

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl **FM** SC **FM** P1 L 29 # 50
 Grow, Robert Self
 Comment Type **E** Comment Status **X**
 Both cx and cz were approved during the March SASB meeting and should be referenced with the year 2023.
 SuggestedRemedy
 Replace "202x" with "2023" here and on page 12.
 Proposed Response Response Status **O**

Cl **FM** SC **FM** P8 L 24 # 52
 Grow, Robert Self
 Comment Type **E** Comment Status **X**
 The WG ballot group is now known, please fill in so that names can be reviewed.
 SuggestedRemedy
 Per comment.
 Proposed Response Response Status **O**

Cl **FM** SC **FM** P4 L 21 # 51
 Grow, Robert Self
 Comment Type **E** Comment Status **X**
 This is not the current front matter.
 SuggestedRemedy
 Replace with current front matter.
 Proposed Response Response Status **O**

Cl **FM** SC **FM** P8 L 42 # 72
 D'Ambrosia, John Futurewei, US Subsidiary of Huawei
 Comment Type **E** Comment Status **X**
 Members of WG Ballot not added
 SuggestedRemedy
 Add WG Balloting List
 Proposed Response Response Status **O**

Cl **FM** SC **FM** P8 L 12 # 71
 D'Ambrosia, John Futurewei, US Subsidiary of Huawei
 Comment Type **ER** Comment Status **X**
 Task Force Leadership not fully recognized
 SuggestedRemedy
 1. Modify
 "Mark Nowell, IEEE P802.3df Task Force Vice Chair"
 to
 Mark Nowell, IEEE P802.3df Task Force Vice Chair, IEEE P802.3df "Optics"Sub-task Force Chair
 2. Add
 Kent Lusted, IEEE P802.3df "Electrical" Sub-task Force Chair
 Mark Gustlin, IEEE P8023df "Architecture and Logic" Sub-task Force Chair
 Proposed Response Response Status **O**

Cl **FM** SC **FM** P12 L 37 # 53
 Grow, Robert Self
 Comment Type **E** Comment Status **X**
 This is not the self description of the approved D3.2 draft. The end of the self description was changed when the original project was split adding P802.3dh. (Publication of IEEE Std 802.3cz-2023 is expected soon.)
 SuggestedRemedy
 "for optical automotive Ethernet using graded-index glass optical fiber."
 Proposed Response Response Status **O**

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl **FM** SC **FM** P **12** L **47** # **54**
 Dudek, Mike Marvell
 Comment Type **E** Comment Status **X**
 IEEE Std 802.3-2022 has been published
 SuggestedRemedy
 Change 202x to 2022
 Proposed Response Response Status **O**

Cl **00** SC **0** P **225** L # **68**
 D'Ambrosia, John Futurewei, US Subsidiary of Huawei
 Comment Type **TR** Comment Status **X**
 As noted in Tables 169-2 and 169-3, 800G AUI variants are optional for both 800G copper and optical PHY types, which means you could have an 800GAUI-8 in the PHY as well as in the extender. This means you would PMA (32:8) and PMA (8:32) to support AUIs - not PMA 8:8
 See Fig 173A-4 as example that a PMA (32:8) is called out for connecting to a 800GAUI-8
 SuggestedRemedy
 The statements regarding the 32:8 and 8:32 PMAs should reflect being present to support 800GAUIs which may not just be in the Extender as currently stated.
 Proposed Response Response Status **O**

Cl **1** SC **1.3** P **30** L **40** # **55**
 Dudek, Mike Marvell
 Comment Type **E** Comment Status **X**
 "One fibre rows" is strange.
 SuggestedRemedy
 Check the reference and correct to "One fibre row" unless the reference does have this in its title.
 Proposed Response Response Status **O**

Cl **1** SC **1.4.135a** P **30** L **49** # **73**
 D'Ambrosia, John Futurewei, US Subsidiary of Huawei
 Comment Type **E** Comment Status **X**
 400GBASE-DR4-2
 The term "lane" is ambiguous when discussing SMF -as a lane may be either a wavelength or a fiber.
 SuggestedRemedy
 Change
 "IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding and 4-level pulse amplitude modulation over four lanes of single-mode fiber, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)"
 to
 IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding and 4-level pulse amplitude modulation over four wavelengths distributed over 4 single-mode fibres, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)
 Proposed Response Response Status **O**

Cl **1** SC **1.4.148i** P **31** L **44** # **24**
 Slavick, Jeff Broadcom
 Comment Type **TR** Comment Status **X**
 Isn't it a 800GMII interface between the RS and either a PCS or Extender and an Extender and a PCS. This definition only lists RS to PCS.
 SuggestedRemedy
 The interface used between the Reconciliation Sublayer (RS), Media Independent Interface Extender Sublayer (XS) and the Physical Coding Sublayer (PCS) for 800 Gb/s operation
 Proposed Response Response Status **O**

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 1 SC 1.4.184b P31 L 6 # 74

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

800GBASE-DR8

The term "lane" is ambiguous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change

"IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of single-mode fiber, with reach up to at least 500 m. (See IEEE Std 802.3, Clause 124.)"

to

IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight wavelengths distributed over 8 single-mode fibers with reaches up to at least 500 m. (See IEEE Std 802.3, Clause 124.)

Proposed Response Response Status O

Cl 1 SC 1.4.184c P31 L 10 # 75

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

800GBASE-DR8-2

The term "lane" is ambiguous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change

"IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of single-mode fiber, with reach up to at least 500 m. (See IEEE Std 802.3, Clause 124.)"

to

IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight wavelengths distributed over 8 single-mode fibers with reaches up to at least 2 km. (See IEEE Std 802.3, Clause 124.)

Proposed Response Response Status O

Cl 1 SC 1.4.184f P31 L 20 # 76

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

800GBASE-SR8

With the introduction of WDM technology over MMF, the term "lane" is ambiguous when discussing MMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change

"IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of multimode fiber, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 167.)"

to

IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight wavelengths distributed over 8 multimode fibers, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 167.)

Proposed Response Response Status O

Cl 1 SC 1.4.184g P31 L 24 # 77

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

800GBASE-VR8

With the introduction of WDM technology over MMF, the term "lane" is ambiguous when discussing MMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change

"IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of multimode fiber, with reach up to at least 50 m. (See IEEE Std 802.3, Clause 167.)"

to

IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight wavelengths distributed over 8 multimode fibers, with reach up to at least 50 m. (See IEEE Std 802.3, Clause 167.)

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 1 SC 1.4.184h P 33 L 37 # 56
 Dudek, Mike Marvell
 Comment Type E Comment Status X
 The editors note has served its purpose
 SuggestedRemedy
 delete the editors note
 Proposed Response Response Status O

Cl 4 SC 4.4.2 P 33 L 32 # 41
 Schreiner, Stephan Rosenberger Hochfrequenztechnik GmbH & Co. KG
 Comment Type E Comment Status X
 in minFrameSize for 2.5 GB/s, 5 GB/s,... is a line break after 512 bits, which might be caused by a different column width
 SuggestedRemedy
 Increase width of column to match the size of the other columns from the MAC data rate
 Proposed Response Response Status O

Cl 1 SC 1.4.184k P 32 L 1 # 47
 Brown, Matt Huawei
 Comment Type E Comment Status X
 No such thing as "800 Gb/s Extender Sublayer". See 171.1.
 SuggestedRemedy
 Change "800 Gb/s Extender Sublayer" to "800GMII Extender Sublayer"
 Also in 1.5, page 32, line 32
 Proposed Response Response Status O

Cl 30 SC 30.5.1.1.2 P 34 L 51 # 78
 D'Ambrosia, John Futurewei, US Subsidiary of Huawei
 Comment Type E Comment Status X
 400GBASE-DR4
 The term "lane" is ambiguous when discussing SMF -as a lane may be either a wavelength or a fiber.
 SuggestedRemedy
 Change
 400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124
 to
 400GBASE-R PCS/PMA over 4 wavelengths distributed over 4 single-mode fibres PMD with reach up to at least 500 m as specified in Clause 124
 Proposed Response Response Status O

Cl 1 SC 1.4.461 P 32 L 18 # 31
 Huber, Tom Nokia
 Comment Type E Comment Status X
 The text has a comma splice
 SuggestedRemedy
 Change "...the PCS distributes data to multiple logical lanes, these logical lanes are called PCS lanes." to "...the PCS distributes data to multiple logical lanes that are called PCS lanes."
 Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 30 SC 30.5.1.1.2 P 35 L 8 # 79

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

800GBASE-DR8

The term "lane" is ambiguous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change

800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124

to

800GBASE-R PCS/PMA over 8 wavelengths distributed over 8 single-mode fibres PMD with reach up to at least 500 m as specified in Clause 124

Makes changes throughout document as appropriate with editorial license

Proposed Response Response Status O

Cl 30 SC 30.5.1.1.2 P 35 L 14 # 61

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

800GBASE-SR8

With the introduction of WDM technology over MMF, the term "lane" is ambiguous when discussing MMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change

800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 100 m as specified in Clause 167

to

800GBASE-R PCS/PMA over 8 wavelengths distributed over 8 multi-mode fibres PMD with reach up to at least 100 m as specified in Clause 167

Makes changes throughout document as appropriate with editorial license

Proposed Response Response Status O

Cl 30 SC 30.5.1.1.2 P 35 L 10 # 80

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

800GBASE-DR8-2

The term "lane" is ambiguous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change

800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 124

to

800GBASE-R PCS/PMA over 8 wavelengths distributed over 8 single-mode fibres PDwith reach up to at least 2 km as specified in Clause 124

Makes changes throughout document as appropriate with editorial license

Proposed Response Response Status O

Cl 30 SC 30.5.1.1.2 P 35 L 16 # 62

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

800GBASE-VR8

With the introduction of WDM technology over MMF, the term "lane" is ambiguous when discussing MMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change

800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 50 m as specified in Clause 167

to

800GBASE-R PCS/PMA over 8 wavelengths distributed over 8 multi-mode fibres PMD with reach up to at least 50 m as specified in Clause 167

Makes changes throughout document as appropriate with editorial license

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 45 SC 45.2.1.7.5 P 40 L 14 # 1

Hajduczenia, Marek Charter Communications

Comment Type E Comment Status X

list uses "." instead of "," in edited list "100GBASE-KR1, 200GBASE-KR2, 400GBASE-KR4, 800GBASE-KR8"

SuggestedRemedy

Change "." to "," before newly added entry. Same on line 19. The same applies to Table 45-12

Proposed Response Response Status O

Cl 45 SC 45.2.1.135 P 45 L 29 # 16

Ran, Adeo Cisco

Comment Type TR Comment Status X

Registers 1.500 through 1.515 and 1.516 through 1.531 are mapped to variables that are used for transmitter equalization (local and remote) with AUI-C2C interfaces at 25 or 50 Gb/s per lane (defined in Annex 120B or 120D respectively). The transmit equalizer has 3 taps and specific sets of tap values (or ratios) with relatively coarse steps.

For 100 Gb/s per lane AUI-C2C, the transmitter equalization is controlled by a different set of variables, as defined in 120F.3.1.7 and 120F.3.2.6. The variables are different from and incompatible with those of Annex 120B/120D - the transmit equalizer has 5 taps and finer step size. The mapping of these variables to MDIO registers is also specified in these subclauses of 120F.

Therefore, Registers 1.500 through 1.531 should be made specific to the AUI-C2C at 25 or 50 Gb/s per lane.

This should have been done in 802.3ck, but if the subclauses of clause 45 are modified by this project, it should be done correctly.

If the suggested remedy is not within scope then, as an alternative, these subclauses of clause 45 should be deleted from 802.3df, since they are irrelevant for 800GAUI-n and thus out of scope.

SuggestedRemedy

In the title and body text of 45.2.1.135, change "50GAUI-n, 100GAUI-2, 200GAUI-n, and 400GAUI-n, and 800GAUI-n" to "50GAUI-n, 100GAUI-2, 200GAUI-8, 200GAUI-4, 400GAUI-16, and 400GAUI-8". Apply the same change in the title of Table 45-107.

Apply similarly in 45.2.1.136, 45.2.1.137 (including Table 45-108), and 45.2.1.138.

Proposed Response Response Status O

Cl 45 SC 45.2.1.135.1 P 45 L 48 # 25

Slavick, Jeff Broadcom

Comment Type TR Comment Status X

With the deletions the paragraph now reads a bit strangely and needs some word smithing. "The value of this bit indicates the value of the variable Request_flag in the lane 0 receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). This indicates whether the chip-to-chip device is issuing a request to change the remote transmitter equalization in the chip-to-chip lane 0 transmitter in the receive direction. If a lane 0 receiver in the receive direction is not present in the package, then the value returned for this bit should be zero."

SuggestedRemedy

Make it so the old paragraph is a full cross out text and replaced with the following paragraph:

"This bit indicates the state of the Request_flag variable of the lane 0 receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). When read as a one, the device is issuing a request to change the transmitter equalization of the transmitter driving lane 0 in the receive direction. If a lane 0 receiver in the receive direction is not present in the package, then the value returned for this bit should be zero."

Proposed Response Response Status O

Cl 45 SC 45.2.1.135.1 P 48 L 44 # 57

Dudek, Mike Marvell

Comment Type E Comment Status X

800GAUI-16 is not being defined in this amendment and therefore 120D and 120B are not used. There is no need to make changes to these sections?

SuggestedRemedy

Remove the changes to sections 45.2.1.135.1 to 45.2.1.135.7 and other equivalent changes. (If 800GAUI-16 is to be included in this amendment then bring in Annex 120D and make appropriate changes (including Title changes)

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 45 SC 45.2.1.135.2 P 46 L 3 # 26
Slavick, Jeff Broadcom
Comment Type TR Comment Status X
The value of these indicates the value is an odd phrase
SuggestedRemedy
Replace "The value of these bits indicates the value of the variable Requested_eq_c1 in the" with "These its indicate the state of the Requested_eq_c1 variable of the"
Proposed Response Response Status O

Cl 45 SC 45.2.1.135.2 P 46 L 3 # 28
Slavick, Jeff Broadcom
Comment Type TR Comment Status X
We're requesting the transmitter that is driving this given receiver to be changed. Not sure this text supports lane reversal between ends of the C2C link or not.
SuggestedRemedy
Replace "for the transmitter equalization in the chip-to-chip lane 0 transmitter in the receive direction." with "for the transmitter equalization of the transmitter driving the lane 0 receiver in the receive direction."
Make the same change in 45.2.1.135.3
Proposed Response Response Status O

Cl 45 SC 45.2.1.135.3 P 46 L 3 # 27
Slavick, Jeff Broadcom
Comment Type TR Comment Status X
The value of these indicates the value is an odd phrase
SuggestedRemedy
Replace "The value of these bits indicates the value of the variable Requested_eq_cm1 in the" with "These its indicate the state of the Requested_eq_cm1 variable of the"
Proposed Response Response Status O

Cl 45 SC 45.2.1.135.4 P 46 L 22 # 29
Slavick, Jeff Broadcom
Comment Type TR Comment Status X
We're providing the transmitter eq that is driving this receiver. Not sure this text supports lane reversal between ends of the C2C link or not.
SuggestedRemedy
Replace "being used in lane 0 of the transmitter in the receive direction" with "being used by the transmitter driving the lane 0 receiver in the receive direction."
Make the same chang in 45.2.1.135.5
Proposed Response Response Status O

Cl 45 SC 45.2.3.25 P 0 L 0 # 17
Slavick, Jeff Broadcom
Comment Type TR Comment Status X
Listing the number of PCS lanes for each PCS type in Clause 45 just adds duplication of information provided in the actual PCS clause. This text is likely to get stale or not updated as new rates or PCS configurations are added.
SuggestedRemedy
Remove the last paragraph that begins with Clause 82
Proposed Response Response Status O

Cl 45 SC 45.2.3.25 P 60 L 1 # 30
Slavick, Jeff Broadcom
Comment Type TR Comment Status X
The second paragraph is not necessary and just make for more work in the future. The first paragraph provides references to all the necessary registers for the maximal width PCS and states the unused lanes for thinner PCS's are to to return 0.
SuggestedRemedy
Remove the last paragraph of 45.2.3.25
Remove the last paragraph of 45.2.4.15
Remove the last paragraph of 45.2.5.15
Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 45 SC 45.2.3.25.1 P0 L0 # 18
 Slavick, Jeff Broadcom
 Comment Type TR Comment Status X
 Including the PCS rate when defining which variable is extraneous information. Just provide the clauses those given variable and the clause numbers.
 SuggestedRemedy
 Change the last sentence to read "This bit reflects the state of am_lock[0] or amps_lock[0] (see 82.2.19.2.2, 119.2.6.2.2, or 172.2.6.2.2)."
 Proposed Response Response Status O

Cl 45 SC 45.2.4.15 P0 L0 # 21
 Slavick, Jeff Broadcom
 Comment Type TR Comment Status X
 Listing the number of PCS lanes for each PCS type in Clause 45 just adds duplication of information provided in the actual PCS clause. This text is likely to get stale or not updated as new rates or PCS configurations are added.
 SuggestedRemedy
 Remove the last paragraph that begins with Clause 119
 Proposed Response Response Status O

Cl 45 SC 45.2.3.48a P0 L0 # 19
 Slavick, Jeff Broadcom
 Comment Type TR Comment Status X
 The clause 45 registers are containers for information the other clauses have. Whether a counter exists is functional Clause dependency not a Clause 45 dependency.
 SuggestedRemedy
 Remove the word "optional" in the second sentence
 Proposed Response Response Status O

Cl 45 SC 45.2.4.16a P0 L0 # 23
 Slavick, Jeff Broadcom
 Comment Type TR Comment Status X
 The clause 45 registers are containers for information the other clauses have. Whether a counter exists is functional Clause dependency not a Clause 45 dependency.
 SuggestedRemedy
 Remove the word "optional" in the second sentence
 Proposed Response Response Status O

Cl 45 SC 45.2.4.15 P0 L0 # 20
 Slavick, Jeff Broadcom
 Comment Type TR Comment Status X
 Including the PCS rate when defining which variable is extraneous information. Just provide the clauses those given variable and the clause numbers.
 SuggestedRemedy
 Change the last sentence to read "This bit reflects the state of amps_lock[0] (see 119.2.6.2.2, or 172.2.6.2.2)."
 Proposed Response Response Status O

Cl 45 SC 45.2.5.16a P0 L0 # 22
 Slavick, Jeff Broadcom
 Comment Type TR Comment Status X
 The clause 45 registers are containers for information the other clauses have. Whether a counter exists is functional Clause dependency not a Clause 45 dependency.
 SuggestedRemedy
 Remove the word "optional" in the second sentence
 Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 45 SC 45.2.5.16a P 81 L 49 # 11

Ewen, John Independent

Comment Type E Comment Status X

Beginning of sentence refers to registers 4.300 to 4.302; however, the subclause is defining registers 5.300 to 5.302

SuggestedRemedy

Change 4.300 - 4.302 to 5.300 - 5.302 respectively in first sentence of second sub-clause paragraph.

Proposed Response Response Status O

Cl 93A SC 93A.1 P 245 L 54 # 9

Lusted, Kent Intel Corporation

Comment Type TR Comment Status X

Table 93A "Physical Layer specifications that employ COM" in the base document, as amended by Std 802.3ck-2022, does not contain entries for the new 800GbE rates.

SuggestedRemedy

Update the table to include the following Physical Layer references and Parameter values:
800GAUI-8 C2C (Annex 120F) | Table 120F-8
800GBASE-CR8 (Clause 162) | Table 162-20
800GBASE-KR8 (Clause 163) | Table 163-11

Proposed Response Response Status O

Cl 116 SC 116.1.3 P 95 L 38 # 63

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

400GBASE-DR4
The term "lane" is ambiguous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change description to:
400 Gb/s PHY using 400GBASE-R encoding over 4 wavelengths distributed over 4 single-mode fibres, with reach up to at least 500 m (see Clause124)

Makes changes throughout document as appropriate with editorial license

Proposed Response Response Status O

Cl 116 SC 116.1.3 P 95 L 41 # 64

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E Comment Status X

400GBASE-DR4-2
The term "lane" is ambiguous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change description to:
400 Gb/s PHY using 400GBASE-R encoding over 4 wavelengths distributed over 4 single-mode fibres, with reach up to at least 2 km (see Clause124)

Makes changes throughout document as appropriate with editorial license

Proposed Response Response Status O

Cl 120 SC 120.5.11.2 P 98 L 13 # 15

Ran, Adeo Cisco

Comment Type T Comment Status X

"All test patterns specified in 120.5.11.2.1, 120.5.11.2.2, 120.5.11.2.3, and 120.5.11.2.4 are defined without precoding."
This should also include 120.5.11.2.a (PRBS9Q test pattern added in 802.3ck).

SuggestedRemedy

Add 120.5.11.2.a.

Proposed Response Response Status O

Cl 120G SC 120G.1 P 255 L 14 # 70

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type ER Comment Status X

The MII's, PCS Sublayers, and AUI's are all distinguished by data rates except the PMA sublayers

SuggestedRemedy

Distinguish PMA sublayers with reference to data rate

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 124 SC 124.1 P 99 L 13 # 66

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type ER Comment Status X

Given progress of 800G in IEEE P802.3dj with the creation of a single lambda solution at 10km, it is assumed that there will be a PCS related to coherent optical signaling that will be different than the PCS for other 800GBASE-R PMDs. Therefore, it is anticipated that there will be multiple PCSs and PMAs at 800G.

SuggestedRemedy

Modify "PCS" to be "800GBASE-R PCS" throughout document in all text and figures with editorial license.
 Modify "PMA" to be "800GBASE-R PMA" throughout document in all text and figures with editorial license.

Proposed Response Response Status O

Cl 124 SC 124.1 P 99 L 36 # 65

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type ER Comment Status X

400GbE will have multiple PCSs and PMAs with the introduction of 400GBASE-ZR. The PCS and PMA associated with 400GBASE-R PMDs has been renamed to 400GBASE-R PCA and 400GBASE-R PMA, respectively.

SuggestedRemedy

Change all instances in text and figures of PCS and PMA in the document that are relevant to 400GBASE-R PMDs to "400GBASE-R PCS" and "400GBASE-R PMA"

Proposed Response Response Status O

Cl 124 SC 124.5.4 P 106 L 10 # 84

Dawe, Piers Nvidia

Comment Type TR Comment Status X

The same modules will be capable of any of 100GBASE-DR, 400GBASE-DR4, 800GBASE-DR8, 100GBASE-FR1, 400GBASE-DR4-2, 800GBASE-DR8-2. Nominal nearly-compliance for a virtually unusable 0.2 dB on an unimportant spec would make the market more complicated and add procedural cost.

SuggestedRemedy

In the longer term, the average launch power (min) for 100GBASE-FR1 should be increased from -3.1 to -2.9 dBm to bring it in line with 100GBASE-DR/400GBASE-DR4. In the meantime: add a recommendation that the SIGNAL_DETECT power criterion for 800GBASE-DR8, 400GBASE-DR4-2 and 800GBASE-DR8-2 (which is: >= average receive power, each lane (min) in Table 124-7) should be -7.1 dBm.
 In practice, module implementers will set it lower than this anyway. See other comments for Tx and Rx specs, and for interoperability text.

Proposed Response Response Status O

Cl 124 SC 124.5.4 P 106 L 10 # 42

Schreiner, Stephan Rosenberger Hochfrequenztechnik GmbH & Co. KG

Comment Type E Comment Status X

Missing Bracket 3x("(" but only 2x")"

SuggestedRemedy

Insert Bracket at the End of Line 11

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 124 SC 124.7.1 P 108 L 23 # 85

Dawe, Piers Nvidia
 Comment Type TR Comment Status X

The minimum OMA for 400GBASE-DR4-2 and 800GBASE-DR8-2 is 0.7 dB higher than for 400GBASE-DR4/100GBASE-DR and 800GBASE-DR8, so setting the average launch power 0.2 dB lower is not helpful. Any transmitter with an extinction ratio lower than 9.8 dB, which is very high, will exceed the 400GBASE-DR4 limit anyway. Modules will be made multi-compliant for convenience in interoperability and breakout - let us document that.
 There is a minor benefit in improving the clearance between Rx min power and Tx off max power, which should be very wide to accomodate better-than-worst receivers and intentional signal detect hysteresis.

SuggestedRemedy

Change Average launch power, each lane (min) from -3.1 to -2.9 dBm
 Change Average receive power, each lane (min) from -7.1 to -6.9 dBm.
 See another commen for interoperability text.

Proposed Response Response Status O

Cl 124 SC 124.7.3 P 110 L 16 # 12

Stassar, Peter Huawei
 Comment Type TR Comment Status X

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, also the same (lower) allocation for MPI penalty can be assumed for DR4/DR8 and DR4-2/DR8-2.

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 3.5 dB.
 Furthermore change Tx min power from x to y and Rx sensitivity from a to b. A supporting presentation will be provided for the comment resolution meeting

Proposed Response Response Status O

Cl 124 SC 124.8.1 P 115 L 8 # 3

Nicholl, Shawn AMD
 Comment Type T Comment Status X

The Pattern column for the Wavelength row contains text "Square wave, 3, 4, 5, 6, or valid 400GBASE-R signal, or 800GBASER signal". Currently, it seems that the word valid is only applied to the 400GBASE-R signal, and not to the 800GBASE-R signal.

SuggestedRemedy

Propose "Square wave, 3, 4, 5, 6, or valid 400GBASE-R signal, or valid 800GBASER signal".

Similar comment for rows pertaining to "Side mode suppression ratio" parameter and to "Average optical power" parameter.

Proposed Response Response Status O

Cl 124 SC 124.11a.1 P 122 L 21 # 86

Dawe, Piers Nvidia
 Comment Type TR Comment Status X

We have a nuisance exception "provided that ... 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" that adds procedural cost for no technical benefit.

SuggestedRemedy

Having made the minimum 400GBASE-DR4-2 transmitter average power the same as for 400GBASE-DR4 (see another comment), 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."
 Similarly in 124.11a.2.

Proposed Response Response Status O

Cl 124 SC 124.12.2 P 123 L 42 # 87

Dawe, Piers Nvidia
 Comment Type E Comment Status X

Missing 124.12.3 Major capabilities/options

SuggestedRemedy

Add major options for the four PMD types

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 124 SC 124.12.4 P 124 L 11 # 45
 Brown, Matt Huawei
 Comment Type E Comment Status X
 In 124.12.4.3a/b/c the PICS item nicknames DR1 and DR2 are repeated. Also, the status variable is not defined and a different variable will need to be defined for each PMD type.
 SuggestedRemedy
 In 124.12.3 create status lable (like "**MD") for each PMD type.
 In 124.12.4.3, 124.12.4.3a, 124.12.4.3b, and 124.12.4.3c...
 - change the item labels such that they are unique
 - update the status with the new status variables
 Proposed Response Response Status O

Cl 124 SC 124.12.4. P 125 L 35 # 92
 Dawe, Piers Nvidia
 Comment Type E Comment Status X
 Need PICS for the 800G MDIs because the IEC connector reference is different to 400G, and there is an interface performance spec.
 SuggestedRemedy
 Per comment
 Proposed Response Response Status O

Cl 124 SC 124.12.4.1 P 124 L 3 # 88
 Dawe, Piers Nvidia
 Comment Type E Comment Status X
 F1 Compatible with 400GBASE-R PCS and PMA
 SuggestedRemedy
 Modify to include 800G
 Proposed Response Response Status O

Cl 124 SC 124.12.4.3a P 124 L 11 # 89
 Dawe, Piers Nvidia
 Comment Type E Comment Status X
 Presumably the "status" criterion in each of these four tables in 124.12.4.3X will be adjusted to the PMD type major options. Also, they could be combined as one table in one subclause: "400GBASE-DR4-2 transmitter meets specifications in" and so on.
 SuggestedRemedy
 Per comment
 Proposed Response Response Status O

Cl 124 SC 124.12.4.4 P 125 L 1 # 90
 Dawe, Piers Nvidia
 Comment Type E Comment Status X
 This subclause title "Optical measurement methods" represents the obsolete thinking that we specify testing, which we don't; we specify parameter limits and explain what the parameters are and how they might be determined by measurement. We started to move away from this in Clause 52, where this subclause was called "Optical measurement requirements", matching 52.9. But 124.8 is called "Definition of optical parameters and measurement methods"
 SuggestedRemedy
 Change "Optical measurement methods" to "Optical parameters and measurement methods".
 Proposed Response Response Status O

Cl 124 SC 124.12.4.4 P 125 L 21 # 91
 Dawe, Piers Nvidia
 Comment Type E Comment Status X
 The status of OM9 to OM12 should depend on the major option for PMD type
 SuggestedRemedy
 Per comment
 Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 162 SC 162.14.4.2 P 139 L 52 # 7

Lusted, Kent Intel Corporation

Comment Type TR Comment Status X

The PICS table for "PMD control function" the base document, as amended by Std 802.3ck-2022, has an incorrect reference to the relevant subclause for the training pattern entries due to the addition of the new item (h) in 3df 162.6.11 and the new sub-clause 162.8.11.1, including Table 162-10a.

SuggestedRemedy

Update 162.14.4.2 PMD Control Function PICS items as follows:

For Item 'PC2':

- update the subclause to be 162.8.11.1
- update value/comment to reference Table 162-10a

For Item 'PC3':

- update the subclause to be 162.8.11.1

Proposed Response Response Status O

Cl 163 SC 163.13.4.2 P 148 L 52 # 8

Lusted, Kent Intel Corporation

Comment Type TR Comment Status X

The PICS table for "PMD control function" the base document, as amended by Std 802.3ck-2022, has an incorrect reference to the relevant subclause for the training pattern entries due to the addition of the new item (h) in 3df 162.6.11 and the new sub-clause 162.8.11.1, including Table 162-10a.

SuggestedRemedy

Update 163.13.4.2 PMD Control Function PICS items as follows:

For Item 'PC2':

- update the subclause to be 162.8.11.1
- update value/comment to reference Table 162-10a

For Item 'PC3':

- update the subclause to be 162.8.11.1

Proposed Response Response Status O

Cl 167 SC 167.1.1 P 151 L 40 # 93

Dawe, Piers Nvidia

Comment Type E Comment Status X

Clause 173 and then Clause 172

SuggestedRemedy

Could be simplified to: Clause 173 then Clause 172

Proposed Response Response Status O

Cl 167 SC 167.8.1 P 159 L 9 # 94

Dawe, Piers Nvidia

Comment Type T Comment Status X

For the transmitter, we aren't talking about an optical signal but the pattern the transmitter is transmitting, which does not depend on V vs. S. It is not stated what "valid" means. One could assume it means the same as compliant, in which case it adds nothing. This table entry has become very long.

We can simplify:

3, 4, 5, 6, or valid 100GBASE-VR1, 200GBASE-VR2, 400GBASE-VR4, 800GBASE-VR8, 100GBASE-SR1, 200GBASE-SR2, 400GBASE-SR4, or 800GBASE-SR8 signal

to

3, 4, 5, 6, or 100GBASE-R1, 200GBASE-R2, 400GBASE-R4 or 800GBASE-R8 signal

Surprisingly, we have not used the term "800GBASE-R8" although in Section 6 we have 100GBASE-R10 and 100GBASE-R4. Such names will be useful for describing PMAs and AUIs, increasingly so as we work on 200G/lane in P802.3dj.

SuggestedRemedy

Change:

3, 4, 5, 6, or valid 100GBASE-VR1, 200GBASE-VR2, 400GBASE-VR4, 800GBASE-VR8, 100GBASE-SR1, 200GBASE-SR2, 400GBASE-SR4, or 800GBASE-SR8 signal

to

3, 4, 5, 6, or 100GBASE-R1, 200GBASE-R2, 400GBASE-R4, 800GBASE-R8, signal

Similarly for Average optical power.

For Stressed receiver sensitivity, just delete "valid". The SRS signal is on the edge of non-compliance anyway, by definition.

Define 100GBASE-R1, 200GBASE-R2, 400GBASE-R4, 800GBASE-R8 in the PMA clauses or introductory clauses 80, 116, 169.

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 167 SC 167.10.3.4 P 165 L 1 # 95

Dawe, Piers Nvidia
 Comment Type TR Comment Status X

A dual-row 24-position connector was recommended for 100GBASE-SR10, long ago. 400GBASE-SR8 has two options: a dual-row twelve-fiber interface (although different positions are used) and a single-row sixteen-fiber interface. Since then, the sixteen-fiber approach has become established. With the higher bandwidth for 800GBASE-SR8 vs. 400GBASE-SR8, the advantage of a single-row angled connector is more important.

SuggestedRemedy

Delete Option A, the dual-row 24-position non-angled connector. Update PICS accordingly.

Proposed Response Response Status O

Cl 167 SC 167.10.3.4 P 165 L 14 # 58

Dudek, Mike Marvell
 Comment Type T Comment Status X

The option B uses the angled interface which is depicted in Figure 167-10 not Figure 167-9

SuggestedRemedy

Change Figure 167-9 to 167-10

Proposed Response Response Status O

Cl 167 SC 167.11.4.6 P 168 L 35 # 59

Dudek, Mike Marvell
 Comment Type E Comment Status X

OC17 appears to be identical to OC16 except in the status column.

SuggestedRemedy

Label one of these with Option A and one with Option B

Proposed Response Response Status O

Cl 169 SC 169.4 P 177 L 27 # 82

Maki, Jeffery Juniper Networks
 Comment Type T Comment Status X

The sum of the sublayer delays of 92.16 ns for 800GBASE-R PMA (up to four PMA stages) and 20.48 ns for 800GBASE-VR8/SR8/DR8/DR8-2 PMD is 112.64 ns, which is less than the observed delay of two PMA stages and the PMD. The concern is that these sublayers delays are specified too small in value to be feasible. Excessive delays of about 50% are seen for optical modules (two PMA stages + PMD).

SuggestedRemedy

Increase Delay values for PMA and PMD to align with prevalent implementation.

Proposed Response Response Status O

Cl 169 SC 169.4 P 177 L 40 # 10

Laubach, Mark IEEE Member / Self
 Comment Type E Comment Status X

"(See 31B.2 for the definition of pause_quanta.)". I see this reference is many places in the standard. Clause 31B.2 defines "pause_time" only and that "The pause_time is measured in units of pause_quanta,..." "pause_quanta" is defined somewhere else, not 31B.2.

SuggestedRemedy

Cross-reference to where pause_quanta is actually defined?

Proposed Response Response Status O

Cl 169 SC 169.5 P 180 L 2 # 81

Li, Mike Intel
 Comment Type TR Comment Status X

The skew numbers in this table 169-5 no longer represent the technology in reality, resulting in lagging skew spec for 800GE which needs to be changed.

SuggestedRemedy

See slide 10 of li_3df_01_0423 (presentation made at the April 26, 2023 "802.3df Architecture and logic ad hoc". Also inserted.

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 169 SC 169.5 P 180 L 9 # 96

Dawe, Piers Nvidia
 Comment Type **TR** Comment Status **X**

As discussed, the Skew and Skew Variation limits were based on a digital clock rate that is slow by modern standards, and CWDM over 40 km which is not going to happen for 800G. Also 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.

SuggestedRemedy

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

Proposed Response Response Status

Cl 169 SC 169.5 P 180 L 31 # 97

Dawe, Piers Nvidia
 Comment Type **E** Comment Status **X**

Table layout

SuggestedRemedy

Adjust column widths

Proposed Response Response Status

Cl 170 SC 170.4.4.2 P 187 L 3 # 98

Dawe, Piers Nvidia
 Comment Type **E** Comment Status **X**

Broken variable name

SuggestedRemedy

Make second column slightly wider

Proposed Response Response Status

Cl 171 SC 171.1 P 189 L 11 # 44

Brown, Matt Huawei
 Comment Type **E** Comment Status **X**

Description of Extender implies it has only one 800GAUI-n, but it can also have two. Also, by definition 800GAUI-n is a physical instantiations so a bit superfluous.

SuggestedRemedy

Change "The 800GMII Extender is composed of a DTE 800GXS at the RS end, and a PHY 800GXS at the PHY end with a physical instantiation of 800GAUI-n between two adjacent PMA sublayers."

To:

"The 800GMII Extender is composed of a DTE 800GXS at the RS end, and a PHY 800GXS at the PHY end with one or two 800GAUI-n between."
 Align definition in 1.4.184j.

Proposed Response Response Status

Cl 171 SC 171.1 P 190 L 22 # 67

D'Ambrosia, John Futurewei, US Subsidiary of Huawei
 Comment Type **TR** Comment Status **X**

The definition of the OSI Physical Layer is incorrect as shown in Fig 171-1. The medium is not part of the Physical Layer

SuggestedRemedy

modify Fig 171-1 to show the Physical Layer bottom border at the bottom of the MDI

Proposed Response Response Status

Cl 171 SC 171.2 P 190 L 46 # 99

Dawe, Piers Nvidia
 Comment Type **TR** Comment Status **X**

I don't see any the modification to the FEC degrade signaling in 171.5. It might be different to the 400GBASE-R PCS, but here we are comparing it to the 800GBASE-R PCS. I thought we sorted this out last time.

SuggestedRemedy

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

Proposed Response Response Status

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 171 SC 171.3 P 192 L 15 # 4

Nicholl, Shawn

AMD

Comment Type TR Comment Status X

Figure 171-2 "Functional block diagram for the PHY 800GXS" shows "Flow <n> Rx" labels in the transmit path of the PHY 800GXS and likewise shows "Flow <n> Tx" labels in the receive path. This introduces confusion.

SuggestedRemedy

Propose one of the following solutions:

* Update the diagram. In the transmit path of the PHY 800GXS (i.e. direction from PMA to 800GMII), use labels "Flow 0 Tx" and "Flow 1 Tx". In the receive path of the PHY 800GXS (i.e. direction from 800GMII to PMA), use labels "Flow 0 Rx" and "Flow 1 Rx". The problem with this proposal is that it contradicts the PICS tables (which for example, indicate that the "171.8.4.1 Transmit function" of the 800GXS includes a 64B/66B to 256B/257B transcoder).

* Update the diagram. Remove the Tx/Rx in the dotted area. Replace "Flow 0 Tx" with "Flow 0". Replace "Flow 1 Tx" with "Flow 1". Replace "Flow 0 Rx" with "Flow 0". Replace "Flow 1 Rx" with "Flow 1". If this solution is chosen, propose to apply similar solution to Figure 172-2 "Functional block diagram".

* Remove the diagram. Since the diagram is effectively an inverted replica of Figure 172-2 "Functional block diagram", rely on the text (in the same manner that 118.1.2 "200GXS/400GXS Sublayer" was able to rely on text without a new diagram).

Proposed Response

Response Status O

Cl 171 SC 171.4 P 193 L 42 # 13

Ran, Adeo

Cisco

Comment Type T Comment Status X

The standard should be explicit about what happens in a PHY connected to an 800GMII Extender when there is no valid input signal.

The precedence is set in 802.3cw: D2.1 states (in 155.2.6.7.2) that the 400GBASE-ZR PCS sends local fault ordered sets to the 400GMII when there is no signal; this means the PHY XS transmits these local fault over the 400GAUI-n toward the DTE XS. There is no provision for "shutting down" the PHY XS output, so the 400GAUI-n in an Extender is never silent.

The behavior of the 800GMII extender should be the same as that of the 400GMII extender as described above.

Note that this behavior is different from existing optical modules that are connected with any AUI-C2M to a PCS (as part of the PHY, not an extender), where it is common to squelch the module electrical output (aka disable the AUI's transmitter) when there is no optical input (PMD:IS_SIGNAL.indication is not_ok); that is indicated to by PCS via PMA:IS_SIGNAL.indication on its adjacent PMA. That would not be compliant behavior when the AUI is within a 800GMII Extender.

The different behavior required from Extender modules may not be obvious and should be mentioned.

Note: if the task force wants to allow squelching the Extender's AUI, it may require more significant changes; as an alternative, an editor's note can be added to capture that intent until a detailed proposal is presented (such as "Editor's note: the behavior of the Extender when there is no input signal from the PHY is to be determined").

SuggestedRemedy

Add the following paragraph at the end of 171.4:

NOTE-link fault signaling generated by the PHY (see 170.3 and 81.3.4) is transmitted to the RS through the 800GMII Extender. Therefore, the electrical interface used within the 800GMII Extender sends valid PHY 800GXS data regardless of the link state of the PHY below the 800GMII.

Proposed Response

Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 171 SC 171.6 P 194 L 26 # 43
 Brown, Matt Huawei
 Comment Type E Comment Status X
 The PMA above the PMD may not be an 800GBASE-R PMA (per Clause 173) and the PMA may not have 8 lanes.
 SuggestedRemedy
 For the PMA immediately above the PMD change "PMA (32:8)" to "PMA".
 Proposed Response Response Status O

Cl 171 SC 171.6 P 194 L 35 # 48
 Brown, Matt Huawei
 Comment Type E Comment Status X
 No such thing as "800 Gb/s Extender Sublayer". See 171.1.
 SuggestedRemedy
 Change "800 Gb/s EXTENDER SUBLAYER" to "800GMII EXTENDER SUBLAYER"
 Proposed Response Response Status O

Cl 171 SC 171.8.4.3 P 201 L 8 # 32
 Huber, Tom Nokia
 Comment Type E Comment Status X
 It is not clear why the coding rules PICS items jump from C6 to C9; the set of items is the same as what is in clause 118, which numbers them sequentially.
 SuggestedRemedy
 Change the numbering of C9 through C11 to C7 through C9, respectively.
 Proposed Response Response Status O

Cl 172 SC 172.1.5 P 204 L 14 # 5
 Nicholl, Shawn AMD
 Comment Type TR Comment Status X
 Figure 172.1.5 "Functional block diagram" contains a functional diagram of the 800G PCS. Currently, the diagram shows "Flow <n> Tx" labels in the transmit path and likewise shows "Flow <n> Rx" labels in the receive path. When/If this diagram is re-used for 800GXS it may cause confusion.
 SuggestedRemedy
 Propose to update the diagram. Remove the Tx/Rx in the dotted area. Replace "Flow 0 Tx" with "Flow 0". Replace "Flow 1 Tx" with "Flow 1". Replace "Flow 0 Rx" with "Flow 0". Replace "Flow 1 Rx" with "Flow 1". See similar comment against Figure 171-2 "Functional block diagram for the PHY 800GXS" in sub-clause 171.3 and apply consistent solution.
 Proposed Response Response Status O

Cl 172 SC 172.2 P 205 L 1 # 100
 Dawe, Piers Nvidia
 Comment Type ER Comment Status X
 This title "Physical Coding Sublayer (PCS)" is as good as the same as the main clause title "Physical Coding Sublayer (PCS), type 800GBASE-R" which can't be right.
 SuggestedRemedy
 Change this to "Functions within the PCS", change 172.2.1 to "Overview of functions within the PCS", "Functions and processes within the PCS" or similar.
 Proposed Response Response Status O

Cl 172 SC 172.2.1 P 205 L 19 # 33
 Huber, Tom Nokia
 Comment Type E Comment Status X
 The word block is overloaded in this paragraph, which discusses 66-, 257-, and 5140-bit blocks, and also uses 'block' to refer to the processes (called functional blocks) in Figure 172-2.
 SuggestedRemedy
 In the second sentence, change "encode and rate matching block" to "encode and rate matching functional block" or "encode and rate matching process".
 Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 172 SC 172.2.1 P 205 L 33 # 34

Huber, Tom Nokia
 Comment Type E Comment Status X

The sentences describing AM lock, reordering, deskewing could be written more clearly.

SuggestedRemedy

Change
 It attains alignment marker lock based on the common marker (CM) portion that is periodically transmitted on every PCS lane. After alignment markers are found on all PCS lanes, the individual PCS lanes are identified using the unique marker portion (UM) and then reordered, reordered and deskewed, and the align_status flag is set..
 to
 It attains alignment marker lock based on the common marker (CM) portion of the alignment markers that are periodically transmitted on every PCS lane and identifies individual PCS lanes using the unique marker portion (UM) or the alignment makers. The PCS lanes are then reordered and deskewed, and the align_status flag is set..

Proposed Response Response Status O

Cl 172 SC 172.2.3 P 206 L 1 # 101

Dawe, Piers Nvidia
 Comment Type E Comment Status X

Same topic, very short subclauses

SuggestedRemedy

Make 172.2.3, 172.2.2.1, or remove this subheading and change the title of 172.2.2 to " 66-bit blocks and the 64B/66B code" or similar.

Proposed Response Response Status O

Cl 172 SC 172.2.4.1.1 P 206 L 29 # 35

Huber, Tom Nokia
 Comment Type E Comment Status X

Per the style guide, a clause should not have a single subclause. It is however useful to have some separation between the general description and this new stateless encoder.

SuggestedRemedy

One option would be to make 172.2.4.1.1 a level-4 heading. The other would be insert a level 5 heading immediately after 172.2.4.1 with an innocuous title like 'Process description' and renumber the existing 172.2.4.1.1 to 172.2.4.1.2. In either case, the cross-reference at line 15 would also need to be updated.

Proposed Response Response Status O

Cl 172 SC 172.2.4.1.1 P 206 L 44 # 103

Dawe, Piers Nvidia
 Comment Type T Comment Status X

If it's OK to combine criteria in the second column it's OK in the third column

SuggestedRemedy

Combine rows 3 and 4, combine rows 5 and 6

Proposed Response Response Status O

Cl 172 SC 172.2.4.4 P 207 L 20 # 106

Dawe, Piers Nvidia
 Comment Type E Comment Status X

Instead of 0 to 31, t might be better to number the lanes 0.0 to 0.15, 1.0 to 1.15

SuggestedRemedy

Per comment

Proposed Response Response Status O

Cl 172 SC 172.2.4.4 P 207 L 27 # 105

Dawe, Piers Nvidia
 Comment Type ER Comment Status X

Please don't make work for your readers

SuggestedRemedy

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.

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 172 SC 172.2.4.4 P 208 L 7 # 104

Dawe, Piers Nvidia
 Comment Type ER Comment Status X

This table is very hard to use. The next is split over two pages

SuggestedRemedy

Make the headings line up with the ~columns, e.g. by inserting spaces.
 Combine the two tables, adjusting the text on the previous page. The PCS lane numbers are unique, but sub-heading rows or another column indication flow 0 and flow 1 can be used.
 Use the orphan rows property to ensure the table is not split.

Proposed Response Response Status O

Cl 172 SC 172.2.4.9 P 210 L 48 # 36

Huber, Tom Nokia
 Comment Type T Comment Status X

It's more clear to say the test pattern is the result of the MII being a continuous stream of Idle characters (which the PCS will then turn into blocks, etc.).

SuggestedRemedy

Change the last sentence of the first paragraph from
 The scrambled idle test pattern is the output of the PCS when the input to the PCS at the 800GMII is a control block with all idle characters.
 To
 The scrambled idle test pattern is the output of the PCS when the input to the PCS at the 800GMII is a contiuous stream of idle characters.

Proposed Response Response Status O

Cl 172 SC 172.2.5.8.1 P 212 L 10 # 37

Huber, Tom Nokia
 Comment Type E Comment Status X

Per the style guide, a clause should not have a single subclause. It is however useful to have some separation between the general description and this new stateless encoder.

SuggestedRemedy

One option would be to make 172.2.5.8.1 a level-4 heading. The other would be insert a level 5 heading immediately after 172.2.5.8 with an innocuous title like 'Process description' and renumber the existing 172.2.5.8.1 to 172.2.5.8.2. In either case, the cross-reference at line 3 would also need to be updated.

Proposed Response Response Status O

Cl 172 SC 172.2.6.3 P 214 L 15 # 38

Huber, Tom Nokia
 Comment Type E Comment Status X

It appears that the only difference between figure 119-3 and figures 172-5 and 172-6 is that figure 119-3 has been split into two parts because the part shown in figure 172-6 is done separately for each flow. It would be helpful if that was more clear in the bullet points that describe the exceptions.

SuggestedRemedy

Change:
 — The PCS synchronization process is depicted in Figure 172–5 and Figure 172–6, instead of in Figure 119–13.
 — The monitor for three consecutive uncorrectable FEC codewords (see Figure 172–6) is done independently within each flow.
 To:
 — The PCS synchronization process is depicted in Figure 172–5 and Figure 172–6, which are derived by splitting Figure 119–13 into two parts to better illustrate that the monitor for three consecutive uncorrectable FEC codewords (see Figure 172–6) is done independently within each flow.

Proposed Response Response Status O

Cl 172 SC 172.7.4.3 P 222 L 21 # 39

Huber, Tom Nokia
 Comment Type E Comment Status X

It appears that Items C7-C9 are omitted here because in clause 119 they are used for EEE-related rules, which are not relevant to 800G - but the remaining items should have been renumbered.

SuggestedRemedy

Change the numbering of C9 through C11 to C7 through C9, respectively.

Proposed Response Response Status O

Cl 173 SC 173.1 P 226 L 26 # 49

Brown, Matt Huawei
 Comment Type E Comment Status X

No such thing as "800 Gb/s Extender Sublayer". See 171.1.

SuggestedRemedy

Change "800 Gb/s EXTENDER SUBLAYER" to "800GMII EXTENDER SUBLAYER"

Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 173 SC 173.3 P 227 L 26 # 60
 Maguire, Valerie Copperopolis
 Comment Type E Comment Status X
 Use a non-breaking space between figures and abbreviations
 SuggestedRemedy
 Use a non-breaking space between "53.125" and "GBd".
 Proposed Response Response Status O

Cl 173 SC 173.4 P 229 L 7 # 107
 Dawe, Piers Nvidia
 Comment Type T Comment Status X
 The grouping into two flows of 16 lanes each is significant to the PMA (although the lane numbers are not).
 SuggestedRemedy
 Instead of one group of 32 input lanes, show two groups of 16, consistent with the PCS figures. Similarly in Figure 173-4.
 Proposed Response Response Status O

Cl 173 SC 173.4 P 231 L 231 # 108
 Dawe, Piers Nvidia
 Comment Type T Comment Status X
 An IC implementing a 8:8 PMA is likely to have signal detect ability in both directions.
 SuggestedRemedy
 Show an optional SIL on the Tx side, that looks at PMA:IS_UNITDATA_0:7.request (there may be no PMA:IS_SIGNAL... primitive). Add MDIO register.
 Proposed Response Response Status O

Cl 173 SC 173.4.2.1 P 232 L 7 # 6
 Nicholl, Shawn AMD
 Comment Type TR Comment Status X
 In 173.4 "Functions within the PMA" the text references the undefined term "restricted bit-multiplexing" and says to "see 173.4.2.1". However, the word "restricted" does not appear in 173.4.2.1 "32:8 PMA bit-level multiplexing".
 SuggestedRemedy
 Propose to update the text in 173.4.2.1 "32:8 PMA bit-level multiplexing". Replace "The multiplexing function has an additional constraint ..." with "This restricted bit-multiplexing function has an additional constraint ..."

Similarly, propose to update the text in 173.4.2.2 "8:32 PMA bit-level multiplexing". Replace "The multiplexing function has an additional constraint ..." with "This restricted bit-multiplexing function has an additional constraint ..."
 Likewise, propose to update the text in 173.4.2.3 "8:8 PMA bit-level multiplexing". Replace "The 4 PCSs received on an input lane shall be mapped ..." with "This restricted bit-multiplexing function has an additional constraint that the 4 PCSs received on an input lane shall be mapped ..."
 Proposed Response Response Status O

Cl 173 SC 173.4.2.1 P 232 L 15 # 2
 Nicholl, Shawn AMD
 Comment Type T Comment Status X
 In 173.4.2.1 "32:8 PMA bit-level multiplexing" the word "contain" is used which is inconsistent with referenced 120.5.2 "Bit-level multiplexing".
 SuggestedRemedy
 Propose to replace "contain" with "carries", so the sentence reads "... each of the 8 output lanes carries two PCSs from ...". Using the word "carries" emphasizes that each lane is carrying a stream of bits.
 Propose to make the same change in 173.4.2.2 "8:32 PMA bit-level multiplexing".
 Proposed Response Response Status O

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 173 SC 173.4.2.3 P 233 L 7 # 83

Nicholl, Gary Cisco Systems

Comment Type T Comment Status X

The multiplexing rules in this section (along with the mutiplexing rules in 173.4.2.1 and 173.4.2.1) were updated based on comment #27 against D1.1 and supporting presentation "https://www.ieee802.org/3/df/public/23_01/0130/ran_3df_01b_230130.pdf".

As captured in slide 3 of ran_3df_01b_230130 the motivation of the proposed change was to avoid the situation "where one of two flows always gets the LSB of the PAM4 symbols"

The changes to the mutiplexing rules for PMA 32:8 (173.4.2.1) and PMA 8:32 (173.4.2.2) achieve this goal.

However the change to the mutiplexing rules for the PMA 8:8 (173.4.3) goes one step futher than the changes to the PMA 32:8 and PMA 8:32. This additional restriction is unnecessary (as the situation this step is trying to avoid can be caused by both the PMA 32:8 and PMA 8:32 anyway), and it any may make some existing 100G PAM4 retimer implementations non-compliant.

The additional step is the requirement that "the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane" This means the PAM4 output must be MSB/LSB aligned to the PAM4 input. It is not clear that this would always be the case, and is something that is not required for the 400GbE generation of PAM4 retimer chips. It is also not fully consistent with the description of the PAM4 Encoding described in 173.4.7.1 (which essentially references the PAM4 encoding rules from Clause 120, which do not require PAM4 outputs to be MSB/LSB aligned to PAM4 inputs).

This step is not required in order to meant the intent captured in slide 3 of ran_3df_01b_230130.pdf.

If the PAM4 input is decoded to a serial bit stream, then in order to meet the intent of ran_3df_01b_230130.pdf, the only rquirement is that the bit stream be sent in the same order (no rearrangement of bits) to the PAM4 output encoder. The output encoder just has to take two bits at a time and encode into a PAM4 symbol (consistent with the description in 173.4.7.1) . There is no need for the PAM4 encoder to be MSB/LSB aligned to the bit stream coming from the PAM4 receiver.

It should also be noted that this section only describes the bit level mutipexing functions of a serial bit stream (in keeping with Figure 173-5), and the PAM4 decoding and encoding rules are described in a different section (173.4.7.1).

SuggestedRemedy

Change from:

"The 4 PCSLs received on an input lane shall be mapped to an output lane 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 173.4.7.1)."

to:

"The 4 PCSLs received on an input lane shall be mapped to an output lane such that the order of PCSLs is maintained from input lane to output lane, except for possible swapping of each bit pair (see 173.4.7.1)."

Proposed Response Response Status

Cl 173 SC 173.4.3.1 P 233 L 26 # 102

Dawe, Piers Nvidia

Comment Type T Comment Status X

On further investigation: this must be output not generate. If there are multiple PMAs they share this limit, as is made clear for the receive direction.

SuggestedRemedy

Per comment

Proposed Response Response Status

IEEE P802.3df D2.0 Initial Working Group ballot comments

Cl 173 SC 173.4.8.3 P 236 L 19 # 14

Ran, Adee Cisco
 Comment Type T Comment Status X
 "Otherwise the SIL reports the signal status as FAIL"

In the case of 8:8 PMA, this FAIL status typically indicates that data is not being received on all 8 input lanes (inst:IS_UNITDATA_0:7.indication). When this happens, the data on the output lanes (PMA:IS_UNITDATA_0:7.indication) cannot be determined from the standard. Apparently it is unspecified, but it isn't stated explicitly.

In optical modules (a common implementation of PMAs similar to this one), the typical behavior is to turn off the electrical output of the AUI-C2M; but this functionality is not specified in the standard, and there is no specification of "output disabled" in 120G.3.2. It can be argued that this common behavior is non-compliant.

With no specification of behavior in this condition, the signal status is not conveyed to the PMA client (host ASIC) in a specified and consistent manner. Moreover, SerDes designers cannot assume what signal appears on the AUI when there is no input, and that is a repeating source of confusion, often leading to bad design or unnecessary over-design.

We need to specify the AUI behavior when signal status is FAIL such that the PMA client can detect this situation. Based on existing module behavior, it is suggested to state that a PMA with a physically instantiated interface disables the transmitters on all lanes of that interface when signal status is FAIL on the other interface, for some minimum time. The PMA client can infer the status by detecting that its input signal corresponds to a disabled transmitter. This requires adding the missing "output disabled" mode in the module output characteristics (120G.3.2).

A possible alternative is to allow the PMA to transmit the PRBS31Q test pattern (120.5.11.2.2), if implemented, instead of disabling the transmitter. The PMA client can then infer the link status by detecting that its input corresponds to a PRBS31Q test pattern. This would not require adding "output disabled" mode, but it is likely not the existing behavior, and would be more disruptive.

Note that this isn't just an 802.3df problem (ambiguity of the module output is a long-standing issue), but since we are defining a new PMA it is a good opportunity to close this gap.

SuggestedRemedy

Add the following paragraph at the end of 173.4.8.3:

"When the signal status is FAIL, an 8:8 PMA shall disable the output on all lanes of its physically instantiated service interface for a minimum time of 50 ms."

Add 120G.3.2 to the draft. Change the first sentence from

"The module output shall meet the specifications given in Table 120G-3"
 to

"When the module output is enabled, it shall meet the specifications given in Table

120G-3. When the module output is disabled, the Differential peak-to-peak output voltage shall be less than 35 mV."<paragraph break>

Change the title of Table 120G-3 to "Module output characteristics in enabled state at TP4"

Proposed Response Response Status O

Cl 173 SC 173.4.11 P 236 L 31 # 69

D'Ambrosia, John Futurewei, US Subsidiary of Huawei
 Comment Type ER Comment Status X

800GAUI-8 has been described elsewhere as an electrical interface in 163.1, but the definition uses "physical instantiation" - use consistent language

SuggestedRemedy

Change description of 800GAUI elsewhere from electrical interface to physical instantiation

Proposed Response Response Status O

Cl 173 SC 173.6.5 P 241 L 15 # 40

Huber, Tom Nokia
 Comment Type E Comment Status X

The status column should be reformatted so the items are not spilling over lines

SuggestedRemedy

Reformat so that the items are not split across lines

Proposed Response Response Status O

Cl 173A SC 173A P 276 L 28 # 46

Brown, Matt Huawei
 Comment Type E Comment Status X

No such thing as "800 Gb/s Extender Sublayer". See 171.1.

SuggestedRemedy

Change "800 Gb/s EXTENDER SUBLAYER" to "800GMII EXTENDER SUBLAYER"
 Also in Figure 173-4, page 277, line 31.

Proposed Response Response Status O