

IEEE P802.3df D3.0 Initial Sponsor ballot comments

CI FM SC FM P1 L29 # I-2

Brown, Matthew Alphawave
 Comment Type E Comment Status A (bucket1)

The order of amendments to IEEE Std 802.3-2022 has been adjusted such that 802.3df precedes 802.3cw, with the former being Amendment 9 and the latter Amendment 10.

SuggestedRemedy

Remove all references to and amendments to 802.3cw and set 802.3df as amendment 9. On the front page, change "Amendment" to "Amendment 9" and remove 802.3cw from the list of preceding amendments. On page 13, remove 802.3cw from the list of amendments. On page 14, add "Amendment 9" at the beginning of the 802.3df description. On page 37 and 41, remove "as modified by IEEE Std 802.3cw-202x)" and adjust changes appropriately. Implement with editorial license.

Response Response Status C
 ACCEPT.

CI FM SC FM P1 L29 # I-1

Hajduczenia, Marek Charter Communications
 Comment Type E Comment Status A (bucket1)

IEEE Std 802.3cy-202x is now approved (2023)

SuggestedRemedy

Update publication year for IEEE Std 802.3cy to 2023 in the whole document.

Response Response Status C
 ACCEPT.

CI FM SC FM P1 L30 # I-20

Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)

802.3df will be published before 802.3cw

SuggestedRemedy

Change "... IEEE Std 802.3cz-2023, IEEE Std 802.3cy-202x, and IEEE Std 802.3cw-202x." to "... IEEE Std 802.3cz-2023, and 802.3cy-202X."

Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-2.

CI FM SC FM P13 L45 # I-19

Marris, Arthur Cadence Design Systems, Inc.
 Comment Type E Comment Status A (bucket1)

802.3df will be published before 802.3cw so references to 802.3cw should be removed

SuggestedRemedy

Delete IEEE Std 802.3cw™-202x entry on line 45 on page 13

On page 1 change "IEEE Std 802.3cy-202x, and IEEE Std 802.3cw-202x" to "and IEEE Std 802.3cy-202x"

Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-2.

CI FM SC FM P13 L45 # I-21

Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)

802.3df will be published before 802.3cw

SuggestedRemedy

Delete the text related to 802.3cw.

Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-2.

CI 0 SC 0 P34 L2 # I-41

Ran, Adeo Cisco Systems, Inc.
 Comment Type E Comment Status A PCS sublayer

"PCS Sublayer" (RAS syndrome) in new text: 1.4.184k, 162.9.5, 163.9.3, 169.2.3, Figure 171-2, 172.1.2, 120F.3.2

SuggestedRemedy

Change "PCS Sublayer" to "PCS" in all instances.

Response Response Status C
 ACCEPT IN PRINCIPLE.

For text added by 802.3df, change instances of "PCS sublayer" to "PCS" except where the word sublayer is added for clarity.

Implement with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 0 SC 0 P104 L12 # I-45

Ran, Adee Cisco Systems, Inc.

Comment Type T Comment Status A delay values

The PMD delay constraint for 800G optical PMDs should be the same in ns terms to those of similar PMDs at the same signaling rate with fewer lanes (viz., 20.48 ns rather than 40.96 ns).

To allow the total delay for 800G modules as has been adopted in response to comment #82 against D2.0, an extra delay of 20.48 ns can be allocated to the PMA instead, to create the same total delay of 87.04 ns (for PMD+PMA). Note that the delay could be added only for the PMA(8:8), but currently, there is no distinction between PMA types.

This comment affects clauses 124, 167, 169, and 173.

SuggestedRemedy

in 124.3.1 and in 167.3.1 Change "32 768 bit times (64 pause_quanta or 40.96 ns)" to "16384 bit times (32 pause_quanta or 20.48 ns)".

In 173.5.4, Change the values in Table 173-1 to "53 248", "104", and "66.56".

Change the corresponding entries in Table 169-4 accordingly.

Response Response Status C

ACCEPT IN PRINCIPLE.
 Implement the changes shown on slide 13 of the following presentation:
https://www.ieee802.org/3/df/public/23_0926/brown_3df_01a_230926.pdf
 Change the delay value for 800GBASE-DR8-2 to 20.48 ns (along with the pause_quanta and bit times) as well.
 Add an informative note explaining that although the PMD service interface definition implies otherwise the delay allocation within a PHY assumes the majority of the SERDES function is in the PMA next to the PMD rather than in the PMD.
 Implement with editorial license.

Cl 0 SC 0 P108 L49 # I-46

Ran, Adee Cisco Systems, Inc.

Comment Type E Comment Status A PCS sublayer

"PCS Sublayer" (RAS syndrome) in existing text - but changes in nearby text may put it in scope for correction: 124.6, 162.4 (twice), 162.9.4, 163.9.2, 167.6, 120F.3.1, 120G.3.1

SuggestedRemedy

Change "PCS Sublayer" to "PCS" in all instances.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-41.

Cl 0 SC 0 P128 L21 # I-47

Ran, Adee Cisco Systems, Inc.

Comment Type TR Comment Status A (bucket1)

The symbol "+" is used on the status column in multiple PICS items, denoting logical-OR. It is not defined in the PICS conventions in clause 21.

SuggestedRemedy

Add Clause 21 to the draft, and amend 21.6.2, adding the sentence:

"<item1>+<item2>: OR-predicate condition, the requirement has to be met if either of the optional items is implemented".

Response Response Status W

ACCEPT IN PRINCIPLE.
 Add Clause 21 to the draft, and amend 21.6.2, adding the sentence:
 "<item1>+<item2>: OR-predicate condition, the requirement has to be met if one or both of the items is implemented"
 Implement with editorial license.

Cl 1 SC 1.1.3.2 P31 L12 # I-37

Ran, Adee Cisco Systems, Inc.

Comment Type ER Comment Status A Interfaces

"While conformance with implementation of this interface is not necessary to ensure communication..."

"Conformance with implementation" does not make sense. The intent is probably "conformance with the specification".

Similarly in the next item, L19.

SuggestedRemedy

Change "conformance with implementation" to "conformance with the specification", twice.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-84 and #I-85.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 1 SC 1.1.3.2 P31 L13 # I-84

Dawe, Piers J G NVIDIA
 Comment Type T Comment Status A Interfaces

This says about the 800GMII: "While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds. The 800GMII is a logical interconnection intended for use as an intra-chip interface. No mechanical connector is specified for use with the 800GMII. The 800GMII is optional." which is much the same as item d, GMII. An exposed 800GMII is much less likely than an exposed GMII. As the current interfaces of choice for "allowing flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds" are AUIs not MIIs, the first sentence quoted is misleading old cruft. 170.1 gives a more convincing reason: "Though the 800GMII is an optional interface, it is used in this standard as a basis for specification".

SuggestedRemedy

Delete the sentence "While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds." or replace it with something like "While conformance with implementation of this interface is not necessary to ensure communication, it is used in this standard as a basis for specification."

Response Response Status C

ACCEPT IN PRINCIPLE.
 On page 31 lines 13 and 18, delete the sentence "While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds."
 See also the response to comment #I-85.

Cl 1 SC 1.1.3.2 P31 L13 # I-38

Ran, Adeo Cisco Systems, Inc.
 Comment Type ER Comment Status A Interfaces

(While conformance... is not necessary...) "it allows flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds"

it's not the conformance that allows flexibility, it's the fact that it's a common service interface.

SuggestedRemedy

Change "it allows" to "it serves as a common logical interface that allows".

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-84.

Cl 1 SC 1.1.3.2 P31 L17 # I-85

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A Interfaces

This text "*The* 800GAUI-n is a physical instantiation of the PMA service interface... While conformance with implementation of *this interface*... *The 800GAUI-n* is intended... For chip-to-chip interfaces and for chip-to-module interfaces, one width of 800GAUI-n is defined: *an eight-lane version* (800GAUI-8) in Annex 120F and Annex 120G. No mechanical connector is specified for use with *the* 800GAUI-n. *The* 800GAUI-n is optional." reads as if there is only one kind of 800GAUI-n, and its specification is spread over two annexes. This is wrong; 800GAUI-n C2M and 800GAUI-n C2C are distinct, not interchangeable, and not intended to interoperate with each other (unlike the original intent for XLAUI). There is not "a version". Also, "the PMA service interface" is inaccurate; there can be more than one PMA service interface per MAC. Note the definition 1.4.184h uses "A" not "The".

SuggestedRemedy

Change the paragraph to: x) 800 Gb/s Attachment Unit Interface (800GAUI-n). An 800GAUI-n is a physical instantiation of a PMA service interface to extend the connection between 800 Gb/s capable PMAs. While conformance with implementation of 800GAUI-n is not necessary to ensure communication, it is recommended, since it allows maximum flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds. 800GAUI-n C2C is intended for use as a chip-to-chip and 800GAUI-n C2M is intended as a chip-to-module interface. One width of 800GAUI-n is defined for chip-to-chip interfaces and one for chip-to-module interfaces: eight-lane 800GAUI-8 C2C in Annex 120F and eight-lane 800GAUI-8 C2M in Annex 120G. No mechanical connector is specified for use with an 800GAUI-n. An 800GAUI-n is optional.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change:

"The 800GAUI-n is a physical instantiation of the PMA service interface to extend the connection between 800 Gb/s capable PMAs. While conformance with implementation of this interface is not necessary to ensure communication, it is recommended, since it allows maximum flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds. The 800GAUI-n is intended for use as a chip-to-chip or a chip-to-module interface. For chip-to-chip interfaces and for chip-to-module interfaces, one width of 800GAUI-n is defined: an eight-lane version (800GAUI-8) in Annex 120F and Annex 120G. No mechanical connector is specified for use with the 800GAUI-n. The 800GAUI-n is optional."

To:

"An 800GAUI-n is an optional physical instantiation of the PMA service interface to extend the connection between 800 Gb/s capable PMAs. 800GAUI-n C2C is intended for use as a chip-to-chip interface and 800GAUI-n C2M as a chip-to-module interface. One width of 800GAUI-n is defined: eight-lane 800GAUI-8 C2C in Annex 120F and 800GAUI-8 C2M in Annex 120G. No mechanical connector is specified for use with an 800GAUI-n."

Implementation with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 1 SC 1.1.3.2 P31 L20 # I-39

Ran, Adee Cisco Systems, Inc.
 Comment Type **TR** Comment Status **A** Interfaces

"since it allows maximum flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds"
 "Maximum flexibility" is questionable, and this is not the motivation of the 800GAUI-n nor of multiple similar AUIs defined for lower data rates.

The motivation of the AUIs is to enable the usage of implemented PCS/PMA sublayers over different media.

SuggestedRemedy

Change "since it allows maximum flexibility in intermixing PHYs and DTEs at 800 Gb/s speeds" to "since it allows links over different media to be used by the same DTE through PHYs that contain medium-dependent components".

Response Response Status **C**

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-85.

Cl 1 SC 1.4.109 P31 L49 # I-40

Ran, Adee Cisco Systems, Inc.
 Comment Type **E** Comment Status **A** (bucket1)

In all other definitions in 1.4 that mention reach (103, 108a, 109a, 135, 135a, 142, 142a, 143, 144, 144a, 184b, 184c, 184f, 184g) there is a comma before "with reach up to". Here there isn't.

SuggestedRemedy

For consistency, add a comma after "in each direction".

Response Response Status **C**

ACCEPT.

Cl 1 SC 1.4.184h P33 L37 # I-86

Dawe, Piers J G NVIDIA
 Comment Type **TR** Comment Status **A** definitions

This says that 800GAUI-n is used for chip-to-chip or chip-to-module electrical interfaces. It says that an eight-lane version when in fact, two versions are defined, that are specified differently and not generally compatible with each other. In the proposed change, the first sentence, shown for context, is unchanged.

SuggestedRemedy

Change: 800 Gb/s Attachment Unit Interface (800GAUI-n): A 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 or chip-to-module electrical interfaces. For chip-to-module interfaces and for chip-to-chip interfaces, one width of 800GAUI-n is defined: an eight-lane version (800GAUI-8). (See IEEE Std 802.3, Annex 120F and Annex 120G.)
 to: 800 Gb/s Attachment Unit Interface (800GAUI-n): A 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 or chip-to-module electrical interfaces. One width of 800GAUI-n is defined for chip-to-chip interfaces and one for chip-to-module interfaces: eight-lane 800GAUI-8 C2C and eight-lane 800GAUI-8 C2M. (See IEEE Std 802.3, Annex 120F and Annex 120G.)

Response Response Status **C**

ACCEPT IN PRINCIPLE.

Change the definition to the following:

"A 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 or chip-to-module electrical interfaces. For chip-to-module interfaces and for chip-to-chip interfaces, one width of 800GAUI-n is defined: eight-lane 800GAUI-8 C2C and 800GAUI-8 C2M. (See IEEE Std 802.3, Annex 120F and Annex 120G.)"

Cl 1 SC 1.4.184k P34 L2 # I-87

Dawe, Piers J G NVIDIA
 Comment Type **E** Comment Status **A** PCS sublayer

Tautology: "PCS Sublayer" and "RS sublayer". 1.4.113 200GXS and 1.4.148 400GXS have the same problem.

SuggestedRemedy

Delete Sublayer and sublayer, or spell out PCS and RS in words, or at least change "PCS Sublayer" to "PCS sublayer".

Response Response Status **C**

ACCEPT IN PRINCIPLE.

Resolve using the response to comments #I-41 and #I-42.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 1 SC 1.4.184k P34 L3 # I-42
 Ran, Adee Cisco Systems, Inc.
 Comment Type E Comment Status A (bucket1)
 "RS Sublayer" (RAS syndrome)
 SuggestedRemedy
 Change to "Reconciliation Sublayer"
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 RS is normally spelled out to distinguish it from the Reed-Solomon FEC (RS-FEC) and this resolves the redundancy.
 Implement the suggested remedy with editorial license.
 [Editor's note: Changed line 34 to 3]

Cl 1 SC 1.4.461 P34 L19 # I-88
 Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)
 Difficult to parse "carried on a physical lane together at the..."
 SuggestedRemedy
 Change to "carried together on a physical lane at the..." or "carried on a single physical lane at the..."
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change "One or more PCS lanes can be multiplexed and carried on a physical lane together"
 To: "One or more PCS lanes can be multiplexed and carried together on a physical lane"

Cl 30 SC 30.5.1.1.2 P36 L45 # I-43
 Ran, Adee Cisco Systems, Inc.
 Comment Type T Comment Status R PMD label (bucket2) (cc)
 Most entries in this list include reach, but some don't, although reach is defined for them. In this project, reach was added for 400GBASE-DR4, but not for other items.
 200GBASE-DR4, 200GBASE-SR4, 400GBASE-SR4.2, 400GBASE-SR8, and 400GBASE-SR16 have reaches included in their definitions in 1.4.
 SuggestedRemedy
 In the 200GBASE-DR4 item, insert "with reach up to at least 500 m" after "PMD".
 In the 200GBASE-SR4 item, insert "with reach up to at least 100 m" after "PMD".
 In the 400GBASE-SR4.2 item, insert "with reach up to at least 150 m" after "PMD".
 In the 400GBASE-SR16 item, insert "with reach up to at least 100 m" after "PMD".
 Response Response Status C
 REJECT.
 The changes already made to these descriptions in Draft 3.0 were justified to remove ambiguity in the definitions for PHYs related to the new PHYs defined in this draft.
 The additional changes proposed in this comment are not addressing innaccuracy or inconsistency.
 There is no consensus to make the proposed changes.

Cl 30 SC 30.5.1.1.2 P37 L34 # I-22
 Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)
 802.3df will be published before 802.3cw
 SuggestedRemedy
 Change the editing instruction to say "Insert the following new entries into "APPROPRIATE SYNTAX" in 30.5.1.1.2 after the entry for 400GBASE-VR4:"
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-2.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 31B SC 31B.3.7 P251 L25 # I-73

Ran, Adeo Cisco Systems, Inc.
 Comment Type ER Comment Status A numbers (bucket1)

"115 840"

The space separator is inconsistent with the format of existing numbers in 31B.3.7 in the base document (e.g., "57920" for 400 Gb/s).

Per the style manual, the use of space as a thousands separator is specified for numbers within tables. There is no need to use it in text and equations, especially where it creates inconsistency.

This comment also applies to 124.3.1 and 167.3.1, where numbers of bit times appear with thousands separators in the text (subject of another comment).

SuggestedRemedy

Change "115 840" to "115840".

Implement similarly for the numbers of bit time in 124.3.1 and 167.3.1 (subject of another comment).

Response Response Status W

ACCEPT IN PRINCIPLE.
 See related comment #I-64.
 For this case the use of a thousands separator is inconsistent with similar equations in 31B.3.7 and the use of the separator makes the equation difficult to read.
 Change "115 840" to "115840" in 31.B.3.7
 Comment #I-45 addresses the changes requested for 124.3.1 and 167.3.1

Cl 45 SC 45.2.1.6 P41 L3 # I-23

Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)

The editing instruction needs to reflect that table 45-7 was modified by 802.3ck-2022, 802.3db-2022, and 802.3cz-2023, and that 802.3cw won't have modified it.

SuggestedRemedy

Change the parenthetical remark in the editing instruction to say "(as modified by IEEE Std. 802.3db-2022, IEEE Std. 802.3ck-2022, and IEEE Std 802.3cz-2023)"

Response Response Status C

ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-2.

Cl 45 SC 45.2.1.6 P41 L3 # I-18

Marris, Arthur Cadence Design Systems, Inc.
 Comment Type E Comment Status A (bucket1)

802.3df is now expected to be published before 802.3cw.

SuggestedRemedy

On page 41 delete "(as modified by IEEE Std 802.3cw-202x)" on line 3
 on page 41 line 24 change "0 1 1 1 1 1 1 = 400GBASE-ZR PMA/PMD" to "0 1 1 1 1 1 1 = reserved"

and in "30.5.1.1.2 aMAUType"
 On page 37 line 35 change "(as modified by IEEE Std 802.3cw-202x)" to "(as modified by IEEE Std 802.3db-2022)"
 Change "after the entry for 400GBASE-ZR" to "after the entry for 400GBASE-VR4"

Response Response Status C

ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-2.

Cl 45 SC 45.2.1.6 P41 L25 # I-24

Huber, Thomas Nokia
 Comment Type T Comment Status A (bucket1)

400GBASE-ZR won't have been defined when 802.3df is approved since 802.3cw is after 802.3df

SuggestedRemedy

Replace "400GBASE-ZR PMA/PMD" with "reserved"

Response Response Status C

ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-2.

Cl 45 SC 45.2.1.7 P42 L16 # I-25

Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)

The period after 400GBASE-KR4 should be a comma, and the punctuation mark should be indicated as text to be inserted

SuggestedRemedy

Change "400GBASE-KR4. 800GBASE-KR8" to "400GBASE-KR4, 800GBASE-KR8" and underline the comma

Response Response Status C

ACCEPT.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 45 SC 45.2.1.7 P42 L21 # I-26
 Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)
 The comma and space following 400GBASE-CR4 should be indicated as text to be inserted
 SuggestedRemedy
 Underline the comma and space.
 Response Response Status C
 ACCEPT.

Cl 45 SC 45.2.1.7.4 P42 L16 # I-16
 Marris, Arthur Cadence Design Systems, Inc.
 Comment Type E Comment Status A (bucket1)
 Replace . with ,
 SuggestedRemedy
 Change "400GBASE-KR4. 800GBASE-KR8" to "400GBASE-KR4, 800GBASE-KR8"
 Response Response Status C
 ACCEPT.

Cl 45 SC 45.2.1.7.4 P42 L16 # I-138
 Dudek, Michael Marvell
 Comment Type E Comment Status A (bucket1)
 *** Comment submitted with the file image.png attached ***
 The separation between 400GBASE-KR4 and 400GBASE-KR4 should be a comma, not a period
 SuggestedRemedy
 Fix it
 Response Response Status C
 ACCEPT.

Cl 45 SC 45.2.1.60b P47 L1 # I-27
 Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)
 The editing instruction should note that 45.2.1.60a was inserted by 802.3cz
 SuggestedRemedy
 Change the editing instruction to say "Insert 45.2.1.60b after 45.2.1.60a (as inserted by IEEE Std. 802.3cz-2023) as follows:"
 Response Response Status C
 ACCEPT.

Cl 45 SC 45.2.3.25.2 P60 L20 # I-17
 Marris, Arthur Cadence Design Systems, Inc.
 Comment Type E Comment Status A (bucket1)
 Delete editor's note as it is no longer needed
 SuggestedRemedy
 Delete editor's note as it is no longer needed
 Response Response Status C
 ACCEPT.

Cl 45 SC 45.2.3.25.2 P60 L20 # I-139
 Dudek, Michael Marvell
 Comment Type E Comment Status A (bucket1)
 *** Comment submitted with the file image.png attached ***
 The editor's note has served its purpose
 SuggestedRemedy
 Delete the note
 Response Response Status C
 ACCEPT.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 73 SC 73 P90 L2 # I-28
 Huber, Thomas Nokia
Comment Type T Comment Status A (bucket1)
 Figure 73-1 (as updated by 802.3ck-2022) should be updated to include 800G MII and 800 Gb/s media
SuggestedRemedy
 Insert clause 73.2, with an editing instruction to replace Figure 73-1 (as replaced by 802.3ck-2022). In the figure itself, change "or 400GMII" to "400GMII, or 800GMII", change "or 400 Gb/s" to "400 Gb/s, or 800 Gb/s", and add "800GMII = 800 Gb/s MEDIA INDEPENDENT INTERFACE" to the legend
Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-140.

Cl 73 SC 73.2 P90 L0 # I-140
 Slavick, Jeff Broadcom Inc
Comment Type TR Comment Status A data rates (bucket1)
 Figure 73-1 does not include 800GMII or 800Gb/s
SuggestedRemedy
 Remove the laundry list of data rates below the MDI
 Change the laundry list of specific MII rates to just be xMII and update the legend accordingly
Response Response Status W
 ACCEPT IN PRINCIPLE.
 Implement the suggested remedy using editorial license.
 In the legend have "xMII = generic Media Independent Interface".

Cl 73 SC 73.6.4 P90 L8 # I-29
 Huber, Thomas Nokia
Comment Type E Comment Status A (bucket1)
 Missing a space in the editing instruction
SuggestedRemedy
 Change "Table73-4" to "Table 73-4".
Response Response Status C
 ACCEPT.

Cl 73 SC 73.7.6 P91 L6 # I-30
 Huber, Thomas Nokia
Comment Type E Comment Status A (bucket1)
 Missing a space in the editing instruction
SuggestedRemedy
 Change "Table73-5" to "Table 73-5".
Response Response Status C
 ACCEPT.

Cl 116 SC 116.1.3 P95 L43 # I-79
 Lusted, Kent Intel
Comment Type T Comment Status D PMD label (bucket2) (cc)
 There is no indication of the supported reach for 200GBASE-SR4 in Table 116-1. An unfamiliar reader may not know of the reach of this specific PHY or be able to differentiate it from the other entries in the table. Note that Table 116-2 for 400 Gb/s PHYs has a description entry for 400GBASE-SR4 that does include "with a reach up to at least 100 m". The reach text is also in the Definitions in 1.4.109 (page 31, line 50)
SuggestedRemedy
 Add "with a reach up to at least 100 m" to the description of 200GBASE-SR4 in Table 116-1.
Proposed Response Response Status Z
 REJECT.
 This comment was WITHDRAWN by the commenter.

Cl 116 SC 116.1.3 P95 L43 # I-44
 Ran, Adeo Cisco Systems, Inc.
Comment Type T Comment Status D PMD label (bucket2) (cc)
 200GBASE-SR4 is defined with a reach (see 1.4.109), but it is the only one for which it is not mentioned in this table.
SuggestedRemedy
 Insert ", with reach up to at least 100 m" after "in each direction".
Proposed Response Response Status Z
 REJECT.
 This comment was WITHDRAWN by the commenter.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 124 SC 124 P115 L16 # I-15
 Stassar, Peter Huawei Technologies Co., Ltd
 Comment Type TR Comment Status A optical budget

This is a resubmission of comment #12 to D2.0.
 Comment #12 was rejected, because it was agreed that the proposed remedy was incomplete.

In clause 124, Table 124-8, for 400G-DR4 and 800G-DR8, the allocation for penalties is 3.5 dB, whereas for 400G-DR4-2 and 800G-DR8-2 it is 3.8 dB. The difference of 0.3 dB seems to originate from the FR4 spec in Clause 151, which is potentially suffering a higher MPI penalty due to larger individual reflections in an FR4 configuration compared to a DR4/DR8 configuration.

Because it was agreed (during the TF phase) to use the same list of requirements for discrete reflectances as shown in in-force Table 124-13, the allocation for penalties for DR4-2/DR8-2 can be lowered by 0.2 dB from 3.8 to 3.6 dB (assuming 0.1 dB for DGD penalty).

SuggestedRemedy

In Table 124-8, in the columns for 400GBASE-DR4-2 and 800GBASE-DR8-2, change the allocation for penalties from 3.8 dB to 3.6 dB.

Furthermore, in Table 124-7 for 400GBASE-DR4-2 and 800GBASE-DR8-2 increase the max Rx sensitivity from -4.5 / -5.9 +TECQ [dbm] to -4.3 / -5.7 +TECQ [dBm].

A supporting presentation with a complete change proposal will be provided for the comment resolution meeting

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement the proposed changes on slide 7 of the following presentation:
https://www.ieee802.org/3/df/public/23_0926/stassar_3df_01a_230926.pdf
 Implement with editorial license.

Cl 124 SC 124.1.1 P103 L3 # I-75
 Ran, Adee Cisco Systems, Inc.
 Comment Type TR Comment Status A BER

For the new 800 Gb/s PMDs the requirement in the second paragraph is that frame loss ratio is less than 3.4e-12, as opposed to 1.7e-12 for 400 Gb/s PMDs

The second paragraph of 124.1.1 in the base standard, which is not modified by this amendment, states that

"If the error statistics are not sufficiently random to meet this requirement, then the BER shall be less than that required to give a frame loss ratio of less than 1.7e-12 for 64-octet frames with minimum interpacket gap".

This statement should also address 800 Gb/s PMDs where the maximum FLR is 3.4e-12.

SuggestedRemedy

Change the second paragraph (currently not in the draft) from:

"If the error statistics are not sufficiently random to meet this requirement, then the BER shall be less than that required to give a frame loss ratio of less than 1.7e-12 for 64-octet frames with minimum interpacket gap"

to:

"If the error statistics are not sufficiently random to meet the specified frame loss ratio for 64-octet frames with minimum interpacket gap, then the BER shall be lower than the value required to meet that frame loss ratio".

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

Cl 124 SC 124.2 P103 L16 # I-31
 Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)

Singular/plural misalignment bewteen subject and verb in the second sentence.

SuggestedRemedy

Change "The service interface for these PMDs are described..." to "The service interface for these PMDs is described..." or "The service interfaces for these PMDs are described..."

Response Response Status C

ACCEPT IN PRINCIPLE.

Change "The service interface for these PMDs are described" to "The service interface for these PMDs is described"

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 124 SC 124.3.1 P104 L13 # I-89
 Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A delay values

The delay for 800GBASE-DR8 or 800GBASE-DR8-2 PMD including 2 m of fiber in one direction should be the same 20.48 ns as 400GBASE-DR4 and all other 200GBASE-R and 400GBASE-R optical PMDs (see tables 116-6 and 7). It was changed "because modern PMDs contain DSP": but that is semantics. We should not have different specification methods for 800GBASE-DR8 and 400GBASE-DR4 PMA/PMD: they are the same modules! For a typical retimed module, the PMA-PMD interface is internal so it doesn't matter (if we say it doesn't matter), but as linear and co-packaged optics become more popular, the interface is accessible, and a spec that has given the time for the A to D to the part that doesn't contain it becomes a problem. See comment against 169.3.3. Also note that a 32:8 or 8:32 PMA is "a SerDes" but an 8:8 PMA may be implemented as two SerDes back to back, with additional delay. See dawe_3df_01a_2307 Module and PMA delay limits, and other comments on delay.

SuggestedRemedy

Revert the PMD allowance to 16,384 bit times (32 pause_quanta or 20.48 ns) for all 8x100G optical, consistent with all 1/2/4x100G optical. With another comment, this gives a module with one PMD and one PMA $20.48+92.16 = 112.64$ ns. vs. D2.1 $40.96+46.08 = 87.04$ ns and 802.3-2018 $20.48 + 92.16/2$ (maybe) = 66.56 ns which seems to be tight for some DSP.

Response Response Status C
 ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-45.

Cl 124 SC 124.7.1 P110 L23 # I-83
 Maniloff, Eric Ciena Corporation
 Comment Type E Comment Status A TX power

Different optical clauses in 802.3 have not maintained consistency in the ER value used to calculate the Minimum Average Launch Power, but unfortunately this is not stated and it is left to the reader to calculate this for each Tx.. Since the different ERs exist in the standard, there should be a footnote added in the Tx tables to provide the value of ER max used to calculate the minimum Tx Power

SuggestedRemedy

Add a footnote to Table 124-6 for Average launch power, each lane (min) based on the final determination of which ER values are used. For example "An ER value of 10dB is used to calculate the Average launch power, each lanea (min)", or if different ER values are used for the different reaches this should be indicated in the footnote.

Response Response Status C
 ACCEPT IN PRINCIPLE.

Resolve using the response to comment I-95.

Also, add the follow note relating to 800GBASE-DR8/DR8-2 and 400GBASE-DR4-2:
 "Average power of -2.9dBm corresponds to an OMA of -0.8dBm with an extinction ratio of approximately 10dB or an OMA of -0.1dBm with an extinction ratio of approximately 16dB." Implement with editorial license.

Straw poll #3 (decision)
 I support adding a note explaining relationship between average power, OMA, and extinction ratio.
 Yes: 24
 No: 8
 Abstain: 18

IEEE P802.3df D3.0 Initial Sponsor ballot comments

CI 124 SC 124.7.1 P110 L23 # I-82

Maniloff, Eric Ciena Corporation

Comment Type TR Comment Status A TX power

The value for Average Launch Power, each lane (min) is calculated using an ER value of 10dB for DR4 and DR8, but using infinite extinction ratio for DR4-2 and DR8-2. There is no rationale presented to have different max ER's for different reaches. The specifications should use a single ER for these values.

SuggestedRemedy

Change the value of Average Launch Power, each lane (min) to -2.2dBm for the 2km reaches.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment I-95.

CI 124 SC 124.7.1 P110 L38 # I-10

Li, Jing YOFC

Comment Type E Comment Status R parameter label (bucket2)
(TECQ) (max)

SuggestedRemedy

(TECQ), each lane (max)

Response Response Status C

REJECT.

There is no consensus to make the proposed change.

CI 124 SC 124.8.1 P117 L8 # I-94

Dawe, Piers J G NVIDIA

Comment Type T Comment Status A (bucket1)

This would be better worded like the base text or Table 167-11 "3, 4, 5, 6, or valid 100GBASE-R, 200GBASE-R, 400GBASE-R, or 800GBASE-R signal".

SuggestedRemedy

Change "3, 4, 5, 6, or valid 400GBASE-R signal or 800GBASE-R signal" to "3, 4, 5, 6, or valid 400GBASE-R or 800GBASE-R signal" (i.e. put "or 800GBASE-R" before the first (pre-existing) "signal" and delete the second one).

Response Response Status C

ACCEPT.

CI 124 SC 124.8.1 P117 L30 # I-76

Ran, Adee Cisco Systems, Inc.

Comment Type TR Comment Status A (bucket1)

In Table 124-10, the subclause reference for the bottom two rows (Stressed receiver conformance test signal calibration, and Stressed receiver sensitivity) is 124.9, but that subclause is "Safety, installation, environment, and labeling" - apparently incorrect.

In the base document, these references are to 124.8.10, which is not part of this draft. If the existing 124.8.10 is adequate for the new PHYs then the reference can simply be corrected.

However, I suspect that other changes are required (for example, 140.7.13 includes a requirement about overshoot and undershoot, which does not exist in 124.8.10, even though these Tx requirements were added in 124.8.5b). If that is the case, then 124.8.10 should be added to this document and amended. I do not have the expertise to propose a detailed solution.

SuggestedRemedy

Change the reference of both table items to 124.8.10.

If it is necessary, add 124.8.10 to this document and make any required changes.

Response Response Status W

ACCEPT IN PRINCIPLE.

In Table 124-10 change the related subclause for Stressed receiver conformance test signal calibration and Stressed receiver sensitivity from 124.9 to 124.8.10.

CI 124 SC 124.8.5.1 P118 L23 # I-32

Huber, Thomas Nokia

Comment Type E Comment Status R (bucket1)

The style guide indicates that there should not be only one subclause at a given level; as such, inserting 124.8.5.1 without also adding a 124.8.5.2 is not appropriate.

SuggestedRemedy

Delete the editing instruction to insert 124.8.5.1 and that new heading. Include the text that would have gone in 124.8.5.1 as part of the changes to be made to 124.8.5.

Response Response Status C

REJECT.

The comment correctly points out that the style manual "recommends" that there should not be lone subclauses. However, this is an exception, where 124.8.5 we wish to retain a similar structure to 121.8.5 and 121.8.5.2 to define the TDECQ test channel characteristics.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 124 SC 124.8.9.2 P120 L17 # I-77
 Ran, Adee Cisco Systems, Inc.
 Comment Type E Comment Status A (bucket1)
 The editorial instruction says "Insert new subclause 124.8.9.2 after Figure 124-4". But that figure might move to another place when a new revision is created.
 The location of the new subclause should be defined by the subclause structure.
 SuggestedRemedy
 Change the instruction to "Insert new subclause 124.8.9.2 after 124.8.9.1".
 Response Response Status C
 ACCEPT.

Cl 124 SC 124.11a P124 L23 # I-95
 Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A TX power
 It would be bad economics to fragment the market for 400GBASE-DR4-2 modules into those that can interoperate with 400GBASE-DR4 and those that say they can't, when there is no cost to being interoperable. D2.0 comment 86, D2.1 comment 19. As 400GBASE-DR4 is well established but 400GBASE-DR4-2 is new, and as having a lower power for the higher performance PMD is counter-intuitive, the draft 400GBASE-DR4-2 should be brought into line. This proposed change will improve paperwork costs and reduce confusion, and have no practical technical effect - it reduces the measurement guard band from 0.9 dB to 0.7 dB at 9.8 dB extinction ratio, which is higher than realistic anyway.

SuggestedRemedy
 Delete "and the 400GBASE-DR4-2 transmitter average power is greater than or equal to the value for average launch power (min) for 400GBASE-DR4 in Table 124-6." In Table 124-6, change the Average launch power, each lane (min) from -3.1 dBm (the value associated with an infinite extinction ratio) to -2.9 dBm, same as 400GBASE-DR4 (associated with an unrealistically high extinction ratio for the same minimum OMA). Similarly for 800GBASE-DR8-2.

Response Response Status C
 ACCEPT IN PRINCIPLE.

The following presentations were reviewed by the CRG:
https://www.ieee802.org/3/df/public/23_0926/stassar_3df_02_230926.pdf
https://www.ieee802.org/3/df/public/23_0926/dawe_3df_02_230926.pdf

Implement the following:
 For 400GBASE-DR4-2 and 800GBASE-DR8-2 increase the minimum average power from -3.1 to -2.9 dBm and update related text with editorial license.
 Also, in Table 124-7 400GBASE-DR4-2 and 800GBASE-DR8-2 change the average receive power, each lane (min) from -7.1 dBm to -6.9 dBm.
 Implement with editorial license.

Straw poll #1 (Choose 1)
 I support the following option as stated in slide 14 of
https://www.ieee802.org/3/df/public/23_0926/stassar_3df_02_230926.pdf.
 A: Option 1
 B: Reject based on either Option 2 or Option 3
 C: Abstain
 A: 10 B: 10 C: 23

Straw poll #2 (decision)
 I support implementing option 1 as stated in slide 14 of
https://www.ieee802.org/3/df/public/23_0926/stassar_3df_02_230926.pdf.
 Yes: 26
 No: 5

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 124 SC 124.12.4.4 P128 L21 # I-96
 Dawe, Piers J G NVIDIA
 Comment Type ER Comment Status A (bucket1)
 This use of + is used in several clauses in this draft. It is not defined in 21.6.2, but it is useful.
 SuggestedRemedy
 In 21.6.2, add: <item1>+<item2>: OR-predicate condition, the requirement has to be met if either or both optional items are implemented
 Response Response Status W
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-47.

Cl 124 SC 124.12.4.4 P128 L21 # I-143
 Slavick, Jeff Broadcom Inc
 Comment Type TR Comment Status A (bucket1)
 PICS don't have a definition for +
 SuggestedRemedy
 For OM9,OM10,OM11,OM12 change the + to a :M and then add a N/A[] in the Support columng
 Response Response Status W
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-47.

Cl 124 SC 124.12.4.6 P128 L10 # I-144
 Slavick, Jeff Broadcom Inc
 Comment Type TR Comment Status A (bucket1)
 PICS don't have a definition for +
 SuggestedRemedy
 Change OC5 Status to be "INS*DR4:M INS*DR42:M"
 Change OC10 Status to be "INS*DR8:M INS*DR82:M"
 Change + to :M in OC3, OC4, OC6, OC7, OC8, OC9
 Response Response Status W
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-47.

Cl 124 SC 124.12.4.6 P129 L14 # I-33
 Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)
 There is a stray : in the Status
 SuggestedRemedy
 Change
 "(DR4+DR42)*INS:M" to
 "(DR4+DR42)*INS:M"
 Response Response Status C
 ACCEPT.

Cl 162 SC 162.1 P130 L20 # I-48
 Ran, Adeo Cisco Systems, Inc.
 Comment Type ER Comment Status A test points (bucket1)
 "Annex 162A provides information on parameters with test points that may not be testable in an implemented system"
 The word "testable" is inappropriate for test points; it is the parameters associated with the test points that might not be testable, because the test points are typically inaccessible.
 SuggestedRemedy
 Change the quoted sentence to
 "Annex 162A provides information on parameters that might not be testable in an implemented system, since the test points they are associated with are typically inaccessible".

Response Response Status W
 ACCEPT.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 162 SC 162.1 P130 L20 # I-97

Dawe, Piers J G

NVIDIA

Comment Type E Comment Status A test points (bucket1)

Bad use of "may not", and contradictory to the meaning at Table 167-6. "The word may is used to indicate a course of action permissible within the limits of the standard (may equals is permitted to)." This issue is fixed in 162A.1. Missing word "associated". Also, see style guide 10.1.2 That and which.

SuggestedRemedy

Change "information on parameters with test points that may not be testable in an implemented system" to "parameters associated with test points which might not be testable in an implemented system", aligning with 162A.1 and 136A.1.

Response Response Status C

ACCEPT IN PRINCIPLE.
Resolve using the response to #I-48.

Cl 162 SC 162.8.1 P137 L8 # I-98

Dawe, Piers J G

NVIDIA

Comment Type TR Comment Status R ground

Ambiguous sentence "The PMDs on both ends of the link have connected ground references." It is not clear whether this is intended to say that:

The PMDs are connected to ground;
the PMDs are connected to each other, and that defines a "ground reference"; or
the lanes in a PMD are connected together to a "ground reference", not necessarily the ground reference for the other PMD.

If this sentence means the PMDs are connected to each other, it is not clear whether it is telling the implementer to arrange such a connection, e.g. through mains earth, or that it is provided, e.g. through the cable assembly. It is not clear whether Signal shield and/or Link shield in Fig 162-2 are involved; "The signal shields are connected to ground contacts in the MDI plug connectors on both ends of the cable assembly" but signal shields are by lane, not by PMD.

It is not clear what "ground reference" (as opposed to "ground") means. It appears in 23.5 and 32.6 (both deprecated clauses) and four times in 802.3ck, reproduced here. The term does not appear in 162.11, Cable assembly characteristics, nor does anything about shields.

SuggestedRemedy

Make clear what is required of 800GBASE-CR8 PHYs and cables. It would be better to use "common" rather than "ground" or ground reference".

When this is clear, a maintenance item for 100GBASE-CR1, 200GBASE-CR2 and 400GBASE-CR4 would be appropriate.

Response Response Status C

REJECT.

Changing the word ground to common in the following sentence does not improve the accuracy or clarity of the draft:

"The PMDs on both ends of the link have connected ground references."

Also, as noted in the suggested remedy this change would apply also to 100GBASE-CR1, 200GBASE-CR2, and 400GBASE-CR4, which are out of scope for this project.

Cl 167 SC 167.1 P156 L13 # I-4

Brown, Matthew

Alphawave

Comment Type E Comment Status A (bucket1)

It is "800GBASE-R PCS" and "800GBASE-R PMA"

SuggestedRemedy

Change "PCS for 800GBASE-R" to "800GBASE-R PCS"

Change "PMA for 800GBASE-R" to "800GBASE-R PMA"

Response Response Status C

ACCEPT IN PRINCIPLE.

Change "PCS for 800GBASE-R" to "800GBASE-R PCS"

Change "PMA for 800GBASE-R" to "800GBASE-R PMA"

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 167 SC 167.7.1 P163 L26 # I-11
 Li, Jing YOFC
 Comment Type E Comment Status R (bucket1)
 4.4|4.4
 SuggestedRemedy
 4.4
 Response Response Status C
 REJECT.
 The value is provided in two column even though it's the same for a reason. The motivation is summarized in comment I-15 in the 802.3db Draft 3.0 final comment report here:
https://www.ieee802.org/3/db/comments/P802d3db_D3p0_comments_final_by_ID_052522.pdf
 The response is repeated here:
 "TDECQ is measured using different fiber emulation filters for VR and SR. TDECQ (max) is specified in separate columns for VR and SR to note this difference even though both PMDs allow the same numerical limit for TDECQ(max) of 4.4 dB."
 In keeping with the decision, relating to this same table, made in the 802.3db task force, the columns should not be merged as proposed.

Cl 167 SC 167.7.1 P163 L30 # I-12
 Li, Jing YOFC
 Comment Type E Comment Status R parameter label (bucket2)
 Overshoot/undershoot (max)
 SuggestedRemedy
 Transmitter overshoot and undershoot (max)
 Response Response Status C
 REJECT.
 The proposed change would result in inconsistent terminology within Clause 167.
 There is no consensus to make the proposed change.

Cl 167 SC 167.7.2 P164 L26 # I-13
 Li, Jing YOFC
 Comment Type E Comment Status R parameter label (bucket2)
 Receiver sensitivity (OMAouter) (max)
 SuggestedRemedy
 Receiver sensitivity, each lane (OMAouter) (max)
 Response Response Status C
 REJECT.
 Per 167.8.2, for receiver sensitivity and stressed receiver sensitivity the interface BER is the average of the BERs of the receive lanes when they are stressed so stating the parameter value applies to "each lane" is incorrect.

Cl 167 SC 167.7.2 P164 L28 # I-14
 Li, Jing YOFC
 Comment Type E Comment Status R parameter label (bucket2)
 Stressed receiver sensitivity (OMAouter)c (max)
 SuggestedRemedy
 Stressed receiver sensitivity, each lane (OMAouter)c (max)
 Response Response Status C
 REJECT.
 Per 167.8.2, for receiver sensitivity and stressed receiver sensitivity the interface BER is the average of the BERs of the receive lanes when they are stressed so stating the parameter value applies to "each lane" is incorrect.

Cl 167 SC 167.11.4.6 P174 L10 # I-145
 Slavick, Jeff Broadcom Inc
 Comment Type TR Comment Status A (bucket1)
 PICS don't have a definition for +
 SuggestedRemedy
 Change + to :M in OC5a, OC16, OC17
 Change OC18 and OC19 to be "INS*VR8:M INS*SR8:M"
 Response Response Status W
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-47.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 169 SC 169.2.1 P178 L3 # I-49

Ran, Adeo Cisco Systems, Inc.

Comment Type TR Comment Status A (bucket1)

The title of this subclause is "Reconciliation Sublayer (RS) and Media Independent Interface (MII)" and the text includes "The Media Independent Interface (MII) specified in Clause 170".

But MII is defined in 1.4.393 (as of 802.3-2022) only with reference to clause 22. Annex 4A (which defines the MAC) does not use MII as a generic term.

For 800G, the term 800GMII (defined in 1.4.184i) should be used.

SuggestedRemedy

Change the title to "Reconciliation Sublayer (RS) and 800 Gb/s Media Independent Interface (800GMII)".

Change the subclause text accordingly

Response Response Status W

ACCEPT IN PRINCIPLE.
Implement the suggested remedy with editorial license.

Cl 169 SC 169.2.6 P178 L53 # I-50

Ran, Adeo Cisco Systems, Inc.

Comment Type ER Comment Status A (bucket1)

"Auto-Negotiation is used by the 800 Gb/s backplane PHY (800GBASE-KR8) and the 800 Gb/s copper PHY (800GBASE-CR8) is specified in Clause 73."

The sentence is incorrect as written (800GBASE-CR8 is not specified in Clause 73).

SuggestedRemedy

Change to "Auto-Negotiation is used by the 800 Gb/s backplane PHY (800GBASE-KR8) and the 800 Gb/s copper PHY (800GBASE-CR8). Auto-Negotiation is specified in Clause 73."

Response Response Status W

ACCEPT IN PRINCIPLE.
To be consistent with wording in other similar clauses implement the following...
Change: "Auto-Negotiation is used by the 800 Gb/s backplane PHY (800GBASE-KR8) and the 800 Gb/s copper PHY (800GBASE-CR8) is specified in Clause 73."
To: "Auto-Negotiation used by the 800 Gb/s backplane PHY (800GBASE-KR8) and the 800 Gb/s copper PHY (800GBASE-CR8) is specified in Clause 73."

Cl 169 SC 169.2.6 P178 L54 # I-34

Huber, Thomas Nokia

Comment Type E Comment Status A (bucket1)

One of the two instances of 'is' in the second sentence was presumably intended to be 'as'.

SuggestedRemedy

Revise the sentence to use the structure of the analogous sentence in clause 80.2.6: Clause 73 auto-negotiation is used by the 800 Gb/s backplane PHY (800GBASE-KR8) and the 800 Gb/s copper PHY (800GBASE-CR8).

Response Response Status C

ACCEPT IN PRINCIPLE.
Resolve using the response to comment #I-50.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 169 SC 169.3.3 P182 L4 # I-90

Dawe, Piers J G

NVIDIA

Comment Type TR Comment Status R PMD SI

Traditionally, the PMD limited a PAM2 signal and the PMA did timing recovery, and might include some PCB. With PAM4, the PMA does Gray mapping too. 116.3.3.2.1, Semantics of the service primitive, says that:

"each of the rx_symbol parameters can either take one of two values: zero or one; or take one of four values: zero, one, two, or three", possibly implying that the PMD makes the decisions (therefore contains any DSP equaliser and associated A to D, as well as analog equalisation). With DSP and soft decision coming to specs related to 802.3df soon, this may need to change or be clarified. We need to be careful where we assume the A to D and DSP functions are when dividing up or combining elements of the delay budget.

For EPoC, 100.2.1.2, PMD_UNITDATA.indication, says:

This primitive defines the transfer of I/Q value pair data from the Clause 100 PMD to the Clause 101 PMA. The semantics of the service primitive are PMD_UNITDATA.indication(I_value, Q_value, ChNum). The data conveyed by PMD_UNITDATA.indication is a continuous stream of I/Q value pairs and received OFDM channel. Both I_value and Q_value are encoded as 32-bit signed integers. ChNum indicates the applicable channel.

P802.3cw 156.2.1.2.1, Semantics of the primitive, says:

The PMD_UNITDATA.indication primitive conveys four *analog* signals, representing... 3cw is not binding here, but EPoC and 3cw are reasonable ways of describing the component parts, that work when more sophisticated signal processing techniques are used. But they put the A to D in different places.

SuggestedRemedy

The "PMD makes the decisions" model will put too much of the PHY in an unrecognisable "PMD sublayer". EPoC's "PMD contains the D to A" model seems un-intuitive, and it would mean that a PMA in an AUI (which obviously can contain an A to D) must have a very different delay allocation to a PMA next to the PMD. P802.3cw's "PMD may provide E/O conversion, gain, and analog EQ" model seems the most promising.

Addressing this question may be needed to set the delay limits of the sublayers.

Add an exception here, that unlike in 116.3.3.2.1, IS_UNITDATA_i.indication(rx_symbol) conveys an analog signal representing a PAM4 signal, possibly with noise and distortion. See other comments on delay.

Response Response Status C

REJECT.

For commonality with 100 Gb/s per lane interfaces for 100 Gb/s, 200 Gb/s, and 400 Gb/s Ethernet the PMD service interface should remain as currently defined.

The proposed changes might be worth considering in a later project, e.g., 802.3dj, for higher signaling rate interfaces.

Cl 169 SC 169.4 P182 L11 # I-99

Dawe, Piers J G

NVIDIA

Comment Type T Comment Status A delay wording

This text "Predictable operation of the MAC Control PAUSE operation ... concatenation of devices." looks like it was copied from 24.6 (for 100BASE-X) when a MAC bit was about 2 m long, the largest nominal reach was 2 km (1000 bits on the line) and there were repeaters. At 800G, a MAC bit is 0.25 mm long and we expect 40 km in P802.3dj (1.6e8 bits on the line, 200,000 ns). So the medium can dominate, and one should not expect all PAUSE implementations to tolerate such long links. And, no-one talks about repeaters now.

In the proposed change, the NOTE is copied from earlier clauses.

SuggestedRemedy

Update and simplify this text, e.g. "The delay limits for each sublayer are relevant to the MAC Control PAUSE operation (Clause 31, Annex 31B).

NOTE—The physical medium interconnecting two PHYs introduces additional delay in a link.

Response Response Status C

ACCEPT IN PRINCIPLE.

The introduction as written is relevant and should not be pared down.

The delays specified for the backplane (KR8) and copper cable (CR8) PMDs include allocation for the medium (14 ns or ~3 m) between Physical Layers. However, for delays specified for optical PMDs (VR8, SR8, DR8, DR8-2) include only 2 m (~10 ns) allocation for the medium between Physical Layers.

Add the following sentence after Table 169-4.

"The physical medium interconnecting two optical PHYs introduces additional delay in a link."

Cl 169 SC 169.4 P182 L13 # I-51

Ran, Adee

Cisco Systems, Inc.

Comment Type E Comment Status A (bucket1)

The sentence "in bit times as specified in 1.4 and pause_quanta as specified in 31B.2 for 800 Gigabit Ethernet" suggests that 31B.2 includes a specification for 800 Gigabit Ethernet - but it does not.

The references to 1.4 and 31B.2 are parenthetic, so corresponding punctuation should be used.

SuggestedRemedy

Change to "in bit times (as specified in 1.4) and pause_quanta (as specified in 31B.2) for 800 Gigabit Ethernet"

Response Response Status C

ACCEPT IN PRINCIPLE.

Implemented the suggested remedy with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 169 SC 169.4 P182 L16 # I-100

Dawe, Piers J G

NVIDIA

Comment Type T Comment Status A delay wording

Instead of "colocated", Clause 45 uses terminology like "instantiated within the same package" and "The definition of the term package is vendor specific and could be a chip, module, or other similar entity." We should use language consistent with Clause 45 if it is the same concept, as it appears to be. I suppose the key here could be whether the sublayers are the responsibilities of different parties or whether the interface between the sublayers is accessible for measurement. Also, this uses the spelling "colocated" (twice) while the base document uses "co-located" (twice in 55B). Spelling should be consistent.

SuggestedRemedy

Change the criterion to say that the delay for the sublayers within a single implementation, which might be a PCB, package, chip or module, is constrained by the sum of constraints for all of the sublayers within it.
If the word "colocated" is kept, reconcile the spelling with the base document.

Response Response Status C

ACCEPT IN PRINCIPLE.

The use of the word "colocated" was (incorrectly) intended to make use of established terminology for this situation. Instead, the term "in/within the same package" is used for this purpose. See examples in 802.3ck-2022 120.1.4, 135.1.4, 162.9.4, 162.9.5.1, 163.9.2, 163.9.3.1, 120F.3.1, 120G.3.1 and 802.3-2022 45.2 (many instances).

Change: "The delay for a set of colocated sublayers may be constrained by the sum of constraints for all of the colocated sublayers."

To: "The delay for a set of sublayers within the same package may be constrained by the sum of constraints for the set of sublayers."

Cl 169 SC 169.4 P182 L18 # I-52

Ran, Adee

Cisco Systems, Inc.

Comment Type TR Comment Status A delay wording

The text says that bit time and pause quanta are "for 800 Gigabit Ethernet".

The title of Table 169-4 has "800GBASE", and footnotes a and b start with "For 800GBASE-R".

Although 800GBASE-R is currently the only defined PHY family, it may not be so in the future; bit time and pause quanta are independent of the PHY type, so the footnotes should not be restricted to one PHY family.

Note that the addition of such footnotes started in Clause 80 in which there were two data rates, so it was required. It isn't required in clauses that define a single data rate, such as Clause 105. If it is anticipated that Clause 169 also introduces 1.6 Terabit Ethernet, then the distinction will be required; otherwise, the data rate can be removed from the footnotes.

The table title should be consistent with the text.

SuggestedRemedy

In the table title, change "800GBASE" to "800 Gigabit Ethernet".

In footnotes a and b, either change "For 800GBASE-R" to "For 800 Gigabit Ethernet", or delete these words.

Response Response Status C

ACCEPT IN PRINCIPLE.

Clause 169 is specific to 800 Gigabit Ethernet or 800GBASE Physical Layer implementations, so the qualifier "(800GBASE)" in the table title is not necessary.

The bit times and pause_quanta, are relevant to any 800GBASE Physical Layer implementation, so the related footnotes should not be specific to 800GBASE-R. However, it is help to be unambiguous that these numbers are specific to 800GBASE in general. In Table 169-4 title delete "(800GBASE)".

In Table 169-4 footnotes 1 and 2, change "800GBASE-R" to "800GBASE".

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 169 SC 169.4 P182 L28 # I-91

Dawe, Piers J G NVIDIA
 Comment Type ER Comment Status A delay values

The delay allowance for an 8:8 PMA is too low, and the allowance for an optical PMD is out of step with other optical PMDs. (The allowance for CR or KR PMD+AN may be wrong too, but it doesn't matter much as they are always combined with PMAs.) See daw_3df_01a_2307 Module and PMA delay limits, and other comments on delay

SuggestedRemedy

Change "800GBASE-R PMA" to "32:8 or 8:32 800GBASE-R PMA". Add a row "8:8 800GBASE-R PMA, 73,728 BT, 144 PQ, 92.16 ns (exactly twice that for the 32:8 or 8:32 PMA). Revert the VR8, SR8, DR8 and DR8-2 PMD allowances to 16,384 BT, 32 PQ, 20.48 ns.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-45.

Cl 169 SC 169.4 P182 L28 # I-101

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A delay wording

It is not clear here whether e.g. a pair of IOs forming an AUI is one PMA sublayer or two. 173.5.4 says "up to four instances of the 800GBASE-R PMA within a Physical Layer", but the relation between instance and sublayer is not given there. 120.5.4, Delay constraints, says "...up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) but it's still ambiguous. In 173.5.4, Delay constraints, "...up to four instances of the 800GBASE-R PMA", and the numbers for the PMA in Table 173-1 (not this table 169-4) apply to an instance not a sublayer. In 173.5.3.5 we have "group of PMAs" which is not explicitly defined: maybe it means any stack of nothing but PMA-things between PMD and PCS, which could be OK for this project but may need more careful definition if an inner FEC is put between or within PMA-things.

SuggestedRemedy

Consolidate the terminology (don't use "sublayer" and instance" for the same thing), and explicitly state somewhere whether a pair of IOs forming an AUI is one PMA sublayer or two. Add cross-references as appropriate, e.g. from the AUI annexes. Write something like "Each instance of a PMA" in the Notes column. Change the heading of the left column to "Sublayer or instance" if appropriate.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change "Table 169-4 contains the values of maximum sublayer delay"
 To "Table 169-4 contains the values of maximum delay for each instance of a sublayer"
 Implement with editorial license.

Cl 169 SC 169.4 P182 L28 # I-137

Maki, Jeffery Juniper Networks, Inc.
 Comment Type TR Comment Status A delay values

800GBASE-R PMA Delay + 800GBASE-DR8 PMD Delay or 800GBASE-DR8-2 PMD Delay is 87.04 ns (the optical module Delay) and is too small in relation to prevalent implementations where values are measured to be as high as 106 ns to 108 ns with the various suppliers reporting values as high as 109 ns to 129 ns.

SuggestedRemedy

Increase the allowed sum to 200 pause_quanta or 128 ns.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-45.

Cl 169 SC 169.5 P185 L34 # I-93

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status R skew variation

D2.0 comment 96: As discussed, the Skew Variation limits were based on a digital clock rate that is slow by modern standards, and they were heavily sandbagged. It is important to sort this out for 800G so that the future 200G/lane-based Ethernet is not locked into decisions made long ago for technology that doesn't apply in this case. This draft has better Skew numbers but Skew Variation needs more investigation.

SuggestedRemedy

Continue the investigation into Skew Variation, revise the numbers according to relevant technology, take out some of the padding.

Response Response Status U

REJECT.

The suggested remedy does not provide sufficient detail to implement.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 169 SC 169.6 P185 L51 # I-102

Dawe, Piers J G

NVIDIA

Comment Type TR Comment Status A FEC degrade

This says "... FEC degrade functionality is identical to that defined ... in 116.6." But 116.6 is just non-normative introduction, it contains no definition and not even any cross-references.

SuggestedRemedy

Change "Optional FEC degrade functionality is identical to that defined for 200 Gigabit Ethernet and 400 Gigabit Ethernet in 116.6." to "Optional FEC degrade functionality is as described for 200 Gigabit Ethernet and 400 Gigabit Ethernet in 116.6. For the 800GBASE-R PCS, it is defined in 172.2.5.3 (see 119.2.5.3), 172.2.5.3 (see 119.2.5.3) and 172.2.6 (see 119.2.6.2). For the 800GMII Extender, see 171.2, 118.2.1, 171.3, 118.2.2, 171.6, and 118.2."

In 116.6, insert a second sentence "For the 200GBASE-R or 400GBASE-R PCS, it is defined in 119.2.5.3, 119.2.5.3, and 119.2.6.2. For the 200GMII Extender and 400GMII Extender, see 118.2.1, 118.2.2, and 118.2."

Response Response Status C

ACCEPT IN PRINCIPLE.

Provide subclauses in Clause 171 and Clause 172 providing a summary of FEC degrade detection and signaling.

In subclause 169.6 point to 116.6 for the overview and then to the new subclauses in Clause 171 and Clause 172.

Implement with editorial license.

Cl 170 SC 170.1 P187 L7 # I-53

Ran, Adee

Cisco Systems, Inc.

Comment Type TR Comment Status A (bucket1)

"This clause defines the characteristics of the Reconciliation Sublayer (RS) and the Media Independent Interface between Ethernet media access controllers and various PHYs"

This clause is specific to 800 Gb/s PHYs. The capitalized "Media Independent Interface" is a different thing, specified for 10M/100M Ethernet in Clause 22 (see 1.4.393).

SuggestedRemedy

Change to "This clause defines the characteristics of the Reconciliation Sublayer (RS) and the 800 Gb/s Media Independent Interface (800GMII) between Ethernet media access controllers and various PHYs".

Response Response Status W

ACCEPT IN PRINCIPLE.

Change to "This clause defines the characteristics of the Reconciliation Sublayer (RS) and the 800 Gb/s Media Independent Interface (800GMII) between Ethernet media access controllers and various 800 Gb/s PHYs".

Cl 170 SC 170.1 P187 L37 # I-54

Ran, Adee

Cisco Systems, Inc.

Comment Type TR Comment Status A (bucket1)

The title of Figure 170-1 has "RS" and "MII", but the labels in the figure are "Reconciliation" and "800GMII".

SuggestedRemedy

Change the title to "Relationship of the Reconciliation Sublayer and 800GMII to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model".

Response Response Status W

ACCEPT.

Cl 170 SC 170.1.1 P188 L9 # I-55

Ran, Adee

Cisco Systems, Inc.

Comment Type T Comment Status A (bucket1)

"The following are the major concepts of the 800GMII:"

But the list discusses both the 800GMII and the RS.

SuggestedRemedy

Change "800GMII" to "800GMII and RS".

Response Response Status C

ACCEPT.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 170 SC 170.1.2 P188 L29 # I-103

Dawe, Piers J G NVIDIA
 Comment Type T Comment Status R MII

This says "This logical interface [the 800GMII] is used to provide media independence so that an identical media access controller may be used with supported PHY types". It's not really media independence; the common PCS and PMA provide that. It would allow an identical media access controller to be used with different PCSs, if the 800GXS were not used. This is unlikely. The real reason has already been stated in 170.1: "Though the 800GMII is an optional interface, it is used in this standard as a basis for specification".

SuggestedRemedy

As it is not inaccurate and not needed, delete the sentence

Response Response Status C

REJECT.
 A similar sentence has been repeated in each xMII Clause starting with 1 Gb/s Ethernet. The intent and presence of the sentence is important as it highlights the common 800GMII allows the MAC and RS to be identical for any PHY type defined or to be defined for 800 Gb/s Ethernet.

Cl 170 SC 170.4.4.1 P191 L19 # I-35

Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)

It seems odd to skip G2. This seems to be copied from clause 117, but it doesn't make any more sense there; if the intent was to align with the numbering in clause 81, the two rows should be G3 and G4 rather than G1 and G3.

SuggestedRemedy

Rather than propagate the presumed typo from clause 117, change G3 to G2

Response Response Status C

ACCEPT.

Cl 170 SC 170.4.4.2 P191 L29 # I-56

Ran, Adeo Cisco Systems, Inc.
 Comment Type T Comment Status R (bucket1)

PICS items PL2 through PL13 refer to 170.1.7 but there is no corresponding text there.

The text in 170.1.7 refers back to 81.1.7 for these functions, with an exception for EEE and LPI, which is not reflected in the PICS.

Having detailed PICS items when the text is just a reference is not helpful. The EEE/LPI exception should be noted.

Similarly for 170.4.4.2 (where multiple items refer to 170.2), and for 170.4.4.4 and 170.4.4.5 (170.3, which has an exception for EEE/LPI),

SuggestedRemedy

Replace PL2 through PL9 with a single item "Primitives mapped as specified in 81.1.7 except for EEE and LPI", 170.1.7, MII:M.

Apply similarly in other tables including the exception where appropriate.

Response Response Status C

REJECT.
 The PICS items are correct and unambiguous as written. The proposed changes do not improve the technical accuracy or clarity of the draft.

Cl 171 SC 171.1.1 P195 L39 # I-104

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)

"Each 800GXS leverages all functions in the 800GBASE-R PCS": this is ambiguous. It might be that an 800GXS uses them, or that its functions are based, more or less, on them but with modification(s). I see the word in 118.1.1; it's not good there but 118 XS functions and 119 PCS functions are not quite identical.

SuggestedRemedy

Change "leverages all functions in" to "has the same functions as".

Response Response Status C

ACCEPT.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 171 SC 171.2 P195 L46 # I-105

Dawe, Piers J G

NVIDIA

Comment Type T Comment Status A FEC degrade

Now that we have agreed that FEC degrade is optional, the same in the XS as in the PCS, there's no difference between the DTE 800GXS and the 800GBASE-R PCS. FEC degrade *signalling* in 118.2.1 (200G and 400G XS) seems to apply, but it's not an exception, and 118.2 is referenced 171.6. We need 172.2.5.3, Reed-Solomon decoder, with the two flows. More references could be useful, somewhere, as the information seems to be scattered between 118, 119, 171 and 172. I wonder if tx_am_sf should get a mention somewhere.

SuggestedRemedy

Delete "with the exception that the FEC degrade signaling is defined in 118.2.1"

Response Response Status C

ACCEPT IN PRINCIPLE.

The response to comment I-102 adds a subclause under Clause 171 to summarize the FEC degrade functionality in the DTE XS.

Change the reference from 118.2.1 to the new subclause.

Implement with editorial license.

Cl 171 SC 171.3 P195 L8 # I-78

Ran, Adee

Cisco Systems, Inc.

Comment Type TR Comment Status A MII rate

The PHY 800GXS is specified identically to the PCS with inverted transmit and receive. The PCS specification includes insertion and deletion of alignment markers. In the transmit direction, after AM insertion the signaling rate is governed by the AUI frequency range, which is +/- 50 ppm. In the receive direction the idles are removed, and _optionally_ (per 172.2.5.10) idles are inserted to compensate.

For the PHY 800GXS, the directions are reversed: it removes AMs in the transmit direction and adds them in the receive direction.

Since the idle insertion in the receive direction by the PCS is optional, and the PHY 800GXS has no exception, the PHY 800GXS is allowed not to insert idles.

The problem is that if the PHY 800GXS does not insert idles to compensate for removal of AMs, the signaling rate at the 800GMII below the PHY 800GXS will be lower than the nominal 800 Gb/s by 49 ppm, and will be different from that of the 800GMII above the DTE 800GXS. It means that the 800GMII Extender changes the rate of the 800GMII. This would be unexpected and architecturally unclean: for example, if stations are connected with synchronous clocking, the frequency difference would accumulate.

Additionally, unless the PCS (below the 800GXS) artificially increases the signaling rate back, this offset consumes 49 out of the 50 ppm that the PMD is allowed to have. This is undesirable.

To prevent the problems above it should be required that a PHY 800GXS inserts idles to compensate for AM removal in the transmit direction. Similarly, an 800GBASE-R PCS that has a PHY 800GXS as its client should be required to insert idles to compensate for AM removal in the receive direction. In both cases, functionally equivalent implementations should be allowed.

SuggestedRemedy

In 171.3, add another item to the list of exceptions:

"A PHY 800GXS is required to maintain the original data rate at the 800GMII despite the deletion of alignment markers in the transmit direction. This is done by Insertion of idle control characters or functionally equivalent behavior".

In 172.2.5.10, add the following paragraph:

"If the client of the PCS is a PHY 800GXS, the PCS is required to maintain the original data rate at the 800GMII despite the deletion of alignment markers in the receive direction. This is done by insertion of idle control characters or functionally equivalent behavior".

Response Response Status C

ACCEPT IN PRINCIPLE.

Add notes in Clause 171 and Clause 172 highlighting the issue described in this comment.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

[Editor's note: CC: 171, 172]

Cl 171 SC 171.3 P196 L8 # [-106]

Dawe, Piers J G

NVIDIA

Comment Type T Comment Status A FEC degrade

Now that we have agreed that FEC degrade is optional, the same in the XS as in the PCS, there's no difference between the DTE 800GXS and the 800GBASE-R PCS. FEC degrade *signalling* in 118.2.2 (200G and 400G XS) seems to apply, but it's not an exception, and 118.2 is referenced 171.6. We need 172.2.5.3, Reed-Solomon decoder, with the two flows. More references could be useful, somewhere, as the information seems to be scattered between 118, 119, 171 and 172. I wonder if tx_am_sf should get a mention somewhere.

SuggestedRemedy

Delete the line "-- FEC degrade signaling is defined in 118.2.2."

Response Response Status C

ACCEPT IN PRINCIPLE.

The response to comment I-102 adds a subclause under Clause 171 to summarize the FEC degrade functionality in the PHY XS.

Change the reference from 118.2.2 to the new subclause.

Implement with editorial license.

Cl 171 SC 171.3.2 P198 L18 # [-57]

Ran, Adee

Cisco Systems, Inc.

Comment Type E Comment Status A PMA name (bucket2)

In most figures in the existing standard the PMAs are designated PMA(m:n) - including in this draft, e.g. Figures 169-2 through 169-5, Figure 171-3, Figure 173-2, and all figures in Annex 173A

However, in the text of clauses 171 and 173 the PMAs are referred to as "32:8 PMA", "8:32 PMA", and "8:8 PMA", and in the PICS (173.7.3) they are listed as "PAM 32:8", "PMA 8:32", and "PMA 8:8".

Consistency is preferable.

SuggestedRemedy

In clauses 171 and 173:

Change 14 instances of "32:8 PMA" to "PMA(32:8)"
 Change 11 instances of "8:32 PMA" to "PMA(8:32)"
 Change 11 instances of "8:8 PMA" to "PMA(8:8)".

Add the missing parentheses in the PICS.

Also, change bare instances of "8:8", "32:8", "8:32" to "PCS(8:8)" etc., where appropriate (e.g. some instances in 173.2 and 173.3).

Response Response Status C

ACCEPT IN PRINCIPLE.
 Change all instances of "PMA(32:8)" to "32:8 PMA".
 Change all instances of "PMA(8:32)" to "8:32 PMA".
 Change all instances of "PMA(8:8)" to "8:8 PMA".
 Also, fix PICS feature for *P328, *P832, *P88.
 Update any other references to PMA types as appropriate.
 Implement with editorial license.

Cl 171 SC 171.3.3 P198 L36 # [-5]

Brown, Matthew

Alphawave

Comment Type E Comment Status A (bucket1)

800GMII is already defined previously in the clause, so no need to spell it out here.

SuggestedRemedy

Change "Media Independent Interface (800GMII)"
 To "800GMII"

Response Response Status C

ACCEPT.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 171 SC 171.6 P12 L12 # I-3

Brown, Matthew Alphawave

Comment Type E Comment Status A FEC degrade

Signaling of FEC degrade (local and remote) as currently defined requires the PHY XS and PCS to snoop signals in the other sublayer rather than using the more conventional method of sending signals using the inter-sublayer service interface. This makes it hard to trace the signaling between sublayers and to abstract that signaling so that different PCS types looks the same to the PHY XS.

SuggestedRemedy

Change the FEC Degrade signaling between sublayers such that it uses common signals on the PCS service interface rather than signals within the other sublayer. A presentation with a full proposal will be provided.

Response Response Status C

ACCEPT IN PRINCIPLE.

The following presentation was reviewed by the CRG:
https://www.ieee802.org/3/df/public/23_0926/nicholl_3df_01_230926.pdf

Implement the proposed changes in the following presentation:
https://www.ieee802.org/3/df/public/23_0926/nicholl_3df_01_230926.pdf
 Implement with editorial license.

Cl 172 SC 172.1.2 P206 L12 # I-58

Ran, Adeo Cisco Systems, Inc.

Comment Type T Comment Status A (bucket1)

Subclause title is "Relationship of 800GBASE-R to other standards" - but the text is specific to the PCS.

SuggestedRemedy

Change the title to "Relationship of the 800GBASE-R PCS to other standards".

Response Response Status C

ACCEPT.

Cl 172 SC 172.1.2 P207 L49 # I-59

Ran, Adeo Cisco Systems, Inc.

Comment Type TR Comment Status A (bucket1)

"Media Independent Interface" is specific to 10M/100M Ethernet.

SuggestedRemedy

Change to "800 Gb/s Media Independent Interface".

Response Response Status W

ACCEPT.

Cl 172 SC 172.2.1 P210 L6 # I-60

Ran, Adeo Cisco Systems, Inc.

Comment Type TR Comment Status A (bucket1)

The first sentence in this subclause states that "The 800GBASE-R PCS is composed of the PCS Transmit and PCS Receive processes"

But the third sentence talks about "transmit channel", and also in line 17 "When the transmit channel is in normal mode" and in line 28 "When the transmit channel is in test-pattern mode"

The term "transmit channel" appears only here while "transmit function" is used elsewhere (5 times for the PCS).

Also, the sentence "The PCS transmit channel can operate in normal mode or test-pattern mode." would be better placed right before these modes are discussed.

SuggestedRemedy

Move the sentence "The PCS transmit channel can operate in normal mode or test-pattern mode." to a separate paragraph after the second paragraph.

Change "transmit channel" to "transmit function", 3 times.

Response Response Status W

ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 172 SC 172.2.4 P211 L10 # I-107

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A encoding example

There is an informative Annex 119A, 200GBASE-R and 400GBASE-R PCS FEC codeword examples.

SuggestedRemedy

As the Clause 172 PCS is subtly different to Clause 119, with partly different alignment markers and the block distribution and synchronised alignment marker groups of the two flow method, there are new opportunities for ambiguity and misunderstanding that 119A won't catch. So, please prepare a similar annex for Clause 172. Add text here and at the beginning of 172 and 169.2.3 mentioning it. Revise the amendment description on page 14.

Please prepare a plain-text file with the large tables for convenient reading into a program, and post it on the project web site for review with future drafts.

Response Response Status C

ACCEPT IN PRINCIPLE.

The following presentation was reviewed by the CRG:
https://www.ieee802.org/3/df/public/23_0926/dawe_3df_04_230926.pdf

Implement the proposed changes in dawe_3df_04_230926.

Implement with editorial license.

Cl 172 SC 172.2.4.1 P211 L10 # I-61

Ran, Adee Cisco Systems, Inc.
 Comment Type E Comment Status A (bucket1)

The subclause title "Encode" does not match the subordinate subclause titles which use "encoder".

Also, "Encode" is also used in 172.2.4.8, a more specific term would better be used here.

Similarly in 172.2.5.9, "Decode".

SuggestedRemedy

Change the title of 172.2.4.1 to "66-bit block encoder".
 Change the title of 172.2.5.9 to "66-bit block decoder".

Response Response Status C

ACCEPT IN PRINCIPLE.
 change title of 172.2.4.1 to "64B/66B encoder"
 change title of 172.2.5.9 to "64B/66B decoder"

Cl 172 SC 172.2.4.1 P211 L11 # I-108

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)

Mixed parts of speech: Encode, State-diagram encoder, Stateless encoder, Rate matching, Block distribution, 64B/66B to 256B/257B transcoder and so on

SuggestedRemedy

Change the odd one out: change Encode to Encoder. Similarly in the title of 172.2.5.9, change Decode to Decoder.

Response Response Status C

ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-61.

Cl 172 SC 172.2.4.1.1 P211 L19 # I-109

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status R Encoder wording

"state-diagram decoder" (a tool to understand state diagrams) is something I would like to have. Would a "state-diagram encoder" turn a state diagram into code? That would be useful. If the alternative encoder needs to know the previous block as well as the one it is encoding, calling it "stateless" is borderline; if it were, we would call the first one "stateful". So these names are not ideal. They could be seen as "original" and "FEC-enabled".

SuggestedRemedy

Change to "Method A", "Method B" as we did for the 10G eye mask, unless someone has a better suggestion.

Response Response Status C

REJECT.
 The proposed changes would not result in an obvious improvement to the overall clarity of the document. The draft is correct as written.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 172 SC 172.2.4.5 P212 L19 # I-110

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A scrambler

"the two scramblers should be set to different states": this is too weak, and readers do not understand the importance of this. The consequence of getting it wrong is much more than the bad spectrum or correlation issues we have seen elsewhere.

SuggestedRemedy

Change should to shall or is.
 Add a sentence: This is because before the link can carry traffic, the 66-bit blocks in the two flows have the same content

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment I-62.

Cl 172 SC 172.2.4.5 P212 L19 # I-62

Ran, Adeo Cisco Systems, Inc.
 Comment Type T Comment Status A scrambler

The recommendation to "set to different states" deserves further explanation.

SuggestedRemedy

Add the following paragraph at the end of 172.2.4.5:

NOTE---if the two scramblers have the same state and the same input (e.g., encoded remote fault signal), their outputs will be identical. With specific choices of PMA lane muxing, this can create atypical sequences on the PMA output".

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

Cl 172 SC 172.2.4.6 P212 L35 # I-111

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A alignment markers (bucket1)

In "and finally a unique pad per PCS lane...", "finally" is unfortunate or incorrect, as the UPs don't come last. As it is only rhetorical, it can be left out.

SuggestedRemedy

Delete "finally"

Response Response Status C

ACCEPT.

Cl 172 SC 172.2.4.6 P212 L36 # I-112

Dawe, Piers J G NVIDIA
 Comment Type T Comment Status R alignment markers

119.2.6 says what to do with the common marker and unique marker portion of the alignment block but doesn't mention the unique pads. As they have so many different values, it is fair to assume they have some purpose. The reader can't know if there is a defect in the spec, or he overlooked something.

More detail: 172.2.4.6, Alignment marker mapping and insertion, incorporates 119.2.4.4, Alignment marker mapping and insertion, with exceptions. 119.2.4.4 is part of 119.2.4, Transmit. It says "The unique pad (UP0 to UP2) within the alignment markers and the PRBS9 pad at the end of the alignment maker group are ignored on receive."

172.2.5, Receive function > 172.2.5.1, Alignment lock and deskew, points to 119.2.5, Receive function. 119.2.5.1, Alignment lock and deskew, uninformatively says "It obtains lock to the alignment markers as specified by the alignment marker lock state diagram shown in Figure 119-12." 119.2.6.2.2, Variables, refers back to 119.2.4.4.

I did not find anything more about the unique pads in the standard. But see anslow_03_0416_logic.

SuggestedRemedy

Please add a few words here explaining why the unique pads are present, such as "The unique pads are remnants of the BIP fields used in the Clause 82 PCS where some PHY types did not use RS-FEC. They are ignored on receive."

Please add a sentence in 172.2.5.1: "Within the alignment block, the common marker (CM) portions are used for synchronising, the unique markers (UM) for identifying PCS lanes, and the unique pads (UP) are ignored."

Response Response Status C

REJECT.

Subclause 172.2.4.6 specifies alignment markers according 119.2.4.4 with some listed exceptions.

Specifications in 802.3 do not typically provide detailed rationale for each of choices made in the specifications. Instead, it provides all of the necessary detail to allow a designer to implement a compliant solution.

The specifications of the alignment markers including the unique pads (UPn) are currently defined with sufficient clarity and accuracy.

The rationale for the unique pad structure is the result of a series of discussions and decisions over several projects. Including this rationale would not improve the technical accuracy and clarity of the standard.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 172 SC 172.2.4.6 P212 L38 # I-113

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status R alignment markers

D2.0 comment 105 (accepted in principle): Add an informative NOTE saying what is common among these lanes, what is the same for the two flows, *and what is the same in 400G*.

SuggestedRemedy

To address the last point, please add something that gives the information in shrikhande_3df_01a_221004 slide 13:
 CM0-CM5 and UP0-UP2 are unchanged from 400GbE CL119
 UM0/UM3 for Flow lanes 0-15 are inverted from 400GbE
 UM1/UM2/UM4/UM5 for Flow lanes 16-31 are inverted from 400GbE
 e.g.:
 NOTE--CM0 to CM5 and UP0 to UP2 are the same as for 400GBASE-R (see Table 119-2). UM1, UM2, UM4, UM5 for flow 0, and UM0 and UM3 for flow 1, are the same as for 400GBASE-R. Other unique markers are bit-wise inversions of the ones in the other flow.

Response Response Status C

REJECT.
 The 802.3 standard does not typically provide rationale for specifications or provide comparisons with other rates or physical layer implementations. This requested changes are not necessary to ensure a interoperable interface for 800 Gb/s Ethernet.
 The proposed changes do not improve the clarity or accuracy of the draft.

Cl 172 SC 172.2.4.6 P213 L8 # I-114

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)

In the text above, CM0 to CM5, UM0, UP0 and so on are in regular text while in the tables, the numbers are subscripts. This should be made consistent. In spite of their use in clauses 82 and 119, the subscripts are inconvenient and not necessary.

SuggestedRemedy

Change the subscripts to regular text in these two figures

Response Response Status C

ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

Cl 172 SC 172.2.4.6 P213 L10 # I-115

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A AM table

*** Comment submitted with the file alignmentMarkerTable.txt attached ***

These table(s) of alignment markers could be put on the web in machine-readable format at <https://standards.ieee.org/downloads/>

SuggestedRemedy

Please publish a plain-text file with the alignment markers (without cell straddling) for convenient reading into a program. One table for all 32 rows x 15 columns, no header or lane number column. Tab delimited, 0x format, as in the uploaded example file. Post it on the project web site for review with future drafts.

Response Response Status C

ACCEPT IN PRINCIPLE.
 A text file provided by the commenter is provided here:
https://www.ieee802.org/3/df/comments/D3p0/I-115_alignmentMarkerTable.txt
 Add text to the draft referencing the text file, using the reference to the SSRPQ file in IEEE Std 802.3-2022 120.5.11.2.3 as a template.
 When the draft is published the text file will be posted on the IEEE web site and the URL updated.

Cl 172 SC 172.2.4.6 P213 L32 # I-63

Ran, Adeo Cisco Systems, Inc.
 Comment Type T Comment Status R AM octet order

Table 172-2 Footnote a states "Each octet is transmitted LSB to MSB".
 The transmitter order of octets should also be stated.
 Similarly in Table 172-3.

SuggestedRemedy

Insert "Octets are transmitted from CM0 to UM5. " at the beginning of the footnote, in both tables.

Response Response Status C

REJECT.
 The AM mapping and insertion is fully described in 172.2.4.6 which references 119.2.4.4. Adding this extra bit of information to the note under the table is not necessary and the suggested remedy is not sufficiently complete.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 172 SC 172.2.4.10 P216 L11 # I-116

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)

This wording causes confusion: "The portion of the figure above the "64B/66B to 256B/257B transcoder" is excluded." Which figure? How can they be excluded, it won't work!

SuggestedRemedy

Change to:
 The 66-bit block distribution of Figure 172-4 feeds the 64B/66B to 256B/257B transcoder of Figure 119-11 in each flow directly, and the portion of Figure 119-11 above the "64B/66B to 256B/257B transcoder" is not used.

Response Response Status C

ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 172 SC 172.2.4.11 P216 L43 # I-117

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A register (bucket2)

"is accessible through the register": which register?

SuggestedRemedy

is accessible through the BASE-R PCS test-pattern control register 3.42.3

Response Response Status C

ACCEPT IN PRINCIPLE.

The referenced text is as follows:

"If a Clause 45 MDIO is implemented, then the tx_test_mode variable is accessible through the register as shown in Table 172-5."

In Table 172-5, the third row provides the MDIO register and bit information for the variable tx_test_mode. No further information is required. Since the mapping of the variable to a register/bit is already provided in the table, the address need not be repeated in 172.2.4.11.

However, a cross-reference directly to the subclause defining PCS registers would be helpful.

Under 172.3, add a cross-reference to 45.2.3.
 Do similar for other clauses in this draft where appropriate.
 Implement with editorial license.

[Editor's note: CC: 172 and others]

Cl 172 SC 172.2.4.11 P216 L44 # I-118

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)

Table 172-5

SuggestedRemedy

This is not a hotlink.

Response Response Status C

ACCEPT IN PRINCIPLE.
 Change the reference to an active cross-reference.

Cl 172 SC 172.2.4.11 P216 L53 # I-142

Slavick, Jeff Broadcom Inc
 Comment Type TR Comment Status A (bucket1)

Clause 119.2.5.1 calls out the explicit amount of skew the PCS must tolerate which is different than the requirement for an 800G system.

SuggestedRemedy

Add a new exception:
 The Skew and Skew Variation requirements are specified in Table 169-5 and Table 169-6.

Response Response Status W

ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #I-119.

Cl 172 SC 172.2.5.1 P216 L54 # I-119

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A (bucket1)

There is a new exception for the alignment lock and deskew process

SuggestedRemedy

The 800GBASE-R PCS receive function shall support a maximum Skew of 152 ns between PCS lanes.
 (Editorial: "support" is lame, this should be tolerate.)

Response Response Status W

ACCEPT IN PRINCIPLE.
 Add a new exception:
 "The maximum Skew is changed from 180 ns to 152 ns"

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 172 SC 172.2.5.2 P217 L3 # I-120

Dawe, Piers J G

NVIDIA

Comment Type T Comment Status A PCS lanes wording

"PCS lanes can be received on different lanes of the service interface from which they were originally transmitted." They aren't usually received on the service interface from which they were originally transmitted, that's loopback. Lanes on lanes doesn't make sense without more explanation. Also, the PCS transmits *to* the PMA service interface beneath it.

SuggestedRemedy

Signals can be received at a PCS with the PCS lanes in a different arrangement in PMA lanes to that at the PMA service interface below the other PCS at which they were originally transmitted.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change: "PCS lanes can be received on different lanes of the service interface from which they were originally transmitted."

To: "The PCS lanes might be received in any order on the PMA service interface lanes."

Cl 172 SC 172.2.5.2 P217 L10 # I-121

Dawe, Piers J G

NVIDIA

Comment Type T Comment Status A (bucket1)

"the original stream of two FEC codewords" - there are many codewords, but two FEC streams per flow.

SuggestedRemedy

Change to: the original two streams of FEC codewords

Response Response Status C

ACCEPT.

Cl 172 SC 172.2.5.9 P217 L49 # I-122

Dawe, Piers J G

NVIDIA

Comment Type T Comment Status A (bucket1)

The receive PCS shall use the decoding method defined in either 172.2.5.9.1 or 172.2.5.9.2.

SuggestedRemedy

The receive PCS shall use one of the two decoding methods that are defined in 172.2.5.9.1 and 172.2.5.9.2.

Response Response Status C

ACCEPT.

Cl 172 SC 172.2.6.2.4 P220 L9 # I-80

Opsasnick, Eugene

Broadcom Inc.

Comment Type TR Comment Status A counters

This section states that the counters for 800GBASE-R PCS use the same values as 119.2.6.2.4 for the 400GBASE-R PCS.

The amp_counter value is used in Figure 119-12 "Alignment marker lock state diagram" to count the appropriate number of FEC codewords between alignment markers. This number is 4096 for 200Gb/s and 8192 for 400Gb/s as specified in 119.2.6.2.4

For 800Gb/s, the spacing between alignment markers should be 16k codewords as shown in the adoped baseline shrikhande_3df_01a_221004.pdf on slide #12.

SuggestedRemedy

Change the wording in 172.2.6.2.4

from:

"The counters are the same as those in specified in 119.2.6.2.4 for the 400GBASE-R PCS."

to:

"The counters are the same as those in specified in 119.2.6.2.4 for the 400GBASE-R PCS with the following exception:
amp_counter

This counter counts the interval of 16,384 FEC codewords containing normal alignment marker payload sequences for the 800GBASE-R PCS."

Response Response Status C

ACCEPT IN PRINCIPLE.

Change:

"The counters are the same as those in specified in 119.2.6.2.4 for the 400GBASE-R PCS."

To:

"The counters are the same as those in specified in 119.2.6.2.4 for the 400GBASE-R PCS with the following exception:
amp_counter

This counter counts the interval of 16 384 FEC codewords containing alignment marker payload sequences for the 800GBASE-R PCS."

Implement with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 172 SC 172.5 P223 L50 # I-64

Ran, Adeo Cisco Systems, Inc.
 Comment Type ER Comment Status R numbers (bucket1)

"640 000"

Per the style manual, the use of space as a thousands separator is specified for numbers within tables. There is no need to use it in text and it adds no clarity.

Adding spaces in numbers within clause creates significant issues in other places of the standard and should be avoided.

SuggestedRemedy

Change "640 000" to "640000".

Response Response Status W

REJECT.

Note that there are other similar instances: "32 768" in 167.3.1, "32 768" in 124.3.1 and 167.3.1"

The guidance from the publication editors is that thousands separator is required with some exceptions, e.g., where the readability is compromised. Readability is not an issue for the cases noted in the comment and in this response above.

The proposed change is not consistent with the editorial guidelines.

See related comment #I-73.

Cl 172 SC 172.7.4 P226 L22 # I-65

Ran, Adeo Cisco Systems, Inc.
 Comment Type E Comment Status A (bucket1)

Many PICS items refer to subclauses in 172 for features that are not explicitly specified there but refer back to clause 119.

SuggestedRemedy

Whenever there are multiple items referring to a subclause that only refers back to clause 119, consider replacing these items with a single item that points to the subclause in clause 172, across the PICS tables.

Response Response Status C

ACCEPT IN PRINCIPLE.

To be consistent, it is an improvement to point to the relevant subclauses in Clause 172 rather than Clause 119. However, it does not seem necessary or practical to collapse multiple items together as proposed nor does this improve the technical clarity or accuracy of the draft.

Change references to subclauses in Clause 119 to the relevant subclauses in Clause 172.

Cl 173 SC 173.1.3 P231 L13 # I-123

Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A squelch (bucket1)

As it is a new observable behaviour, the optional squelch feature should be mentioned here in the overview and in 173.2 PMA service interface. And, the word "squelch" should be used so readers will recognise it.

SuggestedRemedy

In 173.1.3 Summary of functions, add a row:

-- Optionally indicate status by disabling (squelching) a lane or lanes

In 173.2 page 233 line 8, add sentences "The 8:32 PMA optionally provides signal status information to the PMA client by disabling (squelching) a lane or lanes (see

173.5.8.2). "The 8:8 PMA optionally provides signal status information in either direction by disabling (squelching) a lane or lanes (see 173.5.8.3)."

Response Response Status C

ACCEPT IN PRINCIPLE.

The automatic AUI output disable (squelch) is one of many aspects of "Provide signal status information". Providing the extra detail as suggested in this high-level "Summary of functions" in 171.1.3 is not warranted.

However, providing the additional sentences to the service interface definition sections (173.2 and 173.3) would be an improvement to the draft.

In 173.2 page 233 line 7 add the following sentence to end of the third paragraph:

"The 8:32 and 8:8 PMAs may optionally provide signal status information to the PMA client by disabling (squelching) one or more of the PAM4 symbol streams sent to the PMA client (PMA:IS_UNITDATA_0:7.indication), see 173.5.8.2 and 173.5.8.3.

In 173.3 page 233 line 32 add the following new paragraph:

For the 8:8 PMA, if the sublayer below the PMA is another PMA, the 8:8 PMA may optionally provide signal status information by disabling (squelching) one or more of the PAM4 symbol streams sent to the sublayer below (PMA:IS_UNITDATA_0:7.request), see 173.5.8.3.

Implement with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 173 SC 173.4.1 P234 L35 # I-141

Slavick, Jeff Broadcom Inc
 Comment Type T Comment Status A test patterns

The dotted arrows in Figure 173-3, Figure 173-4 and Figure 173-5 aren't accurately placed.

SuggestedRemedy

- In all 3 figures
- Shift the dotted arrow(s) going from test pattern generate to have it go into the PAM4 encode and signal drivers box
- Shift the dotted arrow(s) going into test pattern check to come from the PAM4 decode and CDR box
- Shift the dotted arrow(s) going to the SIL to come from the PAM4 decode and CDR box

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-66.

Cl 173 SC 173.4.1 P234 L35 # I-66

Ran, Adee Cisco Systems, Inc.
 Comment Type T Comment Status A test patterns

The dashed-line arrows in Figure 173-3 are not connected to the right places.

"Test pattern generate" creates bits that are encoded as PAM4 symbols and then driven by the same signal drivers. It should go into the "PAM4 encode/Signal drivers" box.

"Test pattern check" operates on a bit stream, so should take the output of "PAM4 encode/CDR".

The arrow leading to "SIL" denotes information from the CDR. It should be taken from the PAM4 decode/CDR box.

Similarly in Figure 173-4 and Figure 173-5.

SuggestedRemedy

Modified figures will be supplied

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement with editorial license the changes shown on slides 19 to 21 of the following presentation:
https://www.ieee802.org/3/df/public/23_0926/brown_3df_03a_230926.pdf
 except draw the line for the test pattern check directly out of the "PAM4 decode" box instead of after the box.

Cl 173 SC 173.4.3 P237 L46 # I-124

Dawe, Piers J G NVIDIA
 Comment Type T Comment Status R delay wording (bucket1)

While an 8:8 PMA is clear and understandable, it seems that at this speed, with PAM4 and equalisation, implementations are typically back-to-back SerDes. This solves the problem of specifying its maximum delay appropriately.

SuggestedRemedy

If the group sees this as an improvement saying that an 8:8 PMA is specified by assuming that it is back-to back 8:32 and 32:8 PMAs, addressing any conflict between this and 173.5.2.3 restricted bit muxing.

Response Response Status C

REJECT.

The 8:8 PMA is distinctly different from a back-to-back 8:32 PMA and 32:8 PMA. For instance, there is an explicit rule that groups of PCS lanes on each physical lane remain together through the PMA. The latency concern can more easily be addressed, if necessary, by increasing the specified value for the 8:8 PMA.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 173 SC 173.5.2.1 P238 L20 # I-125

Dawe, Piers J G

NVIDIA

Comment Type E Comment Status A (bucket1)

"the function": what or which function? Compare lines 31, 39, 46

SuggestedRemedy

Add words such as "bit-level multiplexing" at least here, the first time, and preferably in 173.5.2.2. e.g. "8:32 bit-level multiplexing" would be better. Also at line 31, but maybe that can be "this function".

Response Response Status C

ACCEPT IN PRINCIPLE.

In 173.5.2.1

- add the following new paragraph before the first paragraph, "The 32:8 PMA provides bit-level multiplexing in both the transmit and receive directions."
- change "In the transmit direction, the function is performed" to "In the transmit direction, the bit-level multiplexing function is performed"
- change "In the receive direction, the function is performed" to "In the receive direction, the bit-level multiplexing function is performed"

In 173.5.2.2:

- add the following new paragraph before the first paragraph, "The 8:32 PMA provides bit-level multiplexing in both the transmit and receive directions. "
- change "In the transmit direction, the function is performed" to "In the transmit direction, the bit-level multiplexing function is performed"
- change "In the receive direction, the function is performed" to "In the receive direction, the bit-level multiplexing function is performed"

In 173.5.2.3:

- add the following new paragraph before the first paragraph, "The 8:8 PMA provides bit-level multiplexing in both the transmit and receive directions."
- change "In the transmit direction, the function is performed" to "In the transmit direction, the bit-level multiplexing function is performed"
- change "In the receive direction, the function is performed" to "In the receive direction, the bit-level multiplexing function is performed"

Implement with editorial license.

Cl 173 SC 173.5.2.1 P238 L23 # I-67

Ran, Adeo

Cisco Systems, Inc.

Comment Type E Comment Status A (bucket1)

"referencing the functional block diagram shown in..." does not sound right.

This appears in 173.5.2.1, 173.5.2.2, and 173.5.2.3, two instances each.

SuggestedRemedy

Change "referencing the functional block diagram shown in" to "as shown in", in all 6 instances.

Response Response Status C

ACCEPT.

Cl 173 SC 173.5.2.1 P238 L28 # I-126

Dawe, Piers J G

NVIDIA

Comment Type TR Comment Status A muxing rules

"with two lanes from ... followed by two lanes from ..." isn't right. Lanes exist continuously, they can be in parallel but cannot follow.

SuggestedRemedy

Bits from the four PCSLs are multiplexed in temporal order with one bit from each of two lanes from PMA client lanes i = 0 to 15 followed by one bit from each of two lanes from PMA client lanes i = 16 to 31. Similarly in 173.5.2.2.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change: "The four PCSLs are multiplexed in temporal order with two lanes from PMA client lanes i = 0 to 15 followed by two lanes from PMA client lanes i = 16 to 31."

To: "Bits from the four PCSLs are multiplexed in temporal order with one bit from each of two lanes from PMA client lanes i = 0 to 15 followed by one bit from each of two lanes from PMA client lanes i = 16 to 31."

Implement similar changes in 173.5.2.2 and elsewhere if appropriate.

Implement with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 173 SC 173.5.2.1 P238 L28 # I-68

Ran, Adeo Cisco Systems, Inc.
 Comment Type T Comment Status A muxing rules

"The four PCSLs are multiplexed in temporal order with two lanes from PMA client lanes i = 0 to 15 followed by two lanes from PMA client lanes i = 16 to 31"

The clarity and accuracy of this sentence can be improved.

SuggestedRemedy

Change to
 "The four PCSLs are multiplexed in temporal order such that two bits received from two of the PMA client lanes with i=0 to 15 are followed by two bits received from two of the PMA client lanes with i=16 to 31".

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-126.

Cl 173 SC 173.5.2.3 P239 L22 # I-81

Ran, Adeo Cisco Systems, Inc.
 Comment Type T Comment Status A muxing rules

The allowed swapping of bit pairs may seem ambiguous to some readers. It can be rephrased to be complete and reduce the risk of misunderstanding.

Alternatively, the option of swapping bits can be removed from the draft; whether it is allowed or not in the standard would not matter in practice. If that solution is chosen, the words "except for possible swapping of each bit pair" should be removed.

SuggestedRemedy

Change from
 "such that the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane, except for possible swapping of each bit pair (see 173.5.7.1)"
 to
 "such that the Gray mapped PAM4 symbol sequence on the output lane is either identical to the Gray mapped PAM4 symbol sequence on the input lane, or is the result of swapping the order of each pair of bits {A, B} to {B, A} in the Gray mapping function (see 173.5.7.1)".

Response Response Status C

ACCEPT IN PRINCIPLE.

Change:
 "such that the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane, except for possible swapping of each bit pair (see 173.5.7.1)"

To:
 "such that the Gray mapped PAM4 symbol sequence on the output lane is either identical to the Gray mapped PAM4 symbol sequence on the input lane, or is the result of reversing the order of each pair of bits {A, B} to {B, A} in the Gray mapping function (see 173.5.7.1)".

Implement the suggested remedy with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 173 SC 173.5.2.3 P239 L22 # I-127

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A muxing rules

"except for possible swapping of each bit pair": discussions have established that bit pairs may not be swapped. Bits within pairs may, but this needs more careful definition because of the Gray mapping. "except for possible" reads like an anti-recommendation in unusual wording contrary to house style, but if the receiver can cope with the bit swapping, there is no point recommending the "identical" method over it.

SuggestedRemedy

Change the item to:
 "The 4 PCSLs received on an input lane shall be mapped to a single output lane. Either the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane, or the sequence on the output lane is the result equivalent to undoing the Gray mapping function (see 173.5.7.1), swapping the bits in each pair of bits {A, B} to {B, A}, and Gray mapping to PAM4."

Response Response Status C
 ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-81.

Cl 173 SC 173.5.3.1 P238 L39 # I-69

Ran, Adeo Cisco Systems, Inc.
 Comment Type E Comment Status A skew wording

"shall produce" here, "shall generate" in 173.5.3.3, "shall deliver" in 173.5.3.5... the title of all three has "skew generation".

In fact, the skew numbers stated are cumulative.

Since the skew at any point is not necessarily generated at that point, the proper requirement seems to be "shall have".

SuggestedRemedy

Change all three "shall" statements in the comment to "shall have".

Response Response Status C
 ACCEPT IN PRINCIPLE.

Implement with editorial license the changes shown on slides 16 and 17 of the following presentation:
https://www.ieee802.org/3/df/public/23_0926/brown_3df_03a_230926.pdf

Cl 173 SC 173.5.3.2 P239 L44 # I-70

Ran, Adeo Cisco Systems, Inc.
 Comment Type T Comment Status A (bucket1)

"the PMA service interface that receives data in the transmit direction ... shall tolerate the maximum amount of Skew Variation"

The PMA has to tolerate skew variation, not its service interface (see also 173.5.3.4 where it's the PMA).

SuggestedRemedy

Delete "service interface".

Response Response Status C
 ACCEPT.

Cl 173 SC 173.5.3.3 P239 L53 # I-128

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A skew wording

In these subclauses, skew is generated, produced or delivered. It is not clear what these terms mean. I believe that all Skew limits are cumulative (unlike for delay) which has a bearing on what the terms mean.

SuggestedRemedy

Write down what generated, produced and delivered mean here and what the differences are.

Response Response Status C
 ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-69.

Cl 173 SC 173.5.4 P240 L35 # I-129

Dawe, Piers J G NVIDIA
 Comment Type T Comment Status A (bucket1)

within a Physical Layer, which is composed of an 800GBASE-R PHY and an optional 800GMII Extender

SuggestedRemedy

within a Physical Layer, which is composed of an 800GBASE-R PHY and, optionally, an 800GMII Extender

Response Response Status C
 ACCEPT.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 173 SC 173.5.4 P240 L35 # I-130

Dawe, Piers J G NVIDIA
 Comment Type T Comment Status A (bucket1)

It would avoid misinterpretation if the words to the effect of delay is the sum of transmit and receive delays, were reinstated. 169.4 says it, but it is not referenced here for definitions and it is borderline non-normative "Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails."

SuggestedRemedy

Insert words: The maximum delay (sum of transmit and receive delays) contributed by each instance ...

Response Response Status C
 ACCEPT.

Cl 173 SC 173.5.4 P240 L46 # I-92

Dawe, Piers J G NVIDIA
 Comment Type TR Comment Status A delay values

This new delay allocation per PMA-instance may be OK where a PMA is packaged with a PCS, XS or PMD, but it is tight for a standalone PMA (e.g. "on-board retimer"). It is unlikely that a PMA will be packaged with an exposed 32x25G PMA interface except in a prototype.

SuggestedRemedy

Double the allowance for the 8:8 PMA only, from 36,864 BT, 72 PQ, 46.08 ns to 73,728 BT, 144 PQ, 92.16 ns. No need to change the delay allocation for 32:8 and 8:32 PMA.

Response Response Status C
 ACCEPT IN PRINCIPLE.

The commenter indicated that this applies to subclause 173.5.4 rather than 173.6.4.

Resolve using the response to comment #I-45.

Cl 173 SC 173.5.5 P240 L51 # I-36

Huber, Thomas Nokia
 Comment Type E Comment Status A (bucket1)

The variable n should be italicized in the first line

SuggestedRemedy

Format the n in "n output lanes" in italics

Response Response Status C
 ACCEPT.

Cl 173 SC 173.5.5 P241 L2 # I-131

Dawe, Piers J G NVIDIA
 Comment Type T Comment Status A (bucket1)

If an output lane's clock is derived from its corresponding input, it's not independent.

SuggestedRemedy

As this is only an example, changing "independent" to "separate" or "its own" would be enough to correct this

Response Response Status C

ACCEPT IN PRINCIPLE.
 Change "each output lane could use an independent clock derived from its corresponding input" to "each output lane could use a separate clock derived from its corresponding input".

Cl 173 SC 173.5.6 P241 L8 # I-71

Ran, Adee Cisco Systems, Inc.
 Comment Type E Comment Status A (bucket1)

"For cases where the interface between the PMA client and the PMA, or between the PMA and the sublayer below the PMA represent a physically instantiated interface, ..."

This sentence is unnecessarily complex and the punctuation is incorrect.

SuggestedRemedy

Change to "When the interface between the PMA client and the PMA, or between the PMA and the sublayer below the PMA, is physically instantiated, ..."

Response Response Status C

ACCEPT IN PRINCIPLE.
 Change: "For cases where the interface between the PMA client and the PMA, or between the PMA and the sublayer below the PMA represent a physically instantiated interface, the PMA provides electrical signal drivers for that interface."
 To: "For a case where there is a physically instantiated interface the PMA provides electrical signal drivers."

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 173 SC 173.5.8.1 P242 L3 # I-72

Ran, Adeo Cisco Systems, Inc.

Comment Type T Comment Status A signal status (bucket2)

The requirement that "data is being sent on all 32 output lanes (PMA:IS_UNITDATA_0:31.indication)" is unique to this PMA (32:8); the other two PMAs set the signal status only based on data being received on the appropriate interface.

In real implementations, an indication to the PCS that data is not being received by the PMA (which may be due to lack of a link partner) would likely be separate from an indication that data is not being transmitted (essentially a local fault). Specifying in the standard that it's the same indication is not helpful for readers.

SuggestedRemedy

Delete the second item in the list.

Consider converting the list to regular paragraph text as in the other two subclauses.

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement with editorial license the changes shown on slide 23 of the following presentation:
https://www.ieee802.org/3/df/public/23_0926/brown_3df_03a_230926.pdf

Cl 173 SC 173.5.8.2 P242 L13 # I-132

Dawe, Piers J G NVIDIA

Comment Type T Comment Status A wording (bucket2)

It is hard work reverse engineering this: "In the *transmit* direction ... The SIGNAL_OK parameter is set to OK when data is being *received*... I believe that less confusing language has been used somewhere. Ingress and egress could be used.

SuggestedRemedy

Change "when data is being received on all 8 input lanes (PMA:IS_UNITDATA_0:7.request)." to "when data is presented to this PMA sublayer by the PMA sublayer above on all 8 transmit lanes (PMA:IS_UNITDATA_0:7.request)". Similarly in 173.5.8.3 8:8, line 23, change "when data is not being received on all 8 input lanes (PMA:IS_UNITDATA_0:7.request)." to "when data is not being presented to this PMA sublayer by the PMA sublayer above on all 8 input lanes (PMA:IS_UNITDATA_0:7.request)".

Response Response Status C

ACCEPT IN PRINCIPLE.

The direction of transmission and the relevant interfaces are clear and unambiguous. The meaning of the word "received" here is clear given the context. The proposed changes are not an improvement to the technical clarity or accuracy of the text. However, it would be helpful to point to the figure that shows the 8:8 PMA interfaces.

In 173.5.8.3...

In the first paragraph add a reference to Figure 173-5.

In the second paragraph, change "Figure 173-4" to "Figure 173-5".

Implement with editorial license.

Cl 173 SC 173.5.8.3 P242 L18 # I-133

Dawe, Piers J G NVIDIA

Comment Type E Comment Status A squelch (bucket1)

Please name this feature by its familiar name so readers can find it. This is a kind of disabling is new to 802.3 but its name is well established in the industry.

SuggestedRemedy

by disabling (squelching) one or more output lanes
 Same (twice) in next subclause

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 173 SC 173.5.8.3 P242 L19 # I-134
 Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)
 Two dumb cross-references, and two more at line 29.
 SuggestedRemedy
 Make them hot links
 Response Response Status C
 ACCEPT.

Cl 173 SC 173.7.3 P246 L12 # I-146
 Slavick, Jeff Broadcom Inc
 Comment Type TR Comment Status A (bucket1)
 PICS don't have a definition for +
 SuggestedRemedy
 Change C2CA and C2MA to be "P832:O/2 P88:O/2"
 Change C2CB, C2MB, PMDE, PMDO to be "P328:O/3 P88:O/3"
 Response Response Status W
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #-I-47.

Cl 173 SC 173.7.3 P246 L32 # I-135
 Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)
 The optional squelch affects how a PMA is used, so it should appear in the PICS major options
 SuggestedRemedy
 Add two major options, for the receive (ingress) direction and for the transmit (ingress) direction, conditionally optional according to PMA type.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 The squelching functions are provided in PICS items SS1 and SS2 in 173.7.9. These are not major functions, but rather one of many minor features that are specified. It is therefore not appropriate to move these to the "major functions" table. However, for SS1 and SS2 the word "squelching" should be added and the subclause references are incorrect. And also PICS items are missing for the general signal status specifications.
 For SS1 and SS2 feature descriptions change "disabling" to "disabling (squelching)"
 For SS1 and SS2 subclause change to 173.5.8.2 and 173.5.8.3, respectively.
 Add new items for signal status for each PMA type per 173.5.8.1, 173.5.8.2, 173.5.8.3"
 Implement with editorial license.

Cl 173 SC 173.7.4 P246 L42 # I-147
 Slavick, Jeff Broadcom Inc
 Comment Type TR Comment Status A (bucket1)
 PICS don't have a definition for +
 SuggestedRemedy
 Change + to a :M in S1, S2, S3, S7, S8, S9
 Response Response Status W
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #-I-47.

Cl 173 SC 173.7.6 P248 L6 # I-148
 Slavick, Jeff Broadcom Inc
 Comment Type TR Comment Status A (bucket1)
 PICS don't have a definition for +
 SuggestedRemedy
 Change + to a :O in T1, T2, T3, T4, T5, T6
 Response Response Status W
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #-I-47.

Cl 173 SC 173.7.7 P248 L37 # I-136
 Dawe, Piers J G NVIDIA
 Comment Type E Comment Status A (bucket1)
 If the two loopback abilities aren't in the major options table as in 120.7.3, there is no point having separate PCS for "PMA local loopback" and "PMA local loopback implemented". Nothing else depends on "LBL".
 SuggestedRemedy
 Move the loopback abilities to the major options, as in 120.7.3, or combine the two pairs
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Remove the *LBL and *LBR rows from the table.
 For LB1 and LB2...
 in the Feature column delete "implemented"
 Change the status to "O".

IEEE P802.3df D3.0 Initial Sponsor ballot comments

Cl 173 SC 173.7.8 P248 L 54 # I-149

Slavick, Jeff Broadcom Inc

Comment Type TR Comment Status A (bucket1)

PICS don't have a definition for +

SuggestedRemedy

Change + to a :M in P1 and + to a :0 in P4

Response Response Status W

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #I-47.

Cl 173A SC 173A P283 L 8 # I-74

Ran, Adeo Cisco Systems, Inc.

Comment Type E Comment Status R (bucket1)

This annex is titled "800 Gb/s PMA sublayer partitioning examples", but it's about Physical layer partitioning examples, not PMA sublayer partitioning. The PMA is not partitioned.

SuggestedRemedy

Change Annex title to "800 Gb/s Physical layer partitioning examples".

Response Response Status C

REJECT.

This annex, like similar ones used for other Ethernet rates, demonstrates variations of a physical layer implementation with different sets of physical instantiations of the PMA service interface (800GAUI-n) and the resulting MMD address to be assigned to each of the PMA sublayers.