

IEEE P802.3df D1.0 1st Task Force review comments

Cl 45 SC 45.2.4 P 47 L 4 # 1

Marris, Arthur Cadence Design Systems

Comment Type T Comment Status A (bucket1)

"45.2.4 PHY XS registers" and "45.2.5 DTE XS registers" subsections need to be brought into the 802.3df draft and modifications made to increase the number of service interface lanes specified from 20 to 32

SuggestedRemedy

Update "Table 45–314—PHY XS registers" and "Table 45–339—DTE XS registers" and relevant sunclauses to address this. This will include an extra "XS alignment status 5" register at location 54, adding extra "XS lane mapping" registers above 415, adding extra "FEC symbol error counter" registers above 631, and add bit 4.801.6 for "Local degraded SER received"

Response Response Status C

ACCEPT.

Cl 172 SC 172.2.5.5 P 168 L 9 # 2

Ran, Adeo Cisco

Comment Type TR Comment Status A (bucket1)

"The alignment marker removal is identical to that of the 400GBASE-R PCS in 119.2.5.5." - but there are 32 AMs, so it can't be identical.

SuggestedRemedy

Make the necessary changes to the text (add exceptions or "for each flow").

Response Response Status C

ACCEPT IN PRINCIPLE.

Change from

"The alignment marker removal is identical to that of the 400GBASE-R PCS in 119.2.5.5"

to
"The alignment marker removal in each flow is identical to that of the 400GBASE-R PCS in 119.2.5.5"

Cl 172 SC 172.2.6.3 P 170 L 21 # 3

Ran, Adeo Cisco

Comment Type E Comment Status A (bucket1)

Numbers above 10 should not be spelled out.

SuggestedRemedy

change "thirty two" to "32".

Response Response Status C

ACCEPT.

Cl 172 SC 172.3.5 P 173 L 31 # 4

Ran, Adeo Cisco

Comment Type ER Comment Status A fec counters

FEC_cw_counter is defined as optional in 161.6.21. Assuming it is optional here too, it should be stated, as in clause 161.

Otherwise, state that it is not optional for this PCS (but I assume it's not the case).

Similarly for 172.3.6 FEC_codeword_error_bin_i.

SuggestedRemedy

Add "(optional)" to the subclause title in 172.3.5 and 172.3.6.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #189.

Cl 173 SC 173.4 P 180 L 6 # 5

Ran, Adeo Cisco

Comment Type E Comment Status A (bucket1)

The concept of restricted bit multiplexing appears in this subclause for the first time. It may be helpful for readers to have a cross reference to the definition of this restriction.

SuggestedRemedy

Add the following paragraphs after each of the three bulleted lists on page 180, respectively:

"Bit multiplexing restrictions for the 32:8 PMA are specified in 173.4.2.1."

"Bit multiplexing restrictions for the 8:32 PMA are specified in 173.4.2.2."

"Bit multiplexing restrictions for the 8:8 PMA are specified in 173.4.2.3."

Response Response Status C

ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

CI 173 SC 173.4.2.1 P 184 L 10 # 6

Ran, Adeo Cisco

Comment Type TR Comment Status R PCSL interleaving (CC)

The restriction for the 32:8 multiplexing is intended to improve the FEC performance with correlated errors. The analysis was done with an AB/CD muxing scheme where one UI has bits from codewords A and B (flow 0) and the following UI has bits from C and D (flow 1). This way, combined with the checkerboard scheme, spreads the errors in a burst across the four codewords with equal probabilities.

The restriction as written does not preclude a different muxing, AC/BD, where one UI has bits from A and C and the following UI has bits from B and D. For example, muxing bits from lanes 0 and 16 as MSB+LSB in one UI and bits from lanes 1 and 17 as MSB+LSB in the next UI.

Since the checkerboard pattern swaps codewords A/B on each pair of lanes in flow 0, and swaps codewords C/D on each pair of lanes in flow 1, this would result in always taking the MSB from either codeword A or B, and the LSB from either codeword C or D. Since the BER for the LSB is twice that of the MSB, this would make flow 1 have an increased BER: it would get 2/3 of the errors (33% higher BER than with the AB/CD muxing).

If this muxing is performed, the result would be an increased FLR (by 1-2 orders of magnitude) compared to 400GBASE-R, just due to sub-optimal muxing - regardless of whether errors are correlated or not!

This degradation can be prevented by adding a restriction that two bits from each flow create one PAM4 symbol.

SuggestedRemedy

Change the second item of the first list in 173.4.2.1 from "The multiplexing function has an additional constraint that each of the 8 output lanes contain two unique PCSLs from PMA client lanes i = 0 to 15 and two unique PCSLs from PMA client lanes i = 16 to 31" to "The multiplexing function has an additional constraint that each of the 8 output lanes contain two unique PCSLs from PMA client lanes i = 0 to 15 encoded as one PAM4 symbol, and two unique PCSLs from PMA client lanes i = 16 to 31 encoded as the subsequent PAM4 symbol (see 173.4.7)."

Make a similar change in the second item of the second list in 173.4.2.2 (which has "service interface lanes" instead of "PMA client lanes").

Also, change the second item of the list in 173.4.2.3 from "The 4 PCSLs received on any input lane shall be mapped together to an output lane. The order of PCSLs from an input lane does not have to be maintained on the output lane." to "The 4 PCSLs received on any input lane shall be mapped together to an output lane, maintaining the bit pairs encoded on each PAM4 symbol. Other than that, the order of PCSLs from an input lane does not have to be maintained on the output lane."

Response Response Status C

REJECT.

The current text and constrained PCSL multiplexing requirement is consistent with the adopted baseline (see slides 17&18 in https://www.ieee802.org/3/df/public/22_10/22_1004/shrikhande_3df_01a_221004.pdf).

The following presentation was reviewed by the task force: https://www.ieee802.org/3/df/public/22_12/ran_3df_01a_2212.pdf

Based on straw polls #1 and #2 there is no consensus to make the proposed changes at this time. The commenter is encouraged to refine the proposal and build consensus.

Straw poll #1 (direction)
I would support the changes proposed on slides 10 in ran_3df_01a_2212.
Yes: 19
No: 6
Need more information: 27

Straw poll #2 (direction)
I would support the changes proposed on slides 11 in ran_3df_01a_2212.
Yes: 13
No: 6
Need more information: 33

CI 169 SC 169.6 P 138 L 49 # 7

Ran, Adeo Cisco

Comment Type TR Comment Status A fec degrade (bucket1)

FEC degrade function is defined as optional in 116.6. Assuming it is optional here too, it should be stated, as in clause 116.

SuggestedRemedy

Add "(optional)" to the subclause title in 169.6.

Response Response Status C

ACCEPT IN PRINCIPLE.

The FEC degrade is indeed intended to be optional for 800 GbE as well. It is not sufficient or necessary to put the word "optional" in the title, rather the word should be included in the text.

Change: "FEC degrade functionality is identical"
To: "Optional FEC degrade functionality is identical"

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.4.4 P 164 L 45 # 8

Ran, Adeo Cisco
 Comment Type TR Comment Status R alignment

"Alignment marker encoding values for flow 1 are specified in Table 172–2 and the variable x in 119.2.4.4.2 takes the values of PCS lane number minus 16"

In 119.2.4.4.2, x is used as part of the variable am_x. We have 32 distinct alignment markers, for lanes 0 through 31, so assigning x to "lane number minus 16" would result in am_0 through am_15 assigned twice, and am_16 through am_31 not assigned at all.

Instead, we should specify that for flow 1, AM are constructed per 119.2.4.4.2 but with x taking values from 16 to 31, and the variable j used in the mapping procedure takes values from 8 to 16 (instead of 0 to 7).

This difference may be listed as another exception, but it seems that it makes it worthwhile to have a new subclause for creating the 32 AMs.

SuggestedRemedy

Replace the reference to 119.2.4.4.2 with a full specification of AM creation and insertion, based on the content (text and equations) of 119.2.4.4.2, but with AMs for lanes 16 to 31 constructed as in the comment.

Response Response Status C

REJECT.

Each Flow is a unique "instance" of the 119.4.4.2 so the fact that there are 2 copies of variable "am_#", one in Flow0 and another in Flow1 that have different values is how it's intended to be specified.

Cl 172 SC 172.2.4.4 P 164 L 51 # 9

Ran, Adeo Cisco
 Comment Type TR Comment Status A AM sync

In the baseline proposal https://www.ieee802.org/3/df/public/22_10/22_1004/shrikhande_3df_01a_221004.pdf, slide 10, it is written that "AM insertion is aligned across the two flows".

I do not see that requirement in clause 172. The text in 172.2.4.4 does not preclude inserting AM blocks independently in each flow.

SuggestedRemedy

If the subclause specifying AM creation is updated to include full text, this requirement can be included in it (a similar statement exists in 119.2.4.4.2 for the 16 lanes).

Otherwise, add this requirement as another exception, with editorial license.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #90.

Cl 172 SC 172.2.4.8 P 166 L 51 # 10

Ran, Adeo Cisco
 Comment Type ER Comment Status A (bucket2)

The functions above the "64B/66B to 256B/257B transcoder" are excluded'

This is confusing - looks as if these functions are not required, but of course they are.

It had to read it several times to understand that they are excluded from the "transmit function" blocks because they are present above them.

SuggestedRemedy

Change from
 The functions above the "64B/66B to 256B/257B transcoder" are excluded
 to
 The functions above the "64B/66B to 256B/257B transcoder" in Figure 119—11 are not included in the transmit function blocks, and instead are located outside of these blocks, as shown in Figure 172—3.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using response to comment #185.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.5.3 P 167 L 52 # 11
 Ran, Adee Cisco
 Comment Type **TR** Comment Status **A** *fec degrade (bucket1)*
 The FEC degrade variables in clause 172 should be stated as optional, as in their original definition in clause 119.
SuggestedRemedy
 Insert "If the optional PCS FEC degraded SER ability is implemented, " at the beginning of the first list item.
Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 It was intended for the FEC degrade to be optional, but as written that is not obvious. Add the following sentence at the end of 172.2.5.3:
 "The FEC degrade functionality is optional."

Cl 172 SC 172.2.5.4 P 168 L 5 # 12
 Ran, Adee Cisco
 Comment Type **TR** Comment Status **A** *(bucket1)*
 "The post-FEC interleave is identical to that specified in 119.2.5.4."
 But 119.2.5.4 talks specifically about two FEC codewords, and we have four.
 In similar subclauses for the transmit functions, the text includes "for each flow".
 Also applies to 172.2.5.6 and 172.2.5.7.
SuggestedRemedy
 Insert "for each flow" after "interleave".
 Make similar changes in 172.2.5.6 and 172.2.5.7, with editorial license.
Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

Cl 124 SC 124.8.5a P 76 L 16 # 13
 Dudek, Mike Marvell
 Comment Type **T** Comment Status **A** *(bucket1)*
 800GBASE-DR4 is not part of this specification
SuggestedRemedy
 Change to 800GBASE-DR8 Also on line 25 and page 77 line 29
Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 Change 800GBASE-DR4 to 800GBASE-DR8

Cl 124 SC 124.11.3.1 P 80 L 34 # 14
 Dudek, Mike Marvell
 Comment Type **T** Comment Status **A** *(bucket1)*
 The optical lane assignments are wrong in figure 124-6.
SuggestedRemedy
 Change them to match Figure 124-6 in the base document.
Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 Figure was intended to be the same as in in-force figure. Probably formatting problem. Check and update figure with editorial license

Cl 124 SC 124.11.3.3 P 81 L 29 # 15
 Dudek, Mike Marvell
 Comment Type **E** Comment Status **A** *(bucket1)*
 Should be plural
SuggestedRemedy
 Change "800GBASE-DR8 and 800GBASE-DR8-2 has" to "800GBASE-DR8 and 800GBASE-DR8-2 have"
Response Response Status **C**
 ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

CI 30 SC 30.5 P 33 L 45 # 16

Dudek, Mike Marvell
 Comment Type T Comment Status A (bucket1)

The base standard and 802.3db all list the "with reach up to at least xxx." to differentiate between the various Phys's. This draft does not.

SuggestedRemedy

Add the reach information to the new Phys.

Response Response Status C

ACCEPT IN PRINCIPLE.

change 400GBASE-DR4 description to:
 "400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124"

change 400GBASE-DR4-2 description to:
 "400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 124"

change 800GBASE-DR4 description to:
 "800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124"

change 800GBASE-DR4-2 description to:
 "800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 124"

change 800GBASE-SR8 description to:
 "800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 100 m as specified in Clause 167"

change 800GBASE-VR8 description to:
 "800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 50 m as specified in Clause 167"

Implement with editorial license.

CI 45 SC 45.2.1.8 P 38 L 13 # 17

Dudek, Mike Marvell
 Comment Type E Comment Status A (bucket1)

In table 45-12 "and" is used in the list for BR but it has been deleted for KR and CR. The table should be consistent for all rows.

SuggestedRemedy

Add the "and" before 800.

Response Response Status C

ACCEPT.

CI 45 SC 45.2.1.23 P 39 L 24 # 18

Dudek, Mike Marvell
 Comment Type T Comment Status A (bucket2)

This is listing register 1.72 but 45.2.1.60b is listing the abilities in Register 1.73

SuggestedRemedy

Change to register 1.72. Also on line39

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #44.

CI 45 SC 45.2.1.161 P 41 L 34 # 19

Dudek, Mike Marvell
 Comment Type T Comment Status A (bucket1)

The mapping of lanes 4-7 is not provided.

SuggestedRemedy

Add the mapping for those lanes. Also in 45.2.1.163 on line 50, 45.2.1.165 and 45.2.1.167

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #45

IEEE P802.3df D1.0 1st Task Force review comments

Cl 162 SC 162.11 P 94 L 51 # 20
 Dudek, Mike Marvell
 Comment Type E Comment Status A (bucket1)
 There are 4 cable assembly types
 SuggestedRemedy
 Change "three" to "four"
 Response Response Status C
 ACCEPT.

Cl 173 SC 173.1.4 P 177 L 28 # 24
 Dudek, Mike Marvell
 Comment Type E Comment Status A (bucket1)
 Should be "a physical instantiation"
 SuggestedRemedy
 Change "an" to "a"
 Response Response Status C
 ACCEPT.

Cl 163 SC 163.3 P 100 L 28 # 21
 Dudek, Mike Marvell
 Comment Type T Comment Status A (bucket1)
 Should be 800GASE-KR8 not KR4
 SuggestedRemedy
 fix it.
 Response Response Status C
 ACCEPT.

Cl 173 SC 173.1.4 P 178 L 33 # 25
 Dudek, Mike Marvell
 Comment Type T Comment Status A (bucket1)
 There are more than just two addresses (1 and 8) available for the MMD. (more are shown in figure 173-2)
 SuggestedRemedy
 Change "1 and 8" to "1,8,9 and 10".
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change from:
 "Manageable Device (MMD) addresses 1 and 8"
 to
 "Manageable Device (MMD) addresses 1,8,9,10 and 11"

Cl 163 SC 163.3 P 100 L 29 # 22
 Dudek, Mike Marvell
 Comment Type T Comment Status A (bucket1)
 should be 800GBASE-CR8 not KR8
 SuggestedRemedy
 Change it.
 Response Response Status C
 ACCEPT.

Cl 124 SC 124.11.3.1.1 P 80 L 32 # 26
 Bruckman, Leon Huawei
 Comment Type E Comment Status A (bucket1)
 In figure 124-6 the labels are all squeezed together
 SuggestedRemedy
 Spread the TX/RX labels to the right position
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #14

Cl 167 SC 167.2 P 110 L 23 # 23
 Dudek, Mike Marvell
 Comment Type E Comment Status A (bucket1)
 "have" should be "has" ("or" makes it singular)
 SuggestedRemedy
 change it.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Replace "PMD have eight" with "PMD has eight".

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.4.9 P 167 L 25 # 27
 Bruckman, Leon Huawei
 Comment Type T Comment Status A test pattern (CC)
 I assume test pattern shall be applied to both flows together
 SuggestedRemedy
 It may be beneficial to note that the test function when activated affects both flows
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 The scrambled idle pattern for 800GBASE-R is different from that for 400GBASE-R and 200GBASE-R.
 Implement with editorial license the changes on slide 18 of the following presentation:
https://www.ieee802.org/3/df/public/22_12/brown_3df_03b_2212.pdf

Cl 172 SC 172.2.6.2.4 P 170 L 15 # 28
 Bruckman, Leon Huawei
 Comment Type T Comment Status R (bucket1)
 From this clause it may be implied that counters are not aggregated, but in the MDIO Table 172-4 shows (and text indicates that) they are aggregated
 SuggestedRemedy
 Add exception indicating that counters are the aggregate of both flows
 Response Response Status C
 REJECT.
 172.2.6.2.4 is defining the counters used in the state diagrams. The definition of these counters is identical to that in 119.2.6.2.4. Therefore, these counters are not aggregated and are not the same as those defined in Table 172-4.

Cl 173 SC 173.2 P 179 L 10 # 29
 Bruckman, Leon Huawei
 Comment Type T Comment Status A PMA SI
 "In the case where the sublayer below the PMA is a PHY 800GXS the PMA does not receive a PHY_XS:IS_SIGNAL.indication as an input to the SIL". Figure 173-4 that describes this interface does include the PHY_XS:IS_SIGNAL.indication
 SuggestedRemedy
 Update Figure 173-4 according to text
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #196.

Cl 45 SC 45.2.1.165 P 42 L 8 # 30
 Huber, Tom Nokia
 Comment Type T Comment Status A (bucket2)
 While the mapping of bits to registers is obvious, it seems incomplete to explicitly describe the mapping for bits 0-3 and say nothing at all about bits 4-7. A simpler statement of how the mapping works for all bits would be better and easier to maintain.
 SuggestedRemedy
 Change "Lane 0 maps to register 1.1320, lane 1 maps to register 1.1321, lane 2 maps to register 1.1322, and lane 3 maps to register 1.1323."
 to
 "Lanes 0-7 map to registers 1.1320 to 1.1327, respectively."
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 The word "to" should be used instead of "-".
 Change text to: "Lanes 0 to 7 map to registers 1.1320 to 1.1327, respectively."

Cl 45 SC 45.2.1.167 P 42 L 23 # 31
 Huber, Tom Nokia
 Comment Type T Comment Status A (bucket2)
 While the mapping of bits to registers is obvious, it seems incomplete to explicitly describe the mapping for bits 0-3 and say nothing at all about bits 4-7. A simpler statement of how the mapping works for all bits would be better and easier to maintain.
 SuggestedRemedy
 Change "Lane 0 maps to register 1.1420, lane 1 maps to register 1.1421, lane 2 maps to register 1.1422, and lane 3 maps to register 1.1423."
 to
 "Lanes 0-7 map to registers 1.1420 to 1.1427, respectively."
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 The word "to" should be used instead of "-".
 Change text to: "Lanes 0 to 7 map to registers 1.1420 to 1.1427, respectively."

IEEE P802.3df D1.0 1st Task Force review comments

Cl 45 SC 45.2.1.168 P 42 L 38 # 32

Huber, Tom Nokia
 Comment Type T Comment Status A (bucket2)

While the mapping of registers to what they control is obvious, it would be better to spell it out a bit more completely to maintain similar structure to the other clauses that are specifying registers per-lane.

SuggestedRemedy

Change "Register 1.1450 controls the PMD training pattern for PMD lane 0; register 1.1451 controls the PMD training pattern for PMD lane 1; etc."
 to
 "Registers 1.1450 to 1.1457 control the PMD training pattern for PMD lanes 0-7, respectively."

Response Response Status C

ACCEPT IN PRINCIPLE.

The word "to" should be used instead of "-".

Change text to: "Registers 1.1450 to 1.1457 control the PMD training pattern for PMD lanes 0 to 7, respectively."

Cl 45 SC 45.2.1.168 P 42 L 41 # 33

Huber, Tom Nokia
 Comment Type E Comment Status R (bucket1)

The text "and 136.8.11.1.3" is in 802.3-2022, so it should not be identified as a change.

SuggestedRemedy

Remove the underlining from this text.

Response Response Status C

REJECT.

The reference to 136.8.11.1.3 is not in the base standard so the underlining should remain.

Cl 45 SC 45.2.1.168 P 42 L 42 # 34

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

The last 3 sentences would be clearer if the order of the last two sentences is swapped, and the (current) last sentence is written more generically to apply to any situation where a polynomial identifier is being reused.

SuggestedRemedy

Replace "The polynomial identifier for each lane should be unique; two physically adjacent lanes having the same identifier could impair operation of the PMD control function. The default identifiers are (binary): for lane 0, 00; for lane 1, 01; for lane 2, 10; for lane 3, 11; for lane 4, 00; for lane 5, 01; for lane 6, 10; for lane 7, 11. For 8-lane use cases different initial seeds should be used where the same polynomial is being reused."
 with

"The polynomial identifier for each lane should be unique; two physically adjacent lanes having the same identifier could impair operation of the PMD control function. If the same polynomial identifier is used for multiple lanes, different initial seeds should be used for each of those lanes. The default identifiers are (binary): for lane 0, 00; for lane 1, 01; for lane 2, 10; for lane 3, 11; for lane 4, 00; for lane 5, 01; for lane 6, 10; for lane 7, 11."

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #122

Cl 45 SC 45.2.3 P 43 L 12 # 35

Huber, Tom Nokia
 Comment Type E Comment Status R (bucket1)

Subclauses 45.2.3.24-26 all exist in 802.3-2022, so they should not be indicated as changes in the table.

SuggestedRemedy

Remove the underlining from 45.2.3.24, 45.2.3.25, 45.2.3.26.

Response Response Status C

REJECT.

Although these clauses are in the base standard, there are no references to them in Table 45-233. Therefore it is appropriate to add them to the table with underlining.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 45 SC 45.2.3 P 43 L 50 # 36

Huber, Tom Nokia
 Comment Type E Comment Status R (bucket1)

Subclause 45.2.3.50 exists in 802.3-2022, so it should not be indicated as a change in the table.

SuggestedRemedy

Remove the underlining from 45.2.3.50

Response Response Status C

REJECT.
 Although this subclause is in the base standard there is no reference to it in the table. Therefore it is appropriate to add it to Table 45-233 with underlining.

Cl 124 SC 124.1 P 59 L 24 # 37

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

Table 124-1 was modified by 802.3ck-2022

SuggestedRemedy

Change the editing instruction to add "(as modified by IEEE 802.3ck-2022)", and insert the rows for Annexes 120F and 120G into the table.

Response Response Status C

ACCEPT IN PRINCIPLE.
 Implement proposed remedy with editorial license

Cl 124 SC 124.1 P 61 L 36 # 38

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

Since there are only two items in the list, they should be separated with and rather than a comma

SuggestedRemedy

Change "400GBASE-DR4, 400GBASE-DR4-2" to "400GBASE-DR4 and 400GBASE-DR4-2"

Response Response Status C

ACCEPT.

Cl 162 SC 162.1 P 85 L 8 # 39

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

Elsewhere in the clause (e.g. in 162.4), 800GAUI-n is used, which seems desirable since it will be more future-proof toward the 200G/lane AUI that will be added.

SuggestedRemedy

Change 800GAUI-8 to 800GAUI-n.

Response Response Status C

ACCEPT.

Cl 169 SC 169.1.2 P 127 L 36 # 40

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

The dashed lines between the OSI layers and the Ethernet layers are not in the correct locations.

SuggestedRemedy

Align the upper two dashed lines with the boundaries of the data link layer in the OSI model.

Response Response Status C

ACCEPT.

Cl 169 SC 169.1.2 P 128 L 4 # 41

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

Singular/plural disagreement in item a)

SuggestedRemedy

Change "when implemented as logical interconnection points" to "when implemented as a logical interconnection point"

Response Response Status C

ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

CI 172 SC 172.1.3 P 161 L 6 # 42

Huber, Tom Nokia

Comment Type E Comment Status A (bucket1)

missing "(to)" in the transcoding description in item b)

SuggestedRemedy

Change "Transcoding from 66-bit blocks to (from 257-bit blocks (25B/257B))" to "Transcoding from (to) 66-bit blocks to (from 257-bit blocks (25B/257B))"

Response Response Status C

ACCEPT IN PRINCIPLE.
Change from "Transcoding from 66-bit blocks to (from) 257-bit blocks (256B/257B))" to "Transcoding from (to) 66-bit blocks to (from) 257-bit blocks (256B/257B))"

CI 45 SC 45.2.1.6 P 36 L 3 # 43

Huber, Tom Nokia

Comment Type E Comment Status A (bucket1)

Since the table includes 400ZR as existing text, the editing instruction should note that the text shown is as modified by 802.3cw.

SuggestedRemedy

Add "(as modified by IEEE 802.3cw-202x)" after "Change Table 45-7"

Response Response Status C

ACCEPT.

CI 45 SC 45.2.1.23 P 39 L 23 # 44

Huber, Tom Nokia

Comment Type T Comment Status A (bucket2)

Register 1.72 is added by 802.3cz; presumably 1.73 is what was intended here

SuggestedRemedy

Change 1.72 to 1.73

Response Response Status C

ACCEPT IN PRINCIPLE.

Change 1.72 to 1.73 on lines 23 and 39

CI 45 SC 45.2.1.161 P 41 L 34 # 45

Huber, Tom Nokia

Comment Type T Comment Status A (bucket2)

While the mapping of bits to registers is obvious, it seems incomplete to explicitly describe the mapping for bits 0-3 and say nothing at all about bits 4-7. A simpler statement of how the mapping works for all bits would be better and easier to maintain.

SuggestedRemedy

Change "Lane 0 maps to register 1.1120, lane 1 maps to register 1.1121, lane 2 maps to register 1.1122, and lane 3 maps to register 1.1123."

to
"Lanes 0-7 map to registers 1.1120 to 1.1127, respectively."

Response Response Status C

ACCEPT IN PRINCIPLE.

The word "to" should be used instead of "-".

Change text to: "Lanes 0 to 7 map to registers 1.1120 to 1.1127, respectively."

CI 45 SC 45.2.1.163 P 41 L 50 # 46

Huber, Tom Nokia

Comment Type T Comment Status A (bucket2)

While the mapping of bits to registers is obvious, it seems incomplete to explicitly describe the mapping for bits 0-3 and say nothing at all about bits 4-7. A simpler statement of how the mapping works for all bits would be better and easier to maintain.

SuggestedRemedy

Change "Lane 0 maps to register 1.1220, lane 1 maps to register 1.1221, lane 2 maps to register 1.1222, and lane 3 maps to register 1.1223."

to
"Lanes 0-7 map to registers 1.1220 to 1.1227, respectively."

Response Response Status C

ACCEPT IN PRINCIPLE.

The word "to" should be used instead of "-".

Change text to:
"Lanes 0 to 7 map to registers 1.1220 to 1.1227, respectively."

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.1 P 163 L 38 # 47

Huber, Tom Nokia
 Comment Type T Comment Status A pcs functions

There is some repetition between the paragraph about the PCS Synchronization process and the paragraph about the PCS Receive process in terms of aligning, reordering, and deskewing. Per the state diagrams, the PCS synchronization process ensures that all the lanes are aligned and deskewed, and the receive process deals with deocding the 66b characters.

SuggestedRemedy

Add a sentence to the end of the penultimate paragraph: "When all 32 lanes are aligned and deskewed, and reordered, the align_status flag is set to indicate that the PCS has obtained alignment."

Revise the first two sentences of the final paragraph as follows: "The PCS Receive process separates the reordered PCS lanes into two sets of 16 PCs lanes..."

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement with editorial license the changes on slide 33 of the following presentation: https://www.ieee802.org/3/df/public/22_12/brown_3df_03b_2212.pdf

Cl 172 SC 172.2.4.1 P 164 L 28 # 48

Huber, Tom Nokia
 Comment Type T Comment Status A OTN reference point

The OTN reference point needs further discussion - it would be preferrable if the mapping point was 257b blocks rather than 66b blocks..

SuggestedRemedy

Supporting presentation to be provided.

Response Response Status C

ACCEPT IN PRINCIPLE.

The following presentation was reviewed by the task force: https://www.ieee802.org/3/df/public/22_12/huber_3df_01_2212.pdf

There is consensus to adopt the proposal on slide 7 in the reviewed presentation.

Implement with editorial license the proposal on slide 7 of the reviewed presentation.

Cl 120F SC 120F.1 P 198 L 25 # 49

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

To maintain parallel structure with the rest of the sentence, the new 800G AUJ should be introduced as 800Gb/s eight-lane

SuggestedRemedy

change "and eight-lane Attachment Unit Interface" to "800 Gb/s eight-lane Attachment Unit Interface"

Response Response Status C

ACCEPT IN PRINCIPLE.

change "and eight-lane" to "and 800 Gb/s eight-lane".

Cl 120F SC 120F.3.1 P 201 L 10 # 50

Huber, Tom Nokia
 Comment Type E Comment Status A rate range

The inserted text is more complex than is necessary.

SuggestedRemedy

Change "800GAUI-8 C2C or for 100GAUI-1, 200GAUI-2, or 400GAUI-4 C2C with" to "100GAUI-1, 200GAUI-2, 400GAUI-4, or 800GAUI-8 C2C"

Response Response Status C

ACCEPT IN PRINCIPLE.

The text intentionally distinguishes between 800GAUI-8, for which the range is always +/- 50 PPM, and the other interfaces, for which it is conditional.

Therefore, the suggested remedy would not be correct. However, the text can be clarified.

In Table 120F-1 change the first sentence in footnote a to the following: "For 100GAUI-1, 200GAUI-2, or 400GAUI-4 C2C with a PMA in the same package as the PCS sublayer or for any 800GAUI-8 C2C."

In Table 120G-1 change the first sentence in footnote a to the following: "For 100GAUI-1, 200GAUI-2, or 400GAUI-4 C2M with a PMA in the same package as the PCS sublayer or for any 800GAUI-8 C2M."

IEEE P802.3df D1.0 1st Task Force review comments

CI 162B SC 162B P 215 L 11 # 51
 Huber, Tom Nokia
 Comment Type E Comment Status A (bucket1)
 The title is missing 'C2M' for 800GAUI-8
 SuggestedRemedy
 Add 'C2M' to the end of the title
 Response Response Status C
 ACCEPT.

CI 173A SC 173A P 226 L 1 # 52
 Huber, Tom Nokia
 Comment Type E Comment Status A (bucket1)
 The text should be referencing figure 173A-3.
 SuggestedRemedy
 Change 173A-4 to 173A-3.
 Response Response Status C
 ACCEPT.

CI 45 SC 45.2.4.4 P 46 L 54 # 53
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 Need to add 800G capability register to PHY XS
 SuggestedRemedy
 Assign a bit in register 4.4 for 800G capable and create a description the same as the 400G bit replacing 400G with 800G
 Response Response Status C
 ACCEPT.

CI 45 SC 45.2.5.4 P 46 L 54 # 54
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 Need to add 800G capability register to DTE XS
 SuggestedRemedy
 Assign a bit in register 5.4 for 800G capable and create a description the same as the 400G bit replacing 400G with 800G
 Response Response Status C
 ACCEPT.

CI 45 SC 45.2.5.15 P 46 L 54 # 55
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 DTE XS AM lock registers need to be updated with 800G references and expanded to 32 AM lanes
 SuggestedRemedy
 Update (see 119.2.6.2.2) to (see 119.2.6.2.2 and 172.2.6.2.2) in 45.2.4.15.* and 45.2.4.16.*
 Add the extra 16 lanes of amps_lock as well as was done for the PCS registers.
 Response Response Status C
 ACCEPT.

CI 45 SC 45.2.5.17 P 46 L 54 # 56
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 DTE XS lane mapping registers need to update with 800G references and expanded to 32 lanes
 SuggestedRemedy
 Bring in and update 45.2.5.17 and 45.2.5.18 adding references to Clause 171 and adding 16 more registers
 Response Response Status C
 ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 45 SC 45.2.5.19 P 46 L 54 # 57
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 DTE XS symbol error counter registers needs update with 800G references and expanded to 32 lanes
 SuggestedRemedy
 Bring in and update 45.2.5.19 and 45.2.5.20 adding references to 172.3.4 and adding 16 more counters
 Response Response Status C
 ACCEPT.

Cl 169 SC 169.3.2 P 133 L 45 # 58
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 800GAUI-n is not listed in the list of acronyms for Figure 169-3
 SuggestedRemedy
 Add 800GAUI-n to list of acronyms in Figure 169-3
 Response Response Status C
 ACCEPT.

Cl 171 SC 171.4 P 151 L 38 # 59
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 There is no am_lock variable in Clause 172
 SuggestedRemedy
 Change am_lock to amps_lock in Table 171-3 and 171-5
 Response Response Status C
 ACCEPT.

Cl 172 SC 172.2.4.4 P 164 L 49 # 60
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A AM sync
 Missing the relationship of the flow 0 257-bit block to the AM group
 SuggestedRemedy
 add "following the alignment marker group" before "in flow 0"
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #90.

Cl 172 SC 172.3.1 P 172 L 35 # 61
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 The variable name is amps_lock not am_lock
 SuggestedRemedy
 Change am_lock to amps_lock in Table 172--4
 Response Response Status C
 ACCEPT.

Cl 45 SC 45.2.3.26a P 44 L 24 # 62
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 Clause 172 (and 119) use a variable named amps_lock[x] for lane alignment lock status. Which was the name used in CI91 and 161 for the FEC sublayers.
 SuggestedRemedy
 Bring in 45.2.3.25.* and 45.2.3.26.*
 For indexes 16 to 32 change the "(see 82.2.19.2.2)." to be "(see 82.2.19.2.2) or amps_lock[16] (see 172.2.6.2.2)"
 For indexes 0 to 15 and change the "(see 82.2.19.2.2)." to be "(see 82.2.19.2.2) or amps_lock[16] (see 119.2.6.2.2 and 172.2.6.2.2)"
 Response Response Status C
 ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.3.5 P 173 L 32 # 63
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A fec counters
 The CW counter is a RS-FEC sublayer counter in MDIO space, not a PCS counter.
SuggestedRemedy
 Copy of the definition of 45.2.1.120a (802.3ck) into a set of PCS registers (45.2.3.###) and replace the Clause 161 references with 172.
 Replace the text in 172.2.3.5 with the same text from 161.6.21 updating the MDIO register references to point to the newly created MDIO registers.
 Update Table 172-4 to point to the newly created MDIO registers.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 172 SC 172.3.6 P 173 L 32 # 64
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A fec counters
 The FEC_codeword_error_bin_i is a RS-FEC sublayer set of counters in MDIO space, not PCS counters.
SuggestedRemedy
 Copy of the definition of 45.2.1.131a (802.3ck) into a set of PCS registers (45.2.3.###) and replace the Clause 161 references with 172.
 Replace the text in 172.2.3.6 with the same text from 161.6.17 updating the MDIO register references to point to the newly created MDIO registers.
 Update Table 172-4 to point to the newly created MDIO registers.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

Cl 45 SC 45.2.3.60.1 P 46 L 54 # 65
 Slavick, Jeff Broadcom
 Comment Type T Comment Status A (bucket1)
 Various clause 45 registers need to some Clause 172 references added.
SuggestedRemedy
 A reference to Clause 172 needs to be added to 45.2.3.49
 A reference to 172.2.5.3 needs to be added to:
 45.2.3.60.1
 45.2.3.60.2
 45.2.4.61.4
 45.2.3.61.6
 45.2.3.64
 45.2.3.65
 45.2.3.66
 45.2.4.21.1
 45.2.4.21.2
 45.2.4.22.2
 45.2.4.22.3
 45.2.4.22.4
 45.2.4.22.5
 45.2.4.25
 45.2.4.26
 45.2.4.27
 45.2.5.21.1
 45.2.5.21.2
 45.2.5.22.2
 45.2.5.22.3
 45.2.5.22.4
 45.2.5.22.5
 45.2.5.25
 45.2.5.26
 45.2.5.27
 A reference to 172.2.6.2.2 needs to be added to:
 45.2.3.61.1
 45.2.3.61.2
 45.2.3.61.3
 45.2.3.61.5
 45.2.4.22.1
 45.2.5.22.1
 A reference to 172.3.2 needs to be added to 45.2.3.62, 45.2.4.23 and 45.2.5.23
 A reference to 172.3.3 needs to be added to 45.2.3.63, 45.2.4.24 and 45.2.5.24

IEEE P802.3df D1.0 1st Task Force review comments

A reference to 172.3.4 needs to be added to 45.2.3.58

Response *Response Status* **C**
ACCEPT.

Cl 45 **SC 45.2.4.15** **P 46** **L 54** # **66**

Slavick, Jeff **Broadcom**

Comment Type **T** *Comment Status* **A** (bucket1)

PHY XS AM lock registers need to be updated with 800G references and expanded to 32 AM lanes

SuggestedRemedy

Update (see 119.2.6.2.2) to (see 119.2.6.2.2 and 172.2.6.2.2) in 45.2.4.15.* and 45.2.4.16.* Add the extra 16 lanes of amps_lock as well as was done for the PCS registers.

Response *Response Status* **C**
ACCEPT.

Cl 45 **SC 45.2.4.17** **P 46** **L 54** # **67**

Slavick, Jeff **Broadcom**

Comment Type **T** *Comment Status* **A** (bucket1)

PHY XS lane mapping registers need to update with 800G references and expanded to 32 lanes

SuggestedRemedy

Bring in and update 45.2.4.17 and 45.2.4.18 adding references to Clause 171 and adding 16 more registers

Response *Response Status* **C**
ACCEPT.

Cl 45 **SC 45.2.4.19** **P 46** **L 54** # **68**

Slavick, Jeff **Broadcom**

Comment Type **T** *Comment Status* **A** (bucket1)

PHY XS symbol error counter registers needs update with 800G references and expanded to 32 lanes

SuggestedRemedy

Bring in and update 45.2.4.19 and 45.2.4.20 adding references to 172.3.4 and adding 16 more counters

Response *Response Status* **C**
ACCEPT.

Cl 45 **SC 45.2.1.161** **P 41** **L 34** # **69**

Lusted, Kent **Intel Corporation**

Comment Type **TR** *Comment Status* **A** (bucket1)

The paragraph provides mapping of registers 1.1120-1.1123 to lanes [0:3] but not the additional lanes of [4:7] used for eight-lane interface types.

SuggestedRemedy

change:

" Lane 0 maps to register 1.1120, lane 1 maps to register 1.1121, lane 2 maps to register 1.1122, and lane 3 maps to register 1.1123."

to:

" Lane 0 maps to register 1.1120, lane 1 maps to register 1.1121, lane 2 maps to register 1.1122, lane 3 maps to register 1.1123, lane 4 maps to register 1.1124, lane 5 maps to register 1.1125, lane 6 maps to register 1.1126, and lane maps to register 1.1127."

Response *Response Status* **C**

ACCEPT IN PRINCIPLE.
Resolve using the response to comment #45

Cl 45 **SC 45.1.2.163** **P 41** **L 50** # **70**

Lusted, Kent **Intel Corporation**

Comment Type **TR** *Comment Status* **A** (bucket1)

The paragraph provides mapping of registers 1.1220-1.1223 to lanes [0:3] but not the additional lanes of [4:7] used for eight-lane interface types.

SuggestedRemedy

change:

" Lane 0 maps to register 1.1220, lane 1 maps to register 1.1221, lane 2 maps to register 1.1222, and lane 3 maps to register 1.1223."

to:

" Lane 0 maps to register 1.1220, lane 1 maps to register 1.1221, lane 2 maps to register 1.1222, lane 3 maps to register 1.1223, lane 4 maps to register 1.1224, lane 5 maps to register 1.1225, lane 6 maps to register 1.1226, and lane maps to register 1.1227."

Response *Response Status* **C**

ACCEPT IN PRINCIPLE.
Resolve using the response to comment #46

IEEE P802.3df D1.0 1st Task Force review comments

CI 45 SC 45.1.2.165 P 42 L 8 # 71

Lusted, Kent Intel Corporation

Comment Type TR Comment Status A (bucket1)

The paragraph provides mapping of registers 1.1320-1.1323 to lanes [0:3] but not the additional lanes of [4:7] used for eight-lane interface types.

SuggestedRemedy

change:

" Lane 0 maps to register 1.1320, lane 1 maps to register 1.1321, lane 2 maps to register 1.1322, and lane 3 maps to register 1.1323."

to:

" Lane 0 maps to register 1.1320, lane 1 maps to register 1.1321, lane 2 maps to register 1.1322, lane 3 maps to register 1.1323, lane 4 maps to register 1.1324, lane 5 maps to register 1.1325, lane 6 maps to register 1.1326, and lane maps to register 1.1327."

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #30

CI 45 SC 45.1.2.167 P 42 L 23 # 72

Lusted, Kent Intel Corporation

Comment Type TR Comment Status A (bucket1)

The paragraph provides mapping of registers 1.1420-1.1423 to lanes [0:3] but not the additional lanes of [4:7] used for eight-lane interface types.

SuggestedRemedy

change:

" Lane 0 maps to register 1.1420, lane 1 maps to register 1.1421, lane 2 maps to register 1.1422, and lane 3 maps to register 1.1423."

to:

" Lane 0 maps to register 1.1420, lane 1 maps to register 1.1421, lane 2 maps to register 1.1422, lane 3 maps to register 1.1423, lane 4 maps to register 1.1424, lane 5 maps to register 1.1425, lane 6 maps to register 1.1426, and lane maps to register 1.1427."

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #31.

CI 162 SC 162.1 P 84 L 35 # 73

Lusted, Kent Intel Corporation

Comment Type TR Comment Status A (bucket1)

In Table 162-3a, the rightmost column heading is incorrect as the table refers to 800GBASE-CR8.

SuggestedRemedy

Change rightmost column heading to "800GBASE-CR8"

Response Response Status C

ACCEPT.

CI 162 SC 162.7 P 89 L 24 # 74

Lusted, Kent Intel Corporation

Comment Type E Comment Status A (bucket1)

With the addition of new sub-note "a", the rest of the sub-notes from the table 162-5 in P802.3ck are re-indexed. (i.e. 'a' becomes 'b', 'b' becomes 'c'). However, the new notes 'b' and 'c' do not have the relevant strikeout text

SuggestedRemedy

Correct as necessary

Response Response Status C

ACCEPT IN PRINCIPLE.

Table footnote are numbered automatically in FrameMaker and cannot be struck out.

Change the editorial instruction from

"Change Table 162-5, Table 162-6, and Table 162-7 as follows:"

to

"Change Table 162-5, Table 162-6, and Table 162-7, including footnotes, as follows:"

CI 162 SC 162.7 P 89 L 49 # 75

Lusted, Kent Intel Corporation

Comment Type E Comment Status A (bucket1)

With the addition of new sub-note "a", the rest of the sub-notes from the table 162-6 in P802.3ck are re-indexed. (i.e. 'a' becomes 'b', 'b' becomes 'c'). However, the new notes 'b' and 'c' do not have the relevant strikeout text

SuggestedRemedy

Correct as necessary

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #74.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 162 SC 162.13.3 P97 L 21 # 76
 Lusted, Kent Intel Corporation
 Comment Type **TR** Comment Status **A** (bucket1)
 Row entry for PMA800 has incorrect status value of "CR4:M". It should be "CR8:M"
 SuggestedRemedy
 Change to "CR8:M"
 Response Response Status **C**
 ACCEPT.

Cl 162 SC 162.13 P96 L 4 # 77
 Lusted, Kent Intel Corporation
 Comment Type **TR** Comment Status **A** (bucket1)
 In P802.3ck, Clause 162.13 is the environmental specifications and Clause 162.14 is the PICS. The 162.13 sub clause is missing from the draft and creates an issue where the PICS became sub clause 162.13.
 SuggestedRemedy
 Fix editing instruction on p96, line 1 to reference the heading of 162.14
 Correct the sub clause number for the PICS to 162.14 in the title and the sub clauses.
 Update all editing instructions as required.
 Implement with editorial license
 Response Response Status **C**
 ACCEPT.

Cl 163 SC 163.3 P100 L 27 # 78
 Lusted, Kent Intel Corporation
 Comment Type **TR** Comment Status **A** (bucket1)
 Text references "CR" PMD types in the PMD service interfaces for Clause 163, which is for backplanes.
 SuggestedRemedy
 Change "100GBASE-CR1, 200GBASE-CR2, 400GBASE-CR4" to "100GBASE-KR1, 200GBASE-KR2, and 400GBASE-KR4"
 Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 The text states that the KR* service interfaces are identical to those of CR*. The addition of "KR8" was erroneous.
 Resolve using the response to comment #22.

Cl 162 SC 162.13 P105 L 4 # 79
 Lusted, Kent Intel Corporation
 Comment Type **TR** Comment Status **D** Withdrawn
 In P802.3ck, Clause 163.13 is the environmental specifications and Clause 163.14 is the PICS. The 163.13 sub clause is missing from the draft and creates an issue where the PICS became sub clause 163.13.
 SuggestedRemedy
 Fix editing instruction on p105, line 1 to reference the heading of 163.14
 Correct the sub clause number for the PICS to 163.14 in the title and the sub clauses.
 Update all editing instructions as required.
 Implement with editorial license
 Proposed Response Response Status **Z**
 PROPOSED REJECT.
 This comment was WITHDRAWN by the commenter.

Cl 169 SC 169.5 P136 L 10 # 80
 Lusted, Kent Intel Corporation
 Comment Type **ER** Comment Status **A** (bucket1)
 Figure 169-4 variable "q" should be italics like 'n' and 'p'. Both in middle and bottom of figure
 SuggestedRemedy
 consider changing 'q' to italics types
 Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 Change font to italic for variable q.
 Implement with editorial license.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 120F SC 120F.1 P 198 L 48 # 81
 Lusted, Kent Intel Corporation
 Comment Type T Comment Status A (bucket1)
 Paragraph omits the eight-lane 800GAUI-8.
 SuggestedRemedy
 Replace the second sentence in the 5th paragraph with "Each 100GAUI-1, 200GAUI-2, 400GAUI-4, or 800GAUI-8 C2C data path contains one, two, four, or eight, respectively, differential lanes, which are AC coupled."
 Response Response Status C
 ACCEPT.

Cl 120F SC 120F.1 P 198 L 52 # 82
 Lusted, Kent Intel Corporation
 Comment Type TR Comment Status A (bucket1)
 The mapping of the differential voltage level to the PAM4 symbol is missing in Annex 120F. It is also not present in Annex 120F in IEEE Std. 802.3ck-202x. The mapping of the differential voltage level to the PAM4 symbol level is important for interoperability.
 SuggestedRemedy
 Add a new sentence to the 5th paragraph: "The highest differential level corresponds to the symbol three and the lowest level corresponds to the symbol zero."
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 In the sixth paragraph, change "The C2C transmitter and the receiver use PAM4 signaling" To:
 "The C2C transmitter and receiver use PAM4 signaling. The highest differential level corresponds to the tx_symbol or rx_symbol value three, and the lowest differential level corresponds to the tx_symbol or rx_symbol value zero."

Cl 162B SC 162B P 215 L 11 # 83
 Lusted, Kent Intel Corporation
 Comment Type E Comment Status A (bucket1)
 The title of Annex 162B is missing "C2M" after the 800GAUI-8 entry.
 SuggestedRemedy
 Add "C2M" after 800GAUI-8
 Response Response Status C
 ACCEPT.

Cl 162 SC 162.8.1 P 91 L 22 # 84
 Opsasnick, Eugene Broadcom
 Comment Type ER Comment Status A (bucket1)
 At top-middle of Figure 162-2, the added text reads "800GBASE-CR4 8x", but "-CR4" should probably be "-CR8".
 SuggestedRemedy
 Replace "800GBASE-CR4 8x" with "800GBASE-CR8 8x".
 Response Response Status C
 ACCEPT.

Cl 163 SC 163.3 P 100 L 27 # 85
 Opsasnick, Eugene Broadcom
 Comment Type ER Comment Status A (bucket1)
 At end of first line of paragraph, 800GBASE-KR4 (wraps to line 28), "-KR4" should probably be "-KR8"
 SuggestedRemedy
 Replace "800GBASE-KR4" with "800GBASE-KR8" and use non-breaking hyphen.
 Response Response Status C
 ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.6.3 P 170 L 19 # 86

Opsasnick, Eugene Broadcom
 Comment Type TR Comment Status A stateless encoder

"State diagrams are identical to those specified in 119.2.6.3 ..."

State diagrams in Figure 119-14 "Transmit state diagram" and Figure 119-15 "Receive state diagram" can cause logic implementation issues at high rate port speeds (i.e. 800GbE) as shown in opsasnick_3df_01a_221005.pdf. A "stateless" encode/decode option to these state diagrams could be allowed since the state diagrams were originally designed for non-FEC interfaces. Interfaces with required FEC should have sufficient protection to allow for the stateless coding. An updated presentation showing the error analysis will be forthcoming.

SuggestedRemedy

To be shown in an updated presentation for December comment resolution meetings.

Response Response Status C

ACCEPT IN PRINCIPLE.

The following presentation was reviewed by the task force:
https://www.ieee802.org/3/df/public/22_12/opsasnick_3df_01a_2212.pdf

There was consensus to adopt the optional stateless encode and decode proposal captured in slides 12 and 13 of the reviewed presentation.

Implement with editorial license the optional stateless encode and decode proposal captured in slides 12 and 13 of the reviewed presentation.

Cl 120G SC 120G.3.2.1 P 209 L 21 # 87

Opsasnick, Eugene Broadcom
 Comment Type ER Comment Status A (bucket1)

In Table 120G-4, four instances of "800GAUI-4" in last two rows of the table should likely be "800GAUI-8"

SuggestedRemedy

Replace "800GAUI-4" with "800GAUI-8"

Response Response Status C

ACCEPT IN PRINCIPLE.
 Change 800GAUI-4 to 800GAUI-8 in the bottom two rows of the table (4 instances).

Cl 124 SC 124.8.5a P 76 L 15 # 88

Opsasnick, Eugene Broadcom
 Comment Type ER Comment Status A (bucket1)

In second line of paragraph, "800GBASE-DR4" should probably be "...-DR8". Same text appears on line 25 in 124.8.5b, and on page 77, line 29, section 124.8.9.2.

SuggestedRemedy

Replace "800GBASE-DR4" with "800GBASE-DR8".

Response Response Status C

ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #13.

Cl 124 SC 124.3.1 P 63 L 13 # 89

He, Xiang Huawei
 Comment Type ER Comment Status A (bucket1)

Looks like a typo. "16834 bit times" should be "16384 bit times"

SuggestedRemedy

Change 16834 to 16384.

Response Response Status C

ACCEPT.

Cl 172 SC 172.1.5 P 162 L 3 # 90

Rechtman, Zvi Nvidia
 Comment Type T Comment Status A AM sync

Figure 172-2—Functional block diagram
 The block diagram includes two flows for TX and Rx.
 Both TX flows are supposed to insert the alignment markers in sync with each other. This does not appear explicitly in the diagram.

SuggestedRemedy

Possible improvement #1:
 Add arrow with the word synchronization between the "Alignment insertion" blocks.
 Possible improvement #2:
 Add a footnote that the two "Alignment insertion" should operate in synchronized manner.

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement with editorial license the changes on slide 38 of
https://www.ieee802.org/3/df/public/22_12/brown_3df_03b_2212.pdf

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.4.4 P 164 L 48 # 91

Rechtman, Zvi

Nvidia

Comment Type T Comment Status A AM sync

"The first 66-bit block of the 257-bit transcoded block .. from the 64B/66B encoder."
This sentence implicitly means that the alignment insertion process of the two flows should be synchronized.
To avoid mistakes, it would be preferable to explicitly state that the two alignment insertion are synchronized

SuggestedRemedy

Add the following sentence before "The first 66-bit... " sentence:
"The marker insertion functions of the two flows must insert their markers at the exact same time (block unit), i.e. in a synchronized manner"

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #90.

Cl 30 SC 30.5.1.1.2 P 33 L 1 # 92

Wang, Haojie

China Mobile

Comment Type ER Comment Status A (bucket1)

There should be "800GBASE-R" other than "400GBASE-R"

SuggestedRemedy

Change "400GBASE-R" to "800GBASE-R"

Response Response Status C

ACCEPT.

Cl 30 SC 30.5.1.1.2 P 33 L 3 # 93

Wang, Haojie

China Mobile

Comment Type ER Comment Status A (bucket1)

There should be "800GBASE-R" other than "400GBASE-R"

SuggestedRemedy

Change "400GBASE-R" to "800GBASE-R"

Response Response Status C

ACCEPT.

Cl 124 SC 124.11.3.1.1 P 80 L 32 # 94

Wang, Haojie

China Mobile

Comment Type ER Comment Status A (bucket1)

The positions of "Rx" in figure 124-6 is inconsistent with the text at line 27, which is depicted as the right-most four positions.

SuggestedRemedy

Plot the four "Rx" at the right-most four positions.

Response Response Status C

ACCEPT IN PRINCIPLE.
Resolve using the response to comment #14.

Cl 124 SC 124.2 P 62 L 16 # 95

Nicholl, Gary

Cisco Systems

Comment Type ER Comment Status A (bucket1)

The space after "these" should be underlined.

SuggestedRemedy

Underline the space after "these"

Response Response Status C

ACCEPT.

Cl 124 SC 124.2 P 62 L 29 # 96

Nicholl, Gary

Cisco Systems

Comment Type ER Comment Status A (bucket1)

The space after "have" should be underlined

SuggestedRemedy

Underline the space after "have"

Response Response Status C

ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

CI 124 SC 124.5.1 P 65 L 13 # 97

Nicholl, Gary Cisco Systems

Comment Type ER Comment Status A (bucket1)

Missing editing instruction to update the title of Figure 124-2 from "Block diagram for 400GBASE-DR4 transmit/receive paths" to "Block diagram for 400GBASE-DR4 or 400GBASE-DR4-2 transmit/receive paths"

SuggestedRemedy

Change the title of Figure 124-2 from "Block diagram for 400GBASE-DR4 transmit/receive paths" to "Block diagram for 400GBASE-DR4 or 400GBASE-DR4-2 transmit/receive paths"

Response Response Status C

ACCEPT.

CI 124 SC 124.5.4 P 65 L 49 # 98

Nicholl, Gary Cisco Systems

Comment Type ER Comment Status A (bucket1)

Missing comma after "400GBASE-DR4-2"

SuggestedRemedy

Add missing comma after " 400GBASE-DR4-2"

Response Response Status C

ACCEPT.

CI 124 SC 124.7.1 P 68 L 44 # 99

Nicholl, Gary Cisco Systems

Comment Type TR Comment Status A table format

Table 124-6. The row for "Outer Optical Modulation Amplitude (OMAouter), each lane (min)b" for 400GBASE-DR4 is different from what we did for 100GBASE-DR in Table 140-6 in 3cu. I am not sure it is correct to add "for TDECQ < 3.4 dB" as the value of OMA (min) is dependent on the value of TDECQ and is not flat across the board at "-0.8dBm"

SuggestedRemedy

I would suggest using the same format for 400GBASE-DR4 that was used for 100GBASE-DR in Table 140-6 of 802.3cu.

Response Response Status C

ACCEPT IN PRINCIPLE.

The row for OMAouter is formatted correctly.

However, in Table 124-6 for Launch power in OMAouter minus TDECQ, each lane (min) value of -2.2 should only apply to 400GBASE-DR4.

Fix the table with editorial license.

CI 124 SC 124.7.1 P 68 L 47 # 100

Nicholl, Gary Cisco Systems

Comment Type TR Comment Status A (bucket1)

Table 124-6. The row "Launch power in OMAouter minus TDECQ, each lane (min)" only applies to 400GBASE-DR4 and not to 800GBASE-DR8.

SuggestedRemedy

Correct this row in accordance with the comment to indicate that is row only applies to 400GBASE-DR\$ and not to 800GBASE-DR8. It should look more like the "TDECQ – 10log10(Ceq)c (max)" row on line 52.

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement proposed remedy with editorial license.

IEEE P802.3df D1.0 1st Task Force review comments

CI 124 SC 124.7.1 P 69 L 15 # 101
 Nicholl, Gary Cisco Systems
 Comment Type ER Comment Status A (bucket1)
 Table 124-6. Why are the rows "Transmitter overshoot and undershoot (max)", Transmitter power excursion (max) and "Transmitter transition time (max)" all in italic ?
 SuggestedRemedy
 Change the font of the text in the rows mentioned in the comment to standard table text font.
 Response Response Status C
 ACCEPT.

CI 124 SC 124.7.1 P 69 L 29 # 102
 Nicholl, Gary Cisco Systems
 Comment Type TR Comment Status A (bucket1)
 Table 124-6. Footnote "b" only applies to 400GBASE-DR4
 SuggestedRemedy
 Update footnote b to make it clear this footnote only applies to 400GBASE-DR4 (see what was done in Table 140-6 in 3cu as an example).
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Implement proposed remedy with editorial license.

CI 124 SC 124.7.2 P 71 L 29 # 103
 Nicholl, Gary Cisco Systems
 Comment Type TR Comment Status R table format
 Table 124-7. The row "Receiver sensitivity (OMAouter), each lane (max)" for 400GBASE-DR4 is different than what was done for 100GBASE-DR in Table 140-7 in 3cu, and I am not sur eit is technically correct.
 SuggestedRemedy
 Remove "for TDECQ < 3.4 dB" for the row for 400GBASE-DR4, to follow the same format that was used for 100GBASE-DR in Table 140-7 in 802.3cu.
 Response Response Status C
 REJECT.
 The format commented to is not broken for 400GBASE-DR4. It was considered an improvement to the format used in the base standard.

CI 124 SC 124.7.3 P 72 L 40 # 104
 Nicholl, Gary Cisco Systems
 Comment Type ER Comment Status A (bucket1)
 The comma after "400GBASE-DR4" should be underlined.
 SuggestedRemedy
 Underline the comma after "400GBASE-DR4".
 Response Response Status C
 ACCEPT.

CI 124 SC 124.11.1 P 79 L 20 # 105
 Nicholl, Gary Cisco Systems
 Comment Type T Comment Status A reflections
 Table 124.11. Why would the optical return loss be any different between DR4/DR8 and DR4-2/DR8-2 ? Don't they both use the same MPO connector. The value of 25dB for DR4-2/DR8-2 appears to have been copied over from 100GBASE-FR1 in 802.3cu, but isn't FR1 using a different optical connector (LC versus MPO).
 SuggestedRemedy
 This is more of a question for clarification.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 In Table 124-11 change 25 dB for Optical return loss (min) to 37 dB.

CI 124 SC 124.11.3.1.1 P 80 L 38 # 106
 Nicholl, Gary Cisco Systems
 Comment Type TR Comment Status A (bucket1)
 Figure 124-6 indicates a different lane assignment for 400GBASE-DR4 than is in Clause 124 of the published version of the 802.3 standard. This would appear to make 400GBASE-DR4 incompatible with the current published standard.
 SuggestedRemedy
 Change the lane assignment in Figure 124-6 in 802.3df D1.0 to match the lane assignment in Figure 124-6 of "P802.3_D3p2".
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #14.

IEEE P802.3df D1.0 1st Task Force review comments

CI 172 SC 172.2.4.1 P 164 L 28 # 107

Nicholl, Shawn

AMD

Comment Type TR Comment Status D OTN reference point

The NOTE says "The stream of 66-bit blocks generated by this process". However, there are two streams generated in the above process. It would be clearer if the end of the sub-clause represented the end of the process and aligned with the OTN reference point in the note.

Also, it would be clearer for the text related to tx_coded<65:0> to coincide with the end of the sub-clause (i.e. for that text to follow any discussion related to rate compensation).

Also, where possible it is helpful to re-use text from 802.3-2022 Clause 119.2.4.1 as it enhances readability (i.e. simplifies compare/contrast between Clause 119 and Clause 172).

Suggested Remedy

Propose the following text:

172.2.4.1 Encode and rate matching

The transmit PCS generates 66-bit blocks based upon the TXD<63:0> and TXC<7:0> signals received from the 800GMII. One 800GMII data transfer is encoded into one 66-bit block. If the transmit PCS spans multiple clock domains, it may also perform clock rate compensation via the deletion of idle control characters or sequence ordered sets or the insertion of idle control characters.

Idle control characters or sequence ordered sets are removed, if necessary, to accommodate the insertion of the alignment markers. See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules, and 172.2.4.5 for more details on alignment markers.

The transmit PCS generates blocks as specified in the transmit state diagram as shown in Figure 119-14. The contents of each 66-bit block are contained in a vector tx_coded<65:0>. tx_coded<1:0> contains the sync header and the remainder of the bits contain the payload.

NOTE: The stream of tx_coded<65:0> 66-bit blocks generated by this process, together with the FEC_degraded_SER and rx_local_degraded bits should be used as the reference signal for mapping to OTN.

172.2.4.1 66B/66B block distribution

The stream of tx_coded<65:0> 66-bit blocks are distributed to the two flows in a round robin fashion by the block distribution function such that the first 66-bit block is sent to flow 0, the second 66-bit block is sent to flow 1, the third 66-bit block is sent to flow 0, and subsequent 66-bit blocks continue in a round robin distribution procedure across the two flows. This forms two streams, tx_coded_flow0<65:0> and tx_coded_flow1<65:0>.

172.2.4.3 64B/66B to 256B/257B transcoder

The 64B/66B to 256B/257B transcoder in each flow is identical to that specified in 119.2.4.2. The transcoder for flow 0 uses the stream of tx_coded_flow0<65:0> 66-bit blocks. The transcoder for flow 1 uses the stream of tx_coded_flow1<65:0> 66-bit blocks.

172.2.4.4 Scrambler

<This Comment Proposes no Changes to Text inside this Sub-Clause>

172.2.4.5 Alignment marker mapping and insertion

<This Comment Proposes no Changes to Text inside this Sub-Clause>

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 172 SC 172.2.4.4 P 164 L 47 # 108

Nicholl, Shawn

AMD

Comment Type TR Comment Status A AM sync

The bullet that says: "The first 66-bit block of the 257-bit transcoded block following the alignment marker ..." may be open to misinterpretation.

Suggested Remedy

Propose the following text:

Let tx_coded_j<65:0> and tx_coded_k<65:0> represent two consecutive blocks in the tx_coded<65:0> stream. Notably, tx_coded_j<65:0> belongs to tx_coded_flow0<65:0> stream. And, tx_coded_k<65:0> belongs to tx_coded_flow1<65:0> stream.

Let tx_coded_j<65:0> represent the first 66-bit block of the 257-bit transcoded block following the alignment marker group in flow 0. It is required that tx_coded_k<65:0> shall be the first 66-bit block of the 257-bit transcoded block following the alignment marker group in flow 1.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #90.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.6.2.2 P 169 L 11 # 109

Nicholl, Shawn AMD
 Comment Type TR Comment Status R (bucket1)

Missing any mention of 800GBASE-R.

SuggestedRemedy

For consistency with 119.2.6.2.2, propose to replace text "with x = 0:31" with text "with x = 0:31 for 800GBASE-R."

Response Response Status C

REJECT.
 The proposed change is not necessary since Clause 172 is only for 800GBASE-R. CL119 specified 200GBASE-R or 400GBASE-R because the same clause includes the PCS for both 200GE and 400GE.

Cl FM SC FM P 1 L 10 # 110

Dawe, Piers Nvidia
 Comment Type E Comment Status R (bucket1)

"Amendment:" - there should be an amendment number here. According to pages 13 and 14, this would be number 10. But 9 amendments before a revision is too many so there should be another roll-up and this could be amendment 1 of 802.3-2023.

SuggestedRemedy

Insert number or placeholder. Also on pages 11 and 27. Add it on page 14. If some amendment numbers including this one are provisional, that can be stated.

Response Response Status C

REJECT.
 As the comment alludes, the amendment number that will be assigned to this amendment is not known at this time with any certainty. An amendments number may be inserted once a number is known with better certainty, likely near the end of WG Ballot.

Cl FM SC FM P 1 L 30 # 111

Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)

Media Access Control Parameters for 800 Gb/s and Physical Layers and Management Parameters for 400 Gb/s and 800 Gb/s Operation. Draft D1.0 is prepared for task force preview

SuggestedRemedy

Media Access Control parameters for 800 Gb/s and Physical Layers and management parameters for 400 Gb/s and 800 Gb/s operation. Draft D1.0 is prepared for Task Force preview

Response Response Status C

ACCEPT IN PRINCIPLE.
 The comment appears to be pointing out that capitalization on some words need s to be corrected.
 Change: "Media Access Control Parameters for 800 Gb/s and Physical Layers and Management Parameters for 400 Gb/s and 800 Gb/s Operation. Draft D1.0 is prepared for task force preview"
 To:
 "Media Access Control parameters for 800 Gb/s and Physical Layers and management parameters for 400 Gb/s and 800 Gb/s operation. Draft D1.0 is prepared for Task Force preview"
 Implement with editorial license.

Cl FM SC FM P 6 L 39 # 112

Dawe, Piers Nvidia
 Comment Type E Comment Status R (bucket1)

The superscript 3 should follow IEEE Xplore, not "contact IEEE."

SuggestedRemedy

Get the template at <https://standards.ieee.org/develop/drafting-standard/resources/> fixed and implement the change.

Response Response Status C

REJECT.
 This footnote location is the same as in the cited template. This text is an official statement copied to the IEEE 802.3 template from the IEEE SA template. According to the 2021 IEEE SA Standards Style Manual , this text "Shall not be altered."

IEEE P802.3df D1.0 1st Task Force review comments

CI **FM** SC **FM** P **10** L **1** # **113**
 Dawe, Piers Nvidia
 Comment Type **E** Comment Status **A** (bucket1)
 "When the IEEE-SA Standards Board": duplicate section
 SuggestedRemedy
 Remove
 Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 The group of text starting with "When the IEEE-SA Standards Board:" is repeated twice.
 Remove one instance.
 Implement with editorial license.

CI **FM** SC **FM** P **27** L **48** # **114**
 Dawe, Piers Nvidia
 Comment Type **E** Comment Status **A** (bucket1)
 3bj and 3bk!! They were approved in 2013 and 2014. 3cy uses 3cx and 3cz as its examples, 3cz uses 3dd, 3cs, 3db, 3ck, 3de and 3cx
 SuggestedRemedy
 Instead of or as well as this bad example, list all the exact amendments and drafts that this draft is built against, as P802.3cz does. Also, say which drafts affect this draft and which are believed not to, preferably clause by clause. The editors must have and agree this information; no reason not to share it with the volunteers who do the review work, and the staff editors.
 Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 The example projects listed are indeed obsolete. This example list from the FrameMaker template needs to be updated for each project and may again change as previous amendments are incorporate into a revision. The examples are not really required so these examples should be deleted here and in the template.
 Delete "(e.g., IEEE P802.3bj and IEEE P802.3bk)"

CI **1** SC **1.4** P **18** L **47** # **115**
 Dawe, Piers Nvidia
 Comment Type **E** Comment Status **R** (bucket1)
 This project is adding another page of definitions to a very long section that doesn't have the usual pdf bookmarks.
 SuggestedRemedy
 To mitigate the deterioration of document structure and usability, divide 1.4 Definitions into subclauses, e.g.
 1.4.1 1 to 8
 1.4.2 A to G
 1.4.3 H to M
 1.4.4 N to S
 1.4.5 T to Z
 If Frame can deliver 1.4.0 ... 1.4.8 1.4.A ... 1.4.Z (some such as 1.4.3 are not needed), that would be even more user-friendly.

Response Response Status **C**
 REJECT.
 The comment is asking for broad changes to the base standard that are not related directly the new content that is being added by this amendment. Such sweeping changes should be addressed using the Base Standard maintenance process.

CI **1** SC **1.5** P **30** L **30** # **116**
 Dawe, Piers Nvidia
 Comment Type **E** Comment Status **R** (bucket1)
 This project is adding to an already long section that lacks the usual level of subdivision (somewhere around one subclause per page would be normal)
 SuggestedRemedy
 To mitigate the deterioration of document structure and usability, divide 1.5 Abbreviations into several subclauses
 Response Response Status **C**
 REJECT.
 The comment is asking for broad changes to the base standard that are not related directly the new content that is being added by this amendment. Such sweeping changes should be addressed using the Base Standard maintenance process.

IEEE P802.3df D1.0 1st Task Force review comments

CI 45 SC 45.2.1.6 P 36 L 20 # 117
 Dawe, Piers Nvidia
 Comment Type T Comment Status A (bucket1)
 Where possible, entries should be in the standard order: slow to fast, short to long, wide to narrow. Here, we have to read upwards because the entries are listed backwards.
 SuggestedRemedy
 Swap VR8 and SR8
 Response Response Status C
 ACCEPT.

CI 45 SC 45.2.1.7.4 P 37 L 23 # 118
 Dawe, Piers Nvidia
 Comment Type T Comment Status A (bucket1)
 Missing entries in transmit fault, receive fault and transmit disable tables
 SuggestedRemedy
 Include rows for
 100GBASE-VR1, 100GBASE-SR1, 200GBASE-VR2, 200GBASE-SR2, 400GBASE-VR4,
 400GBASE-SR4, 800GBASE-VR8, 800GBASE-SR8
 and
 400GBASE-DR4, 400GBASE-DR4-2, 800GBASE-DR8, 800GBASE-DR8-2
 Revise the rubrics.
 Response Response Status C
 ACCEPT.

CI 45 SC 45.2.1.161 P 41 L 34 # 119
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 Lane 0 maps to register 1.1120, lane 1 maps to register 1.1121, lane 2 maps to register 1.1122, and lane 3 maps to register 1.1123.
 SuggestedRemedy
 Lane 0 maps to register 1.1120, lane 1 maps to register 1.1121, and so on, up to lane 7 and register 1.1127.
 Similarly in 45.2.1.163, 45.2.1.165, 45.2.1.167
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #45.

CI 45 SC 45.2.1.168 P 42 L 38 # 120
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 "for PMD lane 1; etc.": a bit terse and informal
 SuggestedRemedy
 Suggested rewording: Register 1.1450 controls the PMD training pattern for PMD lane 0, register 1.1451 controls the PMD training pattern for PMD lane 1, and so on, up to register 1.1457 and PMD lane 7.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #32.

CI 45 SC 45.2.1.168 P 42 L 41 # 121
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 92.7.12 and 136.8.11.1.3
 SuggestedRemedy
 92.7.12, 136.8.11.1.3, or 162.8.11.1 as appropriate
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change the first sentence of the second paragraph of 45.2.1.168 so it reads as "Register bits 12:11 contain a 2-bit identifier that selects the polynomial used for training a particular PMD lane as described in 92.7.12, 136.8.11.1.3, or 162.8.11.1."

IEEE P802.3df D1.0 1st Task Force review comments

CI 45 SC 45.2.1.168 P 42 L 24 # 122

Dawe, Piers

Nvidia

Comment Type TR Comment Status A PRBS seed (bucket1)

This says "The polynomial identifier for each lane should be unique; two physically adjacent lanes having the same identifier could impair operation of the PMD control function."

This is in a section defining the meanings of bits in a memory map. The memory map serves the sublayer, not the other way round. Advice about signal integrity should be in the clause concerned.

With only four polynomials and eight lanes, the polynomials themselves can't all be different, but that's OK. Impairment is very unlikely unless adjacent lanes use the same polynomial AND the PRBS13Qs in the training pattern are aligned in time with each other. We have written generations of PMD and AUI clauses that use the same pattern on multiple lanes, but they should be skewed, e.g. 120G.3.2.2: "For the case where PRBS13Q or PRBS31Q are used with a common clock, there is at least 31 UI delay between the patterns on one lane and any other lane, so that the symbols on each lane are not correlated." The training frame is 98.3% PRBS13Q. In principle, one could incur the risk warned against with a lane carrying "identifier_i" = 0 and an adjacent lane carrying "identifier_i" = 4, with an unlucky timing offset between lanes. As "The PMD shall implement one instance of the PMD control function described in 136.8.11 for each lane", the state machine for each lane can be started and restarted asynchronous to adjacent lanes, so starting the training pattern with a different seed won't solve the issue. The text "For 8-lane use cases different initial seeds should be used where the same polynomial is being reused" recommends a course of action that, on investigation, doesn't address the issue. We should tell the reader what to avoid, not how to avoid it.

Also, the ETC spec has already covered this ground. It uses the same four polynomials and seeds, twice over. No implementation can follow the ETC spec AND this draft (because the default seeds differ) but there is no benefit in the difference.

SuggestedRemedy

1. Put signal integrity recommendations in the spec, not in the register definitions for a memory map!
2. Change "The polynomial identifier for each lane should be unique; two physically adjacent lanes having the same identifier could impair operation of the PMD control function" to "The polynomial identifier for adjacent lanes should be unique to avoid a risk of impairment of the PMD control function".
3. Change "For 8-lane use cases different initial seeds should be used where the same polynomial is being reused." to "For 8-lane use cases, see 162.8.11.1."
4. Make the default seeds in Table 162-10a the same as in the ETC spec (seeds 4 to 7 are the same as seeds 0 to 3).
5. ETC say "it is recommended to ensure that physically adjacent lanes do not use the same polynomial". Recommend this.
6. Also, suggest that when there are more lanes than polynomials to use, significant correlation between any lanes can be avoided by a combination of seed and timing offset. Leave it to the implementer to choose how to do this.

Response

Response Status C

ACCEPT IN PRINCIPLE.

Replace "The polynomial identifier for each lane should be unique; two physically adjacent lanes having the same identifier could impair operation of the PMD control function. The default identifiers are (binary): for lane 0, 00; for lane 1, 01; for lane 2, 10; for lane 3, 11; for lane 4, 00; for lane 5, 01; for lane 6, 10; for lane 7, 11. For 8-lane use cases different initial seeds should be used where the same polynomial is being reused."

with

"The polynomial identifier for adjacent lanes should be unique to avoid a risk of impairment of the PMD control function. If the same polynomial identifier is used for multiple lanes, different initial seeds should be used for each of those lanes. The default identifiers are (binary): for lane 0, 00; for lane 1, 01; for lane 2, 10; for lane 3, 11; for lane 4, 00; for lane 5, 01; for lane 6, 10; for lane 7, 11."

The adopted baseline clearly states what the default seeds in Table 162-10a should be (see: https://www.ieee802.org/3/df/public/22_09/lusted_3df_01a_2209.pdf). A user would be able to change the default values so that the seeds for lanes 4 to 7 match 0 to 3 by writing appropriate seed values to registers 1.1450 through 1.1457. Therefore it is not appropriate to change Table 162-10a.

See also the response to comment #139

CI 124 SC 124 P 59 L 36 # 123

Dawe, Piers

Nvidia

Comment Type T Comment Status R objectives

Wondering if we want 200GBASE-DR2 and 200GBASE-DR2-2. Will people connect 200G-class servers with copper or MMF only until 200GBASE-DR1 is cheaper or they move on to 400G-class servers?

SuggestedRemedy

If people will want to connect 200G-class servers with SMF, perhaps to a CPO switch, before 200GBASE-DR1 is cheaper, then it will happen. If it will happen, it would be best to include it so that it gets official code points.

Response

Response Status C

REJECT.

The comment is proposing the addition of two PMD types which are not in scope for the IEEE P802.3df project.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 124 SC 124.1 P 61 L 36 # 124
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 400GBASE-DR4, 400GBASE-DR4-2
 SuggestedRemedy
 400GBASE-DR4 and 400GBASE-DR4-2
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #38.

Cl 124 SC 124.2 P 62 L 13 # 125
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 six paragraphs 124.2
 SuggestedRemedy
 six paragraphs in 124.2
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change the instruction to:
 "Change the first six paragraphs in 124.2 as follows:"

Cl 124 SC 124.2 P 62 L 40 # 126
 Dawe, Piers Nvidia
 Comment Type TR Comment Status R PCSL interleaving (CC)
 The unlikely case of defective transition density is far more significant than the very modest difference between 2-way and 4-way RS-FEC interleaving. If we are going to break precedent and abandon unrestricted bit-multiplexing, transition density is the first thing to get right, always. With 100G AUI lanes, the Tx silicon can ensure the problem doesn't happen, and we are not mandating 50G/lane AUIs for 800G. We have had some years after this problem was discovered before 800G designs, so it should not be happening now. Let's say so.

SuggestedRemedy
 Change "See NOTE at the end of 120.5.2 concerning the transition density of lanes operating at this nominal signaling rate." to "For 400GBASE-DR4 and 400GBASE-DR4-2, see NOTE at the end of 120.5.2 concerning the transition density of lanes operating at this nominal signaling rate. For 800GBASE-DR8 and 800GBASE-DR8-2, see 173.4.2." Similarly in 124.7.2.
 In 173.4.2, say that unlike in 120, it is the transmit side PCS and PMA's responsibility to avoid the defective transition density, and give some recommendations.
 See other comments.

Response Response Status C
 REJECT.
 Resolve using the response to comment #166.

Cl 124 SC 124.7.2 P 70 L 36 # 127
 Dawe, Piers Nvidia
 Comment Type TR Comment Status R PCSL interleaving (CC)
 The unlikely case of defective transition density is far more significant than the very modest difference between 2-way and 4-way RS-FEC interleaving and we have the opportunity now to exclude it for 800G PMDs (see another comment).

SuggestedRemedy
 As elsewhere: change "See NOTE at the end of 120.5.2 concerning the transition density of lanes operating at this nominal signaling rate." to "For 400GBASE-DR4 and 400GBASE-DR4-2, see NOTE at the end of 120.5.2 concerning the transition density of lanes operating at this nominal signaling rate. For 800GBASE-DR8 and 800GBASE-DR8-2, see 173.4.2."
 In 173.4.2, say that unlike in 120, it is the transmit side PCS and PMA's responsibility to avoid the defective transition density, and give some recommendations.

Response Response Status C
 REJECT.
 Resolve using the response to comment #166.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 124 SC 124.7.2 P71 L 30 # 128
 Dawe, Piers Nvidia
 Comment Type E Comment Status A TECQ
 TDECQ
 SuggestedRemedy
 SECQ (as in 124.8.9.1), three times
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 For parameter, Receiver sensitivity (OMAouter), each lane (max), change first occurrence of TDECQ to SECQ and 2 further occurrences to TECQ.
 Implement with editorial license.

Cl 124 SC 124.8.1 P75 L 4 # 129
 Dawe, Piers Nvidia
 Comment Type E Comment Status A test pattern (CC)
 800G scrambled idle isn't in 119.2.4.9: different rate, different PCS. See another comment.
 SuggestedRemedy
 In Table 124-9, after 119.2.4.9, add "or 172.2.4.9"
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license.

Cl 124 SC 124.8.5 P76 L 5 # 130
 Dawe, Piers Nvidia
 Comment Type E Comment Status A TX test
 This says "The 400GBASE-DR4-2 or 800GBASE-DR8-2 transmitter is tested using an optical channel that meets the requirements for 100GBASE-FR1 in 140.7.5.2" but these PMDs have an optical return loss tolerance of 21.4 while 100GBASE-FR1 uses an optical return loss of 17.1 dB. The cable plant is different (array connectors are angled).
 SuggestedRemedy
 Change
 The 400GBASE-DR4-2 or 800GBASE-DR8-2 transmitter is tested using an optical channel that meets the requirements for 100GBASE-FR1 in 140.7.5.2.
 to
 The 400GBASE-DR4-2 or 800GBASE-DR8-2 transmitter is tested using an optical channel with dispersion and insertion loss as for 100GBASE-FR1 in 140.7.5.2, and optical return loss at the maximum for optical return loss tolerance in Table 124-6.

Response Response Status C
 ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

Cl 124 SC 124.11.1 P79 L 20 # 131
 Dawe, Piers Nvidia
 Comment Type E Comment Status A reflections
 These fiber optic cabling characteristics for 400GBASE-DR4-2 and 800GBASE-DR8-2 are not in the baseline, but are the same as for 100GBASE-FR1. The optical return loss should not follow FR1, as the optical return loss tolerance is significantly different and the table of discrete reflectances is different.
 SuggestedRemedy
 Adjust the optical return loss as necessary to be consistent with the adopted optical return loss tolerance and table of discrete reflectances.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #105.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 124 SC 124.11.2.2 P79 L 43 # 132
 Dawe, Piers Nvidia
 Comment Type T Comment Status R reflections
 Part of the baselines is missing. Both baselines have a table of discrete reflectances above 55 dB
 SuggestedRemedy
 Add this (these) as a new column(s) in Table 124-9
 Response Response Status C
 REJECT.
 The baselines previously agreed were not correct with respect to the table of discrete reflectances for the 2km interfaces. Table 124-13 in the in-force Clause 124 has not been modified (as such not shown in P802.3-df D1.0) and is therefore valid for all 4 DR PMD types.
 No changes to the draft are required.

Cl 124 SC 124.11.2.2 P79 L 43 # 133
 Dawe, Piers Nvidia
 Comment Type E Comment Status R reflections
 It seems odd that the table of discrete reflectances above 55 dB for 800GBASE-DR8 in the baseline is not the same as the existing table for 400GBASE-DR4, but it is the same as 400GBASE-DR4-2 and 800GBASE-DR8-2.
 SuggestedRemedy
 Reconcile the tables for 400GBASE-DR4 and 800GBASE-DR8
 Response Response Status C
 REJECT.
 Resolve using the response to comment #132.

Cl 124 SC 124.11.3.1.1 P 80 L 33 # 134
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 TxTxTxTxRxRxRxRxRx
 SuggestedRemedy
 Should look like the base doc
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #14.

Cl 124 SC 124.11.3.1.2 P 80 L 47 # 135
 Dawe, Piers Nvidia
 Comment Type TR Comment Status A fiber connector (CC)
 This says to use a single-row 16-fiber interface. But this is not in welch_3df_01a_220222, and 8 x100G SMF modules already exist with 2 x 12-way angled connectors.
 SuggestedRemedy
 Change to 2 x 12-way angled connectors.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Members of the task force have indicated that only 16-fiber connectors are being used and the connectors are angled. The editor's note states that this addition connection information was added by the editorial team for completeness.
 In 124.11.3.3 replace the second paragraph with:
 "The MDI adapter or receptacle shall meet the dimensional specifications for interface 7-4-7: MPO adaptor interface – Opposed keyway configuration or interface 7-4-9: MPO active device receptacle, angled interface for 16 fibers, as defined in IEC 61754-7-4. The plug terminating the optical fiber cabling shall meet the dimensional specifications of interface 7-4-1: MPO female plug, down-angled interface for 16 fibers. The MPO-16 female plug connector and MDI are structurally similar to those depicted in Figure 124-7, but with an angled end facet, 16 fibers, an offset keyway, and different pin diameters and locations."
 Implement with editorial license.

Cl 124 SC 124.11.3.1.2 P 80 L 50 # 136
 Dawe, Piers Nvidia
 Comment Type E Comment Status R fiber connector (CC)
 "The transmit optical lanes occupy the leftmost eight positions. The receive optical lanes occupy the rightmost eight positions": as there are only 12 positions, "most" is not really applicable.
 SuggestedRemedy
 Change to "The transmit optical lanes occupy the eight positions on the left. The receive optical lanes occupy the eight positions on the right."
 Response Response Status C
 REJECT.
 The proposed changes do not improve the accuracy or clarity of the draft. There are 16 positions.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 162 SC 162.8.11.1 P 92 L 8 # 137

Dawe, Piers

Nvidia

Comment Type T Comment Status A PRBS seed

the state of the PRBS generator shall be set to a value in the variable - eh? If the variable is a 13-bit seed, it contains 0s and 1s.

SuggestedRemedy

Rewrite for clarity

Response Response Status C

ACCEPT IN PRINCIPLE.

In the second paragraph of 162.8.11.1, change from

"the state of the PRBS generator shall be set to a value in the variable seed_i."

to

"the state of the PRBS generator shall be set to the value of the variable seed_i."

Cl 162 SC 162.8.11.1 P 92 L 9 # 138

Dawe, Piers

Nvidia

Comment Type T Comment Status A PRBS seed

The variable seed_i is not defined. 136.8.11.1.3 says "The default value of seed_i shall be the value given in Table 136-8 for p = i," but neither p nor Table 136-8 apply here. Maybe they should?

SuggestedRemedy

If the seed bits in Table 162-10a are the defaults for seed_i, say so.

Response Response Status C

ACCEPT IN PRINCIPLE.

In the third paragraph of 162.8.11.1, change "the default seed for each lane" to "the default value of seed_i for each lane i".

In Table 162-10a, change the heading of the fourth column from "Default seed bits" to "Default seed_i".

Cl 162 SC 162.8.11.1 P 92 L 29 # 139

Dawe, Piers

Nvidia

Comment Type TR Comment Status R PRBS seed

Dedault seeds 4 to 7 are different to seeds 0 to 3, contrary to the ETC 800G spec. No implementation can follow the ETC spec AND this draft (because the default seeds differ) but there is no benefit in the difference.

We have written generations of PMD and AUI clauses that use the same pattern on multiple lanes, but they should be skewed, e.g. 120G.3.2.2: "For the case where PRBS13Q or PRBS31Q are used with a common clock, there is at least 31 UI delay between the patterns on one lane and any other lane, so that the symbols on each lane are not correlated." The training frame is 98.3% PRBS13Q. In principle, one could incur the risk warned against with a lane carrying "identifier_i" = 0 and an adjacent lane carrying "identifier_i" = 4, with an unlucky timing offset between lanes. As "The PMD shall implement one instance of the PMD control function described in 136.8.11 for each lane", the state machine for each lane can be started and restarted asynchronous to adjacent lanes, so starting the training pattern with a different seed won't solve the issue.

SuggestedRemedy

1. Make the default seeds in Table 162-10a the same as in the ETC spec (seeds 4 to 7 are the same as seeds 0 to 3).
2. ETC say "it is recommended to ensure that physically adjacent lanes do not use the same polynomial". Recommend this.
4. Also, point out that significant correlation between any lanes can be avoided by a combination of seed and timing offset. Leave it to the implementer to choose how to do this.

Response Response Status C

REJECT.

Aligning an IEEE standard with a previously published document may be preferable where possible, but it is not always done.

The default seed values were explicitly set by the adopted baseline proposal https://www.ieee802.org/3/df/public/22_09/lusted_3df_01a_2209.pdf, which included a detailed description, and was approved by unanimous consent.

The seed values are not normative, and using non-default values is permitted, so there is no compliance concern.

The content of item 2 and 4 of the suggested remedy is covered by text in 45.2.1.168 ("should" is a recommendation).

IEEE P802.3df D1.0 1st Task Force review comments

Cl 162 SC 162.9.4 P 93 L 17 # 140

Dawe, Piers Nvidia
 Comment Type E Comment Status A rate range

"For an 800GBASE-CR8 PMD or for a 100GBASE-CR1, 200GBASE-CR2, or 400GBASE-CR4 PMD in the same package as the PCS sublayer": it's very easy to misunderstand this.

SuggestedRemedy

At least put a comma after "CR8 PMD". Also in 163.9.2.

Response Response Status C

ACCEPT IN PRINCIPLE.

The text intentionally distinguishes between 800GAUI-8, for which the range is always +/- 50 PPM, and the other interfaces, for which it is conditional.

The text can be clarified.

In Table 162-11 change the first sentence in footnote a to the following:
 "For 100GBASE-CR1, 200GBASE-CR2, or 400GBASE-CR4 PMD with a PMA in the same package as the PCS sublayer or for any 800GBASE-CR8 PMD."

In Table 163-5 change the first sentence in footnote a to the following:
 "For 100GBASE-KR1, 200GBASE-KR2, or 400GBASE-KR4 PMD with a PMA in the same package as the PCS sublayer or for any 800GBASE-KR8 PMD."

Cl 162 SC 162.9.5 P 93 L 36 # 141

Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)

This text is an informative NOTE in the standard in force, as below. While I can see the reason to make it normative for the transmitter, for the receiver this information about transmitter behaviour is explanation, not something the receiver does.

SuggestedRemedy

Change it from a normative table footnote to an informative table note. Similarly for 163.9.3.

Response Response Status C

ACCEPT.

Cl 167 SC 167.5.1 P 111 L 7 # 142

Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)

Strange to talk about 800G before 100G and 200G: not the usual order (slow MAC to fast MAC).

SuggestedRemedy

The block diagrams for 100GBASE-VR1 and 100GBASE-SR1 are equivalent to Figure 167-2, but for one lane per direction. The block diagrams for 200GBASE-VR2 and 200GBASE-SR2 are equivalent to Figure 167-2, but for two lanes per direction. The block diagrams for 800GBASE-VR8 and 800GBASE-SR8 are equivalent to Figure 167-2, but for eight lanes per direction.

or

The block diagrams for 100GBASE-VR1 and 100GBASE-SR1, for 200GBASE-VR2 and 200GBASE-SR2, and for 800GBASE-VR8 and 800GBASE-SR8 are equivalent to Figure 167-2, but for one, two and eight lanes per direction respectively.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change editing instruction to "Replace the first paragraph in 167.5.1 with the following:" with the text "The PMD block diagram for 400GBASE-VR4 or 400GBASE-SR4 is shown in Figure 167-2. The block diagrams for 100GBASE-VR1 and 100GBASE-SR1 are equivalent to Figure 167-2, but for one lane per direction. The block diagrams for 200GBASE-VR2 and 200GBASE-SR2 are equivalent to Figure 167-2, but for two lanes per direction. The block diagrams for 800GBASE-VR8 and 800GBASE-SR8 are equivalent to Figure 167-2, but for eight lanes per direction."

Cl 167 SC 167.8.1 P 117 L 4 # 143

Dawe, Piers Nvidia
 Comment Type T Comment Status A test pattern (CC)

In Table 167-10, Test patterns, need a new reference for scrambled idle. See another comment.

SuggestedRemedy

Change "82.2.11 and 91, or 119.2.4.9" to "82.2.11 and 91, or 119.2.4.9, or 172.2.4.9"

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement the suggested remedy with editorial license.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 167 SC 167.8.6 P 118 L 9 # 144
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 Font problem
 SuggestedRemedy
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #194.

Cl 167 SC 167.10.3 P 122 L 8 # 145
 Dawe, Piers Nvidia
 Comment Type T Comment Status A fiber connector (CC)
 The two options for 400GBASE-SR8 were defined but we should check if the industry is still split on how to connect 8-lane MMF modules.
 SuggestedRemedy
 Check if Option B, 16-fiber interface, has traction in the industry. If it doesn't, don't include it.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #146.

Cl 167 SC 167.10.3 P 122 L 49 # 146
 Dawe, Piers Nvidia
 Comment Type TR Comment Status A fiber connector (CC)
 This says "While there has not been an adopted baseline for a 16-lane MDI the language in 167.10.3.4 (below) from 400GBASE-SR8 is a good starting point". This material was explicitly EXCLUDED from the baseline murty_3df_01a_220315.pdf "MDI and lane assignments for eight lane MMF links will be taken up in subsequent meetings." It's not as simple as just copy 400GBASE-SR8 because the industry has chosen angled connectors for 8x100G MMF.

SuggestedRemedy
 Add the 2-row x12 angled connector. If appropriate, add the x16 angled connector. If appropriate, delete the one or both "flat" (non-angled) connectors. The text might be like this (references need checking):
 The MDI adapter or receptacle shall meet the dimensional specifications for either interface 7-2-3: MPO adapter interface - opposed keyway configuration or interface 7-2-9: MPO active device receptacle, angled interface, as defined in IEC 61754-7-1. The plug terminating the optical fiber cabling shall meet the dimensional specifications of interface 7-2-1: MPO female plug connector, down-angled interface for 2 to 24 fibres, as defined in IEC 61754-7-1.
 The MDI connection shall meet the interface performance specifications of IEC 63267-1 for performance grade Bm/1m.2

IEC 63267-1 with performance grade 1m specification is available as a Pre-Release Version (PRV) Final Draft International Standard (FDIS); final published version of this specification is expected to be available in 2023.

Response Response Status C
 ACCEPT IN PRINCIPLE.
 Members of the task force have indicated that both 16 and 24 fiber connectors are being used and the 16 fiber connectors are angled and the 24 fiber connectors are flat. The editor's note states that this addition connection information was added by the editorial team for completeness.
 Replace the existing text in 167.10.3.4 with
 "The MDI shall optically mate with the compatible plug on the optical fiber cabling. 800GBASE-VR8 and 800GBASE-SR8 have two optical lane assignment options (see 167.10.3.1a).
 For option A, the MDI adapter or receptacle shall meet the dimensional specifications for interface 7-2-3: MPO adapter interface - opposed keyway configuration or interface 7-2-10: MPO active device receptacle, flat interface, as defined in IEC 61754-7-2. The plug terminating the optical fiber cabling shall meet the dimensional specifications of interface 7-2-4: MPO female plug connector, flat interface for 16 to 24 fibers, as defined in IEC 61754-7-2. The MPO female plug connector and MDI are structurally similar to those depicted in Figure 167-9, but with two rows of fibers. The MDI connection shall meet the interface performance specifications of IEC 61753-1 and IEC 61753-022-2 for performance grade Bm/2m.
 For option B, the MDI adapter or receptacle shall meet the dimensional specifications for interface 7-4-7: MPO adaptor interface – Opposed keyway configuration or interface 7-4-9:

IEEE P802.3df D1.0 1st Task Force review comments

MPO active device receptacle, angled interface for 16 fibers, as defined in IEC 61754-7-4. The plug terminating the optical fiber cabling shall meet the dimensional specifications of interface 7-4-1: MPO female plug, down-angled interface for 16 fibers. The MPO female plug connector and MDI are structurally similar to those depicted in Figure 167-9, but with 16 fibers, an offset keyway, and with different pin diameter and locations. The MDI connection shall meet the interface performance specifications of IEC 63267-1 for performance grade Bm/1m."

Cl 169 SC 169.2.4 P 130 L 33 # 147

Dawe, Piers Nvidia
 Comment Type E Comment Status R PMA description

Wow, this is too mean with the information. Compare 116.2.4: the equivalent of this is missing: "The 200GBASE-R and 400GBASE-R PMAs perform the mapping of transmit and receive data streams between the PCS and PMA via the PMA service interface, and the mapping and multiplexing of transmit and receive data streams between the PMA and PMD via the PMD service interface. In addition, the PMA performs retiming of the received data stream when appropriate, optionally provides data loopback at the PMA or PMD service interface, and optionally provides test pattern generation and checking."

SuggestedRemedy

At least say that a PMA connects the PCS and PMA via the PMA service interface, and the PMA and PMD via the PMD service interface, and that there can be more than one PMA (in series) for one MAC. It performs retiming of the received data stream when appropriate. There are optional defined physical instantiations called AUIs. And/or, at line 35, add "and a summary of its functions is given in 173.1.3".

Response Response Status C

REJECT.
 The description provided in Clause 116 was overly verbose with repeated details that are listed in the reference PMA clause. The PMA description in Clause 169 provides the general function of a PMA with similar detail provided in the other sublayer descriptions and references the relevant PMA subclauses where the reader may find all of the details relevant to each PMA type.

Cl 169 SC 169.2.5 P 130 L 50 # 148

Dawe, Piers Nvidia
 Comment Type E Comment Status A AN

Is a "linked device" defined or explained anywhere? The definition and use of "link" is a delicate area.

SuggestedRemedy

Delete "linked". In the next line, change "the link" to "a link".

Response Response Status C

ACCEPT IN PRINCIPLE.
 The language in this paragraph is consistent with similar subclause 80.2.6 (802.3-2022) and 116.2.5a (802.3ck-2022). However, the term "linked device" rather than just "device" does not seem to provide any useful information. However, the other device is the one on the same link as the local device so "the link" rather than "a link" is correct. Change "linked device" to "device".
 [Editor's note: Page changed from 131 to 130.]

Cl 169 SC 169.3.1 P 132 L 21 # 149

Dawe, Piers Nvidia
 Comment Type T Comment Status R figure lanes

In Figure 116-2, multiple lanes are shown explicitly: PMA:IS_UNITDATA_0.request
 PMA:IS_UNITDATA_1.request ... PMA:IS_UNITDATA_7.request

SuggestedRemedy

As a compromise, follow e.g. Figure 120G-2; add the short diagonal lines "n" to show n lanes, not n requests on one lane with a constant ordering. Several figures, including Fig 172-2 where showing the numbers, 16 and 32, will be helpful.

Response Response Status C

REJECT.
 A single line with an SI parameter with vector notation clearly conveys the fact that there are multiple lanes 0 to n-1. This approach is used to reduce the clutter compared to similar diagrams in Clause 116. This approach is used consistently in various figures in 802.3df. The proposed changes do not improve the accuracy or clarity of the draft.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 169 SC 169.5 P 134 L 53 # 150
 Dawe, Piers Nvidia
 Comment Type E Comment Status R (bucket2)
 116.5 says "Skew (or relative delay) can be introduced between lanes". This says "Skew (or relative delay) can be introduced between PCS lanes" which gives a false impression that PMA and PMD lanes don't get skewed.
 SuggestedRemedy
 Delete "PCS", once.
 Response Response Status C
 REJECT.
 Skew is constrained for each sublayer to limit the net skew between PCS lanes so that the cumulative skew between PCS lanes does not exceed the ability of the specified PCS deskew function.

Cl 169 SC 169.5 P 136 L 27 # 151
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 points for single 800GAUI-n
 SuggestedRemedy
 points for a single 800GAUI-n
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 In Figure 169-4...
 Change "800GBASE-R Skew points for single 800GAUI-n"
 To: "800GBASE-R Skew points for a PHY with a single 800GAUI-n"
 In Figure 169-5...
 Change "Skew points for multiple 800GAUI-n"
 To: "Skew points for a PHY with multiple 800GAUI-n"

Cl 170 SC 170 P 141 L 1 # 152
 Dawe, Piers Nvidia
 Comment Type E Comment Status R (bucket2)
 This has got so little to say it's a waste of a clause number. The 100/200/400/800GMII is like the MAC: almost identical apart from rates, timing and optional EEE.
 SuggestedRemedy
 Merge 170 into 117 or better, merge 170 and 117 into 81.
 Response Response Status C
 REJECT.
 The comment does not provide sufficient justification to support the suggested remedy.
 The current structure of the draft is consistent with the approach taken by previous projects.

Cl 171 SC 171.4 P 152 L 18 # 153
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 activate_t
 hreshold
 SuggestedRemedy
 Make these tables full width, make the right hand columns wider, also in Clause 172. It may be necessary to set break points in these long "words". In maintenance we might change to shorter names, e.g. FEC_degraded_SER_thresh_on
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Improve appearance of the variable names with editorial license.

Cl 171 SC 171.4 P 153 L 11 # 154
 Dawe, Piers Nvidia
 Comment Type T Comment Status R (bucket2)
 Under "MDIO status variable" there is an entry "Lane 0 to 31 aligned" but this isn't a variable that indicates if lanes 0 to 31 are aligned. Table 45-350 has "Name"s Lane 0 aligned, Lane 1 aligned, and so on. Is there such a thing as an "MDIO variable" anyway? Clauses such as PCS have variables, MDIO has registers. The way of talking about such multilane things was solved long ago; e.g. "84.7.5 PMD lane-by-lane signal detect function"
 SuggestedRemedy
 Because a "variable" must be talking about one lane not the pair of registers recording 16 or 32 lanes, change "Lane 0 to 31 aligned" back to how it is in 117: "Lane x aligned" or "Lane i aligned" or better, "Lane aligned". "Lane-by-lane aligned" seems odd, but "DTE XS FEC symbol errors lane 0 to lane 31" below can be "DTE XS FEC symbol errors by lane" Similarly in several tables, also in other clauses such as 172, PCS.
 Response Response Status C
 REJECT.

There is no consensus to make a change at this time. Further work and consensus building is encouraged.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 171 SC 171.4 P 153 L 11 # 155
 Dawe, Piers Nvidia
 Comment Type T Comment Status A (bucket1)
 16 bits for 32 lanes
 SuggestedRemedy
 Need more registers
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Update Table 171-5 to align with register/bit definitions in Clause 45.

Cl 172 SC 172.1.5 P 162 L 12 # 158
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 Transcode
 SuggestedRemedy
 transcode - 4 times Also in this figure: Encode, Decode, Interleave, Lane
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Correct the capitalization with editorial license.

Cl 172 SC 172.1.1 P 160 L 11 # 156
 Dawe, Piers Nvidia
 Comment Type E Comment Status R (bucket2)
 The paragraph of introduction in 119.1.1 is missing: "Both 200GBASE-R and 400GBASE-R are based on a 64B/66B code. The 64B/66B code supports transmission of data and control characters. The 64B/66B code is then transcoded to 256B/257B encoding to reduce the overhead and make room for forward error correction (FEC). The 256B/257B encoded data is then FEC encoded before being transmitted. Data distribution is introduced to support multiple lanes in the Physical Layer. Part of the distribution includes the periodic insertion of an alignment marker, which allows the receive PCS to align data from multiple lanes."
 SuggestedRemedy
 At least refer to 172.1.3 as an introduction.
 Response Response Status C
 REJECT.
 172.1.1 is the scope and the current text is a sufficient description of the scope of the clause. All of the information noted in the comment is provided in 172.1.3 and there is no need to duplicate it in the scope.

Cl 172 SC 172.1.5 P 162 L 23 # 159
 Dawe, Piers Nvidia
 Comment Type T Comment Status A AM sync
 The baseline (shrikhande_3df_01a_221004, see slide 10) shows that the two flows' alignment insertion are connected. 172.2.1 ignores this too, although 172.2.4.4 says what to do, but it should be made obvious in the figure that a linkage is needed.
 SuggestedRemedy
 Show "Alignment insertion" across both flows as in shrikhande_3df_01a_221004, or make the point some other way such as "Synchronization" (used in the ETC 800G spec) or "alignment".
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #180

Cl 172 SC 172.1.5 P 162 L 12 # 157
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 "66B Block distribution": bits not bytes, rogue capital, style
 SuggestedRemedy
 66-bit block distribution
 also 66-bit block collection
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Replace 66B by 66-bit in Fig 172-2 in two places.

Cl 173 SC 173.3 P 179 L 17 # 160
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 another PMA or PMD
 SuggestedRemedy
 a PMD or another PMA
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change from:
 "another PMA or PMD"
 to
 "another PMA or a PMD"

IEEE P802.3df D1.0 1st Task Force review comments

Cl 173 SC 173.3 P 179 L 19 # 161
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 "defined in 169.3" but 173.2 says "defined in 169.3.1"
 SuggestedRemedy
 Reconcile
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change from "169.3" to "169.3.1"

Cl 173 SC 173.4 P 180 L 20 # 162
 Dawe, Piers Nvidia
 Comment Type T Comment Status A PMA SI
 The interface below the PMA (8 lanes) connects with either a PMD or a physically instantiated interface (800GAUI-8).
 SuggestedRemedy
 The interface below the PMA (8 lanes) either connects with a PMD or it is a physically instantiated interface (800GAUI-8) connecting to another 800GAUI-8 PMA interface in another PMA. Similarly twice more.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #196.

Cl 173 SC 173.4 P 180 L 1 # 163
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 Something strange about the page layout; these sections start to the left of the header
 SuggestedRemedy
 Reconcile
 Response Response Status C
 ACCEPT.

Cl 173 SC 173.4 P 180 L 10 # 164
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 32:8 PMA Functional Block Diagram
 SuggestedRemedy
 32:8 PMA functional block diagram - 3 figures
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 In the titles for Figure 173-3, 173-4 and 173-5, change from:
 "Functional Block Diagram"
 to
 "functional block diagram"

Cl 173 SC 173.4.1 P 183 L 44 # 165
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 The next sentence says "at the service interface below the PMA"
 SuggestedRemedy
 So, this one should say "at its service interface"
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Replace the text in 173.4.1 with the following splitting the text into two paragraphs:
 "If the interface between the PMA client and the PMA is physically instantiated as 800GAUI-8, the PMA shall meet the electrical and timing specifications as specified in Annex 120F or Annex 120G as appropriate at the PMA service interface.
 If the interface between the sublayer below the PMA and the PMA is physically instantiated as 800GAUI-8, the PMA shall meet the electrical and timing specifications as specified in Annex 120F or Annex 120G as appropriate at the service interface below the PMA."
 [Editor's note: page was changed from 180 to 183.]

IEEE P802.3df D1.0 1st Task Force review comments

CI 173 SC 173.4.2.1 P 184 L 10 # 166

Dawe, Piers

Nvidia

Comment Type TR Comment Status R PCSL interleaving (CC)

This additional constraint provides a very modest benefit that is judged not necessary in 400G Ethernet. However, the rare but much more harmful "clock content" (transition density) issue that was discovered late in P802.3bs should now be outlawed. There are many easy ways to do this.

SuggestedRemedy

Make this a recommendation "It is recommended that each of the 8 output lanes contain two unique PCSLs from PMA client lanes i = 0 to 15 and two unique PCSLs from PMA client lanes i = 16 to 31".

Add constraint: "The arrangement of lanes and their skew shall ensure that the reduced transition density described at the end of 120.5.2 does not occur."

Response Response Status C

REJECT.

The constrained PCS multiplexing specified in Clause 173 is consistent with slides 17 and 18 in the adopted PCS/PMA baseline (https://www.ieee802.org/3/df/public/22_10/22_1004/shrikhande_3df_01a_221004.pdf).

There is no evidence that clock content is worse than for four-lane 400GBASE-R PMDs lanes. We are not aware of any harmful issues with four-lane 400GBASE-R PMDs due to clock content.

Although some analysis has shown the possibility of reduced clock content, no evidence has been provided to justify further constraints.

CI 173 SC 173.4.2.2 P 184 L 37 # 167

Dawe, Piers

Nvidia

Comment Type TR Comment Status R PCSL interleaving (CC)

This is a PMA. On the receive side, it doesn't know and can't control the PCSLs of the signals it carries.

SuggestedRemedy

Replace this with a practical criterion to ensure that the reduced transition density doesn't happen, if any is needed, e.g. that each of the 8 outputs is derived from four contiguous lanes in the set of 32 incoming PMA lanes. There is negligible benefit in the 4-FEC multiplexing on the receive side because there are only PMAs that can make more errors after this, and their maximum error ratios are far lower than the PMD's.

Response Response Status C

REJECT.

The issue described in the comment is not correct.

Subclause 173.4.2.2 is specifically referring to the 8:32 PMA, which is always co-located with a PHY 800GXS below it (see 173.1.4). In the receive direction, this PMA receives 32 parallel bit streams from the PHY 800GXS. Each one of the 32 bit streams is a specific and known PCSL. The PMA is therefore able to identify the specific PCSLs it is receiving from the PHY 800GXS (from the "PHY_XS:IS_UNITDATA_0:31.indication" service interface primitive) and arrange them appropriately.

This receive direction of the 8:32 PMA is functionally identical to the transmit direction of the 32:8 PMA, where the 32:8 PMA receives 32 parallel bit streams from the 800GBASE-R PCS above it.

The constrained PCSL multiplexing can thus be performed in accordance with slides 17 and 18 in the adopted PCS/PMA baseline (https://www.ieee802.org/3/df/public/22_10/22_1004/shrikhande_3df_01a_221004.pdf).

The clock content mentioned in the suggested remedy are addressed in comments #166, 169, 126, and 127.

CI 173 SC 173.4.2.3 P 185 L 2 # 168

Dawe, Piers

Nvidia

Comment Type E Comment Status A (bucket1)

This can be made clearer.

SuggestedRemedy

Change "lane shall be mapped together to an output lane" to "lane shall be mapped to the same output lane"

Response Response Status C

ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 173 SC 173.4.2.3 P 185 L 3 # 169
 Dawe, Piers Nvidia
 Comment Type **TR** Comment Status **R** PCSL interleaving (CC)
 "The order of PCSLs from an input lane does not have to be maintained on the output lane"
SuggestedRemedy
 Is this enough to exclude the reduced transition density issue? If not, it can be tightened to require the lanes remain in the same or reversed order, not re-ordered about any old how.
 Response Response Status **C**
 REJECT.
 Resolve using the response to comment #166.

Cl 173 SC 173.4.3.5 P 185 L 49 # 170
 Dawe, Piers Nvidia
 Comment Type **E** Comment Status **R** (bucket1)
 "group of PMAs" puzzled me. PMAs are not used in parallel.
SuggestedRemedy
 Change group to series, or sequence
 Response Response Status **C**
 REJECT.
 The text is consistent with subclauses 120.5.3.5 and 83.5.3.6 in the base standard and is accurate as written. The proposed changes do not improve the accuracy or clarity of the text.

Cl 173 SC 173.4.11 P 187 L 20 # 171
 Dawe, Piers Nvidia
 Comment Type **E** Comment Status **A** precoding (CC)
 As I think 120 doesn't address precoding
SuggestedRemedy
 Does 120.5.11.2 need updating or is there a place in 135 that addresses it?
 Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 The base standard is ambiguous about whether precoding should be applied to the PAM4 patterns specified in 120.5.11.2. All patterns other than PRBS31Q are used only in transmitter tests and thus should be used without precoding enabled. The PRBS31Q pattern, which is specified for receiver stress testing, may be used with or without precoding based on AUI or PMD type and the receiver preference.
 Implement the changes on slide 16 of the following presentation:
https://www.ieee802.org/3/df/public/22_12/brown_3df_03b_2212.pdf

Cl 173 SC 173.5 P 187 L 33 # 172
 Dawe, Piers Nvidia
 Comment Type **T** Comment Status **R** (bucket1)
 "Mapping of MDIO control variables to PMA control variables is shown in Table 173–2. Mapping of MDIO status variables to PMA status variables is shown in Table 173–3." But status and control go in opposite directions.
SuggestedRemedy
 Mapping of PMA status variables to MDIO status variables is shown in Table 173–3. Similarly in next sentence.
 Response Response Status **C**
 REJECT.
 The wording is consistent with similar subclauses in multiple clauses in the base standard and is accurate as written. The proposed changes do not improve the accuracy or clarity of the text.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 173 SC 173.5 P 189 L 9 # 173
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket2)
 PRBS Tx pattern testing
 SuggestedRemedy
 PRBS Tx pattern testing error counter
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change "PRBS Tx pattern testing" to "PRBS Tx pattern testing error counter, lane 0 to lane 7"
 Change "PRBS Rx pattern testing" to "PRBS Rx pattern testing error counter, lane 0 to lane 7"

Cl 120F SC 120F P 198 L 8 # 174
 Dawe, Piers Nvidia
 Comment Type E Comment Status A clause name
 This project is lengthening this title but a five-line title is too long. If we had 16 x 100G AUIs it would be even worse.
 SuggestedRemedy
 Name it it the way we name PMD clauses:
 Chip-to-chip 100 Gb/s/lane Attachment Unit Interfaces type 100GAUI-1 C2C, 200GAUI-2 C2C, 400GAUI-4 C2C, and 800GAUI-8 C2C
 Similarly for 120G
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 The titles are indeed long and can be shortened and clarified.
 The suggested remedy introduces the word "Type", which has been used for PHY but not for AUIs. Therefore a slight modification is proposed.
 The same form used for PMD clause titles can be used.
 Change the title of Annex 120F to:
 "Chip-to-chip Attachment Unit Interfaces 100GAUI-1 C2C, 200GAUI-2 C2C, 400GAUI-4 C2C, and 800GAUI-8 C2C"
 Change the title of Annex 120G to
 "Chip-to-module Attachment Unit Interfaces 100GAUI-1 C2M, 200GAUI-2 C2M, 400GAUI-4 C2M, and 800GAUI-8 C2M"
 Change the titles of 120F.5, 120F.5.4, 120G.6, 120G.6.4, the text in 120F.5.1 and 120G.6.1, and the tables in 120F.5.2.2 and 120G.6.2.2, accordingly.
 Change any text affected by these title changes with editorial license.

IEEE P802.3df D1.0 1st Task Force review comments

CI 120F SC 120F.1 P 199 L 9 # 175
 Dawe, Piers Nvidia
 Comment Type E Comment Status A precoding (CC)
 120.5.7.2 doesn't address precoding in C2C
 SuggestedRemedy
 Delete the reference here or change 120.5.7.2
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 The content of 120.5.7.2 is not sufficiently complete for 800 Gb/s Ethernet AUIs and PMDs.
 Implement the changes on slides 14 and 15 of the following presentation:
https://www.ieee802.org/3/df/public/22_12/brown_3df_03b_2212.pdf
 Add references to 173.4.7.2 where appropriate.
 Implement with editorial license.

CI 120G SC 120G.1 P 204 L 44 # 176
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 Each 100GAUI-1, 200GAUI-2, 400GAUI-4 C2M, *and* 800GAUI-8 C2M data path contains one, two, four, *or* eight differential lanes
 SuggestedRemedy
 Change and to or
 Response Response Status C
 ACCEPT.

CI 120G SC 120G.2 P 207 L 8 # 177
 Dawe, Piers Nvidia
 Comment Type E Comment Status A test points
 As dealing with larger numbers of lanes in compliance boards is an engineering issue...
 And by the way, it might have been helpful to show that these are differential.
 SuggestedRemedy
 It would help to add the short diagonal lines showing n lanes. Also Figure 120G-4
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 The test points are separate for each lane.
 However, the clarity of the figure may be improved.
 Add the label "(one per lane)" below TP1a and TP4a in Figure 120G-3, and below TP1 and TP4 in Figure 120G-4.
 In the second and third paragraphs of 120G.2, change "the location of compliance points" to "the location of compliance points for each lane".

CI FM SC FM P 8 L 12 # 178
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 Task Force name Task Force
 SuggestedRemedy
 Task Force 3 times
 Response Response Status C
 ACCEPT.
 Delete "Task force name", three instances
 Also, add list of clause editors.
 Implement with editorial license.
 [Editor's note: The page/line were change from 1/8 to 8/12.]

IEEE P802.3df D1.0 1st Task Force review comments

CI 172 SC 172.2.1 P 163 L 19 # 179

Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)

"distributed in a round-robin fashion into two parallel transmit functions": sort of slang. Where I come from, all robins look round.

SuggestedRemedy

Change "a round-robin fashion" to "an alternating fashion" here; in 172.2.4.1, change "a round robin fashion" to "an alternating fashion". Similarly in 172.2.5.8.

Response Response Status C

ACCEPT IN PRINCIPLE.

This is a new function where 64B/66B blocks are distributed between or combined from two streams or flows, so "alternating" seems more appropriate here than "round-robin". The details of the distribution are not necessary in the summary, but in the detailed functional description "round-robin" should be replaced with "alternating".

In 172.2.1 on page 163 line 19...

Change:

"The 66-bit blocks are then distributed in a round-robin fashion into two parallel transmit functions, referred to as flow 0 and flow 1."

To:

"The 66-bit blocks are then distributed between two parallel transmit functions, referred to as flow 0 and flow 1."

In 172.2.1 on page 163, line 42

Change:

"A 66-bit block collection function merges the 66-bit blocks from the two flows in a round-robin fashion into a single stream of blocks that are then 64B/66B decoded."

To:

"A 66-bit block collection function merges the 66-bit blocks from the two flows into a single stream of blocks that are then 64B/66B decoded."

In 172.2.4.1 on page 164, line 23...

Change:

"The 66-bit blocks are distributed to the two flows in a round robin fashion by the block distribution function such that the first 66-bit block is sent to flow 0, the second 66-bit block is sent to flow 1, the third 66-bit block is sent to flow 0, and subsequent 66-bit blocks continue the round robin distribution procedure across the two flows."

To:

"The 66-bit blocks are distributed to the two flows in an alternating fashion by the block distribution function such that the first 66-bit block is sent to flow 0, the second 66-bit block is sent to flow 1, the third 66-bit block is sent to flow 0, and subsequent 66-bit blocks continue the distribution procedure across the two flows."

In 172.2.5.8 on page 168, line 21

Change: "The block collection reverses the block distribution done in the transmitter (see

172.2.4.1) by combining the 66-bit blocks from the two flows in a round robin fashion to form a single stream of 66-bit blocks."

To: "The block collection reverses the block distribution done in the transmitter (see 172.2.4.1) by combining the 66-bit blocks from the two flows in an alternating fashion to form a single stream of 66-bit blocks."

CI 172 SC 172.2.1 P 163 L 21 # 180

Dawe, Piers Nvidia
 Comment Type T Comment Status A AM sync

"Within each flow, the 66-bit blocks are transcoded to 257-bit blocks, scrambled, and alignment markers are periodically added to the data stream."

SuggestedRemedy

Modify this to say that the insertion of alignment markers is not independent for each flow.

Response Response Status C

ACCEPT IN PRINCIPLE.

The response to comment #90 modifies the text in 172.2.4.4 (Alignment marker mapping and insertion) to clarify the rules for the alignment marker insertion in the two flows, and makes it very clear that the alignment marker insertion is not independent for each flow.

Section 172.2.1 (the subject of this comment) is a high level functional overview of the PCS. It would be useful to introduce that the alignment marker insertion in the two flows is not independent as part of this high level overview.

With editorial license add text to section 172.2.1 to introduce that the alignment marker insertion in the two flows are not independent.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.1 P 163 L 22 # 181
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 The data stream is distributed to two 5140-bit blocks and then FEC encoded. The two FEC codewords are then interleaved before data is distributed to individual PCS lanes.
SuggestedRemedy
 For each flow, the data stream is distributed to two 5140-bit blocks and then FEC encoded. For each flow, the two FEC codewords are then interleaved before data is distributed to individual PCS lanes.
Response Response Status C
 ACCEPT IN PRINCIPLE.
 Replace
 "The data stream is distributed to two 5140-bit blocks and then FEC encoded. The two FEC codewords are then interleaved before data is distributed to individual PCS lanes." with
 "For each flow, the data stream is distributed to two 5140-bit blocks and then FEC encoded. For each flow, the two FEC codewords are then interleaved before data is distributed to individual PCS lanes."

Cl 172 SC 172.2.2 P 163 L 46 # 182
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket1)
 "Use of blocks" - ambiguous: there are 257-bit blocks as well as FEC blocks, even if we call those "codewords". This title dates from 49.2.3 Use of blocks, before 257-bit blocks and FEC.
SuggestedRemedy
 Change "blocks" to "66-bit blocks" here and at line 49.
Response Response Status C
 ACCEPT IN PRINCIPLE.
 Implement the suggested remedy with editorial license

Cl 172 SC 172.2.4.4 P 165 L 8 # 183
 Dawe, Piers Nvidia
 Comment Type E Comment Status R (bucket2)
 The curly brackets must be trying to tell the reader something, but I don't know what.
SuggestedRemedy
 Delete them, or define what they mean, or change to some notation that is defined.
Response Response Status C
 REJECT.
 The curly brackets in Tables 172-1 and 172-2 are consistent with what was used in Table 119-2 in Clause 119. Given that we are striving for consistency between this new and previous PCS specifications, retaining a common format is helpful for comparison.
 The comment does not provide sufficient justification to make the suggested change nor do the proposed changes improve clarity or accuracy of the draft.

Cl 172 SC 172.2.4.4 P 165 L 8 # 184
 Dawe, Piers Nvidia
 Comment Type E Comment Status A (bucket2)
 Two fifths of this table is useless clutter, and it would be good to use spaces in the normal way.
SuggestedRemedy
 Change
 0x9A,0x4A,0x26,0xB6,0x65,0xB5,0xD9,0xD9,0xFE,0x71,0xF3,0x26,0x01,0x8E,0x0C
 to
 9A, 4A, 26, B6, 65, B5, D9, D9, FE, 71, F3, 26, 01, 8E, 0C
 and so on. In the text, say that these are in hex.
 Similarly in Table 172-2.
Response Response Status C
 ACCEPT IN PRINCIPLE.
 Update the format in Table 172-1 and Table 172-2 to use fixed width font.
 There was no consensus to make any other changes.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 172 SC 172.2.4.8 P 166 L 51 # 185

Dawe, Piers

Nvidia

Comment Type T Comment Status A (bucket1)

Careful, "function" has a precise meaning in PCS clauses. This can be more specific and informative.

SuggestedRemedy

Change "The functions ... are" to "the 64B/66B encoder ... is"

Response Response Status C

ACCEPT IN PRINCIPLE.

Change from

"The functions above the "64B/66B to 256B/257B transcoder" are excluded."

to

"The portion of the figure above the "64B/66B to 256B/257B transcoder" is excluded".

Cl 172 SC 172.2.4.9 P 167 L 25 # 186

Dawe, Piers

Nvidia

Comment Type E Comment Status A test pattern (CC)

"Test-pattern generators are identical to that specified in 119.2.4.9" there is only one test pattern, and although it is generated in an analogous way to 119.2.4.9, it's a different PCS and different bits in the pattern.

SuggestedRemedy

Change to "A scrambled idle test pattern can be generated in the same way in the same way as in 119.2.4.9".

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using the response to comment #27.

Cl 172 SC 172.2.5.3 P 168 L 1 # 187

Dawe, Piers

Nvidia

Comment Type E Comment Status A (bucket2)

The relation between hi_ser_0, hi_ser_1 and hi_ser appears later within a state machine variable definition, which is too obscure. More generally, I could not find where the purpose of hi_ser is introduced.

SuggestedRemedy

Add something in regular text (possibly elsewhere) that says that what hi_ser for, and that it is the OR of hi_ser_0 and hi_ser_1.

Response Response Status C

ACCEPT IN PRINCIPLE.

Update draft to clarify the use of hi_ser in in each flow versus at the top level.

Cl 172 SC 172.2.5.8 P 168 L 33 # 188

Dawe, Piers

Nvidia

Comment Type E Comment Status R (bucket1)

This says "See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules" but those subclauses are titled "119.2.3.5 Idle (/I)" and "119.2.3.8 Ordered set (/O)" and the content isn't there so the reader doesn't know to look there, or follow the links from there to 83 to find the deletion and insertion rules.

SuggestedRemedy

Improve the titles of those subclauses: "Idle (/I) and idle insertion and deletion" and "Ordered set (/O) and ordered set deletion"

Response Response Status C

REJECT.

119.2.3.5 and 119.2.3.8 have links to 82.2.3.6 and 82.2.3.9 respectively, which the reader can follow to access the rules for insertion/deletion. Note that this double-reference is common throughout many subclauses in Clause 172. The proposed changes do not improve the accuracy of the draft.

IEEE P802.3df D1.0 1st Task Force review comments

CI 172 SC 172.3.5 P 173 L 31 # 189

Dawe, Piers

Nvidia

Comment Type TR Comment Status A fec counters

I could not find FEC_cw_counter in the base document (802.3-2022 Section 8) or the PCS baseline shrikhande_3df_01a_221004, and in 802.3ck it's for RS-FEC-Int (for 100GBASE-P PHYs 100GBASE-KR1 and 100GBASE-CR1) only. It's not applicable to any 200G or 400G, which is what the 800G PCS is based on. The same applies to 172.3.6 FEC_codeword_error_bin_i, I think.

SuggestedRemedy

Have we had the discussion as to whether we want to copy these features from a feature of a one-speed specialist PCS into a regular PCS feature that applies to any 800GBASE-R PHY?

Response Response Status C

ACCEPT IN PRINCIPLE.

The counters were added even though they are not part of the adopted PCS baseline.

Based on straw poll # 3 there is consensus to retain the FEC bin counters in the draft. Based on straw poll #4 there is more support for making these counters optional rather than mandatory.

There is consensus after the straw polls to leave the counters in the draft and to make them optional.

Reword the draft to indicate that the counters are optional.

Straw poll #3 (directional)

I support keeping the FEC bin counters (FEC_codeword_error_bin_i) and FEC cw counter (FEC_cw_counter) currently defined in 172.3.6 and 172.3.5 respectively.

Y:61

N:1

Straw poll #4 (directional)

I support the FEC bin counters (FEC_codeword_error_bin_i) and FEC cw counter (FEC_cw_counter) currently defined in 172.3.6 and 172.3.5 respectively, being:

- A: Mandatory
- B: Optional
- C: Need more information

A:25, B:36, C:7

CI 173 SC 173.1.4 P 177 L 28 # 190

Dawe, Piers

Nvidia

Comment Type T Comment Status A (bucket1)

"A ... PMA is required to support a physical instantiation of the PMA service interface": doesn't make sense, as the PMA service interface is part of the PMA. an vs. a.

SuggestedRemedy

is used to implement a ...?

Response Response Status C

ACCEPT IN PRINCIPLE.

Change from:

"An 8:8 PMA is required to support a physical instantiation of the PMA service interface (800GAUI-8)"
to
"An 8:8 PMA is required for a physical instantiation of the PMA service interface (800GAUI-8)"

CI 173 SC 173.2 P 178 L 51 # 191

Dawe, Piers

Nvidia

Comment Type T Comment Status A (bucket1)

"The PMA receives": confusing and incomplete.

SuggestedRemedy

In the transmit direction, the PMA receives 32 parallel bit streams, each at the nominal signaling rate of the PCSL. In the receive direction, it delivers 32 parallel bit streams to its client.
Similarly in the next paragraph for an 8-lane interface.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change from

"The PMA receives 32 parallel bit streams, each at the nominal signaling rate of the PCSL."
to
"In the transmit direction, the PMA receives 32 parallel bit streams from either the 800GBASE-R PCS or the DTE 800GXS, each at the nominal signaling rate of the PCSL. In the receive direction, the PMA sends 32 parallel bit streams to the PMA client, each at the nominal signaling rate of the PCSL."

Change from

"The PMA receives PAM4 symbols on each of its input lanes at two times the PCSL rate, each symbol formed from two bits."

to

"In the transmit direction, the PMA receives 8 parallel PAM4 symbol streams from the PMA client, each operating at a nominal signaling rate of 53.125 GBd. In the receive direction, the PMA sends 8 parallel PAM4 symbol streams to the PMA client, each at a nominal signaling rate of 53.125 GBd."

IEEE P802.3df D1.0 1st Task Force review comments

Cl 167 SC 167.7.1 P 114 L 10 # 192

Nicholl, Gary Cisco Systems

Comment Type E Comment Status A (bucket1)

Table 167-7. The order of the PMDs in the 'Signaling rate" row is different from what was done in Clause 124.

SuggestedRemedy

Proposing to reorder the data in this row to put the lower speed and lower lane count PMDs first, i.e.
"Other PMDs"
"800GBASE-VR8, 800GBASE-SR8 PMDs"

Response Response Status C

ACCEPT IN PRINCIPLE.
Change the order and associated parameters as proposed.

Cl 167 SC 167.7.2 P 115 L 12 # 193

Nicholl, Gary Cisco Systems

Comment Type E Comment Status A (bucket1)

Table 167-8. The order of the PMDs in the 'Signaling rate" row is different from what was done in Clause 124.

SuggestedRemedy

Proposing to reorder the data in this row to put the lower speed and lower lane count PMDs first, i.e.
"Other PMDs"
"800GBASE-VR8, 800GBASE-SR8 PMDs"

Response Response Status C

ACCEPT IN PRINCIPLE.
Change the order and associated parameters as proposed.

Cl 167 SC 167.8.6 P 118 L 6 # 194

Nicholl, Gary Cisco Systems

Comment Type E Comment Status A (bucket1)

Table 167-12. The font for the text in the "PMD Type" column looks incorrect. Also the editing instruction is "change this table", but then no underline or strickthrough. Perhaps the editing instruction should have been "Replace Table 167-12 with the following:" ?

SuggestedRemedy

Change the font in the "PMD Type" column to use the standard table font and updt the editing instruction to "Replace Table 167-12 with the following:"

Response Response Status C

ACCEPT IN PRINCIPLE.
Correct the font and text type. Editing instruction is correct, underline new text and correct text alignment.

Cl 171 SC 171.2 P 150 L 4 # 195

Nicholl, Gary Cisco Systems

Comment Type E Comment Status A (bucket1)

800GXS should be 400GXS

SuggestedRemedy

Change
"PCS and 800GXS sublayers specified in 118.2"
to
"PCS and 400GXS sublayers specified in 118.2"

Response Response Status C

ACCEPT.

IEEE P802.3df D1.0 1st Task Force review comments

Cl 173 SC 173.4 P 182 L 38 # 196

Nicholl, Gary Cisco Systems

Comment Type T Comment Status A PMA SI

Figure 173-4 (8:32 PMA) there should be no PMA:IS_SIGNAL.indication towards the PMA (AUI is not able to transfer an out of band status signal) and possibly no "SIL" block in the block diagram.

The same comment applies to the 8:8 PMA in Figure 173-5.

SuggestedRemedy

Remove the PMA:IS_SIGNAL.indication signal and the "SIL" block from Figure 173-4 and Figure 173-5.

Response Response Status C

ACCEPT IN PRINCIPLE.

Comments #196, #197, #162 and #29 (and further review by the editorial team), identified several issues related to the service interface definition (and associated diagrams) for both the PMA and XS sublayers.

Slides 44 to 54 in the following presentation were reviewed by the task force:
https://www.ieee802.org/3/df/public/22_12/brown_3df_03c_2212.pdf

Implement the following changes in the draft with editorial license:

Update Figure 173-3, Figure173-4 and Figure 173-5 with the corresponding figures shown in slides 48-50 of https://www.ieee802.org/3/df/public/22_12/brown_3df_03c_2212.pdf , and update any associated text accordingly.

Update Clause 171 according to slides 51 and 52 of https://www.ieee802.org/3/df/public/22_12/brown_3df_03c_2212.pdf, add DTE 800GXS and PHY 800GXS service interface definitions, and update any associated text accordingly.

Update Figure 169-3 based on slide 53 of https://www.ieee802.org/3/df/public/22_12/brown_3df_03c_2212.pdf, and update any associated text accordingly.

Cl 173 SC 173.4 P 181 L 40 # 197

Nicholl, Gary Cisco Systems

Comment Type E Comment Status A PMA SI

Figure 173-3/4/5/. Need to make it clear if the sublayer above or below is another PMA , that the interface is connected over a physically instantiated AUI (800GAUI-8)

SuggestedRemedy

Update Figure 173-3/4/5 to make it clear if the sublayer above or below is another PMA , that the interface is connected over a physically instantiated AUI (800GAUI-8)

Response Response Status C

ACCEPT IN PRINCIPLE.
 Resolve using the response to comment #196.