

IEEE 802.3dk D3.2 Bidirectional 100Gb/s Optical Access PHYs 2nd Sponsor recirculation ballot comments

Cl 45 SC 45.2.1.6 P20 L13 # R2-17

Ran, Adeo Cisco Systems, Inc.

Comment Type ER Comment Status X

"Editor's note: 101000xx, 1001xxxx, 10001xxx, and 1000011x are reserved by IEEE 802.3dj."
 Editor's notes are common in drafts but are not included in the published standard. This text is not formatted as an editor's note, and it is not stated that it should be removed before publication.

SuggestedRemedy

Move the quoted text into an Editor's note box (Table) and add "(to be removed prior to publication)".

Proposed Response Response Status O

Cl 45 SC 45.2.1.145a P24 L16 # R2-18

Dawe, Piers J G NVIDIA

Comment Type E Comment Status X

Unless one reads the definition of the individual bits, the introductory text for this new register could cause confusion. It is not clear enough that a PMA in 100GAUI-1 has some precoding ability however this register is set, and the PMD-PMA interface for 100GBASE-DR has no precoding ability however this register is set. Also, as this applies to a particular PHY type, there should be a cross-reference to the relevant PMD clause.

SuggestedRemedy

Add a sentence and note (here, because it applies to the whole table): "This register applies to a PMA in 100GBASE-BRx only (see 168.1).
 NOTE--For other PHY types, precoding ability, or the lack of it, is specified by rule (see 135.5.7.2, for example).

Proposed Response Response Status O

Cl 45 SC 45.2.1.145a P24 L25 # R2-19

Dawe, Piers J G NVIDIA

Comment Type E Comment Status X

The language in the table could be better and clearer. "supported" is vague, and the passive voice is undesirable, but Clause 45 was first written using "ability" and "able to" language. This table is to control the state of precoding on the optical path, which is different to PMA precoding in 802.3ck in the context of an optical PHY, where precoding may be used to benefit an AUJ but it is not used on the optical path.

SuggestedRemedy

Change:
 Precoding is supported by Rx input
 Precoding is not supported by Rx input
 Precoding is supported by Tx output
 Precoding is not supported by Tx output
 to
 PMA Rx input is able to remove precoding
 PMA Rx input is not able to remove precoding
 PMA Tx output is able to apply precoding
 PMA Tx output is not able to apply precoding

Proposed Response Response Status O

Cl 45 SC 45.2.1.145a.1 P24 L35 # R2-22

Dawe, Piers J G NVIDIA

Comment Type E Comment Status X

This language could be better and clearer. "supported" is vague, and the passive voice is undesirable, but Clause 45 was first written using "ability" and "able to" language. These bits are to control the state of precoding on the optical path, which is different to PMA precoding in 802.3ck in the context of an optical PHY, where precoding may be used to benefit an AUJ but it is not used on the optical path. Subclause number should not be preceded by "Clause".

SuggestedRemedy

Change:
 When read as a one, bit 1.607.1 indicates that the precoding described in Clause 135.5.7.2 is supported by Rx input. When read as a zero, bit 1.607.1 indicates that the precoding described in Clause 135.5.7.2 is not supported by Rx input.
 to
 When read as a one, bit 1.607.1 indicates that the PMA Rx input is able to remove precoding (decoding, see 135.5.7.2). When read as a zero, bit 1.607.1 indicates that the PMA Rx input is not able to remove precoding.

Proposed Response Response Status O

IEEE 802.3dk D3.2 Bidirectional 100Gb/s Optical Access PHYs 2nd Sponsor recirculation ballot comments

Cl 45 SC 45.2.1.145a.2 P24 L41 # R2-23

Dawe, Piers J G NVIDIA

Comment Type E Comment Status X

This language could be better and clearer. "supported" is vague, and the passive voice is undesirable, but Clause 45 was first written using "ability" and "able to" language. These bits are to control the state of precoding on the optical path, which is different to PMA precoding in 802.3ck in the context of an optical PHY, where precoding may be used to benefit an AUI but it is not used on the optical path.
Subclause number should not be preceded by "Clause".

SuggestedRemedy

Change:
When read as a one, bit 1.607.0 indicates that the precoding described in Clause 135.5.7.2 is supported by Tx output. When read as a zero, bit 1.607.0 indicates that the precoding described in Clause 135.5.7.2 is not supported by Tx output.
to
When read as a one, bit 1.607.0 indicates that the PMA Tx output is able to apply precoding (see 135.5.7.2). When read as a zero, bit 1.607.0 indicates that the PMA Tx output is not able to apply precoding.

Proposed Response Response Status O

Cl 80 SC 80.1.4 P34 L6 # R2-3

Ran, Adeo Cisco Systems, Inc.

Comment Type ER Comment Status X

The title of Table 80-1 is changed from "40 Gb/s and 100 Gb/s PHYs" but it is not marked.

The new Table 80-1a is inserted, not changed, so underlining is not needed. Similarly in Table 80-7a.

SuggestedRemedy

In Table 80-1 include the original table name, strike out "and 100 Gb/s".

In Table 80-1a and Table 80-7a, remove the underline format.

Proposed Response Response Status O

Cl 80 SC 80.2.3 P38 L43 # R2-4

Ran, Adeo Cisco Systems, Inc.

Comment Type ER Comment Status X

The changes in the first paragraph are not marked.

It seems that "100GBASE-LR1, and 100GBASE-ZR PHYs" is being changed to "100GBASE-LR1, 100GBASE-BRx, and 100GBASE-ZR PHYs".

Also in 80.2.5 and maybe other places.

SuggestedRemedy

Apply underline and strikethrough format to the text to show the change from the base standard.

Apply in 80.2.5 and elsewhere as necessary.

Proposed Response Response Status O

Cl 80 SC 80.2.3 P38 L43 # R2-24

Dawe, Piers J G NVIDIA

Comment Type E Comment Status X

The changes vs. the base standard should be shown (in black)

SuggestedRemedy

Underline "100GBASE-BRx,"
Similarly in 80.2.5, possibly more

Proposed Response Response Status O

Cl 80 SC 80.2.5 P38 L47 # R2-6

Ran, Adeo Cisco Systems, Inc.

Comment Type E Comment Status X

Missing closing parentheses

SuggestedRemedy

change "(PMD" to "(PMD)"

Proposed Response Response Status O

EEE 802.3dk D3.2 Bidirectional 100Gb/s Optical Access PHYs 2nd Sponsor recirculation ballot comments

CI 80 SC 80.4 P40 L31 # R2-5

Ran, Adeo Cisco Systems, Inc.

Comment Type ER Comment Status X

The changes to footnotes a and b of Table 80-7 seem to be deletions of the existing text and replacement by new text, but the new text is not underlined.

Also, the table is amended by 802.3ck and 802.3db.

SuggestedRemedy

Apply underline and strikethrough correctly for these footnotes.

Include "as amended by" in the editorial instruction.

Proposed Response Response Status O

CI 135 SC 135.5.7.2 P50 L30 # R2-7

Ran, Adeo Cisco Systems, Inc.

Comment Type E Comment Status X

The second sentence "The PMA may provide 1/(1+D) mod 4 decoding capability on each input lane" is independent of the first one. It would be clearer if placed in a separate paragraph.

SuggestedRemedy

break the quoted sentence into a separate paragraph.

Proposed Response Response Status O

CI 157 SC 157.1.2 P53 L15 # R2-8

Ran, Adeo Cisco Systems, Inc.

Comment Type E Comment Status X

In Figure 157-1, the labels "25GMII, 50GMII, or CGMII" and "25GBASE-R, 50GBASE-R, or 100GBASE-R PCS" are modified from the base standard.

SuggestedRemedy

Apply underlining as necessary

Proposed Response Response Status O

CI 157 SC 157.2.1 P54 L37 # R2-9

Ran, Adeo Cisco Systems, Inc.

Comment Type TR Comment Status X

"50GMII" is still relevant and should not be deleted.

SuggestedRemedy

remove the strikethrough from "50GMII".

Proposed Response Response Status O

CI 157 SC 157.4 P56 L25 # R2-10

Ran, Adeo Cisco Systems, Inc.

Comment Type E Comment Status X

In the draft there are two subclauses, 157.4.1 and 157.4.2, which do not appear in the base standard (802.3-2022).

The insertion of two subclauses should be stated in the editorial instruction, and the titles should be underlined.

Also, the editorial instruction says "(unchanged contents are not shown)" but it seems that the whole content of 157.4 s included.

SuggestedRemedy

Change the editorial instruction to

"Change Clause 157.4 as shown, including additional subclause headings 157.4.1 and 157.4.2."

Format the new headings in underline.

Proposed Response Response Status O

IEEE 802.3dk D3.2 Bidirectional 100Gb/s Optical Access PHYs 2nd Sponsor recirculation ballot comments

Cl 161 SC 161.6.10a P58 L27 # R2-12

Ran, Adeo Cisco Systems, Inc.

Comment Type T Comment Status X

100G_RS_FEC_Int_ability is a variable defined in clause 161 to indicate support of clause 161.

This self reference is problematic: if clause 161 is not implemented then the variable does not exist. If it is implemented, then the value should always be one (constant, not variable).

The only place that 100G_RS_FEC_Int_ability is mentioned is in Table 161-2. There is an MDIO bit in 45.2.1.117.7a indicating RS-FEC-Int ability; that subclause does not refer to the 100G_RS_FEC_Int_ability variable or to 161.6.10a, so the variable is not needed there.

Note that a similar situation exists for the BASE-R FEC, for which there is a FEC capability variable in 74.8.1. However, in that subclause the variable is explicitly defined as belonging to the PHY rather than to the FEC sublayer. This means the variable should have been defined outside of the FEC subclause, e.g. in the PMD subclauses that define the PHY structure. This is an editorial mistake that should not be repeated.

SuggestedRemedy

Delete subclause 161.6.10a.

In Table 161-2, delete "100G_RS_FEC_Int_ability" from the FEC variable column, leaving the cell empty.

Proposed Response Response Status O

Cl 161 SC 161.6.10a P58 L30 # R2-11

Ran, Adeo Cisco Systems, Inc.

Comment Type TR Comment Status X

"When 100G_RS_FEC_Int_ability is asserted, this variable indicates the support of RS-FEC-Int defined in Clause 161"

It is not specified what "asserted" means in a variable.

Compare to 161.6.3 (100G_RS_FEC_Int_enable): "When the 100G_RS_FEC_Int_enable variable is set to one".

(this comment may be overtaken by events if the subclause is deleted, subject of another comment)

SuggestedRemedy

Change "asserted" to "set to one".

Proposed Response Response Status O

Cl 168 SC 168.1 P60 L16 # R2-2

Ran, Adeo Cisco Systems, Inc.

Comment Type TR Comment Status X

The description of the optional precoding in footnote e is lacking crucial detail: As in the case of RS-FEC-Int (footnote d), it has to be supported by both link partners, and configured to the same setting in the transmitter of one link partner and the receiver of the other (and furthermore, different settings can be used in each direction).

Since there is no link training over the medium, this configuration needs to be done by the network operator through station management (e.g. MDIO). This requirement is unique to 100GBAS-BRx, so readers should not be expected to know it; and it is not written explicitly anywhere in this amendment.

Having this information in a footnote to a table of clauses is not the best way to inform readers.

Comment R1-23 suggested adding a new subclause with a detailed description, but its disposition did not add such a subclause.

Similar requirements (but not the same) apply to the use of RS-FEC-Int (without auto-negotiation), and are also worth mentioning explicitly.

These are new requirements, so new text has to be written.

SuggestedRemedy

Add two new subclauses under 168.1 ("Overview") after 168.1.1 ("Bit error ratio") as follows:

<BEGIN>

168.1.2 Reed-Solomon forward error correction codeword-interleaved (RS-FEC-Int) (optional)

100GBASE-BRx PHYs may include the RS-FEC-Int sublayer specified in Clause 161. If this sublayer is included in both link partners, codeword interleaving can be used to reduce the frame loss ratio (see 168.1.1).

If RS-FEC-Int is implemented, it can be either disabled or enabled. The default setting is disabled.

Operation of a 100GBASE-BRx link requires that the RS-FEC-Int setting (enabled or disabled) be the same in the two link partners.

The criteria for using RS-FEC-Int and the means by which both link partners are configured to the same setting are beyond the scope of the standard.

168.1.3 Precoding (optional)

100GBASE-BRx PHYs may provide PAM4 precoding in the PMA, as specified in 135.5.7.2. If this option is provided in both link partners, precoding can be used to reduce the frame loss ratio (see 168.1.1).

Precoding capability is optional in both PMA input and output. If implemented, it can be either disabled or enabled. The default setting is disabled.

The precoding setting is separate for PMA input and output. Operation of a 100GBASE-BRx link requires that the precoding setting (enabled or disabled) be the same in the PMA

EE 802.3dk D3.2 Bidirectional 100Gb/s Optical Access PHYs 2nd Sponsor recirculation ballot comments

output of one link partner and in the PMA input of the other link partner.
 The criteria for using precoding and the means by which the PMAs on both link partners are configured appropriately are beyond the scope of the standard.

<END>

In addition:

In Table 168-1, delete footnotes d and e (or change them to just refer to the new subclauses).

In 168.1.1, change "(Clause 91)" to "(Clause 91 or Clause 161)".

Proposed Response Response Status

Cl 168 SC 168.5.4 P64 L37 # R2-1

Rolfe, Benjamin Blind Creek Associates

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

Regarding R1-18: The SDG failed to address the comment. The comment is: "Not completely sure what this means: "SIGNAL_DETECT shall be a global indicator of the presence of the optical signal" but I think this is a simple statement of consequence when the requirement conditions in Table 168-4 are met."

Citing what is a correct use of "shall" in no way addresses the question raised by the commenter. The group further states they use similar language elsewhere, also not addressing the content of the comment. Either the comment is correct - that this is a statement of fact, not a completely specified requirement, or the comment is incorrect. If the group disagrees with the comment, then an explanation of why would necessary. The group provides no explanation as to why the comment is technically incorrect (such as a reference to where the missing technical specification is to be found).

Suggested Remedy

Change "shall be" to "is".

Proposed Response Response Status

Cl 168 SC 168.6.1 P66 L17 # R2-13

Ran, Adeo Cisco Systems, Inc.

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

"Editor's note: BR20 spec adjustment was done to increase OMAouter tolerance between min and max values."

Editor's notes are common in drafts but are not included in the published standard.

This text is not formatted as an editor's note, and it is not stated that it should be removed before publication.

Looking at the text, it seems to try to explain why BR20 OMA_outer and other power parameters are lower than both BR10 and BR40 (which is counter intuitive, and subject of another comment). If this is true, some kind of explanation should probably be kept in the published standard, but in the form of an informative NOTE.

Suggested Remedy

Either move the quoted text into an "Editor's note" box (table) and add "(to be removed prior to publication)", or delete it now.

If the intent is to explain the unexpected lower power parameters of BR2, add a NOTE at the end of 168.6.2 stating the reasoning clearly - something like

"NOTE-The average launch power (min and max), and OMA_outer (min and max), and transmitter power excursion (max) for 100GBASE-BR20 are lower than the corresponding limits for 100GBASE-BR10 and 100GBASE-BR40. The reason for that is <...>"

Proposed Response Response Status

IEEE 802.3dk D3.2 Bidirectional 100Gb/s Optical Access PHYs 2nd Sponsor recirculation ballot comments

Cl 168 SC 168.6.1 P66 L38 # R2-14

Ran, Adee Cisco Systems, Inc.

Comment Type TR Comment Status X

The average launch power (min and max), and OMA_{outer} (min and max), and transmitter power excursion (max) for 100GBASE-BR20 are lower than the corresponding limits for 100GBASE-BR10 and 100GBASE-BR40.

This is suspicious and inconsistent with Table 160-6 where the values of these limits for 50GBASE-BR20 are between those of 100GBASE-BR10 and 100GBASE-BR40.

I noticed in clause 158 and 159 (10GBASE-BRx and 25GBASE-BRx) there is a similar situation, but at least in clause 158 this may be explained because BR20 uses RS-FEC while BR10 and BR40 don't. For clause 159 I could not find an explanation in the text.

The implication of these values seems to be that a compliant BR20 transmitter may not interoperate with a compliant BR10 receiver over a <10 km distance, because its launch power is allowed to be much lower than that of a BR10 transmitter for the same link budget. While there is no explicit statement that such interoperability is possible, it seems to be expected, because the ranges in Table 168-5 overlap.

The suggested remedy is based on the assumption that these limit values should be in the logical order (as in clause 160). However, if the existing values are a conscious choice, some reasoning should be provided, preferably in an informative NOTE.

SuggestedRemedy

Swap the values of launch power (min and max), and OMA_{outer} (min and max), and transmitter power excursion (max) in the column for BR20 with the ones in the column for BR10.

If compliant BR20 and BR10 are not guaranteed to interoperate over a <10 km distance, please indicate that clearly, preferably both in clause 168 and in clause 157, and consider adjusting the ranges in Table 168-5.

If the values are okay and interoperability is guaranteed, please provide the explanation (which I can't suggest) in an informative NOTE.

Proposed Response Response Status O

Cl 168 SC 168.6.3 P68 L25 # R2-15

Ran, Adee Cisco Systems, Inc.

Comment Type ER Comment Status X

"Editor's note: call for further check of the penalty values."

Editor's notes are common in drafts but are not included in the published standard.

This text is not formatted as an editor's note, and it is not stated that it should be removed before publication.

I assume such further check is not required anymore and the sentence can be deleted.

SuggestedRemedy

Delete the quoted sentence.

Alternatively, make it an Editor's note box (table) and add "(to be removed prior to publication)".

Proposed Response Response Status O

Cl 168 SC 168.7.5 P70 L27 # R2-20

Dawe, Piers J G NVIDIA

Comment Type E Comment Status X

The editor's note at line 27 saying "tap weight limits will be added in alignment with IEEE P802.3 dj" is obsolete now.

SuggestedRemedy

Delete it

Proposed Response Response Status O

IEEE 802.3dk D3.2 Bidirectional 100Gb/s Optical Access PHYs 2nd Sponsor recirculation ballot comments

CI 168 SC 168.7.5 P70 L52 # R2-21

Dawe, Piers J G

NVIDIA

Comment Type TR Comment Status X

The resolution to D3.1 comment R1-37 says:

ACCEPT IN PRINCIPLE.

Add a bullet after line 50 in page 71:

The tap coefficient limits: the *coefficients* of the tap before (pre-cursor) and after (post-cursor) the tap with the largest magnitude tap (cursor) coefficient are less than 0.1. The coefficient of *post-cursor minus pre-cursor* is less than 0.15.

See contribution 3dk_dawe_2511_1.

https://www.ieee802.org/3/dk/public/2511/3dk_dawe_2511_1.pdf

Implement with editorial license.

D3.2 says:

The cursor is the tap with the largest magnitude coefficient *and is used as the basis of normalization*.

The magnitude of the pre-cursor (the tap before the cursor) *normalized* coefficient is less than 0.1.

The magnitude of the post-cursor (the tap after the cursor) *normalized* coefficient is less than 0.1.

The *magnitude of the difference* between the *normalized* coefficients of the post-cursor and pre-cursor is less than 0.15.

This is not correct. In 100G, ECQ coefficients are not normalized to the cursor (although their sum is 1). The post-cursor should not be much more +ve than the pre-cursor but it is acceptable for the post-cursor to be much more -ve than the pre-cursor.

SuggestedRemedy

Change the quoted text to:

The cursor is the tap with the largest magnitude coefficient.

The magnitude of the pre-cursor (the tap before the cursor) coefficient is less than 0.1.

The magnitude of the post-cursor (the tap after the cursor) coefficient is less than 0.1.

The post-cursor coefficient minus the pre-cursor coefficient is less than 0.15.

Proposed Response Response Status

CI 168 SC 168.7.5 P71 L25 # R2-16

Ran, Adee

Cisco Systems, Inc.

Comment Type TR Comment Status X

"Editor's note: tap weight limits will be added in alignment with IEEE P802.3 dj."

Editor's notes are common in drafts but are not included in the published standard.

This text is not formatted as an editor's note, and it is not stated that it should be removed before publication.

Note that P802.3dj is not in SA ballot yet and is expected to be published much later than P802.3dk. This project cannot wait for 802.3dj...

Are the updates made in D3.2 final?

SuggestedRemedy

Delete the quoted text.

Proposed Response Response Status