

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

CI 000 SC 0 P- L- # 50
 Nowell, Mark Cisco

Comment Type T Comment Status A

A number of the specification values that were adopted in the original baselines were colored magenta (or additionally marked with TBC) to represent that they were values which should be considered as TBCs but the current value used was a good starting point unless further analysis suggested changing it. D1.3 represents the 4 review cycle for these values.

If, after the completion of D1.3 Task Force Review, any of these values remaining in magenta font or marked with a "TBC" have not been commented on or modified, then suggest to convert them to black font and/or remove the TBC marking.

The next phase of the process is Working Group ballot and the magenta font has no meaning in this process. If there is a comfort level with having some marking associated with these values, an editors note or footnote could be added but I do not think it is necessary. This will not limit the ability to comment or adjust these values during further reviews or ballots. They will be dealt with consistently with all other specification values in the document.

If an editor's footnote is preferred, I suggest something like: " ** Further work to substantiate these values is anticipated"

SuggestedRemedy

Change all values in magenta fonts that have not been modified at the close of D1.3 comment review to black font. Remove TBC markings. Add editor's footnote to values if deemed necessary.

Response Response Status C
 ACCEPT IN PRINCIPLE.

In 136B.1.1.2, change cross-reference "92.11.3.2" from magenta to forest green (external) text.

In 136B.2.4, change cross-reference in TF7 from "Table 92-13" to "Table 136B-3" and change text color to black.

CI 045 SC 45.2.1.108 P57 L26 # 1
 Anslow, Pete Ciena

Comment Type E Comment Status A <bucket>

"Clause 91" should be a cross-reference and "91.5.2.4" and "91.6.3" should have character tag External applied.

SuggestedRemedy

Make "Clause 91" a cross-reference and apply character tag External to "91.5.2.4" and "91.6.3".

Response Response Status C
 ACCEPT.

CI 045 SC 45.2.1.116m P66 L9 # 55
 Brown, Matt MACOM

Comment Type T Comment Status A

Register bits 1.605.1 and 1.605.0 are status bits indicating a request from the Rx inputs on lane 0 and lane 1 for the connected transmitter. They are currently described as being inputs.

SuggestedRemedy

Change these bits to RO and changed the description appropriately.

Response Response Status C
 ACCEPT IN PRINCIPLE.

Change the register name to:
 "PMA precoder request Rx input status"

Update first sentence in 45.2.1.116m to have the correct register name.

Update the table as follows:

Make the bits read only.
 In footnote remove "R/W = Read/Write,"

Modify the title of table 45.90ai to reflect the new register name.

Change descriptions to:
 45.2.1.116m.1 Lane 1 Rx input precoder request status (1.605.1)
 This bit indicates the Rx input precoder request status for lane 1.

45.2.1.116m.2 Lane 0 Rx input precoder request status (1.605.0)
 This bit indicates the Rx input precoder request status for lane 0

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 045 SC 45.2.1.116n P66 L36 # 56
Brown, Matt MACOM

Comment Type T Comment Status A

Register bits 1.606.1 and 1.606.0 are status bits indicating a request from the Tx inputs on lane 0 and lane 1 for the connected transmitter. They are currently described as being inputs.

SuggestedRemedy

Change these bits to RO and changed the description appropriately.

Response Response Status C

ACCEPT IN PRINCIPLE.

Change the register name to:
"PMA precoder request Tx input status"

Update first sentence 45.2.1.116n to have the correct register name.

Update the table as follows:
Make the bits read only.
In footnote remove "R/W = Read/Write,"

Modify the title of table 45.90aj to reflect the new register name.

Change descriptions to:
45.2.1.116n.1 Lane 1 Tx input precoder request status (1.606.1)
This bit indicates the Tx input precoder request status for lane 1.

45.2.1.116n.2 Lane 0 Tx input precoder request status (1.606.0)
This bit indicates the Tx input precoder request status for lane 0

Cl 069 SC 69.2.3 P81 L34 # 26
Slavick, Jeff Broadcom Limited

Comment Type E Comment Status R

Each PHY type has it's own paragraph in 69, but the 3 new ones have been lumped into a single paragraph.

SuggestedRemedy

Separate the 1 paragraph into 3, starting each paragraph with Backplane Ethernet also specifies <PHY>.
Backplane Ethernet also specifies 50GBASE-KR. The 50GBASE-KR embodiment employs the PCS defined in Clause133, the RS-FEC defined in Clause134, the PMA defined in Clause135, and the PMD defined in Clause137, and specifies 50Gb/s operation using 4-level PAM over one differential path in each direction.

Backplane Ethernet also specifies 100GBASE-KR2. The 100GBASE-KR2 embodiment employs the PCS defined in Clause82, the RS-FEC defined in Clause91, the PMA defined in Clause135, and the PMD defined in Clause137, and specifies 100Gb/s operation using 4-level PAM over two differential paths in each direction.

Backplane Ethernet also specifies 200GBASE-KR4. The 200GBASE-KR4 embodiment employs the PCS defined in Clause 119, the PMA defined in Clause 120, and the PMD defined in Clause137, and specifies 200Gb/s operation using 4-level PAM over four differential paths in each direction.

Response Response Status C

REJECT.

The paragraphs for previous PHYs are separated as these PHYs are either distinctly different or were introduced in different projects.

The text as written is correct and complete.

Cl 073 SC 73.10.2 P89 L21 # 27
Slavick, Jeff Broadcom Limited

Comment Type E Comment Status R

I might be best in Table 73-7 to list the link_fail_inhibit_timer in the same units for all 3 durations.

SuggestedRemedy

Change link_fail_inhibit_timer values for the row with 1.6, 1.7, s to 1600, 1700, ms

Response Response Status C

REJECT.

The suggested remedy does not improve the clarity of the specification.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 131 SC 131.5 P126 L27 # 40
 Dawe, Piers Mellanox

Comment Type TR Comment Status R skew

All 50G PMDs are serial. So the Skew and Skew Variation at SP3 (transmitter MDI), SP4 (receiver MDI) and SP5 (PMD output) can't be different to those at SP2 (PMD input) because there is only one lane from SP2 to SP5. Rebuttal to D1.2 147 did not give a reason (dud reference?) so here is a similar comment again.

SuggestedRemedy

Correct the Skew and Skew Variation limits for 50GBASE-CR, 50GBASE-KR, 50GBASE-SR, 50GBASE-FR and 50GBASE-LR.
 If appropriate, list optional skew values that would apply if there were an 2-lane 50G PMD. But they should not be required - almost all NICs would never see such a PMD even if it existed.
 In Table 131-5, provide columns for serial PHYs and for multi-lane PHYs.

Response Response Status C

REJECT.

The comment is essentially the same as comment #220 against D1.2. That comment was rejected with the following response:

"Based on discussion and comment resolution at the January 2017 task force meeting WRT to the skew specifications for single-lane PMDs the consensus was to implement the specifications consistent with 40G, 100G, and 200G PHYs already specified in IEEE Std 802.3-2015 and P802.3bs.

See the final response for P802.3cd Draft 1.1 Comment #10."

Note that the response to comment #220 against D1.2 shown above contains a typo; the reference to Draft 1.1 Comment #10 should be to comment #80.

There is no new information in the comment to support the suggested change.

Cl 134 SC 134.5.3.3 P151 L49 # 15
 Ran, Adee Intel

Comment Type T Comment Status D <withdrawn>

As shown in a contribution to 802.3bs (see http://www.ieee802.org/3/bs/public/16_09/ran_3bs_01a_0916.pdf), predicting the link performance by the binary event of the average symbol error ratio exceeding some threshold is error prone and would result in problems setting the threshold correctly.

In mass deployment of 802.3cd links, as expected in future data centers, this may result in multiple false alerts or perceived degradations in links that have ample margin for practically error-free operation. The only way to avoid these false alarms is to have a very high margin in all links, but that would increase the cost.

An alternative solution, outlined in http://www.ieee802.org/3/bs/public/16_09/ran_3bs_02a_0916.pdf, is to count codewords with a specific number of symbol errors in separate counters. This information is available from the RS-FEC decoder and would be much more useful for predicting uncorrectable errors and identifying links that have insufficient margin (and the desired margin can be defined after the data is collected).

The proposal above was not accepted, mainly claiming that it is tightly coupled with the PCS FEC which might only be used in an XS while the actual PMD-PMD link would use another FEC. But in 802.3cd there are no XS's and no other FEC is expected, so this method is perfectly adequate.

If information on degradation or prediction of uncorrectable errors is desirable, it should use the relevant information. At the minimum, that information should be available through standard registers. These registers may be queried by management and reported to the partner through higher layer protocols, outside of the scope of 802.3 (or we can add LLDP message in clause 79 later).

SuggestedRemedy

Based on slide 17 of http://www.ieee802.org/3/bs/public/16_09/ran_3bs_02a_0916.pdf:

Define a variable array (16 integers, 12 bits each) for counting received codewords with 1 to 15 symbol errors and uncorrectable codewords. Map these variables to MDIO registers, non-rollover, clear on read.

Add similar variables mapped to the same registers also in clause 91 for the 100G RS-FEC and in clause 119 for the 200G PCS FEC. These should be optional.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

CI 136 SC 136.1 P188 L48 # 37
 Dawe, Piers Mellanox

Comment Type TR Comment Status A BER

These paragraphs imply a requirement for a receiver to give the right BER (FLR) with any compliant transmitter and channel, which usurps the receiver interference tolerance spec and is too vague. We moved off this years ago in favour of clear and specific stressed sensitivity or RITT specs. 136.9.4.1 and 136.9.4.2.5 are now clearer in D1.3 so we don't need this vague double jeopardy any more.

SuggestedRemedy

Create a new subclause 136.1.1 Bit error ratio, as in 138.1.1, 139.1.1, 140.1.1 and 802.3bs; refer to it from 136.9.4.1 and 136.9.4.2.5. Delete "from a compliant input signal" twice. Delete "A compliant input signal is a transmitter output of a compliant PHY that has passed through a compliant cable assembly." Possible other changes for consistency with 138.1.1 and 802.3bs. Similarly in Clause 137.

Response Response Status C

ACCEPT IN PRINCIPLE.

The text in these paragraphs is intended to be descriptive rather than normative, since the normative specification is the interference tolerance test. It declares the expectations from receivers when operating with compliant transmitters and channels.

Adding a "bit error ratio" subclause based on the optical PMD clauses, as suggested, would make the PMD/PMA BER normative, since the text in the optical PMD clauses includes "shall" (these clauses use the BER subclause as a reference for the receiver and transmitter tests).

However, the comment suggests possible incorrect interpretation of this text as a specification. To avoid that, apply the following changes.

Change "are required to detect bits" to "are expected to detect bits", in P188 L50 and in P189 L5.

Apply corresponding changes in clause 137.

CI 136 SC 136.3 P190 L50 # 51
 Nowell, Mark Cisco

Comment Type T Comment Status A <bucket>

Typo on Symbol rate 26.6525 should be 26.5625
 Two instances: Pg 190, line 50 & pg 191 line 3

SuggestedRemedy

fix typo

Response Response Status C

ACCEPT IN PRINCIPLE.

Change 26.6525 to 26.5625 at the indicated locations, and also in other clauses as necessary.

CI 136 SC 136.6.1 P192 L16 # 41
 Dawe, Piers Mellanox

Comment Type TR Comment Status R skew

For 50GBASE-KR, 50GBASE-CR, 50GBASE-SR, 50GBASE-FR and 50GBASE-LR, the Skew at SP3 (transmitter MDI), SP4 (receiver MDI) and SP5 (PMD output) can't be different to those at SP2 (PMD input) because there is only one lane from SP2 to SP5. The draft correctly says there is no Skew Variation at these points.

SuggestedRemedy

Correct the Skew limits for 50GBASE-CR and 50GBASE-KR.

Response Response Status C

REJECT.

See response to comment #40.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 136 SC 136.8.11.5 P206 L9 # 36
 Zvi, Rechtman Mellanox

Comment Type T Comment Status R training

The Coefficient update state machine in figure 136-9 defines the transmitter behavior upon peer receiver requests. While the requestor flow is not explicitly defined in the clause.

SuggestedRemedy

Add a new first paragraph to "136.8.11.5" with explicitly definition of receiver request flow (based on 72.6.10.2.3.3 text):

"A coefficient update request is assigned to a 2-bit field describing a requested update. Four request encodings are defined: no-equalization, increment, decrement, and hold. The default state for a given tap is hold, which corresponds to no change in the coefficient. The no-equalization, increment or decrement encodings are transmitted to request that the corresponding coefficient be reset, increased or decreased. The amount of change implemented by the transmitter in response to the coefficient update request shall meet the requirements as define in the algorithm below. An no-equalization, increment or decrement request shall continue to be transmitted until the update status for that coefficient, indicates: updated, coefficient at limit, coefficient not supported, equalization limit or coefficient limit and equalization limit . At that point, the outgoing requests for that tap shall be set to hold.

A new request to no-equalization, increment or increment any coefficient or a new request for Initial condition request , shall not be sent before the incoming status message revert to not_updated."

Response Response Status C

REJECT.

The commenter has pointed out a potential issue that may need to be addressed.

The suggested remedy is not complete and it seems that the required changes are not simple.

The commenter is invited to resubmit this comment against a future draft with a complete proposal and consensus.

Cl 136 SC 136.9 P218 L30 # 57
 Hegde, Raj Broadcom Ltd.

Comment Type T Comment Status D training <late>

The transmit equalizer coefficient values corresponding to the 'preset 2' and 'preset 3' settings of the variable ic_req in Table 136-12 provide -6dB of de-emphasis and +6dB of pre-shoot respectively. This level of de-emphasis and pre-shoot is a lot higher than the average de-emphasis/pre-shoot values needed over the range of Clause 136 channels.

SuggestedRemedy

Reduce the initial coefficient de-emphasis/pre-shoot settings to +/- 2dB.
 Change c(1) corresponding to preset 2 in Table 136-12 to -0.1 +/- 0.05
 Change c(-1) corresponding to preset 3 in Table 136-12 to -0.1 +/- 0.05

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

Cl 136 SC 136.9.3 P215 L33 # 43
 Dawe, Piers Mellanox

Comment Type ER Comment Status A <bucket>

Align with 802.3bs 120D. Clause 94 should be deprecated and we should not refer to it in new clauses. The same definitions and figure as in 94.3.12.3 are in 93.8.1.3.

SuggestedRemedy

Change the references to 94.3.12.3 (five here, one in PICS 136.14.4.3, one in PICS 137.12.4.3) to 93.8.1.3.

Response Response Status C

ACCEPT IN PRINCIPLE.

Align references with Annex 120D in 802.3bs, in Table 136-11 (cells and footnotes), and in PICS tables in 136.14.4.3 and 137.12.4.3.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 136 SC 136.9.3 P216 L6 # 16
Ran, Adee Intel

Comment Type T Comment Status A

SNDR is set to 30.5 dB in magenta. This is close to the SNR_ISI (30 dB) and their RSS would create a transmitter noise equivalent to SNR_TX=27 dB. However SNR_TX is 32.5 in COM. So channel spec assumes better than transmitter spec, and the budget is broken.

I have submitted a comment to 802.3bs asking to change SNDR to 31.8 dB, SNR_ISI to 35.3 dB, and measure only in unequalized state, with 2 dB allocated to signal loss due to equalization.

The SNR_TX in COM in annex 120D is 31 dB. The proposal is to "dial back" 1.5 dB from both SNDR and SNR_ISI, in order to get SNR_TX of 32.5 dB in COM.

These may not be the final values for clause 136 but they would at least result in a closed budget.

SuggestedRemedy

Change SNDR to 33.3 dB.
Change SNR_ISI to 36.8 dB, measured only in unequalized setting.

Response Response Status C

ACCEPT IN PRINCIPLE.

In order to close the budget, implement the suggested remedy and change the SNDR and SNR_ISI text to black.

Add an editor's note stating "The values for SNDR, SNR_ISI, and SNR_TX require confirmation and may change."

Cl 136 SC 136.9.3 P216 L11 # 39
Dawe, Piers Mellanox

Comment Type T Comment Status R jitter <NSR>

Is it reasonable to expect the same jitter at TP2 (this clause) as at TP0a (Cl.137)? Won't the connector crosstalk make them different.

SuggestedRemedy

Ensure that the jitter limits in 137.9.2 / Table 120D-1, the mated CB crosstalk specs, and the jitter limits here, are compatible.

Response Response Status C

REJECT.

There is insufficient information in the comment and in the suggested remedy to implement any change to the draft.

The commenter is invited to resubmit the comment against a later draft with a complete proposal and consensus.

Cl 136 SC 136.9.3 P216 L11 # 38
Dawe, Piers Mellanox

Comment Type TR Comment Status R jitter

J4 (all but 1e-4 of the edges, or 1e-4*0.75 of the number of UI, divided between early and late, so 3.75e-5 per UI or 1.875e-5 per bit) is overkill for the spec BER of 2.4e-4, and J3 (1.875e-4 per bit) is a good match to the spec BER - just as J4 is a good match to the BER of 1e-5 for 120D. Also, not all edges cause errors. We can make the spec better (more accurate, less performance left on the table) and reduce test time.

SuggestedRemedy

Change J4 to J3, adjusting the limit by use of eq 136-6 and 7 (dual-Dirac theory). In Eq 136-6 change Q4=3.8906 to Q3=3.2905, Q(Q3) = 5 x 10^-4.

Response Response Status C

REJECT.

The suggested does not provide a complete proposal to implement.

The commenter is welcome to build consensus and submit a comment on a future draft.

Cl 136 SC 136.11 P223 L18 # 28
Tracy, Nathan TE Connectivity

Comment Type T Comment Status A

This paragraph is outdated due to the additional MDI added at the Vancouver meeting. Replace lines 18 through 22 with the suggested remedy

SuggestedRemedy

replace with "50GBASE-CR, 100GBASE-CR2 and 200GBASE-CR4 cable assemblies are defined in Annex 136C. There are five possible MDIs which are defined in Annex 136D including the single lane SFP MDI, the four lane QSFP and microQSFP MDIs and the eight lane OSFP and QSFP-DD MDIs. This results in multiple possible implementations including the opportunity to have multiple 50GBASE-CR types in a multi-lane MDI such as QSFP, microQSFP, OSFP and QSFP-DD, multiple 100GBASE-CR2 types in a multi-lane MDI such as QSFP, microQSFP, OSFP and QSFP-DD and multiple 200GBASE-CR4 types in a multi-lane MDI such as OSFP and QSFP-DD.

Response Response Status C

ACCEPT IN PRINCIPLE.

Replace the paragraph starting with "Since 50GBASE-CR has two specified..." with the following paragraph:

"50GBASE-CR, 100GBASE-CR2, and 200GBASE-CR4 cable assembly form factors are described in Annex 136C. There are five possible MDIs which are defined in Annex 136D."

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 136 SC 136.11.7 P224 L38 # 33
 Hidaka, Yasuo Fujitsu Lab of America

Comment Type E Comment Status D COM

As the result of resolution for comment #-34 against P802.3bs D3.0, the package impedance Z_c was added to Table 93A-1 as a COM parameter, because the package Z_c was not a COM parameter.

Similarly, we should add the host PCB trace impedance Z_c (=109.8 ohm) in Table 92-12 referred from 136.11.7.1, the channel PCB signal path length z_p (=151mm) in 136.11.7.1.1, and channel PCB crosstalk path length z_p (=72mm) in 136.11.7.1.2 should be added to the COM parameter table for the copper cable PMD such as Table 136-15. It is very inconvenient and error prone to scatter these parameters in various locations, because they are important parameters for signal integrity.

Alternatively, we may just take out the host PCB trace impedance Z_c from Table 92-12 and add it to Table 93A-1, and leave the channel PCB signal/crosstalk path length z_p as descriptions in 136.11.7.1.1 and 136.11.7.1.2.

SuggestedRemedy

Add the host PCB trace impedance Z_c, channel PCB signal path length z_p, and channel PCB crosstalk path length z_p to Table 93A-1 as COM parameters which are valid only for copper cable PMDs.

Define default values for those new COM parameters in Annex 93A.

Add those new COM parameters to Table 136-15.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

Cl 136 SC 136.11.7 P224 L47 # 17
 Ran, Adee Intel

Comment Type E Comment Status A COM, magenta

Capacitance and impedance values are in magenta.

The existing values (from the baseline proposal) have persisted for three drafts, with no claims that they are not feasible nor a specific proposal to change them to other values. Such proposals can still be brought forward, but there is no need to "remind us that there is no consensus".

Note that these values are now also aligned with annex 120D.

SuggestedRemedy

Apply normal text format for C_d, C_p and Z_c.

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement the suggested remedy.

Also, add an editor's note as follows:

"The value for Z_c requires confirmation and may change."

Cl 136 SC 136.11.7 P224 L51 # 11
 Ran, Adee Intel

Comment Type E Comment Status A COM, magenta

The value for spectral density is in magenta.

The values (from the baseline proposal) has persisted for three drafts, with no claims that it is not feasible nor a specific proposal to change it to another value. Such proposals can still be brought forward, but there is no need to "remind us that there is no consensus".

Just for information, calculation shows that the PSD of 1.64e-8 V²/GHz is equivalent to 0.54 mV RMS of noise at a flat bandwidth of 16 GHz. This is 7.5 dB above the thermal noise floor.

SuggestedRemedy

Apply normal text format for eta_0.

Response Response Status C

ACCEPT.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 136 SC 136.11.7.2 P227 L41 # 2
 Anslow, Pete Ciena

Comment Type E Comment Status A <bucket>

The format of Table 136-16 is not according to the 802.3 template:
 The heading row is not bold.
 The internal ruling is not according to the template (the heading row should have a thicker line at the bottom than it does).
 The font used in the table is incorrect.

SuggestedRemedy

Re-create the table using the "IEEE" table format.

Response Response Status C

ACCEPT.

Cl 136 SC 136.14.3 P231 L14 # 3
 Anslow, Pete Ciena

Comment Type T Comment Status A <bucket>

Item FFQS has a subclause entry of "136.11 and 136.11"

SuggestedRemedy

Change to "136.11"

Response Response Status C

ACCEPT.

Cl 136B SC 136B.1.1.6 P373 L14 # 49
 Dawe, Piers Mellanox

Comment Type TR Comment Status R

Just as for the QSFP connector, we will need better crosstalk to support PAM4 with the SFP connector.

SuggestedRemedy

Tighten NEXT from 1.8 mV rms towards 1.5 as feasible, by changing "shall meet the specification in Table 110B-1." to "shall be less than 1.6 mV."

Response Response Status C

REJECT.

The commenter has not provided sufficient justification to change SFP28 mated test fixture integrated near-end crosstalk to "shall be less than 1.6 mV".

Cl 136B SC 136B.1.1.6 P373 L33 # 9
 Anslow, Pete Ciena

Comment Type E Comment Status A <bucket>

Extra space before the "." in "in Table 136B-2 ."

SuggestedRemedy

Delete the space.

Response Response Status C

ACCEPT.

Cl 136C SC 136C.1 P377 L15 # 34
 Hidaka, Yasuo Fujitsu Lab of America

Comment Type E Comment Status A <bucket>

Here, references for SFP28 and QSFP28 are 110.11.1 and 92.12.1.1, respectively, but there are 136D.2.1 and 136D.2.2 in this same annex 136C. References within the same annex 136C are better, because 136D.2.1 and 136D.2.2 have reference to 110.11.1 and 92.12.1.1.

SuggestedRemedy

Change the reference for SFP28 to 136D.2.1 and the reference for QSFP28 to 136D.2.2.

Response Response Status C

ACCEPT.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

CI 136C SC 136C.2.3 P377 L54 # 29
 Tracy, Nathan TE Connectivity

Comment Type T Comment Status A
 Suggest to include the description of a breakout cable as is provided for the QSFP host form factor (136C2.2)

SuggestedRemedy
 Additional text: "A microQSFP form factor host can also form up to four 50 Gb/s links to either another microQSFP form factor host, using a microQSFP to microQSFP form factor cable assembly (see 136C.3.2), or four separate SFP28 form factor hosts (see 136C.2.1) using a microQSFP to 4xSFP28 form factor cable assembly (see 136C.3.3)."

Response Response Status C
 ACCEPT IN PRINCIPLE.

Change: 136C.2.2 QSFP28 host form factor first paragraph;

To: A host may use the QSFP28 receptacle specified in 136D.2.2 as the MDI for one, two, three or four 50GBASE-CR PHYs, one or two 100GBASE-CR2 PHYs, or one 200GBASE-CR4 PHY. This is referred to as a QSFP28 host form factor.

In 136C.2.3 microQSFP host form factor
 (1)L54 delete sentence end of first paragraph...This MDI may also be used as a single-lane MDI.

(2)Add second paragraph
 A microQSFP form factor host can also form up to four 50 Gb/s links to either another microQSFP form factor host, using a microQSFP to microQSFP form factor cable assembly (see 136C.3.2), or four separate SFP28 form factor hosts (see 136C.2.1) using a microQSFP to 4xSFP28 form factor cable assembly (see 136C.3.3).

In 136C.2.3 microQSFP host form factor
 (1)L54 delete sentence end of first paragraph...This MDI may also be used as a single-lane MDI.

(2)Add second paragraph
 A microQSFP form factor host can also form up to four 50 Gb/s links to either another microQSFP form factor host, using a microQSFP to microQSFP form factor cable assembly (see 136C.3.2), or four separate SFP28 form factor hosts (see 136C.2.1) using a microQSFP to 4xSFP28 form factor cable assembly (see 136C.3.3).

CI 136D SC 136D.2.3 P386 L54 # 10
 Anslow, Pete Ciena

Comment Type E Comment Status A <bucket>
 The footnotes at the foot of pages 386, 387, and 300 should be formatted more helpfully.

SuggestedRemedy
 Change the footnote on page 386 to: "microQSFP specifications are available at <http://www.microqsfp.com>"
 Change the footnote on page 387 to: "QSFP-DD specifications are available at <http://www.qsfp-dd.com>"
 Change the footnote on page 388 to: "OSFP specifications are available at <http://www.osfpmsa.org>"

Response Response Status C
 ACCEPT.

CI 137 SC 137.1 P243 L2 # 4
 Anslow, Pete Ciena

Comment Type E Comment Status A <bucket>
 "." missing at the end of the sentence

SuggestedRemedy
 Add the "."

Response Response Status C
 ACCEPT.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

CI 137 SC 137.9.2 P244 L22 # 12
 Ran, Adee Intel

Comment Type T Comment Status D <withdrawn>

SNDR is set to 32.5 dB and SNR_ISI is set to an unmeasurably high value (43 dB) so that their RSS would create a transmitter noise equivalent to SNR_TX=32.5 dB used in COM.

I have submitted a comment to 802.3bs asking to change SNDR to 31.8 dB, SNR_ISI to 35.3 dB, and measure only in unequalized state, with 2 dB allocated to signal loss due to equalization.

The SNR_TX in COM in annex 120D is 31 dB. The proposal is to "dial back" 1.5 dB from both SNDR and SNR_ISI, in order to get SNR_TX of 32.5 dB in COM.

These may not be the final values for clause 137 but they would at least result in a closed budget.

SuggestedRemedy

Change SNDR to 33.3 dB.
 Change SNR_ISI to 36.8 dB, measured only in unequalized setting.

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

CI 137 SC 137.9.3 P241 L30 # 31
 Mellitz, Richard Samtec

Comment Type TR Comment Status R RITT

Clause 93 in Table 93-6 includes a set of test values for no-FEC even though FEC is required in the PMD. There is not a corresponding test in clause 137.

SuggestedRemedy

And items to list in 137.9.3
 o. Add columns for Test 3 (NO-FEC, low loss)
 .. PCS FEC Symbol error ratio max = 1e-11
 .. Insertion loss at 13.2813 GHz min=11.5 dB and max=12dB
 .. RSS_DFE4 min=0.5
 .. COM including effects of broadband noise target=3dB
 In addition the COM parameter DER0 is set to 1e-12

Response Response Status C

REJECT.

100GBASE-KR4 uses the Clause 91 RS-FEC with an option to bypass error correction.

The text above Table 93-6 includes the following sentence:
 "Tests 1 and 2 are for the case when error correction is bypassed in the RS-FEC sublayer (see 91.5.3.3) and for these cases COM is computed with a DER0 value of 10^-12."

For the PHYs defined in clause 137, there is no option to bypass the RS-FEC correction, and no expectation that they can operate over any channel with the required FLR and MTTFPA unless errors are corrected. Therefore the proposed additional tests are not required.

CI 137 SC 137.9.3.1 P241 L46 # 48
 Dawe, Piers Mellanox

Comment Type TR Comment Status R RL

I doubt that the low frequency RL at 14.25 dB is significant for signal integrity compared with the 8.7 dB at 6 GHz. This RL is much tighter than CEI-56G-MR at low (and high) frequency but looser between 4 and 9 GHz. I've made a similar comment against 802.3bs 120D.

SuggestedRemedy

Change 14.25 - f to 12 -0.625f

Response Response Status C

REJECT.

Several existing RL specifications in 802.3 are not fully consistent with other standards. This is not necessarily a problem.

The existing RL equation is related to the COM package model (See http://www.ieee802.org/3/cd/public/Mar17/mellitz_3cd_01b_0317.pdf).

Further evidence should be provided to justify the suggested change.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

CI 137 SC 137.10 P243 L20 # 13
 Ran, Adee Intel

Comment Type E Comment Status A COM, magenta

Capacitance and impedance values are in magenta.

The existing values (from the baseline proposal) have persisted for three drafts, with no claims that they are not feasible nor a specific proposal to change them to other values. Such proposals can still be brought forward, but there is no need to "remind us that there is no consensus".

Note that these values are now also aligned with annex 120D.

SuggestedRemedy

Apply normal text format for C_d, C_p and Z_c.

Response Response Status C

ACCEPT IN PRINCIPLE.

Implemented the suggested remedy.

Also, add an editor's note as follows:

"The value for Z_c requires confirmation and may change."

CI 137 SC 137.10 P244 L24 # 14
 Ran, Adee Intel

Comment Type E Comment Status A COM, magenta

The value for spectral density is in magenta.

The values (from the baseline proposal) has persisted for three drafts, with no claims that it is not feasible nor a specific proposal to change it to another value. Such proposals can still be brought forward, but there is no need to "remind us that there is no consensus".

Just for information, calculation shows that the PSD of $1.64e-8 V^2/GHz$ is equivalent to 0.54 mV RMS of noise at a flat bandwidth of 16 GHz. This is 7.5 dB above the thermal noise floor.

SuggestedRemedy

Apply normal text format for eta_0.

Response Response Status C

ACCEPT.

CI 137 SC 137.10. P243 L23 # 32
 Hidaka, Yasuo Fujitsu Lab of America

Comment Type TR Comment Status R

This is a comment to follow-up comment #145 against D1.2.

As explained in hidaka_3cd_01a_0317, the current COM spec has a hole between channel test and Rx ITT. Suppose some chip vendor designed Rx using the Rd and Zc values of the COM parameter as their nominal target value. For Rx ITT, the chip vendor may choose a test channel that shows bad matching with their Rx. Suppose they barely passes Rx ITT with this bad-matching test channel. If they chose a good matching channel, they would likely have failed Rx ITT, because more broadband noise would have to be injected to push higher COM (due to good matching) down to 3dB by calibration of test channel. On the other hand, suppose some system vendor may have some channel that shows good matching with the reference Rd and Zc of the COM parameter, but just barely passes the 3dB COM criteria. Although this channel and the above Rx are both compliant to the nominal spec of the standard, they will not operate with each other.

In other words, the current standard has zero tolerance for variation of the device impedance. The size of the hole is about 0.9dB COM for 30dB loss channels when the device impedance has +/- 10% variation as shown in hidaka_3cd_01a_0317.

The straw poll indicated that the concern of low yield due to pessimistic spec assuming the worst case is greater than the concern of low interoperability due to optimistic spec of zero tolerance for variation. On the other hand, the concern of low interoperability is not that negligible small as well. A practical solution may be some compromise between the current zero tolerance for variation and the worst-case scenario of variation.

A problem of option 1 in hidaka_3cd_01a_0317 was how to make this compromise, because the method in hidaka_3cd_01a_0317 was merely based on option 2 by allowing overlap of distribution. There was no good justification to make a compromise for option 1 in this way. Besides, option 2 did not get a good support.

A better and more straight forward way to make a compromise for option 1 is to re-use the same framework as option 1 but reduce the amount of variation from +/- 10% down to +/- 5% (or maybe to +/- 7%).

Note that this amount of variation does not represent the amount of variation of real devices. It is OK for real devices to have larger variation. Also, it may be actually possible to achieve +/- 5% variation with real devices using leading-edge technologies.

In a sense, this is in direction of option 1 in hidaka_3cd_01a_0317 which had more objection (15 people) than support (11 people) in straw poll #2. However, this is also a new method that was not clearly covered in hidaka_3cd_01a_0317. Hopefully, this new method to make a compromise may get more support than option 1 in hidaka_3cd_01a_0317.

SuggestedRemedy

Implement option 1 in hidaka_3cd_01a_0317 to clause 137 with the following parameters:

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Rd : 50 ohm +/- 5%
 Package Zc : 93 ohm +/- 5%
 Av : 0.415 V +/- 5% (same variation as Tx Rd)
 Afe : 0.415 V +/- 5% (same variation as Tx Rd)
 Ane : 0.611 V +/- 5% (same variation as Tx Rd)

Apply similar changes to clause 136 with the following additional parameter:

Host PCB trace Zc : 100 ohm +/- 5% (override the value in Table 92-12 in 136.11.7.1.)

I will prepare a presentation for discussion.

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

There is no consensus to make the suggested changes at this time.

A complete proposal with consensus is required. It is recommended that further work take place in the ad hoc meetings.

The commenter is invited to resubmit a new comment against a later draft.

Cl 138 **SC 0** **P261** **L22** # **30**
 Maguire, Valerie Siemon
Comment Type **TR** *Comment Status* **A** om5

OM5 has been adopted as the named for the "wideband MMF" optical fiber defined in TIA-492AAAE.

SuggestedRemedy

Change all occurrences of "wideband MMF (TIA-492AAAE)" to "OM5".

Do a selective search to update all other references that may lack or use other terminology. For example, on line 23 of page 261, replace,

"A compliant PMD operates on 50/125 um multimode fibers, type A1a.2 (OM3), type A1a.3.(OM4), or fiber compliant to TIA-492AAAE,"

with,

"A compliant PMD operates on 50/125 um multimode fibers, type A1a.2 (OM3), type A1a.3.(OM4), or fiber compliant to TIA-492AAAE (OM5)."

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

[Editor's note: Clause # changed from 000 to 138]

Resolve using response to comment #23.

Cl 138 **SC 138** **P252** **L1** # **42**
 Dawe, Piers Mellanox

Comment Type **TR** *Comment Status* **R**

This -SRn draft is still a baseline, no presentations at the last meeting - where is the indication that these numbers, or others, actually work with technical and economic feasibility?

SuggestedRemedy

While in Task Force review, show some evidence: eyes, receiver waterfall plots, TDECQ measurements and so on. Adjust the draft as appropriate.

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

The suggested remedy offers no changes to the draft.

Cl 138 **SC 138.2** **P255** **L41** # **52**
 Nowell, Mark Cisco

Comment Type **T** *Comment Status* **A** <bucket>

Typo on Symbol rate 26.6525 should be 26.5625
 Two instances: Pg 255, line 41 & line 46

SuggestedRemedy

fix typo

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

Cl 138 **SC 138.7** **P261** **L23** # **20**
 Kolesar, Paul CommScope

Comment Type **T** *Comment Status* **A** om5

The IEC specification containing the equivalent of TIA-492AAAE was approved for publication during the week of 24 April 2017 by SC86A WG1. The equivalent fiber is called A1a.4.

SuggestedRemedy

Consider replacing "or fiber compliant to TIA-492AAAE" with "or type A1a.4". Note that while approval of the IEC CDV ballot allowed OM5 content to remain in ISO 11801-1, the approval of the OM5 term is pending completion of ISO's FDIS ballot. Then "(OM5)" can be added to the description as well.

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

Resolve using response to comment #23.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 138 SC 138.7 P261 L35 # 21
Kolesar, Paul CommScope

Comment Type T Comment Status A om5

The IEC specification containing the equivalent of TIA-492AAAE was approved for publication during the week of 24 April 2017 by SC86A WG1. The equivalent fiber is called A1a.4.

SuggestedRemedy

Consider replacing "(TIA-492AAAE)" with "(type A1a.4)". Note that the OM5 cabling name is very likely to be approved with the FDIS ballot of 11801-1. Then the cell entry can be simplified to "0.5 m to 100 m for OM5".

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using response to comment #23.

Cl 138 SC 138.7 P261 L39 # 60
King, Jonathan Finisar

Comment Type TR Comment Status A <late>

In Table 138-8, 138-9, and 138-10, the values for OMAouter, OMAouter minus TDECQ, and TDECQ, the values for SRS, receiver sensitivity and SEC, and the values for link budget and link penalties, are all marked TBC. There have been no contributions to confirm or counter these values.

SuggestedRemedy

Replace the TBC's within the tables 138-8, 138-9 and 138-10 with editors notes below each table which say 'The values for OMAouter, OMAouter minus TDECQ, and TDECQ, require confirmation and may change' and 'The values for SRS, receiver sensitivity and SEC, require confirmation and may change' and 'The values for link budget and allocation for penalties require confirmation and may change'.

Response Response Status C

ACCEPT.

[Editor's note: This comment was received after the task force review closed.]

Cl 138 SC 138.7.2 P262 L48 # 5
Anslow, Pete Ciena

Comment Type E Comment Status A <bucket>

Minus signs should be an en-dash (Ctrl-q Shft-p)

SuggestedRemedy

in "-7", change the minus sign to an en-dash (Ctrl-q Shft-p)

Response Response Status C

ACCEPT.

Cl 138 SC 138.7.3 P263 L9 # 59
King, Jonathan Finisar

Comment Type TR Comment Status D <withdrawn>

In Table 138-10, the OM3 column, reach insertion loss and additional insertion loss are in magenta text. The OM3 reach of 70m and channel insertyion loss of 1.8 dB is entirely consistent with the 100 m OM4 reach and insertion loss. These values are consistent with previous 100m OM4 capable PMDs e.g. clause 95, and do not need to be in magenta text

SuggestedRemedy

Change the magenta text in the OM3 column of Table 138-10 .

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

[Editor's note: This comment was received after the task force review closed.]

Cl 138 SC 138.7.3 P263 L9 # 19
King, Jonathan Finisar

Comment Type TR Comment Status A

In Table 138-10, the OM3 column, reach insertion loss and additional insertion loss are in magenta text. The OM3 reach of 70m and channel insertyion loss of 1.8 dB is entirely consistent with the 100 m OM4 reach and insertion loss. These values are consistent with previous 100m OM4 capable PMDs e.g. clause 95, and do not need to be in magenta text

SuggestedRemedy

Change the magenta text in the OM3 column of Table 138-10 .

Response Response Status C

ACCEPT IN PRINCIPLE.

Change the magenta text to black in the OM3 column of Table 138-10.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

CI 138 SC 138.8.1.1 P265 L4 # 45
 Dawe, Piers Mellanox

Comment Type T Comment Status A

Instead of just giving an arbitrary (usually excessive) requirement on the test, give the reason, as in 121.8.5.1 and 122.8.5.1, so that implementers can't easily evade the spirit of the spec but can use their common sense to make affordable test rigs.
 The 1 ns (about 27 UI) of Skew that is called out in footnote a to Table 116-7 is mostly Skew that the host might make, not Skew between module input and PMD circuitry.
 giannakopoulos_01_0508 said:
 'PMA to PMD connection
 - Traces should in any case be carefully laid out
 - Should be less than 1" (per direction), which is 0.45 ns (RX and TX)'
 When testing from the optical side, there could be mismatch between optical paths, but 1 ns = 20 cm is not likely.
 The point is that the lanes should not be correlated in the module, and as both the input and output signals are available, the tester can find out what is really needed if he wishes.

SuggestedRemedy

After "there is at least 31 UI delay between the test pattern on one lane and the pattern on any other lane", add "so that the symbols on each lane are not correlated within the PMD".

Response Response Status C
 ACCEPT.

CI 138 SC 138.9.2 P267 L44 # 58
 King, Jonathan Finisar

Comment Type TR Comment Status D <withdrawn>

Hazard level is TBD

Analysis shows the max average launch power per lane specs for 50GBASE-SR 100GBASE-SR2 and 200GBASE-SR4 can be consistent with Hazard level 1M requirements.

SuggestedRemedy

In 138.9.2 and 138.9.7, change 'TBD' to '1M', similarly for the PICS ES2 .

Proposed Response Response Status Z
 PROPOSED REJECT.

This comment was WITHDRAWN by the commenter because it is a duplicate.

[Editor's note: This comment was received after the task force review closed.]

CI 138 SC 138.9.2 P267 L44 # 18
 King, Jonathan Finisar

Comment Type TR Comment Status A hazard level

Hazard level is TBD

Analysis shows the max average launch power per lane specs for 50GBASE-SR 100GBASE-SR2 and 200GBASE-SR4 can be consistent with Hazard level 1M requirements.

SuggestedRemedy

In 138.9.2 and 138.9.7, change 'TBD' to '1M', similarly for the PICS ES2 .

Response Response Status C
 ACCEPT.

See king_3cd_02_0517.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 138 SC 138.10.1 P268 L28 # 23
Kolesar, Paul CommScope

Comment Type T Comment Status A om5

The IEC specification containing the equivalent of TIA-492AAAE was approved for publication during the week of 24 April 2017 by SC86A WG1. The equivalent fiber is called A1a.4.

SuggestedRemedy

Consider replacing "(TIA-492AAAE)" with "(IEC type A1a.4)" in three instances within the paragraph. Note that while approval of the IEC CDV ballot allowed OM5 content to remain in ISO 11801-1, the approval of the OM5 term is pending completion of ISO's FDIS ballot. When the FDIS is approved, the phrase "wideband MMF (TIA-492-AAAE)" can be replaced with "OM5".

Response Response Status C

ACCEPT IN PRINCIPLE.

Implement the following changes with editorial license.

In Table 138-7 replace 'wideband MMF (TIA-492AAAE)' with 'OM5'.

In 138.7 replace 'fiber compliant to TIA-492AAAE' with 'type A1a.4 (OM5)'.

Underneath Table 138-8, in note c, replace 'If measured into type A1a.2 or type A1a.3, 50 µm fiber, in accordance with IEC 61280-1-4, or fiber compliant to TIA-492AAAE.' with 'If measured into type A1a.2, type A1a.3, or A1a.4, 50 µm fiber, in accordance with IEC 61280-1-4.'

In Table 138-10, replace 'Wideband MMF (TIA-492AAAE)' with 'OM5'.

In Table 138-10.1 replace 'wideband MMF (TIA-492AAAE)' with 'OM5', in three places.

In Table 138-14, replace 'Wideband MMF (TIA-492AAAE)' with 'OM5'.

In Table 138-15, replace 'Wideband MMF' with 'OM5', and in note c replace 'TIA-492AAAE' with 'IEC 60793-2-10 type A1a.4'.

Cl 138 SC Table 138-10 P263 L13 # 22
Kolesar, Paul CommScope

Comment Type T Comment Status A om5

The IEC specification containing the equivalent of TIA-492AAAE was approved for publication during the week of 24 April 2017 by SC86A WG1. The equivalent fiber is called A1a.4.

SuggestedRemedy

Consider replacing "(TIA-492AAAE)" with "(IEC type A1a.4)". Note that while approval of the IEC CDV ballot allowed OM5 content to remain in ISO 11801-1, the approval of the OM5 term is pending completion of ISO's FDIS ballot. Then the heading can be simplified to "OM5" to match the others.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using response to comment #23.

Cl 138 SC Table 138-14 P268 L38 # 24
Kolesar, Paul CommScope

Comment Type T Comment Status A om5

The IEC specification containing the equivalent of TIA-492AAAE was approved for publication during the week of 24 April 2017 by SC86A WG1. The equivalent fiber is called A1a.4.

SuggestedRemedy

Consider replacing "(TIA-492AAAE)" with "(IEC type A1a.4)". The name "OM5" is likely to approved with the FDIS ballot of 11801-1. Then the complete heading can be simplified to "OM5" to match the other headings.

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using response to comment #23.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

CI 138 SC Table 138-15 P269 L31 # 25
Kolesar, Paul CommScope

Comment Type T Comment Status A om5

The IEC specification containing the equivalent of TIA-492AAAE was approved for publication during the week of 24 April 2017 by SC86A WG1. The equivalent fiber is called A1a.4.

SuggestedRemedy

Consider replacing "TIA-492AAAE" with "IEC 60793-2-10 type A1a.4". Note that the OM5 name is likely to be approved with the FDIS ballot of 11801-1. Then the heading in the table at line 16 that currently says "Wideband MMF" can be simplified to "OM5".

Response Response Status C

ACCEPT IN PRINCIPLE.

Resolve using response to comment #23.

CI 139 SC 139.2 P279 L11 # 53
Nowell, Mark Cisco

Comment Type T Comment Status A <bucket>

Typo on Symbol rate 26.6525 should be 26.5625
Two instances: Pg 279, line 11 & line 16

SuggestedRemedy

fix typo

Response Response Status C

ACCEPT.

CI 139 SC 139.6.1 P283 L45 # 44
Dawe, Piers Mellanox

Comment Type TR Comment Status R extinction ratio

Following up on D1.2 comments 138 and 200. Requiring an extinction ratio of 4.5 dB restricts the range of transmitter technologies, pushing up cost and, depending on technology, worsening distortion.

SuggestedRemedy

Reduce the extinction ratio limit from 4.5 dB to 3.5 dB. Add requirements:
FR OMA-TDECQ+0.15*ER>=2.3
LR OMA-TDECQ+0.2*ER>=-1.1
(quantities in dBm or dB).
In Table 139-8, change (for maximum TDECQ) to (for maximum TDECQ and 4.5 dB extinction ratio), twice.

Response Response Status C

REJECT.

While there was some support for this comment, there is not consensus to implement the suggested remedy.

The commenter is invited to submit a comment on this topic against a future draft.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

Cl 139 SC 139.7.4 P286 L36 # 6
 Anslow, Pete Ciena

Comment Type E Comment Status A

Comment i-20 against P802.3bs D3.0 made changes to the equivalent text to this in the P802.3bs draft.
 See
http://www.ieee802.org/3/bs/comments/P802d3bs_D3p0_comments_final_a_ID.pdf#page=7
 This comment therefore proposes to make the equivalent change in the P802.3cd draft.
 Also, the draft says a "test pattern specified for extinction ratio" rather than "test pattern specified for OMAouter".
 Same issues in 140.7.4

SuggestedRemedy

In 139.7.4, change:
 "The OMAouter shall be within the limits given in Table 139-6 for 50GBASE-FR and 50GBASE-LR if measured using a test pattern specified for extinction ratio in Table 139-10. The OMAouter is defined as the difference between ." to:
 "The OMAouter shall be within the limits given in Table 139-6 for 50GBASE-FR and 50GBASE-LR. The OMAouter is measured using a test pattern specified for OMAouter in Table 139-10 as the difference between ."
 In 140.7.4, change:
 "The OMAouter shall be within the limits given in Table 140-6 if measured using a test pattern specified for extinction ratio in Table 140-10. The OMAouter is defined as the difference between ." to:
 "The OMAouter shall be within the limits given in Table 140-6. The OMAouter is measured using a test pattern specified for OMAouter in Table 140-10 as the difference between ."

Response Response Status C
 ACCEPT.

Cl 139 SC 139.7.9.1 P290 L3 # 7
 Anslow, Pete Ciena

Comment Type T Comment Status A <bucket>

Prompted by the Pre-ballot Mandatory Editorial Coordination (MEC) review of the P802.3bs draft, comment i-39 against P802.3bs D3.0 made changes to the equivalent text to this in the P802.3bs draft.
 See
http://www.ieee802.org/3/bs/comments/P802d3bs_D3p0_comments_final_a_ID.pdf#page=14
 This comment therefore proposes to make the equivalent change in the P802.3cd draft.

SuggestedRemedy

Change: "Baseline wander and overshoot and undershoot should be minimized." to:
 "Baseline wander, overshoot, and undershoot should be negligible."

Response Response Status C
 ACCEPT.

Cl 139 SC 139.7.9.3 P291 L12 # 8
 Anslow, Pete Ciena

Comment Type T Comment Status A <bucket>

Prompted by the Pre-ballot Mandatory Editorial Coordination (MEC) review of the P802.3bs draft, comment i-40 against P802.3bs D3.0 made changes to the equivalent text to this in the P802.3bs draft.
 See
http://www.ieee802.org/3/bs/comments/P802d3bs_D3p0_comments_final_a_ID.pdf#page=14
 This comment therefore proposes to make the equivalent change in the P802.3cd draft.

SuggestedRemedy

Change:
 "Care should be taken to minimize the noise/jitter introduced by the O/E, filters, and oscilloscope and/or to correct for this noise." to:
 "The noise/jitter introduced by the O/E, filters, and oscilloscope should be negligible or the results should be corrected for its effects."

Response Response Status C
 ACCEPT.

Cl 140 SC 140 P300 L1 # 46
 Dawe, Piers Mellanox

Comment Type TR Comment Status R

This is a "cutting edge" proposal.

SuggestedRemedy

Show technical and economic feasibility for these draft numbers, or change them (e.g. if better receiver sensitivity is possible).

Response Response Status C
 REJECT.

Insufficient evidence in comment.
 No changes to the draft proposed in suggested remedy.

The current specification in D1.2 was based on the baseline proposal in [traverso_3cd_03a_0916.pdf](#) adopted in the September 2016 interim meeting by motion #7, Y: 62 N: 0 A: 15.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

CI 140 SC 140.2 P302 L11 # 54
 Nowell, Mark Cisco

Comment Type T Comment Status A
 Typo on Symbol rate 26.6525 should be 26.5625
 Two instances: Pg 302, line 11 & line 16

SuggestedRemedy
 fix typo

Response Response Status C
 ACCEPT IN PRINCIPLE.

Change "26.6525" to "53.125" in two places.

CI 140 SC 140.6.1 P306 L40 # 47
 Dawe, Piers Mellanox

Comment Type TR Comment Status R
 Following up on D1.2 comments 139 and 211. Requiring an extinction ratio of 5 dB restricts the range of transmitter technologies, pushing up cost and, depending on technology, worsening distortion.

SuggestedRemedy
 Reduce the extinction ratio limit from 5 dB to 3.5 dB. Add a requirement that $OMA-TDECQ+0.25*ER > -0.1$ (quantities in dBm or dB).
 In Table 140-8, change (for max TDECQ) to (for maximum TDECQ and 5 dB extinction ratio), twice.

Response Response Status C
 REJECT.

While there was some support for this comment, there is not consensus to implement the suggested remedy.

The commenter is invited to submit a comment on this topic against a future draft along with a complete proposal with consensus.

CI 140 SC 140.7.5 P309 L38 # 35
 David Lewis Lumentum

Comment Type TR Comment Status D
 Data will be presented at the next interim meeting in support of changing the TDECQ reference equalizer for 100GBASE-DR transmitters.

Although the TDECQ reference equalizer does not imply any particular receiver equalizer implementation, there will be unnecessary margin in the link budget if the penalty based on TDECQ is overstated. Feedback from those developing 53 GBd PAM4 receiver ICs is that for the foreseeable future, the receiver's ADC will acquire 1 sample per symbol and the equalizer will have a minimum of 7 T-spaced FFE taps. It is therefore reasonable to specify a TDECQ/SECQ reference equalizer with 5 T-spaced FFE taps for 100GBASE-DR.

TDECQ testing of high quality 53 GBd PAM4 transmitters is failing the 2.5 dB limit in Table 140-6.

Experimental results show that increasing the reference equalizer length from $5*T/2$ to $7*T$ or longer reduces TDECQ to below 2.5 dB.

Short equalizers such as $5*T/2$ or $3*T$ result in higher TDECQ compared to longer equalizers such as $5*T$ or $7*T$. See 802.3bs smf ad hoc presentations lecheminant_01_1016_smf page 4 and mazzini_01a_0317_smf page 8.

SuggestedRemedy

Change from:
 The TDECQ shall be within the limits given in Table 140-6 if measured using the methods specified in 121.8.5.1, 121.8.5.2, and 121.8.5.3 using a reference equalizer as described in 121.8.5.4, with the following exceptions:
 The optical return loss of the transmitter compliance channel is 15.5 dB
 -- The signaling rate of the test pattern generator is as given in Table 140-6.
 -- There are no interfering lanes and therefore the delay requirement of at least 31 UI between test pattern on one lane and any other lane, as specified in 121.8.5.1, is redundant. The combination of the O/E converter and the oscilloscope has a fourth-order Bessel-Thomson filter response with a bandwidth of 38.68 GHz.

Change to:
 The TDECQ shall be within the limits given in Table 140-6 if measured using the methods specified in 121.8.5.1, 121.8.5.2, and 121.8.5.3, with the following exceptions:
 The optical return loss of the transmitter compliance channel is 15.5 dB.
 -- The signaling rate of the test pattern generator is as given in Table 140-6.
 -- There are no interfering lanes and therefore the delay requirement of at least 31 UI between test pattern on one lane and any other lane, as specified in 121.8.5.1, is redundant. The combination of the O/E converter and the oscilloscope has a fourth-order Bessel-Thomson filter response with a bandwidth of 38.68 GHz.

The reference equalizer is a 5 tap, T spaced, feed-forward equalizer (FFE), where T is the symbol period.

IEEE P802.3cd 50 Gb/s, 100 Gb/s, 200 Gb/s Ethernet 4th Task Force review comments

NOTE -- This reference equalizer is part of the TDECQ test and does not imply any particular receiver equalizer implementation.

Proposed Response *Response Status* **Z**

REJECT.

This comment was WITHDRAWN by the commenter.