

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 000 SC 0 P L # 66
 Froroth, Ingvar Marvell

Comment Type E Comment Status A bucket

The PDF Document Properties are not filled in completely:

Title: IEEE P802.3xx name of Task Force
 Subject: IEEE P802.3xx amendment
 Author: IEEE P802.3xx Task Force
 Keywords: P802.3xx,

SuggestedRemedy

Fill in the PDF Document Properties with 802.3by and relevant details.

Response Response Status C

ACCEPT.

CI 000 SC 0 P L # 105
 Lusted, Kent Intel

Comment Type T Comment Status A no hyphen, CC

The dash "-" in 25G-AUI and 25G-MII does not follow the convention in the base standard.

SuggestedRemedy

Consider changing "25G-AUI" to "25GAUI" and "25G-MII" to "25GMII" in the draft.

Response Response Status C

ACCEPT IN PRINCIPLE.

See comments #193 and #194.

CI 000 SC 0 P L # 192
 Dawe, Piers Mellanox

Comment Type E Comment Status D CC

Presumably this is an amendment to what is presently 802.3bx, which is further towards ratification than this amendment. The Frame source for 802.3bx is available so the links in green to the base document can be made active (or for copied text, could be / could have been left active). Doing so will both reduce the number of stale and incorrect cross-references and ease the process of review and checking.

If the links are live there is no need for them to be green, because they won't need special attention later. For material copied and modified from early clauses, this would save the editors' time spent turning them green. It may be that it would be very onerous to make all the links live - if so, a mix of live links and green dummy links would still be an improvement.

SuggestedRemedy

Use live links to the base document where practical. Leave the live links black as convenient. Update the note on page 24.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 000 SC 0 P L # 108
 Lusted, Kent Intel

Comment Type TR Comment Status D withdrawn

there is confusion related to the fact that we use the term 25GBASE-CR to refer to a PMD, a PHY type, a cable assembly label, a host, an MDI, and a link. 25GBASE-CR-S makes sense with some of these, but not with others.

SuggestedRemedy

See presentation.

Proposed Response Response Status Z

PROPOSED REJECT.

This comment was WITHDRAWN by the commenter.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 000 SC 0 P35 L1 # 20
 Anslow, Pete Ciena

Comment Type E Comment Status A

The first editing instruction for 45.2.1 includes:
 "as modified by IEEE Std 802.3bn and IEEE Std 802.3bw"

When referencing other amendments to 802.3, the amendment name should be followed by a year.

SuggestedRemedy

Change to:
 "as modified by IEEE Std 802.3bn-201x and IEEE Std 802.3bw-201x"
 Make an equivalent change to all other such references in the draft.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change all occurrences including those at the the following locations:
 page 35, line 1 (two occurrences), CL45
 page 36, line 25, CL45
 page 39, line 3, CL45
 page 40, line 23, CL45
 page 45, line 3 and 24, CL45

CI 000 SC 000 P L # 197
 Matt Brown

Comment Type E Comment Status A

Keep text in legacy clauses synchronized with P802.3bx.

SuggestedRemedy

Update legacy clauses based upon approved comments against 802.3bx D3.0, as appropriate.

Response Response Status C

ACCEPT IN PRINCIPLE.

[This comment was added post-review by the editor for formal task force approval.]

Update based upon approved comments against 802.3bx D3.0, as appropriate.

CI 000 SC 000 P25 L25 # 193
 Dawe, Piers Mellanox

Comment Type ER Comment Status A no hyphen, CC

We don't put a hyphen the G in XAUI, XLAUI or CAUI-n. Nor do we put a hyphen after the G in 10GBASE-SR and so on, GMII, XGMII, XLGMII or CGMII. This is no different.

SuggestedRemedy

Change 25G-AUI to 25GAUI throughout.

Response Response Status C

ACCEPT.

[The editor changed to Clause/Subclause from 001/1.4.64a to 000/000 since this applies to many parts of the document.]

See comment #105 and #194.

See strawpoll #2.
 24 supported this change
 3 did not support
 5 abstained

CI 000 SC 000 P25 L29 # 194
 Dawe, Piers Mellanox

Comment Type ER Comment Status A no hyphen, CC

We don't put a hyphen after the G in GMII, XGMII, XLGMII or CGMII. Nor do we put a hyphen after the G in 10GBASE-SR and so on, XAUI, XLAUI or CAUI-n. This is no different.

SuggestedRemedy

Change 25G-MII to 25GMII throughout.

Response Response Status C

ACCEPT.

[The editor changed to Clause/Subclause from 001/1.4.64a to 000/000 since this applies to many parts of the document.]

See comments #105 and #193.

See strawpoll #2.
 24 supported this change
 3 did not support
 5 abstained

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 000 SC 000 P37 L13 # 53
 Booth, Brad Microsoft

Comment Type T Comment Status A CR/KR nomenclature, CC

Noticed that the draft uses "25GBASE-CR or 25GBASE-CR-S" and "25GBASE-KR or 25GBASE-KR-S". This could be simplified with a definition like we did with "25GBASE-R".

SuggestedRemedy

Add two new definitions:

25GBASE-C: A family of Physical Layer entities for 25 Gb/s operation over one lane of twinaxial copper cable. (See IEEE Std. 802.3, Clause 110.)

25GBASE-K: A family of Physical Layer entities for 25 Gb/s operation over one lane of an electrical backplane. (See IEEE Std. 802.3, Clause 111.)

Replace the "or" statements in draft with the corresponding "-C" and "-K" names. Replace could also apply to "and" statements used in the 110 annexes.

Response Response Status C

ACCEPT IN PRINCIPLE.

[The editor changed the clause/subclause from 045/45.2.1.6 to 000/000 as this comment applies to multiple clauses.]

See proposed responses to comment #116 and #139 which address ambiguities in the terminology.

CI 001 SC 1.1.3 P25 L4 # 106
 Lusted, Kent Intel

Comment Type TR Comment Status A

need an entry in 1.1.3 Compatibility interfaces for 25G-MII

SuggestedRemedy

Insert where appropriate into the list under P802.3bx D3.1 Clause 1.1.3, the following text: "25 Gigabit Media Independent Interface (25G-MII). The 25G-MII is designed to connect a 25 Gb/s capable MAC to a 25 Gb/s PHY. While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 25 Gb/s speeds. The 25G-MII is a logical interconnection intended for use as an intra-chip interface. No mechanical connector is specified for use with the 25G-MII. The 25G-MII is optional."

Response Response Status C

ACCEPT IN PRINCIPLE.

Similar text is required for the 25G-AUI as well.

For the 25G-MII, insert the following text where appropriate in 1.1.3.2:

"25 Gigabit Media Independent Interface (25G-MII). The 25G-MII is designed to connect a 25 Gb/s capable MAC to a 25 Gb/s PHY. While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 25 Gb/s speeds. The 25G-MII is a logical interconnection intended for use as an intra-chip interface. No mechanical connector is specified for use with the 25G-MII. The 25G-MII is optional."

For the 25G-AUI, insert the following text where appropriate in 1.1.3.2:

"25 Gigabit Attachment Unit Interface (25G-AUI). The 25G-AUI is a physical instantiation of the PMA service interface to extend the connection between 25 Gb/s capable PMAs. While conformance with implementation of this interface is not necessary to ensure communication, it is recommended, since it allows maximum flexibility in intermixing PHYs and DTEs at 25 Gb/s speeds. The 25G-AUI is intended for use as a chip-to-chip or a chip-to-module interface. No mechanical connector is specified for use with the 25G-AUI. The 25G-AUI is optional."

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 004 SC 4.4.2 P27 L42 # 107
Lusted, Kent Intel

Comment Type TR Comment Status A

Add reference to 25GMII in Note 4. XGMII is listed but not the 25G version.

SuggestedRemedy

consider changing "XGMII" to "XGMII or 25GMII"

Response Response Status C

ACCEPT IN PRINCIPLE.

Change "XGMII" to "XGMII or 25G-MII"

Also, see comment #105.

Cl 030 SC 30.5.1.1.4 P30 L6 # 94
Rannow, Randy k APIC

Comment Type E Comment Status R

Para 30.5.1.1.4 Page 30, line 6 the first two sentences appear confusing

"At power-up or following a reset, the value of this attribute will be "unknown" for AUI, 10BASE5, 10BASE2, 10BROAD36, and 10BASE-FP MAUs. For these MAUs loopback will be tested on each transmission during which no collision is detected."

SuggestedRemedy

At power-up or following a reset, the attribute value of the following MAUs will be "unknown" for AUI, 10BASE5, 10BASE2, 10BROAD36, and 10BASE-FP. For these, MAUs loopback will be tested on each transmission during which no collision is detected.

Response Response Status C

REJECT.

This text is out of scope for P802.3by. If the commenter is concerned about this text the commenter may consult with the 802.3 maintenance committee.

Cl 031B SC 31B.3.7 P196 L40 # 37
Ran, Adee Intel

Comment Type T Comment Status A

The value 60 pause_quanta is insufficient for PHYs that operate in RS-FEC mode and consume the maximum delays allowed in table 105-3: 16 for RS, 7 for PCS, 48 for RS-FEC, 8 for PMA and 1 for PMD sum up to 80 pause_quanta.

Also, plural for pause_quantum is pause_quanta, and it's a time value, as defined in the base document, so "bit times" is an incorrect addition (I have submitted a comment on that to 802.3bx).

Also, the editor's note is not required for the next draft.

SuggestedRemedy

Change "60 pause_quantum bit times" to "80 pause_quanta".

On page 197, change max_overrun formula for 25G to 5120+frame_length, and delete editor's note.

Response Response Status C

ACCEPT IN PRINCIPLE.

Adopt suggested remedy and change "pause_quantum bit times" to "pause_quanta" elsewhere in 31B.3.7

Cl 031B SC 31B.3.7 P197 L11 # 25
Marris, Arthur Cadence Design Syste

Comment Type E Comment Status A

Delete editor's note as it is no longer needed.

bucket

SuggestedRemedy

Delete editor's note as it is no longer needed.

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 045 SC 2 P35 L22 # 82
 Nowell, Mark Cisco

Comment Type E Comment Status R

The use of "single-lane" is used throughout the clause as a replacement for 10G and meant to simplify saying 10GBASE-R and 25GBASE-R I believe. I'm concerned we don't have a specific definition for "single-lane" and therefore the reader may not understand what it implies. While 1G doesn't use FEC it is also a single lane PHY.

Occurrences of this are:

- Table 45-3 Page 35 line 22
- Table 45-3 Page 35 line 25
- 45.2.1.94 Page 41 line 31,34 & 36
- Table 45-74 Page 41 line 40
- 45.2.1.95 Page 42 line 1,4 & 7
- Table 45-75 Page 42 line 11

SuggestedRemedy

Need a discussion on approach - either create a definition (but we use single-lane elsewhere in the document around MDI connectors in 110)

Change to a "10G/25GBASE-R" format to just be explicit rather than "Single-lane PHY BASE-R", which was the original intention of the change I believe.

Response Response Status C

REJECT.

The name "Single-lane PHY BASE-R FEC corrected blocks counter" for register 1.172 was chosen to differentiate the register from the "BASE-R FEC corrected blocks counter, lanes 0 through 19" at 1.300 through 1.339.

Using "Single-lane" makes it clear that the register at 1.172 should be used rather than the one at 1.300 for 10 and 25G.

If a PHY does not use FEC then the reader will ignore this register.

CI 045 SC 45.2.1 P35 L2 # 21
 Anslow, Pete Ciena

Comment Type E Comment Status A

The first editing instruction for 45.2.1 includes:
 "which will insert new registers at addresses 1.17 and 1.18"

However, when IEEE Std 802.3by-201x is published, the other amendments will have been published (otherwise they shouldn't be mentioned).

Same issue in 45.2.1.4, 45.2.1.10, and 45.2.3.7

SuggestedRemedy

- In the first editing instruction for 45.2.1 change to:
 "which inserted new registers at addresses 1.17 and 1.18"
- In the editing instruction for 45.2.1.4 change to:
 "which inserted a row for bit 1.4.10"
- In the editing instruction for 45.2.1.10 change to:
 "which inserted a row for bit 1.11.11"
- In the editing instruction for 45.2.3.7 change to:
 "which inserted a row for bit 3.8.6"

Response Response Status C

ACCEPT.

CI 045 SC 45.2.1.101 P42 L30 # 23
 Anslow, Pete Ciena

Comment Type E Comment Status A bucket

In Table 45-79, "R0" should be "RO". i.e., what appears to be a zero should be a capital "o"

SuggestedRemedy

Change "R0" to "RO"

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 045 SC 45.2.1.101.1 P43 L50 # 175
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status D

Bypass indication added clause 108 to the "see" list, but Bypass correction removed the references to clause 91.

SuggestedRemedy

Remove the "see 91.X" references from the following sections:

- 45.2.1.101.1 RS-FEC Bypass Indication enable
- 45.2.1.101.2 RS-FEC Bypass Correction enable
- 45.2.1.102.7 RS-FEC High SER
- 45.2.1.102.8 FEC Bypass Indication ability
- 45.2.1.102.9 FEC Bypass Correction ability

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 045 SC 45.2.1.101.2 P43 L5 # 5
 Anslow, Pete Ciena

Comment Type ER Comment Status A

In 45.2.1.101.2 "FEC bypass correction enable" the two cross-references to 91.5.3.3 have been removed rather than adding additional cross-references to clause 108. This makes it much, much harder for the reader to figure out which clauses use the "FEC bypass correction enable" feature. It may be fairly obvious when looking at the 802.3by amendment, but it becomes much more difficult when the amendment gets rolled up into the base standard.

Same issue in 45.2.1.102.7, 45.2.1.102.8, and 45.2.1.102.9

SuggestedRemedy

Rather than removing the two cross-references to 91.5.3.3, add two additional cross-references to 108.5.3.2.

Make equivalent changes in 45.2.1.102.7, 45.2.1.102.8, and 45.2.1.102.9.

Response Response Status C

ACCEPT.

CI 045 SC 45.2.1.101.a P42 L42 # 9
 Anslow, Pete Ciena

Comment Type T Comment Status A bucket

It is not clear exactly what happens when the 25GBASE-R Reed-Solomon FEC is disabled. Is this just the decoding? If the encoding is not performed, what is the signal format?

The answer to these questions can be found in 108.6.3 "25G RS-FEC Enable". Please make this information easier to obtain by adding a cross-reference to 108.6.3

SuggestedRemedy

Add: "(see 108.6.3)" to the end of the last sentence.

Response Response Status C

ACCEPT.

CI 045 SC 45.2.1.102.1 P43 L15 # 24
 Anslow, Pete Ciena

Comment Type E Comment Status A bucket

In the added text in 45.2.1.102.1, "Clause 91" should be green and "Clause 108" should be a cross-reference.

Same issues in 45.2.1.102.2

SuggestedRemedy

In the added text in 45.2.1.102.1, apply character tag "External" to "Clause 91" and make "Clause 108" a cross-reference.

Make the same changes in 45.2.1.102.2.

Response Response Status C

ACCEPT.

CI 045 SC 45.2.1.102.1 P43 L15 # 28
 Marris, Arthur Cadence Design System

Comment Type E Comment Status A bucket

Fix cross references in 45.2.1.102.1 and 45.2.1.102.2

SuggestedRemedy

Mark Clause 91 and Clause 108 as cross references

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 045 SC 45.2.1.102.1 P43 L16 # 1
 Anslow, Pete Ciena
 Comment Type E Comment Status A bucket
 "alignment markers lock" should be "alignment marker lock"
 SuggestedRemedy
 Change "alignment markers lock" to "alignment marker lock"
 Response Response Status C
 ACCEPT.

CI 045 SC 45.2.1.103 P44 L0 # 174
 Slavick, Jeff Avago Technologies
 Comment Type T Comment Status A
 45.2.1.103, 45.2.1.104 and 45.2.1.106 (corrected, uncorrected, lane 0 RS-FEC codeword and symbol error counters) contain references to Clause 91 for their definitions. But no reference to clause 108
 SuggestedRemedy
 Add Clause 108 as a location that can define the error counters.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

CI 045 SC 45.2.1.2.3 P36 L14 # 95
 Rannow, Randy k APIC
 Comment Type E Comment Status R
 Para 45.2.1.2.3 Page 36, line 14 appears as a run-on sentence
 "Fault is a global PMA/PMD variable. When read as a one, bit 1.1.7 indicates that either (or both) the PMA or the PMD has detected a fault condition on either the transmit or receive paths."
 SuggestedRemedy
 Recommended:
 Fault is a global PMA/PMD variable. When read as a one, bit 1.1.7 indicates that either (both) the PMA or (and)the PMD has (have) detected a fault condition on either the transmit or receive path.
 Response Response Status C
 REJECT.
 The text is out of scope of P802.3by.

CI 045 SC 45.2.1.7.4 P38 L13 # 22
 Anslow, Pete Ciena
 Comment Type E Comment Status A bucket
 Tables 45-9, 45-10, and 45-12 already contain rows with multiple PMD types. For instance: "10GBASE-S, 10GBASE-L, 10GBASE-E 52.4.8"
 These do not use "or" between the PMD types
 SuggestedRemedy
 Use the already established format for multiple PMD types in Tables 45-9, 45-10, and 45-12.
 Replace " or" with "," in Table 45-9 (2 instances), Table 45-10 (2 instances), and Table 45-12 (2 instances),
 Response Response Status C
 ACCEPT.

CI 045 SC 45.2.1.94 P0 L0 # 176
 Slavick, Jeff Avago Technologies
 Comment Type T Comment Status D
 Sections 45.2.1.94 and 45.2.1.95 are labeled as 10GBASE-R but used in 25GBASE-R operations as well.
 SuggestedRemedy
 Retitle these sections as Single Lane BASE-R similar to how Table 74-1 has been updated.
 Proposed Response Response Status Z
 REJECT.
 This comment was WITHDRAWN by the commenter.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 045 SC 45.2.1.96 P42 L18 # 155
 Maki, Jeffery Juniper Networks, Inc.

Comment Type ER Comment Status A

Draft needs to include this subclause with an expanded subclause title, first sentence, and Table 45-76 title, and first sentence of 45.2.1.96.1 that includes 25G-AUI.

SuggestedRemedy

Replace "CAUI-4" with "CAUI-4 and 25G-AUI" throughout 45.2.1.96 including 45.2.1.96.1.

Response Response Status C

ACCEPT IN PRINCIPLE.

Import 45.2.1.96 and Table 45-76 from the base standard and change "CAUI-4 chip-to-module" to "CAUI-4 C2M and 25GAUI C2M":
 In the 45.2.1.96 heading text.
 In the first sentence in 45.2.1.96.
 in the Table 45-76 table title.
 In the first sentence in 45.2.1.96.1.

CI 045 SC 45.2.3.7.3a P45 L41 # 2
 Anslow, Pete Ciena

Comment Type E Comment Status A bucket

In the editing instruction, "new subclauses 45.2.3.7.3a" should be "new subclause 45.2.3.7.3a"

SuggestedRemedy

Change "new subclauses 45.2.3.7.3a" to "new subclause 45.2.3.7.3a"

Response Response Status C

ACCEPT.

CI 045 SC 45.2.7.12 P47 L11 # 36
 Ran, Adee Intel

Comment Type T Comment Status A

We need a new bit in register 7.48 to indicate that RS-FEC was negotiated, similar to bit 4 which indicates BASE-R FEC.

For good order, there should be a variable in the clause 73 that this bit can be mapped to, so that it becomes clear how it is set. This also applies to "BASE-R FEC negotiated" which does not have an associated variable.

SuggestedRemedy

In clause 45:

Use bit 7.48.7(currently reserved) for "RS-FEC negotiated".

Rename 45.2.7.12.1 from "BASE-R FEC negotiated" to "FEC negotiated" and append text for RS-FEC:

"Bit 7.48.7 indicates that RS-FEC operation has been negotiated. This bit is set only if a PHY type with optional RS-FEC operation has also been negotiated."

In clause 73:

Add two variable definitions in 73.10.1:

an_baser_fec_control - Indicates whether usage of BASE-R FEC has been negotiated for PHYs that have optional BASE-R FEC. Values: false - BASE-R FEC not negotiated or negotiated PHY does not have optional BASE-R FEC (default); true - BASE-R FEC negotiated. NOTE-This variable is set by this variable definition; it is not set explicitly in the state diagrams.

an_rs_fec_control - Indicates whether usage of RS-FEC has been negotiated for PHYs that have optional RS-FEC. Values: false - RS-FEC not negotiated or negotiated PHY does not have optional RS-FEC (default); true - RS-FEC negotiated. NOTE-This variable is set by this variable definition; it is not set explicitly in the state diagrams.

In 73.6.5, add a new paragraph: "The choice of FEC resulting from this subclause is indicated by the variables an_baser_fec_control and an_rs_fec_control."

Append a row to table 73-6:

an_baser_fec_control | 74.48.4 BASE-R FEC negotiated
 an_rs_fec_control | 74.48.7 RS-FEC negotiated

Response Response Status C

ACCEPT IN PRINCIPLE.

Make bit 7.48.7 "RS-FEC FEC negotiated" and implement bit description with editorial license.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Implement the response to comment i-67 against draft 3.0 of 802.3bx and add new variable an_rs_fec_control with editorial license

Cl 045 SC 45.2.7.12.2 P47 L32 # 3
Anslow, Pete Ciena

Comment Type E Comment Status A

"these bits in register 7.48" has been changed to "the bits in register 7.48".
However, not all of the bits in register 7.48 indicate the negotiated port type, so it was better as it was.

SuggestedRemedy

Remove the change so that the text reads: "these bits in register 7.48" which then only refers to the bits in the subclause title.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change "the" back to "these".

Cl 069 SC 69.1.1 P50 L14 # 96
Rannow, Randy k APIC

Comment Type E Comment Status R

Para 69.1.1 Page 50, line 14 appears verbose and confusing (shall operator vs may operate vs can operate?)

"For 25 Gb/s operation, there is 25GBASE-KR and 25GBASE-KR-S that operate over one lane. For 40 Gb/s operation, there is 40GBASE-KR4 that operates over four lanes. For 100 Gb/s operation, the 100GBASE-R family is extended to include 100GBASE-KR4 and 100GBASE-KP4 that operate over four lanes."

SuggestedRemedy

Recommended:
For 25 Gb/s operation, there is 25GBASE-KR and 25GBASE-KR-S that operate over one lane. For 40 Gb/s operation, 40GBASE-KR4 operates over four lanes. For 100 Gb/s operation, the 100GBASE-R family is extended to include 100GBASE-KR4 and 100GBASE-KP4 that operate over four lanes.

Response Response Status C

REJECT.

The suggested remedy does not improve the text.

Also, the suggested remedy changes text that is out of scope for this project. The commenter may address this through the maintenance process.

Cl 069 SC 69.2.3 P52 L24 # 100
Butter, Adrian IBM

Comment Type TR Comment Status A embodiment

The following statement neglects to capture mandatory inclusion of the BASE-R FEC: 'These embodiments employ the PCS defined in Clause 107, the RS-FEC defined in Clause 108, the PMA defined in Clause 109, and the PMD defined in Clause 111 and specifies 25Gb/s operation over one differential path in each direction.'

SuggestedRemedy

Update to include BASE-R FEC: 'The 25GBASE-KR embodiment employs the PCS defined in Clause 107, the BASE-R FEC defined in Clause 74, the RS-FEC defined in Clause 108, the PMA defined in Clause 109, and the PMD defined in Clause 111 and specifies 25Gb/s operation over one differential path in each direction. The 25GBASE-KR-S embodiment employs the PCS defined in Clause 107, the BASE-R FEC defined in Clause 74, the PMA defined in Clause 109, and the PMD defined in Clause 111 and specifies 25Gb/s operation over one differential path in each direction.'

Response Response Status C

ACCEPT IN PRINCIPLE.

Change to the following:

"The 25GBASE-KR embodiment employs the PCS defined in Clause 107, the BASE-R FEC defined in Clause 74, the RS-FEC defined in Clause 108, the PMA defined in Clause 109, and the PMD defined in Clause 111, and specifies 25 Gb/s operation over one differential path in each direction. The 25GBASE-KR-S embodiment employs the PCS defined in Clause 107, the BASE-R FEC defined in Clause 74, the PMA defined in Clause 109, and the PMD defined in Clause 111, and specifies 25 Gb/s operation over one differential path in each direction."

See comment #135.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 069 SC 69.2.3 P52 L25 # 135
 Dudek, Mike QLogic

Comment Type T Comment Status A embodiment

The RS-FEC defined in Clause 108 is not used by 25GBASE-KR-S

SuggestedRemedy

Change "These embodiments employ the PCS defined in Clause107, the RS-FEC defined in Clause108, the PMA defined in Clause109, and the PMD defined in Clause111 and specifies 25Gb/s operation over one differential path in each direction."

to
 "These embodiments employ the PCS defined in Clause107, the PMA defined in Clause109, and the PMD defined in Clause111 and specifies 25Gb/s operation over one differential path in each direction. In addition the Backplane Ethernet 25GBASE-KR embodiment employs the RS-FEC defined in Clause 108."

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #100.

Cl 073 SC 73.6.4 P56 L5 # 132
 Dudek, Mike QLogic

Comment Type E Comment Status A

It is strange to provide a reason for why separate bits are used for backplane and copper cable when at 25G we are using the same bits and the same argument could apply. It would be better to just state what should be done without providing a confusing reason.

SuggestedRemedy

Delete "as the MDI and physical medium are different".

Response Response Status C

ACCEPT.

Cl 073 SC 73.6.5 P56 L10 # 4
 Anslow, Pete Ciena

Comment Type E Comment Status A bucket

There is no editing instruction associated with 73.6.5

SuggestedRemedy

Add "Change 73.6.5 as follows:"

Response Response Status C

ACCEPT.

Cl 073 SC 73.6.5 P56 L15 # 136
 Dudek, Mike QLogic

Comment Type T Comment Status A

I think the F0 and F1 bits are used for 10Gb/s per lane, and we should be more precise.

SuggestedRemedy

Consider changing the descriptions to
 F0 is the 10Gb/s per lane FEC ability
 F1 is the 10Gb/s per lane FEC requested.

In two other places replace "for other speeds of operation" with "for 10Gb/s per lane operation."

Response Response Status C

ACCEPT.

Cl 073 SC 73.6.5 P56 L20 # 62
 Dove, Daniel Dove Networking Solut

Comment Type TR Comment Status D

This section of text does not match what I expected from Dudek_3by_01a_0315. It doesn't explain how to operate between link partners that are CR on one side, CR-S on the other. One can assume a management agent would change a CR PHY to CR-S, but this is not shown.

SuggestedRemedy

Proposed:A presentation suggesting the change will be provided.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 073 **SC 73.6.5** **P56** **L29** # **63**
 Dove, Daniel Dove Networking Solut

Comment Type **TR** **Comment Status** **R**
 Incorrect Statement: do not support RS-FEC operation."

SuggestedRemedy
 Should say "are not required to support RS-FEC operation."

Response **Response Status** **C**
 REJECT.

As far as auto-negotiation is concerned 25GBASE-KR-S and 25GBASE-CR-S PHYs do not support RS-FEC operation.

The statement "This is because 25GBASE-KR-S and 25GBASE-CR-S PHYs do not support RS-FEC operation" is correct.

Cl 073 **SC 73.6.5** **P56** **L36** # **137**
 Dudek, Mike QLogic

Comment Type **T** **Comment Status** **A**
 The paragraph starting at line 36 only applies "for other speeds". It is not obvious that this is true.

SuggestedRemedy
 Create two sub-sections.

Insert sub-section heading "For 25G PHYs" at line 21
 Insert sub-section heading "For other speeds of operation" at line 31. (or title "For 10Gb/s per lane PHYs" if my other comment is accepted.

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

Cl 074 **SC 1** **P59** **L21** # **84**
 Nowell, Mark Cisco

Comment Type **E** **Comment Status** **R**
 In the sentence:
 "The 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR and 25GBASE-KR-S PHYs described in Clause 110 and Clause 111 are required to implement the FEC sublayer and may use it with links with a BER of 10-8 or better."

Remove the "may" - it isn't optional.

Also, use phrasing around BER levels consistent with following sentence:
 "The 40GBASE-CR4 and 100GBASE-CR10 PHYs described in Clause 85 optionally use the FEC sublayer to improve the BER performance beyond 10-12."

SuggestedRemedy
 Change to:

"The 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR and 25GBASE-KR-S PHYs described in Clause 110 and Clause 111 are required to implement the FEC sublayer to improve the BER performance beyond 10-8"

Response **Response Status** **C**
 REJECT.

The FEC sublayer being referred to here is exclusively referring to the Clause 74 BASE-R FEC and does not include the RS-FEC of Clause 108.

The use of the word "may" is appropriate as the 25GBASE-CR and 25GBASE-KR PHYs might use the Clause 108 FEC instead of the BASE-R FEC.

Also for the case of the 25GBASE-CR-S and 25GBASE-KR-S PHYs they are not required to (ie may not) use FEC if the link can deliver a BER of 10-12 without FEC.

The original objective of Clause 74 was to improve BER beyond 10-12 on already compliant channels. 25G is different in that Clause 74 FEC can be used to achieve a BER of 10-12 on an otherwise failing channel. Therefore use of the word "beyond" is not appropriate for 25G.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 074 SC 74.1 P59 L20 # 10
 Anslow, Pete Ciena

Comment Type E Comment Status A bucket

The IEEE Editorial style manual contains:
 "In a series of three or more terms, use a comma immediately before the coordinating conjunction (usually and, or, or nor)."

SuggestedRemedy

Change:
 "The 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR and 25GBASE-KR-S PHYs" to:
 "The 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, and 25GBASE-KR-S PHYs"

Response Response Status C

ACCEPT.

CI 074 SC 74.7.4.5 P65 L36 # 58
 Baden, Eric Broadcom

Comment Type T Comment Status D

To match the hi_ber parameters (97 bad SHs over 2ms), the FEC should corrupt the SHs in all 32 blocks of the code word to indicate an error.

SuggestedRemedy

Change text to indicate all 32 blocks of the code word have the SHs corrupted.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 074 SC 74.7.4.5.1 P66 L32 # 59
 Baden, Eric Broadcom

Comment Type T Comment Status D

To match the hi_ber parameters, indicate to corrupt all 32 sets of SHs in the code word.

SuggestedRemedy

Indicate for 25GBASE-R to corrupt all 32 sets of SHs in the code word.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 074 SC 74.7.4.8 P67 L51 # 79
 Cober, Don CoMIRA Solutions Inc

Comment Type E Comment Status A

Sub Clause 74.7.4.8 will need to be modified to mention the Clause 107 PCS.

SuggestedRemedy

Add:
 "If the optional EEE deep sleep capability is supported, then a Clause 107 PCS sublayer will be encoding encodes /I/ during the wake state and /LI/ during the refresh state, which produces the two types of deterministic FEC blocks."

Response Response Status C

ACCEPT IN PRINCIPLE.

Bring subclause 74.7.4.8 into P802.3by and add the following text.

"If the optional EEE deep sleep capability is supported, then a Clause 107 PCS sublayer encodes /I/ during the wake state and /LI/ during the refresh state, which produces the two types of deterministic FEC blocks."

CI 074 SC 74.8.1 P68 L34 # 11
 Anslow, Pete Ciena

Comment Type E Comment Status A

The title of 74.8.1 in the base standard is "FEC capability", but in the draft it is shown as "25GBASE-R FEC capability" which is not appropriate

SuggestedRemedy

Correct the title of 74.8.1

Response Response Status C

ACCEPT.

CI 074 SC 74.8.1 P68 L35 # 138
 Dudek, Mike QLogic

Comment Type T Comment Status A

subclause 74.8.1 in the base document contains information not related to 25GBASE-R. It shouldn't be put in a paragraph with 25GBASE-R in the title.

SuggestedRemedy

Change the title fo 74.8.1 back to FEC capability.

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 074 SC 74.8.1 P68 L36 # 12
 Anslow, Pete Ciena
 Comment Type E Comment Status A bucket
 "Clause 73" should be a cross-reference
 SuggestedRemedy
 Make "Clause 73" a cross-reference
 Response Response Status C
 ACCEPT.

CI 074 SC 8 P68 L15 # 85
 Nowell, Mark Cisco
 Comment Type E Comment Status R
 Similar comments to my Clause 45 comments on use of "single-lane". Concern over lack of definition of what "single-lane" encompasses
 Appears twice in Table 74-1
 SuggestedRemedy
 Suggest changing "Single-lane PHY BASE-R FEC uncorrected blocks counter register"
 to:
 "10G/25GBASE-R FEC uncorrected blocks counter register"
 Response Response Status C
 REJECT.
 See response to comment #82

CI 078 SC 78.1.1 P70 L23 # 30
 Marris, Arthur Cadence Design Syste
 Comment Type T Comment Status A
 Make it clearer where LPI PMA signalling is defined.
 SuggestedRemedy
 Change:
 "Coding defined in Clause83 and Clause 109 also"
 to:
 "Coding defined in 83.5.11"
 Delete "(See 83.5.11.1.)" on line 26
 Response Response Status C
 ACCEPT.

CI 078 SC 78.2 P72 L24 # 71
 Koehler, Daniel MorethanIP
 Comment Type T Comment Status A
 The value for Tr in Table 78-2 is inconsistent with values from Table 107-1. Table 107-1 specifies wake time Twl as 10.9-11.1 which is larger than Tr of this Table 78-2 would allow.
 SuggestedRemedy
 Change Table 78-2 values for 25G Tr values considering values from Table 107-1. Possibly the original 10G values of Tr being 16.9 to 17.5 should be used.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 See response to comment #72

CI 078 SC 78.2 P72 L24 # 72
 Cober, Don CoMIRA Solutions Inc
 Comment Type T Comment Status A
 Tr (Duration of the refresh signal) value for 25G is incorrect.
 This value is equal to the length of the tx alert + tx wake + tx sleep state:
 For the table in Clause 107 this is:
 Min: 1.1 + 10.9 + 4.9 = 16.9us
 Max: 1.3 + 11.1 + 5.1 = 17.4us
 SuggestedRemedy
 Change Tr min to 16.9
 Change Tr max to 17.4
 for all 25G types
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change Tr min to 16.9
 Change Tr max to 17.5
 for all 25G types

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 078 SC 78.5 P73 L27 # 73
 Cober, Don CoMIRA Solutions Inc

Comment Type T Comment Status A

All of the timing parameter values are incorrect for 25G deep sleep modes.

These timing values are derived from the counter values listed in Clause 107.
 The Clause 107 timing parameters should match the clause 49 timing parameters

For Case 3 of the 25G (RSFEC mode) the values should be the same as case 1 because the scr_bypass state is skipped.

SuggestedRemedy

For the 25G deep sleep modes, these values should be copied from the 10GBASE-KR
 For Case 3 of the 25G (RSFEC mode) the values should be the same as case 1.

Response Response Status C
 ACCEPT.

Cl 105 SC 1.3 P76 L45 # 117
 Goergen, Joel Cisco Systems, Inc.

Comment Type ER Comment Status R 25gbase-cr-n, CC

if the 2M no fec solution stays in the draft, then the -N version should be listed here

SuggestedRemedy

include 25GBASE-CR-N as a defined physical implementation

Response Response Status C
 REJECT.

Clause 110 defines only two PMD types: 25GBASE-CR and 25GBASE-CR-S. A CA-N cable is compatible with both a 25GBASE-CR and 25GBASE-CR-S PHY operating without an FEC.

See comment #116 and #118.

Cl 105 SC 105.1 P77 L42 # 115
 Goergen, Joel Cisco Systems, Inc.

Comment Type T Comment Status R 25gbase-cr-n, CC

If the 2M -N cabling stays in the draft, then -N needs to be listed in table 105-1

SuggestedRemedy

list the -N phy type in the table 105-1

Response Response Status C
 REJECT.

Clause 110 defines only two PMD types: 25GBASE-CR and 25GBASE-CR-S. A CA-N cable is compatible with both a 25GBASE-CR and 25GBASE-CR-S PHY operating without an FEC.

See comment #117.

Cl 105 SC 105.2 P78 L14 # 6
 Anslow, Pete Ciena

Comment Type T Comment Status A

Table 105-2 calls out Annex 109A as optional for all PHY types, but Annex 109B is not mentioned.

SuggestedRemedy

Add a column for Annex 109B and show it as optional for the 25GBASE-SR PHY (as per Table 112-1)

Response Response Status C
 ACCEPT.

Cl 105 SC 105.4.1 P80 L13 # 56
 Baden, Eric Broadcom

Comment Type E Comment Status A bucket

the word 'of' is missing between transfer and a (define the transfer of a stream of data)

SuggestedRemedy

Add the work of between the words transfer and a (stream)

Response Response Status C
 ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 105 SC 105.5 P87 L10 # 64
 Dove, Daniel Dove Networking Solut

Comment Type TR Comment Status A

Text "25GBASE-CR FEC" incomplete.

SuggestedRemedy

Replace with "25GBASE-CR BASE-R FEC"

Response Response Status C

ACCEPT IN PRINCIPLE.

[It appears that the commenter used coordinates from the compare version. Page and line number have been changed by the editor from 95/30 to 87/10.]

Change "25GBASE-R FEC" to "25G BASE-R FEC".

Cl 105 SC 2 P78 L10 # 110
 Nicholl, Gary Cisco Systems

Comment Type ER Comment Status A

Table 105-2 does not include a column for the 25G-AUI C2M Annex 109B.

SuggestedRemedy

Add a column to reference 25G-AUI C2M Annex 109B into Table 105-2. The column should be optional for all rows.

Response Response Status C

ACCEPT IN PRINCIPLE.

The 25G-AUI C2M should be optional only for the optical PHY 25GBASE-SR. It is otherwise not applicable. See IEEE 802.3-2012 Table 80-2.

See comment # 6.

Cl 105 SC 2 P78 L24 # 158
 Andrewartha, Mike Microsoft

Comment Type E Comment Status R

In Table 105-2, the row for 25GBASE-CR and column for 25GBASE-CR-S PMD should be marked M since all the capabilities of CR-S are required for CR. Likewise the row for 25GBASE-KR and column for 25GBASE-KR-S PMD should be marked M.

SuggestedRemedy

Add an 'M' in the cells at:
 row 25GBASE-CR and column 25GBASE-CR-S PMD
 and
 row 25GBASE-KR and column 25GBASE-KR-S PMD

Response Response Status C

REJECT.

The 25GBASE-CR and 25GBASE-CR-S are different PHYs. A 25GBASE-CR PHY is not a 25GBASE-CR-S PHY. The common and unique capabilities are specified in the Clause 110.

A 25GBASE-CR implementation might be configured as a 25GBASE-CR PHY, but this is outside the scope of the standard.

Cl 105 SC 2 P78 L27 # 119
 Goergen, Joel Cisco Systems, Inc.

Comment Type TR Comment Status R 25gbase-cr-n, CC

It has become evident that -N is only incorporated into clause 110. This needs to be fully addressed in clause 105 if it is to stay in the document

SuggestedRemedy

incorporate the -N cable so it is defined properly within the spec - or remove -N option completely

The remedy here is to remove the -N 2M cable option completely from clause 110. Change the 3M -S phy type to no fec with base-r fec optional. adjust the COM margin to 2dB from 3dB. I will present a presentation on this.

Response Response Status C

REJECT.

Clause 110 specifies only two PHY types: 25GBASE-CR and 25GBASE-CR-S. The CA-N is one of three medium classes that is supported by the 25GBASE-CR and 25GBASE-CR-S PHY types. There is no reason to specify the medium in Clause 105. Configuration for operation with a CA-N cable is accomplished directly through the management interface or as a consequence of auto-negotiation.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 105 SC 4.3.2.3 P83 L29 # 162
Andrewartha, Mike Microsoft

Comment Type ER Comment Status A

Table 105-3 shows the PMA layer twice.

SuggestedRemedy

Remove the duplicate PMA layer and associated text.

Response Response Status C

ACCEPT IN PRINCIPLE.

[The commenter indicated that the comment should refer to Figure 105-3 not Table 105-3.]

There should indeed be only one PMA sublayer between the FEC and PMD sublayers.

In Figure 105-3, remove the one PMA sublayer and associated service interface and text.

Cl 106 SC 106.1 P88 L10 # 13
Anslow, Pete Ciena

Comment Type E Comment Status A bucket

"Clause 46" should be shown in green

SuggestedRemedy

Apply the "External" character tag to "Clause 46"

Response Response Status C

ACCEPT.

Cl 106 SC 106.1.7.1 P90 L32 # 133
Dudek, Mike QLogic

Comment Type E Comment Status A

Poor grammar.

SuggestedRemedy

Change "The RS maps the primitive PLS_DATA.request to the 25G-MII signals TXD<31:0>, TXC<3:0>, and TX_CLK in the same way as for XGMII is mapped as specified in 46.1.7.1." to

Either

"The RS maps the primitive PLS_DATA.request to the 25G-MII signals TXD<31:0>, TXC<3:0>, and TX_CLK in the same way as for XGMII, as specified in 46.1.7.1."

Or

"The RS maps the primitive PLS_DATA.request to the 25G-MII signals TXD<31:0>, TXC<3:0>, and TX_CLK in the same way as for XGMII mapping, as specified in 46.1.7.1."

Make the similar change in 106.1.7.2 and 106.1.7.5

Response Response Status C

ACCEPT IN PRINCIPLE.

See response to comment #14.

Cl 106 SC 106.1.7.1 P90 L32 # 14
Anslow, Pete Ciena

Comment Type E Comment Status A

"in the same way as for XGMII is mapped as specified" does not make sense. Same issue in 106.1.7.2 and 106.1.7.5

SuggestedRemedy

Change to: "in the same way as for the XGMII as specified". Make the same change in 106.1.7.2 and 106.1.7.5.

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 107 SC 1.2 P94 L23 # 111
 Nicholl, Gary Cisco Systems

Comment Type T Comment Status D
 Where did the number 97 in 2ms come from ? I don't see that in
http://www.ieee802.org/3/by/public/Mar15/baden_3by_02_0315.pdf
 and I don't see how 16 errors in 125us at 10G scales to 97 errors in 2ms at 25G ? It should scale to 40 errors in 125us (or some multiple thereof). The current hi-ber count proposal appears to equate to a ber of ~6.4e-5

SuggestedRemedy

Proposed Response Response Status Z
 REJECT.

This comment was WITHDRAWN by the commenter.

CI 107 SC 1.2 P94 L23 # 164
 Andrewartha, Mike Microsoft

Comment Type T Comment Status R
 The definition of hi_ber given differs substantially from the definition in 49.2.13.3 with no justification or explanation elsewhere in the draft. Why are the time period and thresholds different? Needs further explanation.

SuggestedRemedy

Not clear based on lack of information about why the threshold and count are so different. Add appropriate explanation or change the values.

Response Response Status C
 REJECT.

The original hi_ber parameters defined in Clause 49 are not compatible with the error event due error marking when RS-FEC correction bypass is enabled. The parameters were changed to match the window size for 40G/100G.

See baden_3by_02_0315.

See comment #111.

CI 107 SC 1.2 P94 L23 # 112
 Nicholl, Gary Cisco Systems

Comment Type T Comment Status D
 baden_3by_02_0315 talks about disabling the PCS hi-ber count when the RS-FEC is used, and using RS-FEC symbol error counts to monitor the link quality instead. I don't see this captured anywhere in Clause 107 ?

SuggestedRemedy

Proposed Response Response Status Z
 REJECT.

This comment was WITHDRAWN by the commenter.

CI 107 SC 107 P97 L0 # 173
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status A
 Generation of scrambled idle patterns is not defined in clause 107 outside of the 107.1.2 item 2 comment regarding the differences between 107 and 49.

SuggestedRemedy

Copy 82.2.11 into clause 107 and update appropriately for clause 107 usage.

Add MDIO register control to select between test pattern prbs, square wave, and scrambled idle

Response Response Status C
 ACCEPT IN PRINCIPLE.

Create new subclause 107.2.3 based on "82.2.11 Test-pattern generators".

Reference and change "45.2.3.17 BASE-R PCS test-pattern control register (Register 3.42)" to control the test patterns.

CI 107 SC 107.2 P96 L1 # 57
 Baden, Eric Broadcom

Comment Type E Comment Status A bucket
 Change the word codes to encodes for better readability.

SuggestedRemedy

Change the word codes to encodes.

Response Response Status C
 ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 107 SC 107.3 P96 L50 # 60
 Dove, Daniel Dove Networking Solut
 Comment Type ER Comment Status A
 The word "and" seems incorrectly placed
 SuggestedRemedy
 Replace with "but" or end split the sentence into two.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 See response to comment #78

CI 107 SC 107.3 P96 L53 # 78
 Cober, Don CoMIRA Solutions Inc
 Comment Type E Comment Status A
 Rework this subclause to make clear that the LPI state diagrams do apply in deep sleep mode.
 SuggestedRemedy
 Add:
 "If the 25GBASE-R PCS is part of a PHY configured for EEE deep sleep operation, the PCS shall follow the state diagrams specified in Figure 49-12 and Figure 49-13."
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Change:
 "If the 25GBASE-R PCS is part of a PHY configured for EEE fast wake operation, the PCS shall encode and decode LPI when indicated and the state diagrams specified in Figure 49-12 and Figure 49-13 do not apply."
 To:
 "If the 25GBASE-R PCS is part of a PHY configured for EEE deep sleep operation, the PCS shall follow the state diagrams specified in Figure 49-12 and Figure 49-13.
 If the 25GBASE-R PCS is part of a PHY configured for EEE fast wake operation, the PCS shall encode and decode LPI when indicated but the state diagrams specified in Figure 49-12 and Figure 49-13 do not apply. Management functions may use MDIO register bit LPI_FW 3.20.0 to select fast wake operation (see 45.2.3.9.11)."

CI 107 SC 107.3 P97 L33 # 98
 Butter, Adrian IBM
 Comment Type TR Comment Status A
 For 25GBASE-R links, there are timing parameter discrepancies between those values shown in Table 78-2 (on p. 72), and those values shown in Table 107-1 (on p. 97) and Table 107-2 (on p. 98).
 SuggestedRemedy
 Update timing parameter values to be consistent among these tables.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 See response to comments #72 and #73

CI 107 SC 107.3 P97 L50 # 172
 Slavick, Jeff Avago Technologies
 Comment Type T Comment Status A
 For fast wake operations the definitions for what do exists, but there is no enable register to control whether to do Deep Sleep or Fast Wake LPI operations.
 SuggestedRemedy
 Add a LPI_FW MDIO register for Clause 107 that enables Fast Wake operations
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 This is already supported by MDIO see "45.2.3.9.11 LPI_FW (3.20.0)".
 Add cross reference to 45.2.3.9.11 in 107.3 to make this clearer.
 See response to comment #78

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 107 SC 107.3 P97 L52 # 170
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status A
 "these state diagrams" is unnecessary, what we really are stating is to use the timer values in the tables instead of the clause 49 tables.

SuggestedRemedy
 Change:
 The LPI functions shall use timer values for these state diagrams as shown in Table 107-1 for transmit and Table 107-2 for receive.
 To:
 The LPI functions shall use the timer values in Table 107-1 and Table 107-2 for EEE deep sleep operation.

Response Response Status C
 ACCEPT.

Cl 108 SC 108.1.1 P101 L10 # 29
 Marris, Arthur Cadence Design Syste

Comment Type E Comment Status A bucket
 Clause 108 is a single specification for the 25G RS_FEC so it should be singular.

SuggestedRemedy
 Change:
 "The specifications are closely related to those of the RS-FEC sublayer for 100GBASE-R PHYs"
 To:
 "The specification is closely related to that of the RS-FEC sublayer for 100GBASE-R PHYs"

Response Response Status C
 ACCEPT.

Cl 108 SC 108.2.7 P106 L8 # 75
 Cober, Don CoMIRA Solutions Inc

Comment Type T Comment Status D RS-FEC LPI signaling, timer
 It is not clear whether the unscrambling of data starts immediately after entry into the TX_WAKE or sometime after.

The exact distance between the unscrambled -> scrambled transition and the CWM is not clear.
 SuggestedRemedy

Proposed Response Response Status Z
 REJECT.

This comment was WITHDRAWN by the commenter.

Cl 108 SC 108.3 P103 L1 # 31
 Marris, Arthur Cadence Design Syste

Comment Type T Comment Status A
 With 25G-AUI the PMA is always a client of the the RS_FEC so it is not "may be".

SuggestedRemedy
 Change:
 "The 25GBASE-R PMA sublayer may be a client of the 25GBASE-R RS-FEC sublayer, when 25G-AUI C2C is used between a device that includes a PCS and a device that includes the RS-FEC."
 To:
 "When 25G-AUI C2C is used between a device that includes a PCS and a device that includes the RS-FEC, the 25GBASE-R PMA sublayer is the client of the 25GBASE-R RS-FEC sublayer."

Response Response Status C
 ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 108 SC 108.5.1 P104 L14 # 99
Butter, Adrian IBM

Comment Type TR Comment Status A

In Figure 108-2 on the left (transmit) side, the arrow between the 'Rate compensation for CW markers' and 'CW markers insertion' blocks points in the wrong direction.

SuggestedRemedy

Change the arrow to point from 'Rate compensation for CW markers' to 'CW markers insertion'.

Response Response Status C

ACCEPT IN PRINCIPLE.

The arrow between the "codeword marker insertion" and "rate compensation" blocks shows the direction in which information about insertion is passed.

The block diagram is intended to show the functions and not the way in which they are implemented.

On the left side, remove the arrow between "rate compensation for codeword marker" and "codeword marker insertion"

On the right side, remove the arrow between "rate compensation for codeword marker" and "codeword marker removal".

Cl 108 SC 108.5.2.2 P103 L34 # 61
Dove, Daniel Dove Networking Solut

Comment Type ER Comment Status A bucket

Inaccurate phrasing

SuggestedRemedy

Replace "periodic" with "periodically occurring"

Response Response Status C

ACCEPT.

Cl 108 SC 108.5.2.2 P103 L38 # 38
Ran, Adeed Intel

Comment Type T Comment Status D RS-FEC LPI signaling, idles

scrambler_bypass as currently defined has the effect of sending unscrambled control codes over the channel. This occurs during refresh and wake cycles, so the PCS input data can be sequences of either // or /L/ characters.

// characters are translated to control code of seven "0" bits, so an unscrambled block would contain a 8-bit block type (possibly shortened to 4-bit by transcoding) and then 56 "0" bits. A repeating pattern of these blocks has very low transition density and is strongly unbalanced, so is unsuitable for AC-coupled signaling and for CDRs.

A simple remedy is to specify that the effect of scrambler_bypass includes a deterministic manipulation of the 56 bits that hold the payload, which will keep the pattern easy to synchronize to - such as inverting every other group of 7 characters (for example, for block type 0x1e, invert C0, C2, C4 and C6) . The receiver will reverse this effect.

SuggestedRemedy

Change

"When scrambler_bypass is true, the descrambled data is passed to the transcoder, rather than the data from the scrambler output" to

"When scrambler_bypass is true, the data passed to the transcoder is created by applying a bitwise exclusive-or with the fixed 64-bit value 0x00FE03F80FE03F80 to each block of descrambled data, rather than using the data from the scrambler output".

In 108.5.3.6, change

"When descrambler_bypass is true, the received data is used without descrambling" to

"When descrambler_bypass is true, bitwise exclusive-or with the fixed 64-bit value 0x00FE03F80FE03F80 is applied to each block instead of regular descrambling".

In 108.5.3.6, change

"This causes the rate compensation function to use the receive data without descrambling (see 108.5.3.6)" to

"This enables the rate compensation function (see 108.5.3.6) to operate correctly with unscrambled data sent from the remote RS-FEC transmit function (see 108.5.2.2)".

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 108 SC 108.5.2.4 P104 L48 # 171
 Slavick, Jeff Avago Technologies

Comment Type T Comment Status R remove CWM, BTI

Codeword markers are a complicated method for framing the fec codewords.

SuggestedRemedy

Remove codeword markers and post scramble the codeword to prevent mis-alignment.
 See slavick_03by_01_0515.pdf

Response Response Status C

REJECT.

Per Motion #3, there is not consensus to make this change.

Cl 108 SC 108.5.2.4 P105 L7 # 7
 Anslow, Pete Ciena

Comment Type T Comment Status A

The convention used in 802.3 for ranges of bits shown within "<>" marks is for the highest number to appear first.

In 108.5.2.4 1), "tx_cwm<0:23>" should be "tx_cwm<23:0>".
 Swap the order of the numbers for items 1 through 16.

Same issue in 108.5.4.2 (5 instances)

SuggestedRemedy

Swap the order of the numbers within "<>" marks for 108.5.2.4 items 1 through 16.
 In 108.5.4.2 Swap the order of the numbers within "<>" marks (3 instances) and change:
 "bits 0:23 and 32:55" to "bits 23:0 and 55:32"

Response Response Status C

ACCEPT IN PRINCIPLE.

Change:
 "constructed as follows"

To:
 "constructed from the eight octets M0, M1, M2, BIP3, M4, M5, M6, and BIP7 (bits 65 to 2) of the alignment markers defined in Table 82-2 with the bit order shown in Figure 82-9, with BIP3 set to the constant value 0x33 and BIP7 set the constant value 0xCC, as follows".

Change list items 1-17 to the following items:

- 1)tx_cwm<63:0> are set to bits <65:2> of the alignment marker of PCS lane 0.
- 2)tx_cwm<127:64> are set to bits <65:2> of the alignment marker of PCS lane 1.
- 3)tx_cwm<191:128> are set to bits <65:2> of the alignment marker of PCS lane 2.
- 4)tx_cwm<255:192> are set to bits <65:2> of the alignment marker of PCS lane 3.
- 5)tx_cwm<256> is set to 0.

Cl 108 SC 108.5.2.4 P105 L8 # 68
 Baden, Eric Broadcom

Comment Type TR Comment Status D CWM

The CWMs should be comprised of AMs from the 40G specification and not from the 100G specification. The justification for the change in the previous draft from 40G AM0 to 100G for AM0 can be shown to be inconsequential. A presentation will be provided in support of this.

SuggestedRemedy

Change the reference to table 82-2 to instead reference table 82-3 on lines 8 and 12.

Proposed Response Response Status Z

PROPOSED REJECT.

This comment was WITHDRAWN by the commenter.

Cl 108 SC 108.5.2.7 P106 L4 # 69
 Koehler, Daniel MorethanIP

Comment Type T Comment Status A ?S-FEC LPI signaling, RCWM

The use of scrambler_bypass seems not necessary to allow the receiver to achieve a fast lock. Instead the already existing functions for inserting CWMs using rapid CWMs could be used for a simpler scheme (for the receiver see my comment on 108.5.3.7).

SuggestedRemedy

I am suggesting not to use scrambler_bypass at any time during EEE. Instead replace lines 4 to 17 as follows:

a) The variable tx_rapid_cwm (new variable) is set to true until 1µs before tx_tw_timer_done becomes true.

b) While tx_rapid_cwm is true insert a CWM at every RS-FEC codeword start. The CWM is identical to the CWM used during normal operation.

c) When tx_rapid_cwm transitions from true to false insert one more CWM at the next codeword start and then enter normal operation inserting CWMs every 1024 codewords.

As a result of the transmit function behavior at least 48 codewords with CWMs are sent. This enables rapid synchronization of the remote RS-FEC receive function and determination of the block types and receive LPI state by the remote PCS.

Response Response Status C

ACCEPT IN PRINCIPLE.

Addressed by the adopted remedy for comment #179.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 108 SC 108.5.2.7 P106 L5 # 179
 Wertheim, Oded Mellanox Technologies

Comment Type TR Comment Status A ?S-FEC LPI signaling, RCWM

During the WAKE periods, when FEC:IS_TX_MODE.request primitive from ALERT to DATA the transmitter sends unscrambled Idles/LPIs for a time period of 0.9us to 1.1us.

This results in:

1. Sending a non-DC balanced signal without enough transition density, which is unsuitable for the electrical signaling.
2. Sending unscrambled LPIs/Idles enables the peer port to detect transcoding block boundaries. However it does not provide a way to differ between the 20 transcoding blocks with 4 LPIs/Idles that each FEC codeword contain and thus detect the FEC codeword boundaries.

Suggested Remedy

Instead of sending unscrambled data, send rapid codeword markers (RCWMs) to enable the peer port to rapidly achieve codeword lock.

Sending Rapid codeword markers provides the peer port a simple mechanism for fast codeword lock while the transmitter sends a DC balanced output with sufficient transition density (scrambled Idles / LPIs).

Detailed remedy:

1. Modify 108.5.2.7 RS-FEC encoding for rapid codeword lock (EEE deep sleep) (page 105)
 - Replace a), b) (lines 5-9) with:
 - a) Set down_count to 16 and send 16 rapid codeword markers (RCWMs). This causes the Codeword marker insertion function (108.5.2.4) to insert a RCWM in the beginning of each of the following 16 FEC codewords.
 - b) The first regular codeword marker is inserted at the beginning of the 1024th RS-FEC codeword after the RCWM with down_count = 1.

- Remove lines 14-17 "As a result . by the remote PCS"

2. Add at the bottom of 108.5.2.4 Codeword marker insertion: (page 105, line 37)

For the optional EEE capability, a rapid method of FEC alignment is used when operating in the deep sleep low power state using Rapid Codeword Markers (RCWMs). RCWMs are inserted at the beginning of 16 codewords following the transmitter tx_mode transition from ALERT to DATA.

RCWMs are identical to regular CWMs with the exception that the constant value of 0x33 in offsets 24:31 is replaced with a down_count value, and the constant value of 0xCC in offsets 56:63 is replaced with the bit-wise inversion of the down_count. The down_count is decremented each time a RCWM is sent.

3. In 108.5.3.6 Rate compensation for codeword markers in the receive direction, (page 109, line 9)
 - Replace: "Insert idle characters, according to the rules in 49.2.4.7, to fill in as necessary for any deleted codeword markers."
 - With: "Insert idle or low power idle (LPI) characters, according to the rules in 49.2.4.7, to fill

in as necessary for any deleted codeword markers or rapid codeword markers."

4. In 108.5.3.7 Rapid codeword lock for EEE deep sleep: (page 109)
 - Remove a) in line 29.
 - Replace c) in line 33 with:

c) Enable the RS-FEC rapid codeword lock mechanism, which attempts to detect rapid codeword markers sent by the remote RS-FEC transmit function (see 108.5.2.7). When two sequential rapid codeword markers are detected, the start location of the RS-FEC codeword is set to the start location of the rapid codeword markers. The next codeword marker position is set to 1024 codewords following the rapid codeword marker with down_count = 1.

- Remove the sentence in line 44: "When the decoding .. Set to false"
- Replace 1) in line 50 "Two 64B/66B .. true to false" with:

1) The RS-FEC codeword monitor state diagram (Figure 108-6) reaches the CW_GOOD state.

5. In 108.5.4.2 State variables:

- Remove page 110, lines 35-40: descrambler_bypass
- Remove page 110, lines 52-54: scrambler_bypass
- Remove page 112, lines 1-3: "optional EEE . always false."

- Add a new variable to page 100, lines 35-40:

down_count
 A counter that is used in rapid codeword markers and is decremented each time a RAM is sent. The counter initial value is set by the RS-FEC transmit function when the tx_mode parameter of the FEC:IS_TX_MODE.request primitive from ALERT to DATA.

Response Response Status C
 ACCEPT IN PRINCIPLE.

Implement changes in cober_3by_01a_0515 slides 15 to 18 with editorial license.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 108 SC 108.5.2.7 P8 L7 # 180
 Wertheim, Oded Mellanox Technologies

Comment Type TR Comment Status A RS-FEC LPI signaling, RCWM

The transmitter inserts the first codeword marker at the beginning of the second full codeword (after the first full codeword has been transmitted).

The receiver may not succeed to identify the codeword boundaries in time and miss the codeword marker.

The transmitter sends unscrambled data for 0.9us - 1.1us. The second full codeword is sent 0.2usec - 0.4usec after the scrambler_bypass was set to TRUE.

SuggestedRemedy

Instead of sending unscrambled data, send rapid codeword markers (RCWMs) to enable the peer port to rapidly achieve codeword lock.

See the remedy in comment #1

Response Response Status C

ACCEPT IN PRINCIPLE.

Addressed by the adopted remedy for comment #179.

CI 108 SC 108.5.3.2 P108 L1 # 97
 Rannow, Randy k APIC

Comment Type E Comment Status R

Para 108.5.3.2 Page 108, line 1 appears confusing, first 1st

"is not supported or not enabled), it shall ensure that, for every other 257-bit block within the codeword starting with the first (1st, 3rd, 5th, etc.),"

SuggestedRemedy

Suggested:
 is not supported or not enabled), it shall ensure that, for every other 257-bit block within the codeword starting with the first (e.g. 1st, 3rd, 5th, etc.),

Response Response Status C

REJECT.

This text is consistent with similar text in 91.5.3.3 (IEEE Std 802.3bj).

CI 108 SC 108.5.3.2 P108 L23 # 80
 Cober, Don CoMIRA Solutions Inc

Comment Type T Comment Status A

The HiSER monitor should be bypassed when in LPI mode

SuggestedRemedy

Add the following to 108.5.3.2:
 "For the optional EEE deep sleep capability, the error monitor employed when FEC_bypass_indication_enable is asserted shall be disabled when rx_lpi_active=true. The next block of 8192 codewords considered by the error monitor shall begin on the codeword boundary following the transition of rx_lpi_active from true to false."

Add the following to 108.5.4.2:
 "The following variables are only used for the optional EEE deep sleep capability. If this capability is not supported, the values of rx_lpi_active are set to false.

Rx_lpi_active
 A Boolean variable that is set to true when the RS-FEC sublayer infers that the Low Power Idle is being received from the link partner and is set to false otherwise."

Response Response Status C

ACCEPT IN PRINCIPLE.

Add the following to 108.5.3.2:
 "For the optional EEE deep sleep capability, the error monitor employed when FEC_bypass_indication_enable is asserted shall be disabled when rx_lpi_active=true. The next block of 8192 codewords considered by the error monitor shall begin on the codeword boundary following the transition of rx_lpi_active from true to false."

Add the following to 108.5.4.2:
 "The following variable is only used for the optional EEE deep sleep capability. If this capability is not supported, the value of rx_lpi_active is set to false.

Rx_lpi_active
 A Boolean variable that is set to true when the RS-FEC sublayer infers that the Low Power Idle is being received from the link partner and is set to false otherwise."

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 108 SC 108.5.3.2 P108 L5 # 35
 Ran, Adee Intel

Comment Type T Comment Status A

Several sentences in this subclause include descriptions of the behavior of the PCS and AN, which are not the subject of this clause. These are informative explanations, so should better be put as NOTES (which would make them informative).

Also applies in 108.5.3.3.

SuggestedRemedy

In 108.5.3.2, move the text "This causes the PCS to discard all frames 64 bytes and larger that are fully or partially within the codeword" to a NOTE, rephrasing as necessary..

Also in 108.5.3.2, move the text "As a result, the PCS sets hi_ber (see Figure 49-15) to true, which inhibits the processing of received packets. When Auto-Negotiation is supported and enabled, assertion of hi_ber causes Auto-Negotiation to restart" to a NOTE, rephrasing as necessary.

In 108.5.3.3, delete the text ", leading to hi_ber being set by the PCS. When Auto-Negotiation is supported and enabled, this event causes Auto-Negotiation to restart.". Instead, add a note:

NOTE--marking multiple 64B/66B blocks as bad causes a PCS to lose its block lock. When AN is enabled, this event restarts the AN".

Response Response Status C
 ACCEPT.

Cl 108 SC 108.5.3.4 P108 L34 # 177
 Wertheim, Oded Mellanox Technologies

Comment Type E Comment Status A bucket

The sunclause defines codeword maerker removal and not alignment marker removal

SuggestedRemedy

Change the subclause to: 108.5.3.4 Codeword marker removal

Response Response Status C
 ACCEPT.

Cl 108 SC 108.5.3.4 P108 L34 # 163
 Andrewartha, Mike Microsoft

Comment Type ER Comment Status A bucket

The heading for 108.5.3.4 should say "Codeword marker removal" instead of "Alignment marker removal"

SuggestedRemedy

Change heading for 108.5.3.4 to "Codeword marker removal"

Response Response Status C
 ACCEPT.

[Editor changed subclause from 5.3.4 to 108.5.3.4]

See #177.

Cl 108 SC 108.5.3.5 P108 L48 # 168
 Andrewartha, Mike Microsoft

Comment Type TR Comment Status A P802.3bx

Text refers to the Figure 82-4 for the block types used by 25GBASE-R PCS but Figure 82-4 deals with PCS Receive bit ordering. Should refer to Figure 82-5.

SuggestedRemedy

change reference to Figure 82-4 to Figure 82-5

Response Response Status C
 ACCEPT IN PRINCIPLE.

[Editor changed subclause from 5.3.5 to 108.5.3.5]

The correct reference is to Figure 82-5 as it appears in IEEE Std 802.3-2012. However, P802.3bx (as of D3.0) has renumbered this figure and it is now Figure 82-4.

Align Figure # with P802.3bx Draft 3.1.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 108 SC 108.5.3.6 P109 L9 # 77
Cober, Don CoMIRA Solutions Inc

Comment Type E Comment Status A CWM

In b) where the idle insertion is described it should be made clear that "idle character" means Idle // and Low Power Idle /LI/

SuggestedRemedy

Change:
"b) Insert idle characters, according to the rules in 49.2.4.7 ..."
to:
"b) Insert // and /LI/ characters, according to the rules in 49.2.4.7 ..."

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement based on suggested remedy but change "and" to "or".

CI 108 SC 108.5.3.7 P109 L22 # 178
Wertheim, Oded Mellanox Technologies

Comment Type T Comment Status A RS-FEC LPI signaling, RCWM

The rapid codeword lock for EEE deep sleep does not define the mechanism to identify the codeword marker.

The rapid codeword lock describes provides a way to determine the start location of RS-FEC codewords. However it does not describe how to identify the codeword marker which is inserted by the transmitter in the second full codeword.

SuggestedRemedy

Enable the receiver to rapidly lock on the codeword marker using rapid codeword markers. See comments #1.

Response Response Status C

ACCEPT IN PRINCIPLE.

Addressed by the adopted remedy for comment #179.

CI 108 SC 108.5.3.7 P109 L24 # 169
Slavick, Jeff Avago Technologies

Comment Type T Comment Status A LPI signaling, remove CWM

If we disable scrambler during bypass scrambler time the data stream produced is not very random.

SuggestedRemedy

Use EEE signaling methods described in slavick_03by_01_0515.pdf

Response Response Status C

ACCEPT IN PRINCIPLE.

Per Motion #3, there is not consensus to adopt the suggested remedy.

Addressed by the adopted remedy for comment #179.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 108 SC 108.5.3.7 P109 L29 # 70
 Koehler, Daniel MorethanIP

Comment Type T Comment Status A RS-FEC LPI signaling, RCWM

The use of descrambler_bypass seems not necessary as CWMs are already defined and its detection function exists that could be re-used with rapid CWMs. Using both, CWMs and unscrambled Idle/LPI detection seems redundant and adding unnecessary complexity.

SuggestedRemedy

Replace Lines 29 to 52 with the following:

- a) Set rapid_cwm (new variable) to true. This enables fast lock based on rapid CWMs for the FEC Synchronization statemachine (Fig. 108-5) as well as the codeword marker removal function following the RS-FEC decoder. A rapid CWM is identical to a normal CWM but occurs in every RS-FEC codeword start position while the transmitter LPI state (Fig. 49-12) is TX_WAKE.
 - b) While rapid_cwm is true the FEC synchronization statemachine (Fig 108-5) is used with a cwm_counter_done occurring at every codeword distance. A fast implementation is required (implementation dependent, out of scope of this standard) to minimize (eliminate) necessary SLIPs and reliably detect two consecutive CWMs within less than 6 codewords (i.e. within~1.2µs).
 - c) Once FEC_align_status becomes true the CWM removal function at the output of the RS-FEC decoder becomes active and removes the CWM in every codeword until it detects no more. Once it finds no more CWM at a codeword start, it sets rapid_cwm to false and enters normal operation removing the CWMs at nominal distance (every 1024 codewords).
- Note - as this step operates on corrected data the missing CWM can unambiguously be identified as start of normal marker distance operation.
- d) Together with changing rapid_cwm to false the FEC sublayer asserts signal_ok to enable the normal PCS operations.

Further modifications from this imply:

- remove variables descrambler_bypass and scrambler_bypass from 108.5.4.2
- remove lines 38-42 of 108.5.2.2 page 103
- change 108.5.2.7 (see comment on 108.5.2.7)
- remove lines 14-17 of 108.5.3.6 page 109

Response Response Status C

ACCEPT IN PRINCIPLE.

Addressed by the adopted remedy for comment #179.

Cl 108 SC 108.5.3.7 P109 L32 # 76
 Cober, Don CoMIRA Solutions Inc

Comment Type T Comment Status A RS-FEC LPI signaling, timer

The holdoff timer started in b) does not match the PCS counterpart.
 The value is mistakenly assuming the PCS is going through the scrambler bypass state.

SuggestedRemedy

Change 108.5.3.7 to read:
 " b) Start a hold-off timer whose duration is greater than or equal to 11.5 us."

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 108 SC 108.5.3.7 P109 L39 # 74
 Cober, Don CoMIRA Solutions Inc

Comment Type T Comment Status D RS-FEC LPI signaling, timer

the RSFEC EEE mechanism for determining the transition from unscrambled to scrambled will not detect correctly and is vulnerable to errors.

Firstly, the unscrambled blocks will not be transcoded correctly until step e2 of the transcoding is bypassed (this is the reverse scrambling of the block type nibble for lookup)

Secondly, searching for an errored block to find scrambled transition will miss random data that shows up as a start of frame or ordered set. Searching for a value not equal to I or LI is more reliable.

Thirdly, even searching for not // /LI/ will fail if the link has uncorrectable errors. It would be more reliable to check an entire codeword of 80 blocks, and consider the codeword unscrambled if any of the 80 blocks is filled with // or /LI/. This would require that the transition between scrambled and unscrambled happen on a codeword boundary.

SuggestedRemedy

Modify 108.5.2.7 (pg 106, ln 5) to read:

"a) The variable scrambler_bypass is set to TRUE for a period of 0.9 ls to 1.1 ls. This causes the ratecompensation function (108.5.2.2) to generate unscrambled data. This variable is only changed on codeword boundaries, such that any codeword will contain all scrambled or all unscrambled blocks. "

Modify 108.5.3.5 (pg 108, ln 54) to read:

"If descrambler_bypass is enabled, then step e2) is bypassed and g<i> = f_c<i> for i=0 to 3"

Modify 108.5.3.7 (pg 109, ln 44) to read:

"When the decoding in item a) of 108.5.3.6 generates a set of 80 blocks from a codeword, none of which are a control block filled with // or /LI/ characters while codeword monitor is in CW_GOOD state and descrambler_bypass is true, it is an indication that the remote RS-FEC transmitter has re-enabled scrambling, and descrambler_bypass is set to false "

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 108 SC 108.5.4.2 P110 L31 # 26
 Marris, Arthur Cadence Design System

Comment Type E Comment Status A bucket

Delete editor's note as it is no longer needed.

SuggestedRemedy

Delete editor's note as it is no longer needed.

Response Response Status C

ACCEPT.

CI 108 SC 108.6.3 P116 L2 # 45
 Ran, Adeo Intel

Comment Type E Comment Status A bucket

The editor's note about RS-FEC enable/disable capability is not needed in the next draft.

SuggestedRemedy

Delete editor's note.

Response Response Status C

ACCEPT.

CI 109 SC 1.3 P123 L6 # 159
 Andrewartha, Mike Microsoft

Comment Type E Comment Status A per input lane

reference to "per-input lane" is unnecessary since only a single lane is defined.

SuggestedRemedy

Change a) to read: "Provide clock and data recovery"

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 109 **SC 1.3** **P123** **L6** # **109**
 Nicholl, Gary Cisco Systems

Comment Type **E** **Comment Status** **A** *per input lane*

This is a single lane project, so there is no need to use terms like "per-lane" in the text.

SuggestedRemedy
 Change "Provide per input-lane clock and data recovery."
 to
 "Provide input lane clock and data recovery"

Response **Response Status** **C**
 ACCEPT IN PRINCIPLE.

See comment #159.

Cl 109 **SC 109.1.3** **P123** **L24** # **52**
 Ran, Adee Intel

Comment Type **T** **Comment Status** **A**

"local loopback" label in Figure 109-2 has a footnote c, "Optional". But 109.4.3 has a mandatory requirement for local loopback in a PMA adjacent to PMDs except for 25GBASE-SR.

This footnote conflicts with the clause text.

I have submitted a comment to 802.3bx on the similar issue in clause 83.

SuggestedRemedy
 Add a new footnote d to "Local loopback, with the text "Local loopback is required for PMAs adjacent to some PMDs, and optional for other PMAs. See 109.4.3".

Response **Response Status** **C**
 ACCEPT IN PRINCIPLE.

The text proposed in the suggested remedy is overly verbose for a figure. A reference to the subclause provides the full context.

Add a new footnote:
 "d Conditional (see 109.4.3)."

Change the label from:
 "local loopback <c>"
 To:
 "local loopback <d>"

Cl 109 **SC 109.2** **P125** **L27** # **39**
 Ran, Adee Intel

Comment Type **E** **Comment Status** **A**

The PMA sends a bit stream to the PMA client. Here it says "one stream", which is unclear.

SuggestedRemedy
 Change "one stream" to "a bit stream".

Response **Response Status** **C**
 ACCEPT IN PRINCIPLE.

Change:
 "sends one stream to the PMA client"
 To:
 "sends a bit stream to the PMA client"

Cl 109 **SC 109.2** **P125** **L30** # **32**
 Ran, Adee Intel

Comment Type **T** **Comment Status** **A**

This text about "the SIGNAL_OK parameter of the PMA:IS_SIGNAL.indication primitive" seems incorrect, in view of the newly added text discussing this primitive, four paragraphs below.

SuggestedRemedy
 Delete this paragraph.

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

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 109 SC 109.2 P126 L8 # 40
 Ran, Adee Intel

Comment Type E Comment Status A

The statements starting with "The ability to support transition..." and "Transition to the low power state..." use "register" and "direction" in an inconsistent order, which reduces their legibility.

I have submitted a comment to 802.3bx on the similar issue in clause 83.

SuggestedRemedy

Change "The ability to support transition to a low power state in the ingress direction is indicated by register 1.1.9 (PMA Ingress AUI Stop Ability, PIASA) and register 1.1.8 for the egress direction (PMA Egress AUI Stop Ability, PEASA)."

To "The ability to support transition to a low power state in the ingress direction is indicated by register 1.1.9 (PMA Ingress AUI Stop Ability, PIASA). The ability to support transition to a low power state in the egress direction is indicated by register 1.1.8 (PMA Egress AUI Stop Ability, PEASA)."

Change "Transition to the low power state is enabled in the ingress direction by register 1.7.9 (PMA Ingress AUI Stop Enable, PIASE) and register 1.7.8 for the egress direction (PMA Egress AUI Stop Enable, PEASE)."

To "Transition to the low power state in the ingress direction is enabled by register 1.7.9 (PMA Ingress AUI Stop Enable, PIASE). Transition to the low power state in the egress direction is enabled by register 1.7.8 (PMA Egress AUI Stop Enable, PEASE)."

Response Response Status C

ACCEPT.

Cl 109 SC 109.3 P126 L23 # 41
 Ran, Adee Intel

Comment Type E Comment Status A bucket

The final part of the sentence, "(e.g., another PMA, FEC, or PMD)", repeats an identical itemization in the previous sentence.

SuggestedRemedy

Delete the parthesized text

Response Response Status C

ACCEPT.

Cl 109 SC 109.4.1 P127 L3 # 152
 Brown, Matthew APM

Comment Type E Comment Status A

Subclause 109.4.1 "Delay Constraints" is under the subclause 109.4 "Functions within the PMA". "Delay Constraints" is a performance metric, not a function. This should be a subclause with a level 2 heading similar to clauses 110, 111, and 112.

SuggestedRemedy

Change the heading level of subclause 109.4.1 to a level 2 heading and place the subclause just prior to the current subclause 109.5.

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 109 SC 109.4.1 P127 L5 # 33
 Ran, Adee Intel

Comment Type T Comment Status A

Referece to the delay of "up to four PMA stages" appears only here. The PICS just state the delay without referring to multiple stages. 105.5 does not refer to multiple PMAs either.

A solution could be to limit the delay of a single PMA (to 1/4 of the current value) in the PMA clause, and account for four such delays in clause 105. However, this would limit implemetations with fewer than 4 PMAs. Also, there is a precedence in clause 83 that we may want to follow.

If we keep the current specification, then the fact that this is the total delay is not obvious, and should be consistently stated elsewhere.

I have submitted a comment to 802.3bx on a similar issue in clause 83.

SuggestedRemedy

Change the PICS and clause 105 to match clause 109 as follows:

In Table 105-3, row "25GBASE-R PMA", prepend to the Notes: "Cumulative value for up to four PMA instances at one end of the link".

In 109.6.4.2, item PC1, append to Feature: ", cumulative value for up to four PMA instances".

Response Response Status C

ACCEPT IN PRINCIPLE.

In Table 105-3, add a table footnote to "25GBASE-R PMA" as follows (using language from 109.4.1):
 "Cumulative round-trip delay contributed by up to four PMA stages in a PHY."

In 109.6.4.2, item PC1, replace feature with:
 "Cumulative round-trip delay contributed by up to four PMA stages in a PHY."

CI 109 SC 109.4.3 P127 L36 # 42
 Ran, Adee Intel

Comment Type E Comment Status A

The first sentence in this paragraph, "The PMA sublayer may provide a local loopback function", isn't necessary. The previous paragraph has already stated required/optional conditions. For PMAs that are required to provide local loopback, "may" is incorrect.

SuggestedRemedy

Delete "The PMA sublayer may provide a local loopback function."

Response Response Status C

ACCEPT IN PRINCIPLE.

The text suggested for removal is redundant. The sentence that follows needs to be modified as the subject is unclear.

Delete: "The PMA sublayer may provide a local loopback function."

In the following sentence.

Change: "The function involves"

To: "The PMA local loopback function involves"

CI 109 SC 109.4.5.1 P128 L31 # 27
 Marris, Arthur Cadence Design Syste

Comment Type E Comment Status A bucket

Delete editor's note as it is no longer needed.

SuggestedRemedy

Delete editor's note as it is no longer needed.

Also on page 129 line 21

Also on page 130 line 5

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 109 SC 109.4.5.2 P128 L46 # 153
Brown, Matthew APM

Comment Type T Comment Status A

In 109.4.5.2 and 109.4.5.6, there is an incorrect reference to "link status" when referring to the service interface below the PMA. The service interface subclause 109.2 only refers to "status". "link status" has a particular connotation in some subclauses.

SuggestedRemedy

Change "link status" to "status" in two places.
page 128, line 46, subclause 109.4.5.2
page 130, line 18, subclause 109.4.5.6

Response Response Status C

ACCEPT.

Cl 109 SC 109.4.5.7 P130 L41 # 151
Brown, Matthew APM

Comment Type T Comment Status A

The text incorrectly refers to the "transmit process".

SuggestedRemedy

Change:
"If the optional Clause 45 MDIO is implemented, the PMA receive process maps the Square_wave_ability and Square_wave_enable_0 variable to the registers and bits defined in 109.5"
To:
"If the optional Clause 45 MDIO is implemented, the PMA transmit process maps the Square_wave_ability and Square_wave_enable_0 variable to the registers and bits defined in 109.5"

Response Response Status C

ACCEPT.

Cl 109 SC 109.6.4.1 P134 L41 # 34
Ran, Adee Intel

Comment Type T Comment Status A

PF3, PF6, PF7 and PF9 refer to transmit test pattern generation, and receiver test pattern checking. All these items have status "optional" and are conditional on 25G-AUI below (PIB). But the test patterns can also be used (and may be necessary for testing) in 25GBASE-CR and 25GBASE-KR PHYs, where the PMA is co-located with a PMD, without an AUI.

SuggestedRemedy

Can be corrected by changing the definition of the "PIB" condition to include ", or adjacent to PMD" as done in definition of JTP2 in 83.7.5, or by adding "KRCR:O" to the status.

Response Response Status C

ACCEPT IN PRINCIPLE.

In 109.6.4 "Major Capabilities" add a new row below *PIB.
Item: *PMB
Feature: PMD below.
Subclause: 109.3
Status:O
Support: Yes [] No []

For PF3, PF6, PF7, and PF9 add the following to the status column: "PMB:O"

Cl 109 SC 3 P126 L33 # 160
Andrewartha, Mike Microsoft

Comment Type E Comment Status A bucket

Typo: "The service interface below the PMA has and input and output" should be "The service interface below the PMA has an input and output"

SuggestedRemedy

Replace the word 'and' with 'an' betwen 'has' and 'input'

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 109 SC 3 P126 L38 # 161
 Andrewartha, Mike Microsoft

Comment Type E Comment Status A
 Sentence fragment or possible missing text in this paragraph.
 There seems to be text missing between 'interface' and 'is'

SuggestedRemedy
 Insert correct words or edit as appropriate to convey intended meaning.

Response Response Status C
 ACCEPT IN PRINCIPLE.

Change "when data received via"
 To "data received via"

Cl 109B SC 109B.1 P207 L14 # 195
 Dawe, Piers Mellanox

Comment Type ER Comment Status R 25g-aii nomenclature, CC
 Make the abbreviation match the phrase it's abbreviating: "chip-to-module 25 Gigabit Attachment Unit Interface". Note that all the adjectives come before the noun in this phrase, not a mix.

SuggestedRemedy
 Change "25G-AUI C2M" to "C2M 25G-AUI" throughout the document.
 Similarly for 25G-AUI C2C.

Response Response Status C
 REJECT.

This comment is a restatement of a Draft 0.1 comment #138 by the same commenter. The comment was rejected by the task force.

Cl 109B SC 109B.1 P207 L40 # 121
 Dawe, Piers Mellanox

Comment Type E Comment Status A bucket
 Entries in key should be in alphabetical order.

SuggestedRemedy
 Move FEC entry to its place in alphabetical order.

Response Response Status C
 ACCEPT.

Cl 109B SC 109B.1 P207 L50 # 122
 Dawe, Piers Mellanox

Comment Type E Comment Status A
 Equation (83E-1) doesn't depict a typical 25G-AUI C2M application with loss budget per section.

SuggestedRemedy
 Delete "and Equation (83E-1)".

Response Response Status C
 ACCEPT.

Cl 109B SC 109B.1 P208 L3 # 123
 Dawe, Piers Mellanox

Comment Type E Comment Status R
 The 25G-AUI C2M interface is even more similar to chip-to-module CAUI-4 than to CEI-28G-VSR, and it helps the reader to know that.

SuggestedRemedy
 Change sentence to "The 25G-AUI C2M interface is almost identical to a single lane of the CAUI-4 chip-to-module interface, and is defined using a specification and test methodology that is similar to that used for CEI-28G-VSR defined in OIF-CEI-03.1 [B56]."

Response Response Status C
 REJECT.

In Draft 1.0, there is a new relaxed requirement for 1E-6 BER which is significantly dissimilar to either CAUI-4 or CEI-28G-VSR. The references to Annex 83E provide ample recognition of the existing commonalities.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 109B SC 109B.1.1 P208 L25 # 128
 Dawe, Piers Mellanox

Comment Type T Comment Status A

These "shalls" are not actionable here: neither host nor module can speak for the other party. There are separate shalls for host and module BER performance that may refer back to here.

SuggestedRemedy

Change
 bit error ratio (BER) shall be less than 10-15 with any errors sufficiently
 to
 bit error ratio (BER) specification is less than 10-15 with any errors sufficiently
 Change
 bit error ratio shall be less than 10-6 with any errors sufficiently
 to
 bit error ratio specification is less than 10-6 with any errors sufficiently
 Delete PICS row, item BER.

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement suggested remedy with editorial license.

CI 109B SC 109B.1.1 P208 L29 # 8
 Anslow, Pete Ciena

Comment Type T Comment Status A

A new paragraph has been added:
 "For a PHY with the RS-FEC sublayer (Clause 108), the 25G-AUI C2M bit error ratio shall be less than 10-6 with any errors sufficiently uncorrelated to ensure an acceptably high mean time to false packet acceptance (MTTFPA) assuming 64B/66B and RS-FEC encoding."
 With RS-FEC in operation the MTTFPA is protected by "The probability that the decoder fails to indicate a codeword with 8 or more symbol errors as uncorrected is expected to be lower than 10-6." in 108.5.3.2. The issue with correlated errors is whether the expected FLR at the MAC is met.

SuggestedRemedy

Change to: "For a PHY with the RS-FEC sublayer (Clause 108), the 25G-AUI C2M bit error ratio shall be less than 10-6 with any errors sufficiently uncorrelated to ensure a frame loss ratio (see 1.4.223) of less than 6.2×10^{-10} for 64-octet frames with minimum inter-packet gap when processed according to Clause 108."

Response Response Status C

ACCEPT.

CI 109B SC 109B.1.1 P208 L31 # 188
 Dawe, Piers Mellanox

Comment Type E Comment Status D

A point that is easily forgotten: what FEC options an implementation supports for the 25GBASE-CR family has no bearing on these options.

SuggestedRemedy

Add NOTE--The MDI for 25GBASE-CR and 25GBASE-CR-S is not 25G-AUI C2M, and this Annex does not apply to it.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 109B SC 109B.2 P208 L35 # 142
 Dudek, Mike QLogic

Comment Type T Comment Status A

The SFP MCB/HCB should be called out in addition to the QSFP test fixture referenced in 83E.

SuggestedRemedy

Change "25G-AUI C2M compliance points are defined in 83E.2." to "25G-AUI C2M compliance points are defined in 83E.2 with the exception that the single lane compliance boards specified in annex 110B can be used as alternates to the multi-lane compliance boards specified in clause 83E.4.1

Response Response Status C

ACCEPT IN PRINCIPLE.

Add text in 109B.2 as follows:

The HCB and MCB for a four-lane connector (e.g., QSFP) are described and specified in 92.11.2.

The HCB and MCB for a single-lane connector (e.g., SFP28) are described and specified in Annex 110B."

Modify the text in 109B.3.2.1 as follows.

Change:

"Figure 83E-11 depicts an example module output eye test configuration. The module output eye is measured at TP4, as shown in Figure 83E-5, using compliance boards specified in 92.11.2 (QSFP) or 110B.1.2 (SFP28)."

To:

"Figure 83E-11 depicts an example module output eye test configuration. The module output eye is measured at TP4, as shown in Figure 83E-5."

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl **109B** SC **109B.3.1** P**208** L**43** # **130**
 Dawe, Piers Mellanox

Comment Type **TR** Comment Status **D**

Do we want to give a 25G-AUI C2M host the same relief that we give to the module?

SuggestedRemedy

Consider modifying the host output and input specs in the same way as done for the module.

Proposed Response Response Status **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

Cl **109B** SC **109B.3.2** P**209** L**12** # **143**
 Dudek, Mike QLogic

Comment Type **T** Comment Status **A**

Is PHY the correct name here.

SuggestedRemedy

Consider changing "PHY that includes" to "Module used for a PHY that requires" on line 12 and "PHY that does not include" to "Module used for a PHY that does not require" Also the equivalent changes on page 209 line 53 and page 210 line 1.

Response Response Status **C**

ACCEPT IN PRINCIPLE.

The module is part of the PHY.

In keeping with the style of the subsequent subclauses make changes as follows.

Change:

"For a PHY that includes an RS-FEC sublayer (Clause 108), the eye opening shall meet the eye opening requirements in either 109B.3.2.1.1 or 109B.3.2.1.2."

To:

"For a PHY that includes an RS-FEC sublayer (Clause 108), the module output eye opening shall meet the eye opening requirements in either 109B.3.2.1.1 or 109B.3.2.1.2."

Change:

"For a PHY that does not include an RS-FEC sublayer, the eye opening shall meet the eye opening requirements in 109B.3.2.1.1."

To:

"For a PHY that does not include an RS-FEC sublayer, the module output eye opening shall meet the eye opening requirements in 109B.3.2.1.1."

Cl **109B** SC **109B.3.2** P**209** L**24** # **124**
 Dawe, Piers Mellanox

Comment Type **E** Comment Status **A**

I wondered why there were two references to define PRBS31. It turns out neither do, but they point towards it.

SuggestedRemedy

Insert reference to 49.2.8. Also in 109B.3.2.1.2.

Response Response Status **C**

ACCEPT IN PRINCIPLE.

The current references are to the PMA test pattern generators and checkers which are not relevant in this context. The reference in the suggested remedy is correct.

In 109B.3.2.1.1 and 109B.3.2.1.2...

Change:

"(109.4.5.1, 109.4.5.2)"

To:

"(49.2.8)"

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 109B SC 109B.3.2.1.2 P209 L28 # 125
 Dawe, Piers Mellanox

Comment Type E Comment Status A

Could give this a more specific name, especially as it's the basic, mainstream requirement for 25G-AUI C2M.

SuggestedRemedy

Change "alternate measurement method" to "25G-AUI C2M measurement method". Similarly for 109B.4.1 Alternate eye width, eye height, and eye closure measurement method.

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #144.

Change the 109B.3.2.1.1 heading name to: "Eye opening using measurement method A"

Change the 109B.3.2.1.2 heading name to: "Eye opening using measurement method B"

Change the Table 109B-2 heading name to: "Module stressed input parameters for measurement method B"

In 109B.3.4, label each of the two test scenarios "method A" and "method B" with editorial license.

In 109B.4.1...

Change "Eye width eye height" to "Eye width and eye height"

CI 109B SC 109B.3.2.1.2 P209 L35 # 126
 Dawe, Piers Mellanox

Comment Type E Comment Status A

a valid 25GBASE-R encoding with RS-FEC encoding.

SuggestedRemedy

a valid RS-FEC encoded 25GBASE-R signal.

Response Response Status C

ACCEPT IN PRINCIPLE.

Retain a similar style to the text for the scrambled idle pattern previously in the same sentence and for 25GBASE-R in the previous subclause.

Change:

"a valid 25GBASE-R encoding with RS-FEC encoding"

To:

"a valid 25GBASE-R signal with RS-FEC encoding"

CI 109B SC 109B.3.4.1 P210 L17 # 18
 Anslow, Pete Ciena

Comment Type E Comment Status A bucket

In Table 109B-1, "Table 88-13" should be green.

Same issue for "83E.4.2" in 109B.4.1

SuggestedRemedy

In Table 109B-1, apply character tag "External" to "Table 88-13"
 Do the same for "83E.4.2" in 109B.4.1

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 109B SC 109B.3.4.1 P210 L4 # 144
 Dudek, Mike QLogic

Comment Type T Comment Status A

The Title of this section is too broad, as the methodology in 83E.3.4.1 is also used for 25G_AUI C2M module stressed input test.

SuggestedRemedy

In Table 109B-1 change "the title to "Alternate 25G_AUI C2M module stressed input test.

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #125.

CI 109B SC 109B.3.4.1 P210 L6 # 145
 Dudek, Mike QLogic

Comment Type T Comment Status A

For the module stressed input test for use with an RS-FEC module the key difference is that the required BER is 1e-6 (not 1e-15). Also for clarity it would be good to make it clear that the eye height and eye width here should be EH8 and EW8

SuggestedRemedy

Add a row to Table 109B-1. Parameter Required BER Value <1e-6.

Add (EW8) to the Eye width parameter, and (EH8) to the Eye height parameter.

Response Response Status C

ACCEPT IN PRINCIPLE.

Table 109-1 specifies the characteristics of the input stressed signal not the result of the test. However, the referenced test methodology in 83E.3.4.1.1 does not specify a goal of the stressed test.

In Table 109-1.

Change "Eye width" to "Eye width (EW8)"
 Change "Eye height" to "Eye height (EH8)"

Add a sentence based on text in 83A.3.5.1 in 109B.3.4.1 as follows:

"The module receiver shall operate with a BER of better than 10⁻⁶ in the presence of a compliant stressed input signal."

Also, to add some clarification.

Change:

"The input shall satisfy the input tolerance defined in Table 109B-1."

To:

"The stressed input signal shall satisfy the input tolerance defined in Table 109B-1."

CI 109B SC 109B.3.4.1 P210 L7 # 129
 Dawe, Piers Mellanox

Comment Type TR Comment Status A

Need to tie the module stressed input test back to 109B.1.1 Bit error ratio.

SuggestedRemedy

Change

with the exception that the input eye height and eye width are measured according to the method in 109B.4.1.

to

with the following exceptions:

- a) The input eye height and eye width are measured according to the method in 109B.4.1.
 - b) The module 25G-CAUI-4 receiver under test shall meet the BER requirement for a PHY with the RS-FEC sublayer given in 109B.1.1, using three Recommended_CTLE_value values for both the high loss test and low loss test.
- Adjust PICS to reflect this.

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #145.

CI 109B SC 109B.3.4.1 P210 L8 # 156
 Maki, Jeffery Juniper Networks, Inc.

Comment Type T Comment Status A different CTLE lane

Text needs to be added to make clear that Recommended_CTLE_value is per 25G-AUI lane regardless of whether the module supports a single lane or multiple lanes of 25G-AUI. A common module should not be required or implied to use the same Recommended_CTLE_value for each 25G-AUI.

SuggestedRemedy

Add text to the end of the paragraph, "If a Clause 45 MDIO is implemented, the variable Recommended_CTLE_value is accessible through register 1.169 (see 45.2.1.96) and is to be understood to be for a single 25G-AUI regardless of whether using a single or multi-port 25 Gb/s module. Each 25G-AUI may be different in a multi-port 25 Gb/s module."

Response Response Status C

ACCEPT IN PRINCIPLE.

The commenter is concerned that the word module is ambiguous as to whether it refers to a single lane or multiple lanes.

Add the following paragraph on Page 207 line 49:

"For module form factors with more than one 25GAUI C2M lane, each lane comprises a separate 25GAUI C2M interface with independent management."

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 109B SC **109B.4.1** P**211** L**14** # **183**
 Dawe, Piers Mellanox
 Comment Type **E** Comment Status **A**
 see Equation (109B-2)
 SuggestedRemedy
 is the eye height defined in Equation (109B-2).
 Response Response Status **C**
 ACCEPT.

Cl 109B SC **109B.5.2.2** P**210** L**46** # **186**
 Dawe, Piers Mellanox
 Comment Type **E** Comment Status **A**
 from CDFL
 from the CDFR
 ...
 from CDF1
 from CDF0
 SuggestedRemedy
 Change "from the CDFR" to "from CDFR".
 Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 [The editor changed the page/line from 212/37 to 210/46.]
 A similar discrepancy occurs in the text above.
 Page 210 line 30 change "of the CDFL and CDFR" to "of CDFL and CDFR"
 page 210 line 46 change "from the CDFR" to "from CDFR" (per suggested remedy)

Cl 109B SC **109B.5.2.2** P**212** L**37** # **184**
 Dawe, Piers Mellanox
 Comment Type **E** Comment Status **A** bucket
 Clause 109B
 SuggestedRemedy
 Annex 109B
 Response Response Status **C**
 ACCEPT.

Cl 109B SC **109B.5.2.2** P**212** L**50** # **185**
 Dawe, Piers Mellanox
 Comment Type **E** Comment Status **A** bucket
 Orphan heading
 SuggestedRemedy
 Keep with table on next page.
 Response Response Status **C**
 ACCEPT.

Cl 109B SC **109B.5.3** P**213** L**11** # **191**
 Dawe, Piers Mellanox
 Comment Type **ER** Comment Status **A** bucket
 MM1 and MM2 are not major options because knowledge of their states does not allow for any different action to knowledge of the major option above, RSFEC.
 SuggestedRemedy
 Move these two PICS to 109B.5.4.2 Module output.
 Response Response Status **C**
 ACCEPT.

Cl 109B SC **109B.5.3** P**213** L**9** # **190**
 Dawe, Piers Mellanox
 Comment Type **ER** Comment Status **A**
 A 25G-AUI host or module might support both a RS-FEC port type and a non-RS-FEC port type, if one existed that one would use with 25G-AUI C2M. I believe that for the 25G-AUI, PHY support of 25G RS-FEC is effectively mandatory, because it's a subset of (easier than) PHY support of 25G non-RS-FEC.
 SuggestedRemedy
 Change "PHY support of 25G RS-FEC" to "No PHY support of 25G without RS-FEC".
 There may be other ways to build the PICS logic.
 Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 Change: "PHY support of 25G RSFEC"
 To: "PHY with 25G RS-FEC"

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl **109B** SC **109B.5.3** P**213** L**9** # **181**
 Dawe, Piers Mellanox
 Comment Type **E** Comment Status **A** bucket
 Wrong subclause. The two choices are first laid out in 109B.1.1.
 SuggestedRemedy
 Change "109B.3.2.1" to "109B.1.1".
 Response Response Status **C**
 ACCEPT.

Cl **109B** SC **109B.5.4.2** P**214** L**19** # **182**
 Dawe, Piers Mellanox
 Comment Type **E** Comment Status **A** bucket
 Signal rate
 SuggestedRemedy
 Signaling rate
 Response Response Status **C**
 ACCEPT.

Cl **109B** SC **109B.5.4.4** P**215** L**15** # **127**
 Dawe, Piers Mellanox
 Comment Type **ER** Comment Status **D** withdrawn
 PICS RM1, 25G-AUI module input characteristics, doesn't agree with the text in 109B.3.4.
 SuggestedRemedy
 Remedy to follow.
 Proposed Response Response Status **Z**
 PROPOSED REJECT.
 This comment was WITHDRAWN by the commenter.

Cl **109B** SC **109B.5.4.4** P**215** L**15** # **189**
 Dawe, Piers Mellanox
 Comment Type **ER** Comment Status **A**
 PICS RM1, 25G-AUI module input characteristics, and RM2, BER requirement, don't agree with the text in 109B.3.4.
 SuggestedRemedy
 Change "83E.3.4" to "109B.3.4".
 Change "Table 83E-7" to "83E.3.4 except module stressed input test."
 Create PICS options for module stressed input test requirements in 83E.3.4.1 or 109B.3.4.1, dependent on RSFEC major option, as done for MM1 and MM2. These could replace RM2. For the CAUI-4 method, subclause 83E.3.4.1 and value/comment "As 83E.4.1.1 with settings associated with Recommended_CTLE_value, Table 83E-8". For the 25G-AUI/alternate method, subclause 109B.3.4.1 and value/comment "As 109B.3.4.1 with settings associated with Recommended_CTLE_value, Table 109B-1".
 Response Response Status **C**
 ACCEPT IN PRINCIPLE.
 Implement suggested remedy with editorial license.

Cl **109B** SC **109B.5.4.4** P**215** L**18** # **154**
 Maki, Jeffery Juniper Networks, Inc.
 Comment Type **ER** Comment Status **A** bucket
 Draft 1.0 was not updated properly to reflect the final response made for Comment #110 against Draft 0.1.
 SuggestedRemedy
 Value/Comment for Item RM2 should refer to 83E.3.4.1.1, not 83E.4.1.1. There is no subclause 83E.4.1.1.
 Response Response Status **C**
 ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl **109B** SC **3.2.1.2** P**209** L**28** # **113**
 Nicholl, Gary Cisco Systems

Comment Type **T** Comment Status **D**

It is not clear to me how you measure (guarantee) the module electrical output to a ber of 1e-8 based on an optical input that is only spec'ed to a ber of 1e-6, without having to add a PRBS generator in the PMA function within the module (i.e. generate the electrical output signal independently from the optical input signal). I thought that being able to avoid the PRBS generator was one of the main drivers for making this change and moving away from the standard CAUI-4 module output electrical ber spec of 1e-15 ?

SuggestedRemedy

Proposed Response Response Status **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

Cl **109B** SC **3.4.1** P**210** L**7** # **114**
 Nicholl, Gary Cisco Systems

Comment Type **T** Comment Status **D**

It is not clear to me how you measure (guarantee) the module electrical input to a ber of 1e-8 based on an optical output that is only spec'ed to a ber of 1e-6, without having to add a PRBS checker in the PMA function within the module (i.e. directly detect errors on the electrical input signal, without having to use the optical output to monitor the errors externally). I thought that being able to avoid the PRBS checker was one of the main drivers for making this change and moving away from the standard CAUI-4 module input electrical ber spec of 1e-15 ?

SuggestedRemedy

Proposed Response Response Status **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

Cl **109C** SC P**218** L**26** # **67**
 Froroth, Ingvar Marvell

Comment Type **T** Comment Status **A**

Figure caption at Figure 109C-4 says "Separate SERDES for optical module interface" but the Figure itself does not indicate which parts provide the SERDES functionality, nor is there any accompanying text stating this.

SuggestedRemedy

Although this deficit is consistent with Figure 83C-5 of 802.3-2012, my suggestion is to add an indication in Figure 109C-4, at the parts representing the SERDES.

Response Response Status **C**

ACCEPT IN PRINCIPLE.

Change "separate SERDES" to "Intermediate PMA device".

Cl **110** SC **11** P**153** L**33** # **157**
 Andrewartha, Mike Microsoft

Comment Type **T** Comment Status **D** MDI

Need to state the requirement for AC coupling in the plug connector. 110.11 refers to 92

SuggestedRemedy

Proposed Response Response Status **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 110 SC 11 P153 L33 # 165
Andrewartha, Mike Microsoft

Comment Type **TR** Comment Status **A** MDI

Need to define the requirement for AC coupling in the plug connector. This section refers to 92.12.1.1 but the requirement for AC coupling in the plug connector is in 92.12.1. This requirement is explicitly called out in 110.11.1.

SuggestedRemedy

Either add text similar to 110.11.1, page 154, lines 1-4 or add a reference to 92.12.1 to incorporate the requirement there.

Response Response Status **C**

ACCEPT IN PRINCIPLE.

Move the following paragraph from 110.11.1 to 110.11 (just before heading 110.11.1)...
"For 25GBASE-CR plug connectors, the receive lanes are AC-coupled; the AC-coupling shall be within the plug connectors. It should be noted that there may be various methods for AC-coupling in actual implementations. The low-frequency 3 dB cutoff of the AC-coupling shall be less than 50 kHz. It is recommended that the value of the coupling capacitors be 100 nF. The capacitor limits the inrush charge and baseline wander."

CI 110 SC 11 P153 L42 # 167
Andrewartha, Mike Microsoft

Comment Type **T** Comment Status **D** MDI

The text refers to 92.12.1.1 to the exclusion of 92.12.1.2. Is the intent to exclude use of the style 2 QSFP28? If so, why?

SuggestedRemedy

Resolution depends on the intent. If style 2 has been deprecated in the industry, a note to that effect is in order. If not, then a reference to 92.12.1.2 makes sense.

Proposed Response Response Status **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

CI 110 SC 110.1 P137 L2 # 15
Anslow, Pete Ciena

Comment Type **E** Comment Status **A** bucket

The base standard is almost consistent in using "interpacket" rather than "inter-packet"
Same issue in 111.1 and 112.1.1 (2 instances)

SuggestedRemedy

Change "inter-packet" to "interpacket"
Make the same change in 111.1 (page 163 line 54) and 112.1.1 (page 181 lines 42 and 43)

Response Response Status **C**

ACCEPT IN PRINCIPLE.

Implement the suggested remedy, also including:
page 76, line 16, clause 105

Affects clauses: 105, 110, 111, 112

CI 110 SC 110.10 P149 L35 # 92
Mellitz, Richard Intel Corporation

Comment Type **TR** Comment Status **D** Cable reach

A base-R FEC cable assembly can support a cable up to at least 4 meters
See mellitz_3by_01_0515.pdf

SuggestedRemedy

Change:
b) Cable assembly short (CA-S): Cable assembly that supports links between two PHYs that operate in BASE-R FEC mode, with cable length up to 3 m.
To
b) Cable assembly short (CA-S): Cable assembly that supports links between two PHYs that operate in BASE-R FEC mode, with cable length up to 4 m.

Proposed Response Response Status **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 110 SC 110.10 P149 L38 # 93
 Mellitz, Richard Intel Corporation

Comment Type TR Comment Status R Cable reach

A no-FEC cable assembly can support a cable up to at least 3 meters with a COM limit of 2.5dB
 See mellitz_3by_01_0515.pdf

SuggestedRemedy

Change
 c) Cable assembly no-FEC (CA-N): Cable assembly that supports links between two PHYs that operate in no-FEC mode, with cable length up to 2 m.
 To
 c) Cable assembly no-FEC (CA-N): Cable assembly that supports links between two PHYs that operate in no-FEC mode, with cable length up to 3 m.

Response Response Status C

REJECT.

There is no consensus to make the suggested change.

CI 110 SC 110.10 P151 L10 # 140
 Dudek, Mike QLogic

Comment Type T Comment Status A

The CA-N cable is expected to be shorter than the CA-S cable (and shorter than 4 meters). There is no need to have a smaller frequency step for the CA-N cable.

SuggestedRemedy

Change the maximum frequency steop for CA-N cable to be the same as CA-S is 0.01GHz.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change max. step for CA-N to 0.01.

Move the footnote in the CA-L column to the text "Maximum frequency step".

CI 110 SC 110.10.2 P150 L24 # 86
 Mellitz, Richard Intel Corporation

Comment Type TR Comment Status R Cable reach

A base-R FEC cable assembly can support a cable up to at least 4 meters and a no-FEC cable assembly up to a least 3 meters
 See mellitz_3by_01_0515.pdf

SuggestedRemedy

Change
 The measured insertion loss at 12.8906 GHz of the CA-S cable assembly shall be less than or equal to 16.48 dB. The measured insertion loss at 12.8906 GHz of the CA-N cable assembly shall be less than or equal to 12.98 dB
 To
 The measured insertion loss at 12.8906 GHz of the CA-S cable assembly shall be less than or equal to 19.48 dB. The measured insertion loss at 12.8906 GHz of the CA-N cable assembly shall be less than or equal to 15.98 dB

Response Response Status C

REJECT.

There is no consensus to make the suggested change.

CI 110 SC 110.10.7 P151 L1 # 89
 Mellitz, Richard Intel Corporation

Comment Type TR Comment Status R Cable COM

For low a DER of 1e-12, COM may be somewhat pessimistic which may result in 3 meter cables not passing the COM of 3dB.

SuggestedRemedy

Change:
 "COM for any channel within the cable assembly shall be greater than or equal to 3 dB for each test."
 To
 "COM for any channel within the CA-S and CA-L cable assembly shall be greater than or equal to 3 dB for each test. COM for any channel within the CA-N cable assembly shall be greater than or equal to 2.5 dB for each test."

This solution is least disruptive to schedule and change creep. See presentation mellitz_3by_01_0515.pdf

Response Response Status C

REJECT.

There is no consensus to make the suggested change.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 110 SC 110.10.7.1.2 P152 L17 # 49
 Ran, Adee Intel

Comment Type E Comment Status A bucket

The editor's note about implementation of comment #103 against D0.1 (scattering parameters) is not needed in the next draft.

SuggestedRemedy

Delete editor's note.

Response Response Status C

ACCEPT.

CI 110 SC 110.11 P154 L6 # 54
 Booth, Brad Microsoft

Comment Type TR Comment Status R MDI

I'm a bit concerned that we're reference specific connectors (SFP28 and QSFP28) for the 25GBASE-CR and CR-S port types to be IEEE 802.3by compliant. I believe the goal should be to require compliance to the electrical parameters and not to the mechanical requirements.

SuggestedRemedy

Text commonly used has been:
 "When the MDI is a connector plug and receptacle connection, it shall meet the interface performance specifications of the following:"
 And any mechanical information is described as:
 "These connectors are depicted (for informational use only) in..."

Response Response Status C

REJECT.

The mechanical interface (MDI) is specified to ensure plug compatibility at the MDI. The MDI couples the PMD (110.7 and 110.8) to the cable assembly (110.10). The specification is consistent with Clause 92 and Clause 85.

CI 110 SC 110.6 P140 L1 # 16
 Anslow, Pete Ciena

Comment Type E Comment Status A bucket

The IEEE Editorial style manual contains:
 "In a series of three or more terms, use a comma immediately before the coordinating conjunction (usually and, or, or nor)."

SuggestedRemedy

Change:
 "in RS-FEC, BASE-R FEC or no-FEC mode." to:
 "in RS-FEC, BASE-R FEC, or no-FEC mode."

Response Response Status C

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 110 SC 110.6 P140 L1 # 116
 Goergen, Joel Cisco Systems, Inc.

Comment Type T Comment Status A CC

There needs to be clarity in the clause 105 introduction to the operation of the phy types with respect to -L/-S/ -N. I wrote three comments already on the -N option as it is confusing the the phy types themselves can each support the 3 cable types.

SuggestedRemedy

address the cable assembly operational modes in clause 105 by addressing teh cable types within each phy type.

Response Response Status C

ACCEPT IN PRINCIPLE.

[Editor changed subclause from 6 to 110.6]

Clause 105 is an introduction to 25G Ethernet, which covers several media types.

The cable assembly types are specific to Clause 110. They are defined in 110.10 and their compatibility with the two PHY types is listed.

However, the difference between 25GBASE-CR and 25GBASE-CR-S in terms of cable assembly types is described only in subclause 110.6, although these terms are used many times earlier in the text.

It seems worthwhile to note the difference as part of the overview subclause.

Insert a new paragraph after Table 110-1:

"A 25GBASE-CR PHY supports operation over cable assemblies of types CA-N, CA-S and CA-L (see 110.10). A 25GBASE-CR-S PHY supports operation over cable assemblies of types CA-N and CA-S, but not CA-L. A 25GBASE-CR-S PHY interoperates with a 25GBASE-CR PHY."

Cl 110 SC 110.6 P140 L7 # 166
 Andrewartha, Mike Microsoft

Comment Type T Comment Status A CC

Based on the current definition of AN, selecting no-FEC operation is likely to require management intervention since the decision is based on the logical 'OR' of the requested bits. As such, the determination of FEC mode should also be under management control. The current text of this paragraph only mentions selection via AN.

SuggestedRemedy

Change "The FEC mode is determined using AN (Clause 73) and is used" to

"The FEC mode is determined using AN (Clause 73) or management control and is used"

Response Response Status C

ACCEPT IN PRINCIPLE.

[Editor changed subclause from 6 to 110.6]

Change:

"The FEC mode is determined using AN (Clause 73) and is used"

To:

"The FEC mode is determined using AN (Clause 73) if AN is enabled or management control if AN is disabled and is used"

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 110 SC 110.7 P140 L19 # 118
Goergen, Joel Cisco Systems, Inc.

Comment Type ER Comment Status A

figure 110-2 defines the cable type from center of connector to center of connector, yet several terms for the cable assembly are actually defined as TP1 to TP4. The term "cable assembly" is used interchangeably between the two definitions, causing confusion on the allocated budgets.

even though this is called out in 110.10 page 149 line 27. Here is is refered to "cable assembly specification"

SuggestedRemedy

Change the term "cable assembly" to "cable assembly specification" in the figure. or modify the following values to include TP1-TP4 or identify as cable assembly specification. table 110-9 assets TP1-TP4, but clearly discusses the cable assembly and no the specification points.

page 220 line 4/5 addresses cable assembly but line 3 clearly lists points. so do I use the picture definition of cable assembly or the point definiton?

page 220 line 23/24 same thing

page 220 line 28/29 same thing

Figure 110A-1 again defines cable assembly as between the two connector sets. yet all definitions appear to reference cable assembly specification as between TP1 and TP4. so same thing here. change cable assembly or add the cable assembly specification clearly in the figure.

Response Response Status C

ACCEPT IN PRINCIPLE.

[Editor changed subclause from 7 to 110.7]

In "110.10 Cable assembly characteristics" (P149, L27) test points and test fixtures are identified. Figure 110A-1 depicts TP1 and TP4 as cable assembly measurment reference.

Make the following changes...

Page 220 L4 add (TP1-TP4) after "maximum cable assembly insertion loss" and add (TP1-TP4) after "is the minimum cable assembly insertion loss".

Cl 110 SC 110.8.4.2 P144 L47 # 46
Ran, Adee Intel

Comment Type E Comment Status A bucket

The editor's note about required block error ratio in BASE-R FEC mode is not needed in the next draft.

SuggestedRemedy

Delete editor's note.

Response Response Status C

ACCEPT.

See also comment #146.

Cl 110 SC 110.8.4.2 P144 L48 # 146
Dudek, Mike QLogic

Comment Type TR Comment Status A RX test

The editor's note suggests that Block error ratio of 1e-8 was accepted. That is not correct. A BER of 1e-8 was the accepted value. A BASE-R FEC block error ratio of 4.7 e-10 is a much more stringent value. (The block error ratio should be the block length x the BER. The BASE-R FEC block length is 2112 bits and therefore the block error ratio should be 2.1 e-5).

SuggestedRemedy

In table 110-6

Either Change the BASE-R block error ratio required row to Bit error Ratio required with a value of <1e-8 and delete footnote a.

Or change the Block error ratio to <2.1 e-5

Also delete the editor's note.

Response Response Status C

ACCEPT IN PRINCIPLE.

Comment #50 against D0.1 suggested "Replace TBD with 10⁻⁸ (.) [in] Page 150 line 17". That TBD was the block error ratio. The comment was accepted, but the value suggested is unsuitable for the block error ratio (although it was suitable in two other places).

The block error ratio required for BASE-R FEC to achieve FLR=6.2e-10 is calculated in slide 3 of

http://www.ieee802.org/3/by/public/adhoc/architecture/ran_020415_25GE_adhoc.pdf as 4.71e-10.

To achieve this block error ratio (as uncorrectable blocks), the corrected block error ratio is 2.1e-5, as the commenter suggests.

Change the block error ratio maximum to 2.1e-5.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 110 SC 110.8.4.2 P145 L28 # 91
 Mellitz, Richard Intel Corporation

Comment Type TR Comment Status A Cable COM, RX test

In 110.8.4.2.3 Test channel calibration, page 147, the channel is specified between reference points in as figure 110-4 (page 147).. This include a cable assembly and approximately of 6.9dB instrument cabling from the test fixture to the instrument as deduced from table 10-5(29.44dB-22.5dB). The data from mellitz_040815_25GE_adhoc.pdf and shanbhag_020415_25GE_adhoc_v2.pdf suggest that many 3 meter cables have a COM of approximately 4 dB. The data in mellitz_3by_01_0515.pdf will show that a 4.2 meter cable is closer to the length that can support at least 3dB COM. The values for a1, a2, and a4 should reflect a this cable plus the instrumented cable.

SuggestedRemedy

In table 10-7 change COM (max) to 2.5 dB and use a1,a2, and a4 suggested in mellitz_3by_01_0515.pdf

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #149.

CI 110 SC 110.8.4.2 P145 L40 # 149
 Dudek, Mike QLogic

Comment Type TR Comment Status A RX test

In Table 110-6 the high loss case 2 should be based on the performance of the CA-S cable. The difference in loss between the CA-S cable and the CA-L cable is 6dB therefore the max loss for the Base-R FEC mode should be 23.44dB. The values of the "a" components in the suggested response have been scaled from the RS-FEC case to achieve this loss.

SuggestedRemedy

Change the Test 2 values in Table 110-6 as below.

a1 from 3.96 to 3.42
 a2 from 0.18 to 0.46
 a4 no change.
 Approximate fitted loss at 12.89 GHz from 21.04dB to 23.44dB.

Response Response Status C

ACCEPT.

See comment #91.

CI 110 SC 110.8.4.2 P145 L45 # 88
 Mellitz, Richard Intel Corporation

Comment Type TR Comment Status A Cable COM, RX test

In 110.8.4.2.3 Test channel calibration, page 147, the channel is specified between reference points in as figure 110-4 (page 147).. This include a cable assembly and approximately of 6.9dB instrument cabling from the test fixture to the instrument as deduced from table 10-5(29.44dB-22.5dB). The data from mellitz_040815_25GE_adhoc.pdf and shanbhag_020415_25GE_adhoc_v2.pdf suggest that many 3 meter cables have a COM of approximately 4 dB. The data in mellitz_3by_01_0515.pdf will show that a 4.2 meter cable is closer to the length that can support at least 3dB COM. The values for a1, a2, and a4 should reflect a this cable plus the instrumented cable.

SuggestedRemedy

In table 10-7 change COM (max) to 2.5 dB and use a1,a2, and a4 suggested in mellitz_3by_01_0515.pdf

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #150.

CI 110 SC 110.8.4.2 P146 L1 # 90
 Mellitz, Richard Intel Corporation

Comment Type TR Comment Status A Cable COM, RX test

In 110.8.4.2.3 Test channel calibration, page 147, the channel is specified between reference points in as figure 110-4 (page 147). This include a cable assembly and approximately of 6.9dB instrument cabling from the test fixture to the instrument as deduced from table 10-5 (29.44dB-22.5dB). The data from mellitz_3by_01_0315.pdf should be refined to add in the instrumented cable. In addition if the COM limit of 2.5dB is accepted 110.10.7, the data for a1, a2, and a4 should reflect a 3 meter cable plus the instrumented cable.

SuggestedRemedy

In table 110-7 change COM (max) to 2.5 dB and use a1,a2, and a4 suggested in mellitz_3by_01_0515.pdf

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #149.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 110 SC 110.8.4.2 P146 L12 # 150
 Dudek, Mike QLogic

Comment Type TR Comment Status A RX test

In Table 110-7 the high loss case 2 should be based on the performance of the CA-N cable. The difference in loss between the CA-N cable and the CA-L cable is 9.5dB therefore the max loss for the no-FEC mode should be 19.94dB. The values of the "a" components in the suggested response have been scaled from the RS-FEC case to achieve this loss.

SuggestedRemedy

Change the Test 2 values in Table 110-7 as below.

a1 from 3 to 2.91
 a2 from 0.29 to 0.39
 a4 from 0.02 to 0.03.
 Approximate fitted loss at 12.89 GHz from 21.04dB to 19.94dB.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change the Test 2 values in Table 110-7 as below.

a1 from 3 to 2.91
 a2 from 0.29 to 0.39
 a4 from 0.02 to 0.027.
 Approximate fitted loss at 12.89 GHz from 17.57dB to 19.96dB.

CI 110 SC 110.8.4.2 P148 L5 # 131
 Dudek, Mike QLogic

Comment Type E Comment Status A RX test

Related to Comment #53 to draft 0.1 and the editor's note on page 148. This original comment was related to the Receiver interference tolerance test not the jitter tolerance test and this reference should be pointing to the Interference tolerance parameter jitter in tables 110-5, 110-6 and 110-7 as appropriate. (in table 92-8 of 802.3bj, not the Jitter tolerance jitter in table 92-9 of 802.3bj. Note that table 110-8 is identical to table 92-9, not to table 92-8.)

SuggestedRemedy

On page 148 line 5 change Table 110-8 to Table 110-5, Table 110-6 or Table 11-7 as appropriate.

Delete the editor's note on page 148 line 36.

Response Response Status C

ACCEPT IN PRINCIPLE.

Draft 0.1 comment #53 wasn't implemented correctly.

Implement the suggested remedy with editorial licence.

CI 110 SC 110.8.4.2.1 P146 L42 # 147
 Dudek, Mike QLogic

Comment Type TR Comment Status A RX test

In order to calibrate COM, noise needs to be added to the signal.

SuggestedRemedy

Add a summing junction and "Channel noise source" box (similar to that shown in Figure 93C-2 between the pattern generator and Test reference in Figures 110-3, or re-label the Pattern Generator box "Pattern Generator with noise injection".

Response Response Status C

ACCEPT IN PRINCIPLE.

Re-label the Pattern Generator box "Pattern Generator with noise injection".

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 110 SC 110.8.4.2.3 P147 L46 # 134
 Dudek, Mike QLogic
 Comment Type E Comment Status A bucket
 This is a good solution to the Comment #52.
 SuggestedRemedy
 Delete the editor's note.
 Response Response Status C
 ACCEPT.
 See comment #47.

Cl 110 SC 110.8.4.3 P148 L36 # 48
 Ran, Adee Intel
 Comment Type E Comment Status A bucket
 The editor's note about implementation of comment #53 against D0.1 (jitter tolerance tables) is not needed in the next draft.
 SuggestedRemedy
 Delete editor's note.
 Response Response Status C
 ACCEPT.

Cl 110 SC 110.8.4.2.3 P147 L46 # 47
 Ran, Adee Intel
 Comment Type E Comment Status A bucket
 The editor's note about implementation of comment #52 against D0.1 (modified COM parameters) is not needed in the next draft.
 SuggestedRemedy
 Delete editor's note.
 Response Response Status C
 ACCEPT.

Cl 110 SC 110.8.4.3 P148 L26 # 17
 Anslow, Pete Ciena
 Comment Type E Comment Status A bucket
 "test 2 as specified of Table 110-5" should be "test 2 as specified in Table 110-5"
 SuggestedRemedy
 Change "test 2 as specified of Table 110-5" to "test 2 as specified in Table 110-5"
 Response Response Status C
 ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 110 SC 110.9 P149 L10 # 139
 Dudek, Mike QLogic

Comment Type T Comment Status A

This paragraph applies to 25GBASE-CR-S as well.

SuggestedRemedy

Change the order of the two paragraphs so that the fact that "25GBASE-CR channel" is generically used for both 25GBASE-CR and 25GBASE-CR-S appears first.

Response Response Status C

ACCEPT IN PRINCIPLE.

25GBASE-CR and 25GBASE-CR-S are PHY types. There is no similar separation of channels or cable assemblies. To reduce ambiguity, cable assemblies and channels should not be labeled "25GBASE-CR".

Apply the following:

Remove the labels "25GBASE-CR" in figures 110-3 and 110-4.

Change "The 25GBASE-CR channel" to "The channel", in 110.7.1 (page 141 line 6), table 110-4 (page 141 line 23), and 110.9 (page 149 lines 9 and 14).

In 110.10, change "25GBASE-CR cable assemblies" to "Cable assemblies defined in this subclause".

In 110.10.7.2, change "Several 25GBASE-CR cable assembly form factors are available" to "Cable assemblies have several form factors".

On page 149 line 31, delete "for 25GBASE-CR"

On page 143 lines 44/50, delete "25GBASE-CR"

In Annex 110C,
 Change "25GBASE-CR cable assembly" to "cable assembly".
 Change "25GBASE-CR link" to "25 Gb/s link".
 Change "25GBASE-CR host" to "host".

Implement with editorial license.

Cl 110A SC 110A.5 P220 L37 # 87
 Mellitz, Richard Intel Corporation

Comment Type TR Comment Status R

A base-R FEC cable assembly can support a cable up to at least 4 meters and a no-FEC cable assembly up to a least 3 meters
 See mellitz_3by_01_0515.pdf

SuggestedRemedy

In Table 110A-1
 Change
 IL_Chmax for CA-S
 From 29 to 31
 IL_Camax for CA-S
 From 16.48 to 19.48

 IL_Chmax for CA-N
 From 25.5 to 28
 IL_Camax for CA-N
 From 12.48 to 15.48

And on page 227 line 40ff table 110C-a
 Change CA-S references for RS-FEC, BASE-R FEC from 3m to 4m and
 Change CA-N references for RS-FEC, BASE-R FEC, no FEC from 2m to 4m

Response Response Status C

REJECT.

There is no consensus to make the suggested change.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

CI 110A SC 5 P220 L35 # 120
Goergen, Joel Cisco Systems, Inc.

Comment Type TR Comment Status R

CA-N definition for a 2M cable has a camin listed at 8dB. this is not low enough for a sub 2M cable used in a stacking environment that could be less then 8dB.

SuggestedRemedy

Since a 2M solution is not of value, but perhaps within a rack, delete the -N assembly, set the -S assembly to optional FEC using the base-r FEC, reduce the COM margin to 2dB.

On could also limit the style of connector to single stack in the 3M no fec solution, there by giving .62dB times 2 back to the over all margin. the 1dB margin debated in COM could come from here. -S no fec solution could be defined as a non stacking device. I would prefer to see the co margin lowered to 2dB and the -N assembly removed from teh document.

I will present something on this.

Response Response Status C

REJECT.

Commentor proposes options to delete -N on basis that "2M solution is not of value, but perhaps within a rack" but the support for 2m reach is largely based on "within a rack" e.g., server-TOR. See andrewartha_3by_01a_0115.pdf.

Commentor does not provided sufficient information to implement deletion of -N in draft.

No presentation was submitted or presented.

Further, regarding minimum IL, please note the minimum IL (8 db @12.8906 GHz) applies to all of the 802.3by and 802.3bj cable assemblies i.e., the measured insertion loss of the CA-L, CA-S and CA-N cable assembly shall be greater than or equal to the minimum cable assembly insertion loss given in Equation (92-26) and illustrated in Figure 92-12.

CI 110B SC 110B.1 P222 L14 # 101
Lusted, Kent Intel

Comment Type ER Comment Status A

Type "QFP28"

SuggestedRemedy

Change "QFP28" to "QSFP28"

Response Response Status C

ACCEPT IN PRINCIPLE.

Correct all occurences of "QFP28" to "QSFP28".

CI 110B SC 110B.1 P222 L18 # 44
Ran, Adee Intel

Comment Type E Comment Status A bucket

Cable assembly form factor is a new concept. Form factors mentioned here should have references to the definitions.

SuggestedRemedy

Insert "(See 110C.3.1)" after "SFP28-SFP28 form factor".
Insert "(See 110C.3.2)" after "QSFP28-QSFP28 form factor".
Insert "(See 110C.3.3)" after "QSFP28-4xSFP28 form factor".

Response Response Status C

ACCEPT.

CI 110B SC 110B.1 P222 L18 # 102
Lusted, Kent Intel

Comment Type ER Comment Status A

Type "QFP28"

SuggestedRemedy

Change "QFP28" to "QSFP28"

Response Response Status C

ACCEPT IN PRINCIPLE.

See comment #101.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 110B SC 110B.1.1 P222 L29 # 103
 Lusted, Kent Intel
Comment Type ER Comment Status A test fixture heading
 Subclause title does not follow convention from 802.3bj-2014.
SuggestedRemedy
 consider changing title to "TP2 or TP3 Test fixture" to align with P802.3bx D3.1 CI 92.11.1
Response Response Status C
 ACCEPT IN PRINCIPLE.
 The D1.0 subclauses were named to explicitly identify with SFP28 therefore exact alignment with 92.11.1 (TP2 or TP3) is not sufficient.
 Change:
 "110B.1.1 SFP28 Host test fixture"
 To:
 "110B.1.1 SFP28 TP2 or TP3 test fixture".
 In paragraph below change "The host test fixture" to "The test fixture".
 Apply to relevant references in other Clauses/Annexes.

Cl 110B SC 110B.1.2 P222 L44 # 104
 Lusted, Kent Intel
Comment Type ER Comment Status R test fixture heading
 Subclause title does not follow convention from 802.3bj-2014.
SuggestedRemedy
 consider changing title to "TP2 or TP3 Cable Assembly Test fixture" to align with P802.3bx D3.1 CI 92
Response Response Status C
 REJECT.
 In IEEE Std 802.3bj-2014 cable assembly test fixture is 92.11.2 Cable assembly test fixture.
 The subclauses were named to explicitly to identify with SFP28 therefore exact alignment with 92.11.2 is not sufficient.

Cl 110B SC 110B.1.3.6 P224 L7 # 141
 Dudek, Mike QLogic
Comment Type T Comment Status A
 For SFP mated test fixtures there is no Far end aggressor.
SuggestedRemedy
 Delete equation 110B-2 and the Far end aggressor amplitude and risetime in Table 110B-2. Also delete "and Fft" and "and Tft respectively" in the sentence on line 15.
Response Response Status C
 ACCEPT.

Cl 111 SC 111.6 P167 L5 # 65
 Dove, Daniel Dove Networking Solut
Comment Type TR Comment Status R
 I'm concerned about only a "recommendation" given that there is no guarantee of interoperability if the FEC mode does not match the channel type.
SuggestedRemedy
 Replace "It is recommended" with "In order to ensure interoperability, it is required"
Response Response Status C
 REJECT.
 Information about channel type may not be available in some backplane Ethernet devices, so this cannot be a normative requirement.
 Backplane systems should be configured for correct operation, e.g. by programming switches to request FEC mode in AN based on channel knowledge.
 AN does not ensure interoperability over an unknown channel. (compare: 100GBASE-CR4 across a 40GBASE-CR4 cable).

Cl 111 SC 111.8.3 P170 L6 # 50
 Ran, Adee Intel
Comment Type E Comment Status A bucket
 The editor's note about implementation of comment #59 against D0.1 (receiver characteristics in no-FEC and BASE-R FEC modes) is not needed in the next draft.
SuggestedRemedy
 Delete editor's note.
Response Response Status C
 ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 111 SC 111.8.3.1 P171 L17 # 148
 Dudek, Mike QLogic

Comment Type **TR** Comment Status **A** RX test

The BASE-R FEC block error ratio required for BASE-R FEC is incorrect. The block error ratio should be the block length x the BER. The BER required is 1e-8. The BASE-R FEC block length is 2112 bits and therefore the block error ratio should be 2.1 e-5

SuggestedRemedy

In table 111-5
 Either Change the BASE-R block error ratio required row to Bit error Ratio required with a value of <1e-8 and delete footnote c.
 Or change the Block error ratio to <2.1 e-5

Response Response Status **C**

ACCEPT IN PRINCIPLE.

See comment #146.

In Table 111-5, change the block error ratio maximum to 2.1e-5.

Cl 111 SC 111.9 P172 L30 # 51
 Ran, Adee Intel

Comment Type **E** Comment Status **A** bucket

The editor's note about implementation of comment #59 against D0.1 (channel characteristics for 25GBASE-KR-S) is not needed in the next draft.

SuggestedRemedy

Delete editor's note.

Response Response Status **C**

ACCEPT.

Cl 112 SC 112.10 P189 L18 # 196
 Dawe, Piers Mellanox

Comment Type **T** Comment Status **A**

As for Table 112-5, 25GBASE-SR operating range and 112.9 Fiber optic cabling model, it helps the user to understand that the characteristics of the fiber optic cabling (channel) are compatible with 100GBASE-SR4.

SuggestedRemedy

Add a sentence:
 The requirements for the optical fiber, connection insertion loss and maximum discrete reflectance are the same as 100GBASE-SR4 (See Clause 95).

Response Response Status **C**

ACCEPT.

Cl 112 SC 112.3 P182 L25 # 43
 Ran, Adee Intel

Comment Type **E** Comment Status **A** bucket

The delay constraint values are in magenta.

SuggestedRemedy

Change to normal black font.

Response Response Status **C**

ACCEPT.

Cl 112 SC 112.3 P182 L27 # 55
 King, Jonathan Finisar

Comment Type **TR** Comment Status **A**

Reference to system delay constraints should be to 105.5

SuggestedRemedy

change "105.4 and its references" to "105.5 and its references"

Response Response Status **C**

ACCEPT.

Cl 999 SC P10 L11 # 19
 Anslow, Pete Ciena

Comment Type **E** Comment Status **A** bucket

As the P802.3bw project is entering Sponsor ballot, it is reasonable to assume that the 802.3bw amendment will be published before the 802.3by amendment.

SuggestedRemedy

Add the summary for IEEE Std 802.3bw to the frontmatter above that for IEEE Std 802.3by:

IEEE Std 802.3bwTM-201x
 This amendment includes changes to IEEE Std 802.3-201x and adds Clause 96. This amendment adds 100 Mb/s Physical Layer (PHY) specifications and management parameters for operation on a single balanced twisted-pair copper cable.

Response Response Status **C**

ACCEPT.

IEEE 802.3by D1.0 25 Gb/s Ethernet 2nd Task Force review comments

Cl 999 SC P12 L9 # 81
 Nowell, Mark Cisco

Comment Type E Comment Status A bucket

Table of contents entry for 45.2.1.94 & 45.2.1.95 both state "Single-lane PHY 10GBASE-R..."

The intention of the changes in 45.2.1.94 & 95 is to replace "10GBASER-R" with "Single-lane PHY BASE-R"

Additional comments to follow on the use of "Single-lane"...

SuggestedRemedy

Correct table of contents entry for both 45.2.1.94 & 45.2.1.95 to state "Single-lane PHY BASE-R..."

Response Response Status C

ACCEPT.

Cl 999 SC P18 L46 # 83
 Nowell, Mark Cisco

Comment Type E Comment Status A bucket

Table of contents entry for 110.11.1 needs a space after clause number or else it looks like 110.11.125

SuggestedRemedy

Add space to Table of contents entry for 110.11.1

Response Response Status C

ACCEPT IN PRINCIPLE.

The space is missing for all level-3 headings in the TOC. Fix the formatting appropriately.

Cl 999 SC 99 P15 L14 # 187
 Dawe, Piers Mellanox

Comment Type E Comment Status A bucket

Formatting /alignment problem?

SuggestedRemedy

Fix

Response Response Status C

ACCEPT.