 Gb/s and 800 Gb/s data rates, place the new sub-clause 172.2.5.8.1 (PCS stateless decoder) into Clause 119 directly.

SuggestedRemedy

Propose to create a new sub-clause 119.2.5.8.1 containing the current text of 172.2.5.8.1 (PCS stateless decoder), except replace "800GMII vector" with "MII vector". Or replace with "rx_raw vector" instead.

In sub-clause 119.2.5.8 (Decode and rate matching), change "... state diagram as shown in Figure 119-15." to "... state diagram as shown in Figure 119-15 or (for 400GBASE-R PCS or 800GBASE-R PCS) by the stateless decoder specified in 119.2.5.8.1."

In sub-clause 172.2.5.8 (Block collection, decode, and rate matching), change "stateless decoder specified in 172.2.5.8.1." to "stateless decoder specified in 119.2.5.8.1."

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.5.8.1** *P* **204** *L* **18** # 23

Ran, Adee — Cisco

Comment Type **TR** *Comment Status* **X**

Table 172–4 has "reset" as the first column, but reset is not defined in clause 172.

Similarly, LBLOCK_R, EBLOCK_R, R_TYPE, and the block types E, S, D, T, C, DECODE, and rx_raw are not defined anywhere in this draft.

SuggestedRemedy

Add text pointing to the definitions of LBLOCK_R and EBLOCK_R in 119.2.6.2.1, reset and rx_raw in 119.2.6.2.2, and R_TYPE and DECODE in 119.2.6.2.3.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.5.8.1** *P* **204** *L* **23** # 24

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **X**

In Table 172-4, row 3, column "R_TYPE (rx_coded_i)", the value is "S or D or T or C".

The possible R_TYPE values (based on 119.2.6.2.3) are C, LI, S, T, D, and E; LI is not valid for clause 172 (per 172.2.3, EEE and low power idle are not supported). Therefore, "S or D or T or C" is equivalent to "not E". This excludes only the combination "E | E".

However, the combination "E | E" matches the second row, and therefore results in the same rx_raw, EBLOCK_R. So having R_TYPE(rx_coded_i-1)=E with any value of R_TYPE(rx_coded_i) would result in EBLOCK_R.

This means the table can be simplified and made more readable.

SuggestedRemedy

Change the third row to the following contents:
"0 | E | any block type | EBLOCK_R".

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.6.1** *P* **204** *L* **38** # 81

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

"its value is to be incremented": by how much? Does it depend on the circumstances?

SuggestedRemedy

Add "by one", or whatever is meant.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.2.6.2.2** *P* **205** *L* **21** # 82

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

this variable mapped per Table

SuggestedRemedy

this variable is mapped per Table
Also at line 28

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.3.3** *P* **209** *L* **20** # 83

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

Without the information in 119.3.3, the title is ambiguous or misleading. This isn't a count of uncorrected codewords which would include the ones that didn't have errors and didn't need correcting; it's a count of errored codewords that were not corrected.

SuggestedRemedy

Add sentence: This counter counts FEC codewords that contain errors that were not corrected.

Proposed Response *Response Status* **O**

Cl **172** *SC* **172.3.6** *P* **209** *L* **34** # 25

Ran, Adee Cisco

Comment Type **T** *Comment Status* **X**

The PMA lane muxing is specified with restrictions intended to ensure that all codewords are represented on each physical lane (and ideally have the same BER).

In practice, devices might use muxing that does not meet these restrictions, and the PCS has to work with any muxing scheme. In some schemes, the four FEC decoders may have different BER and different codeword bin counts. This information can be important for link performance analysis and prediction.

It is suggested to have separate counters for each flow. This is sufficient because, within each flow, the BER seen by the two codewords is inherently the same, due to the checkerboard pattern. Also, FEC_cw_counter in 172.3.5 is the same for both flows and need not be duplicated.

SuggestedRemedy

Replace the FEC_codeword_error_bin_i variables with two sets of variables, flow<j>_FEC_codeword_error_bin_i, where j goes from 0 to 1.

Add MDIO addresses for these variables and update variable mapping tables as appropriate.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.1.3** *P* **212** *L* **51** # 84

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

Adapt the PCSL (PCS lane) formatted signal to the appropriate number of abstract or physical lanes

SuggestedRemedy

Adapt the PCSL (PCS lane) formatted signal to the appropriate number and grouping of abstract or physical lanes

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.1.3** *P* **213** *L* **10** # 85

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

In common cases (800GAUI-8) receive link status information may be used but isn't forwarded.
"Provide receive link status information in the receive direction": do we need another bullet, that when connected to a PHY XS, it provides link status information in the transmit (egress) direction?

SuggestedRemedy

Per comment

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.1.3** *P* **213** *L* **11** # 86

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

173.4 says "Three forms of the 800GBASE-R PMA are defined: 32:8, 8:32, and 8:8" but that information is needed earlier, in 173.1.4, 173.2 and 173.3

SuggestedRemedy

Insert a sentence here, saying that.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.3** *P* **215** *L* **43** # 127

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

"For the 8:32 PMA ... In this case a PHY_XS:IS_SIGNAL.indication primitive is not received from the PHY 800GXS". Why not? The module knows if its incoming signal is good or not, so it can pass that information to the 8:32 PMA, which can e.g. squelch appropriately. This would be a normal behaviour for non-XS modules.

SuggestedRemedy

Discuss

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.3** *P* **215** *L* **49** # 26

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

"The PHY_XS:IS_SIGNAL.request primitive is generated through a set of SIL that reports signal health"

"SIL" is defined in 173.2 as a function, not a set.

SuggestedRemedy

Change the quoted sentence to "The PHY_XS:IS_SIGNAL.request primitive is generated through a signal indication logic (SIL) function that reports signal health".

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4** *P* **217** *L* **6** # 87

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

PMA:IS_UNITDATA_0:31.request would be better shown as PMA:IS_UNITDATA_0:15.request and PMA:IS_UNITDATA_16:31.request as in Figure 172-2. The PMA doesn't really know lane numbers, it doesn't read alignment markers, but it needs to know the two groups to apply the restricted bit muxing rules.
The output lanes can stay as one group.

SuggestedRemedy

Show two groups of 16 input lanes, PMA:IS_UNITDATA_0:15.request and PMA:IS_UNITDATA_16:31.request.
Similarly for the 32 PHY_XS:IS_UNITDATA_0:31.indication lanes in Figure 173-4, 8:32 PMA functional block diagram.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.2** *P* **220** *L* **1** # 88

Dawe, Piers — Nvidia

Comment Type **TR** *Comment Status* **X**

Ensure that the restricted bit multiplexing rules exclude combinations of lanes and Skew that suffer the "clock content" (transition density) issue mentioned at the end of 120.5.2.

SuggestedRemedy

Per comment

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.2.1** *P* **220** *L* **15** # 27

Ran, Adee — Cisco

Comment Type **TR** *Comment Status* **X**

As observed in comment #6 against D1.0, the existing restrictions enable a muxing scheme where one of the two PCS flows is always assigned to the LSBs of the PAM4 symbols, while the other flow is always assigned to the MSB.
This scheme (labeled "option B" in ran_3df_01a_2212) will cause an increase of x34 in the frequency of uncorrectable errors in the link partner, compared to the scheme that was assumed for the baseline proposal, which splits the LSBs equally between the two flows ("option A").

Comment #6 suggested restricting the muxing further to prevent using "option B" in the transmitter. The receiver is required to tolerate any muxing order, so transmitters using "option B" would be interoperable, but they should not be considered compliant.

Straw polls taken during the resolution of comment #6 had inconclusive results indicating need for additional information. In discussions since then, no specific examples of applications that would break by the additional restrictions have been found. These restrictions are therefore suggested again. If there is no consensus to have them as mandatory requirements, they can be added as recommendations.

A presentation providing further explanations and justification for the suggested restrictions will be provided.

SuggestedRemedy

In 173.4.2.1 and 173.4.2.2, change the second list item to
"The multiplexing function has an additional constraint that each of the 8 output lanes contain two unique PCSLs from PMA client lanes i = 0 to 15 followed by two unique PCSLs from PMA client lanes i = 16 to 31".

In 173.4.2.3, change the second list item to
"The 4 PCSLs received on an input lane shall be mapped to an output lane such that the Gray-coded PAM4 symbol sequence on the output lane is identical to the Gray-coded PAM4 symbol sequence on the input lane (see 173.4.7.1)."

Modify wording and/or add illustrations with editorial license.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.2.1** *P* **220** *L* **16** # 89

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **X**

Avoid the bad "option B" bit muxing that Adee has described.
Fixing this is more useful than applying any restricted muxing on the XS.
I doubt that the language of lanes containing lanes will stretch to the ordering restriction needed, so wordsmithing to "constructed from".

SuggestedRemedy

Change
The multiplexing function has an additional constraint that each of the 8 output lanes contain two unique PCSLs from PMA client lanes i = 0 to 15 and two unique PCSLs from PMA client lanes i = 16 to 31
to
The multiplexing function has an additional constraint that each of the 8 output lanes is constructed from two PCSLs from PMA client lanes i = 0 to 15 and two PCSLs from PMA client lanes i = 16 to 31, arranged so that after PAM4 encoding, the first bits of the pairs used to form PAM4 symbols are taken alternately from one of the two PCSLs from PMA client lanes i = 0 to 15, and one of the two PCSLs from PMA client lanes i = 16 to 31.
Similarly in 173.4.2.2, or delete the restricted muxing rule from the 8:32 PMA, as the XS AUI shouldn't make enough errors to trouble the FEC.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.2.1** *P* **220** *L* **17** # 90

Dawe, Piers Nvidia

Comment Type **E** *Comment Status* **X**

I doubt that one can have two unique anythings. Unique means one of a kind, so if there are two, they aren't unique. I think we mean different, but as it is obvious enough from 120.5 that each PCS lane is used just once, there is no need for any such word.

SuggestedRemedy

Delete "unique", twice

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.2.3** *P* **221** *L* **9** # 91

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

"The 4 PCSLs received on any input lane shall be mapped to the same output lane" is ambiguous: this could mean the same lane number (which seems unnecessary) or merely that the PCSLs are kept together. (I know this text is based on my comment - apologies.)

SuggestedRemedy

Clarify. And see next comment.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.2.3** *P* **221** *L* **10** # 92

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

"The order of PCSLs from an input lane does not have to be maintained on the output lane": but to avoid a rogue 8:8 PMA turning the benign properly bit-muxed "option A" into the defective "option B", we can't allow all possible re-ordering.

SuggestedRemedy

As there is no practical reason not to, require that the PMA output the streams of PAM4 symbols that it receives (but without requiring preservation of lane number).

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.3.1** *P* **221** *L* **27** # 93

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **X**

This says "the PMA ... shall produce no more than" while 173.4.3.3 says "the PMA ... shall generate no more than"

SuggestedRemedy

If there is a difference between produce and generate, as I suspect there is, explain. If there isn't, use one word not two.
See another comment that the limits are higher than needed now.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.3.3** *P* **221** *L* **43** # 94

Dawe, Piers — Nvidia

Comment Type **T** *Comment Status* **X**

Not clear "as well" as what.

SuggestedRemedy

Please explain.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.5** *P* **222** *L* **38** # 95

Dawe, Piers — Nvidia

Comment Type **E** *Comment Status* **X**

This says that the clock architecture is identical to that specified in 120.5.5.
Clocking architecture not clock architecture
Rates in 120.5.5 are based on bit rates, here bit rate is not mentioned.
120.5.5 addresses cases of 200GBASE-R and 400GBASE-R, not 800G.
120.5.5 says "... rearrangement of PCSLs between input lanes and output lanes (although rearrangements are allowed)" but this clause has rules forbidding some rearrangements.

SuggestedRemedy

Add material to define what the clocking architecture for this clause is

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.7.2** *P* **223** *L* **1** # 28

Ran, Adee — Cisco

Comment Type **ER** *Comment Status* **X**

The title "Precoding for PAM4 encoded lanes" is used in clause 120, but in clause 173 all lanes are PAM4 encoded.

SuggestedRemedy

Change the title to "Precoding".

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.7.2** *P* **223** *L* **3** # 29

Ran, Adee — Cisco

Comment Type **T** *Comment Status* **X**

The first paragraph of this subclause effectively excludes 800GAUI-8 C2M, making precoding impossible over this interface.

Precoding can also be beneficial for C2M in certain cases, and it is likely implemented as part of the SerDes in many products. Therefore, it would be good to allow it as an optional feature that, if available, can be enabled as required by the application.

This would only apply in the interface lanes connected to the AUI, and not to those that are connected to the PMD, so the optical signal will not be affected.

The fact that this option is not explicitly defined for 400GAUI-4 C2M etc. does not preclude it from being defined in this project.

SuggestedRemedy

With editorial license, make both precoding and decoding optional for PMAs lanes that are part of a 800GAUI-8 C2M link (this may affect both Clause 167 and annex 120G).

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.8** *P* **223** *L* **30** # 129

Dawe, Piers — Nvidia

Comment Type **T** *Comment Status* **X**

This says that the PMA link status functions identically to that specified in 120.5.8. 120.5.8 says "the PMA shall provide link status information to the PMA client using the PMA:IS_SIGNAL.indication primitive." That's too simple; this primitive is not carried over the AUI, and for the 8:32 PMA, link status

SuggestedRemedy

Please write out what actually happens

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.4.11** *P* **223** *L* **47** # 30

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

120.5.11.2 is now included in this draft.

SuggestedRemedy

Make 120.5.11.2 an active cross reference.

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.5** *P* **224** *L* **10** # 96

Dawe, Piers Nvidia

Comment Type **T** *Comment Status* **X**

This says MMDs 8, 9, and 10 while 173.1.4 says 1, 8, 9, 10, and 11

SuggestedRemedy

Reconcile 11

Proposed Response *Response Status* **O**

Cl **173** *SC* **173.6.5** *P* **229** *L* **20** # 31

Ran, Adee Cisco

Comment Type **ER** *Comment Status* **X**

120.5.11.2.2 is now included in this draft.

SuggestedRemedy

Make all instances of 120.5.11.2.2 in this table active cross references.

Proposed Response *Response Status* **O**