Cl **137** *SC* **137.9.3.1** *P* **250** *L* **1** # 37

Dawe, Piers — Mellanox

Comment Type **TR** *Comment Status* **R** *Return loss*

Now that COM is defined with a near-neutral termination and package impedance, receiver mismatch is the receiver designer's concern, not the standard's, unless it is very extreme, because the receiver interference tolerance test finds its effect combined with other receiver attributes. And we don't expect transmitter return loss to align to the COM model any more. This RL is much tighter than CEI-56G-LR at low (and high) frequency (although apparently looser between 4 and 9 GHz). At low frequencies it is tighter than the channel RL. The effect of (good) RL at low frequency is much less than the less good RL at higher frequencies anyway. So we can go back to what we had a few drafts ago.

SuggestedRemedy

Change "shall meet
Equation (137-1)" to "shall meet Equation (93-3)" and delete Eq 137-1 and Fig 137-3.
Or, change 14.25 - f to 12 -0.625f, revise the figure.

Response *Response Status* **U**

REJECT.

The presentation dawe_3cd_01a_0917 was reviewed. Further information was requested by the task force on the system implications of the proposed return loss relaxation.

There was no consensus to implement the proposed changes.

Cl **138** *SC* **138.7.1** *P* **270** *L* **10** # 38

Dawe, Piers — Mellanox

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

It seems that it is possible to make a bad transmitter (e.g. with a noisy or distorted signal), use emphasis to get it to pass the TDECQ test, yet leave a realistic, compliant receiver with an unreasonable challenge, such as high peak power, high crest factor, or a need to remove emphasis from the signal, contrary to what equalizers are primarily intended to do. With some of the changed low-bandwidth TDECQ being used to equalize the reference receiver's own bandwidth, this issue becomes more apparent. Note the receiver is tested for a very slow signal only, not for any of these abusive signals. This is an issue for all the PAM4 optical PMDs, although it may be worse for MMF because of the high TDECQ limit.

SuggestedRemedy

1. To screen for noisy or distorted signals with heavy emphasis
Define TDECQrms = 10*log10(A_RMS/(s*3*Qt*R)) where A_RMS is the standard deviation of the measured signal after the 13.28125 GHz filter response, Qt and R are as already in Eq 212-12. s is the standard deviation of a fast clean signal with OMA=2 and without emphasis, observed through the 13.28125 GHz filter response (around 0.7 - can be calculated when the filter bandwidth is stable). Set limit for TDECQrms according to what level of dirty-but-emphasised signal we decide is acceptable, add max TDECQrms row to the table. Alternatively, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be in the TDECQ procedure 121.8.5.3 as proposed in bs comment(s). Similarly in clauses 139, 140.
2. To protect the TIA input, consider a peak power spec as in Clause 86.
3. To protect the TIA and any AGC and TIA from unreasonable signals, consider a crest factor spec.
4. To protect the equalizer from having to support unnecessary settings, require that the cursor is one of the first three taps.
5. To protect the receiver from having to "invert" heavily over-emphasised signals, set a minimum cursor weight.

Response *Response Status* **U**

REJECT.

This comment is related to unsatisfied comments i-140 and r02-35 against 802.3bs draft 3.2.

The resolution to P802.3bs comment r02-35 was:
"REJECT
Insufficient evidence of the claimed problem and that the proposed remedy fixes the problem. The commenter is invited to provide a contribution that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed additional requirement prevents this issue from occurring."

Insufficient evidence was provided of the claimed problem and that the suggested remedy fixes the problem. A contribution is invited that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that

the proposed additional requirements prevent this issue from occurring.

Cl **139** *SC* **139.6.1** *P* **291** *L* **36** # 40

Dawe, Piers Mellanox

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

The discussion around D2.0 comment 152 implied that there is receiver margin to spare in 50GBASE-FR.

SuggestedRemedy

reduce all the optical power levels for 50GBASE-FR (except Rx damage) by 1 dB.
Bring more evidence for what optical power levels and TDECQ limits are right, including TDECQ measurements with SSPRQ, and correlation to actual receiver performance.
Review the TDECQ limit.

Response *Response Status* **U**

REJECT.

This comment is a follow up comment to comment #152 to D2.0.

The current values are based on the adoption of a baseline proposal in http://www.ieee802.org/3/cd/public/May16/cole_3cd_01_0516.pdf during the May 2016 meeting in Whistler by a motion with the following results. Y: 54 N: 0 A: 25.

It is known that there are margins in both transmitter and receiver specifications when the baseline proposal was adopted.

No analysis has been provided that changing the current values by 1 dB would enable lower cost solutions and/or better performance.

Cl **140** *SC* **140.6.1** *P* **314** *L* **33** # 42

Dawe, Piers Mellanox

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

D2.0 comment 128: PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen too little experimental information showing technical and economic feasibility. As measurements with the new TDECQ method and with new receiver designs become available, it may be that optical power levels can be reduced and the spec as in this draft would be uneconomic.

SuggestedRemedy

Reduce all the optical power levels for 100GBASE-DR by 0.5 dB.
Bring more evidence for what optical power levels and TDECQ limits are right; in particular, TDECQ measurements with SSPRQ, and correlation to actual receiver performance.
Review the TDECQ limit.

Response *Response Status* **U**

REJECT.

No analysis has been provided that changing the current values by 0.5 dB would enable lower cost solutions and/or better performance.

Furthermore the existing values for 100GBASE-DR are intentionally consistent with the values for one lane in 400GBASE-DR4 in P802.3bs.

A presentation (dawe_3bs_03_0917) containing similar proposals pertaining to 400GBASE-DR4 in P802.3bs D3.3 was not accepted.

Cl **136** *SC* **136.9.3** *P* **224** *L* **10** # 43

Dawe, Piers Mellanox

Comment Type **TR** *Comment Status* **R** *<NSR>*

As explained before, J4u should be changed to J3u. The equivalent J3u is known (D2.0 comment 144) but we need an estimate of the difference in jitter between TP0a and TP2 so that we can choose more appropriate limits for the two test points (D2.0 comment 143).

SuggestedRemedy

Change J4u to J3u, here and in 137. Choose the limit at TP2 considering jitter limit at TP0a, the mated compliance board crosstalk specs, and the slower edges at TP2.

Response *Response Status* **U**

REJECT.

The suggested remedy is not specific and cannot be used to apply a change in the draft.

More consensus around a specific remedy is required.

Cl **136** *SC* **136.11.2** *P* **232** *L* **28** # 44

Dawe, Piers — Mellanox

Comment Type **TR** *Comment Status* **R** *Cable assembly*

Where did 17.16 dB come from? the limit should be consistent with other 3 m cables: 16.48 or 15.5 dB (CA-25G-S CA-25G-N), adjusted for Nyquist frequency. Setting it too high is objective creep.

SuggestedRemedy

Set the max loss to be no more than consistent with CA-25G-S. Set the RITT losses accordingly.

Response *Response Status* **U**

REJECT.

The value 17.16 dB is included in the resolution of comment #124 against D2.0, based on palkert_3cd_01b_0717 and the task force discussion following the presentation. This number makes the channel IL the same as for Clause 137.

No further changes are required to close the budget.

Cl **136** *SC* **136.11.7** *P* **233** *L* **18** # 47

Dawe, Piers — Mellanox

Comment Type **TR** *Comment Status* **R** *Cable assembly*

The COM impedances should be moved towards neutral, as explained in D2.0 comment 71 and 113.

SuggestedRemedy

Make changes proposed in D2.0 comment 71 and hidaka_3cd_01_0717 - except don't change the parameter name unless it is coordinated with the name used in Annex 93A.

Response *Response Status* **U**

REJECT.

Comment #71 against D2.0 suggested changing COM parameters to use well-matched impedances: terminations of 50 Ohm, package impedance of 95 Ohm and board impedance of 100 Ohms.

D2.0 comment #71 was rejected due to lack of consensus.

The related changes suggested in D2.0 comment #113 were also not in consensus.

The comment does not provide any new information, nor address any concerns that prevented the prior comments from being adopted.

Cl **137** *SC* **137.9.2** *P* **249** *L* **28** # 49

Dawe, Piers — Mellanox

Comment Type **TR** *Comment Status* **R** *Tx specs*

Transmitter output residual ISI SNR_ISI (min) 36.8 dB (Clause 136) and 43 dB (Clause 137) is still too high see dawe_3bs_04_0717 and dawe_3cd_02a_0717 - can barely measure the IC through the test fixture. The warning NOTE in 120D.3.1.7 shows the issue, but doesn't solve it. D2.0 comment 140

SuggestedRemedy

See presentation.

Response *Response Status* **U**

REJECT.

The task force reviewed rysin_3cd_02_0917.

There was no consensus for implementing the proposed changes.

Cl **137** *SC* **137.9.2** *P* **249** *L* **29** # 50

Dawe, Piers — Mellanox

Comment Type **TR** *Comment Status* **A** *Tx specs*

Signal-to-noise-and-distortion ratio (min), increased to 33.3 dB (Clause 136) and to 32.5 dB (Clause 137) for all Tx emphasis settings, is too high: see dawe_3bs_04_0717 and dawe_3cd_02a_0717 - can barely measure the IC through the test fixture. It seems SNDR depends on emphasis, while COM assumes the spec limit at all emphasis settings which is pessimistic and not realistic. Also I suspect there is double counting of jitter in SNDR and as jitter, in COM. D2.0 comment 139.

SuggestedRemedy

Apply a SNDR limit that accounts for the way Pmax varies with emphasis: SNDR0+20log10(Pmax_equalized/Pmax_unequalized), or apply the SNDR spec for no emphasis only.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

The response to comment #139 against D2.0 was:
"REJECT.
dawe_3cd_02_0717 was presented.
The comment highlights some issues in the current draft, but there was no consensus for adopting any of the proposed solutions.
The commenter is encourged to build consensus and bring a new proposal."

The suggested remedy is a new proposal.

The commenter points out an issue and proposed solution that need further investigation.

There is no consensus to make the proposed changes.

Cl **137** *SC* **137.9.2** *P* **249** *L* **29** # 51

Dawe, Piers — Mellanox

Comment Type **TR** *Comment Status* **R** *Tx specs*

COM SNR_TX is defined at the TX output. SNDR is measured thru package and TF by real (imperfect) test equipment therefore is lower than SNR_TX, causing some double counting in COM. D2.0 comment 139.

SuggestedRemedy

Reduce the SNDR specification to 29 dB for both Clause 136 and 137 to account for the degradation caused by the package and test fixture as well as by the measurement impairments.

Response *Response Status* **U**

REJECT.

The task force reviewed rysin_3cd_01_0917.

The package and test fixture effects are linear, so are effectively de-embedded in the linear fit procedure.

The claim that measured SNDR is lower than "real" SNDR is not substantiated.

Creating a difference of 3.5 dB between the COM parameter (SNR_TX) and the corresponding TX parameter (SNDR) would break the budget. Bad transmitters may pass the Tx specs but cause their partner's receiver to fail.

There is no consensus to make the proposed changes.

Cl **137** *SC* **137.9.2** *P* **249** *L* **30** # 52

Dawe, Piers — Mellanox

Comment Type **TR** *Comment Status* **R** *Return loss*

Now that COM is defined with a near-neutral termination and package impedance, we don't expect transmitter return loss to align to the COM model any more. This RL is much tighter than CEI-56G-LR at low (and high) frequency (although apparently looser between 4 and 9 GHz). At low frequencies it is tighter than the channel RL. The effect of (good) RL at low frequency is much less than the less good RL at higher frequencies anyway, and there is less concern about end-to-end reflections than in C2C because the loss is higher when the receiver is challenged. So we can go back to what we had a few drafts ago.

SuggestedRemedy

If bs doesn't fix this, add another exception and create new equation for Tx RL that is similar to the Cl.93 and the channel RL at low frequencies; 12 -0.625f, 8.7-0.075f. Add figure to illustrate. Refer to new equation instead of existing 137-1.
If 137-1 is revised as above for the receiver, can continue to point to it.

Response *Response Status* **U**

REJECT.

The presentation dawe_3cd_01a_0917 was reviewed. Further information was requested by the task force on the system implications of the proposed return loss relaxation.

There was no consensus to implement the proposed changes.

Cl **136** *SC* **136.11** *P* **223** *L* **42** # 20113

Dudek, Mike — Cavium

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

Equation 92-27 for the differential return loss gives 5.3dB return loss at 13.28GHz. This is not the 6dB listed and is a relatively poor value and could lead to significant differences between system performance with a real host and the COM calculated with the single 110 Ohm host board trace equivalent. Work on backplanes and C2C (e.g. Hidaka_3cd_01a_0317, Dudek_3bs_02_0517) has shown that this affect is significant and it would be better to test COM with nominal impedances and have a guard band between the channel COM and the Interference tolerance COM.

SuggestedRemedy

Change 6 to 5.3 Change the COM value to 3.5dB. In table 136-15 change the value of Rd to 50 Ohm, the value of Zc to 95 Ohm, On page 224 line 40 change the value of COM to 3.5dB. Change the impedance of the test trace from TP0 to TP1 and TP4 to TP5 to 100 Ohm by changing on page 226 line 41 from "using zp = 151 mm in length, representing an insertion loss of 6.42 dB at 13.28 GHz on each PCB." to "using Zc = 100 Ohm and zp = 151 mm in length, representing an insertion loss of 6.42 dB at 13.28 GHz on each PCB." Also change to 3.5dB in PICS CA8.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

[Editor's note: This D2.0 comment was unsatisfied.]

In Table 136-14, Change "Minimum differential return loss at 13.28 GHz" from 6 dB to 5.3 dB.

The rest of the suggested remedy requires more consensus building.

See also #71

Cl **138** *SC* **138.7.1** *P* **262** *L* **18** # 20127

Dawe, Piers — Mellanox

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

It seems that it is possible to make a bad transmitter (e.g. with a noisy or distorted signal), use emphasis to get it to pass the TDECQ test, yet leave a realistic, compliant receiver with an unreasonable challenge (up to 4/2 dB worse than the SRS test?) With some of the changed low-bandwidth TDECQ being used to equalize the reference receiver's own bandwidth, this issue becomes more apparent.
This is an issue for all the PAM4 optical PMDs, although it may be worse for MMF because of the high TDECQ limit.

SuggestedRemedy

Define TDECQrms = 10*log10(A_RMS/(s*3*Qt*R)) where A_RMS is the standard deviation of the measured signal after the 13.28125 GHz filter response. s is the standard deviation of a fast clean signal with OMA=2 and without emphasis, observed through the 13.28125 GHz filter response (around 0.7 - can be calculated when the filter bandwidth is stable).
Set limit for TDECQrms according to what level of dirty-but-emphasised signal we decide is acceptable, add max TDECQrms row to the table. Alternatively, if the same relative limit is acceptable for all PAM4 optical PMDs, the limit could be in the TDECQ procedure 121.8.5.3 as proposed in P802.3bs D3.2 comment r02-35.
Similarly in clauses 139, 140.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

A similar comment was made to P802.3bs D3.2 via comment r02-35, which was rejected.

Insufficient evidence of the claimed problem and that the proposed remedy fixes the problem.
A contribution is invited that demonstrates the problem (a waveform that passes TDECQ but cannot be decoded by a reasonable receiver implementation) and that the proposed additional requirement prevents this issue from occurring.

Cl **140** *SC* **140.6.1** *P* **306** *L* **33** # 20128

Dawe, Piers — Mellanox

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

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen too little experimental information showing technical and economic feasibility. As measurements with the new TDECQ method and with new receiver designs become available, it may be that optical power levels can be reduced and the spec as in this draft would be uneconomic.

SuggestedRemedy

Bring more evidence for what optical power levels and TDECQ limits are right; in particular, TDECQ measurements with SSPRQ, and correlation to actual receiver performance.
Based on evidence, reduce all the optical power levels for 100GBASE-DR by 0.5 or 1 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

The suggested remedy does not propose a specific change to the draft.

Cl **139** *SC* **139.7.7** *P* **289** *L* **15** # 20133

Dawe, Piers — Mellanox

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

With the lower receiver bandwidth, measuring RIN in approximately the signaling rate (twice as much) seems too much; 1/2 to 3/4 would be better. A T-spaced equalizer cannot independently adjust for good ISI and RIN filtering, so can an adequate estimate of RIN can be obtained as a by-product of the TDECQ procedure? While a T/2-spaced equalizer could enhance the RIN, it would not choose to do so if RIN were a problem, so a T-spaced reference equalizer and a T/2-spaced product equalizer are compatible from this point of view, I think. As 52.9.6 says, this RIN method is intended for components (TOSAs) not a "system level test" suitable for a complete optical module.
This is much the same as P802.3bs D3.2 comment r02-39.

SuggestedRemedy

Review; reduce the bandwidth and simplify RIN measurement to a Qsq measurement (see 68.6.7) or eliminate as appropriate. Remove 135.5.10.2.4 Square wave (quaternary) test pattern and any associated registers.
Similarly in 140.7.9.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

The suggested remedy suggests 2 different approaches to change the draft.
Changing the RIN measurement to a Qsq measurement has not been demonstrated to provide the same safeguards that are expected from the RIN requirement.
Eliminating the RIN measurement was discussed in the response to comment #130 against D2.0 of P802.3bs on the basis that "The transmitter RINxOMA spec is intended to screen out potentially bad transmitters even if the noise correction required by the TDECQ test is not very accurate."

Cl **140** *SC* **140.7.9** *P* **310** *L* **28** # 20134

Dawe, Piers — Mellanox

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

The lack of consistency between the low frequency jitter specs in 802.3bs affects 802.3cd also. Here is P802.3bs D3.2 comment r02-40 for those who have not been following this issue. Depending how this inconsistency is fixed, there may be little or no explicit change in the P802.3cd draft.
Following up on P802.3bs D3.0 comment 153 and D3.1 comment 55: if the jitter corner frequency for 26.5625 GBd (NRZ and PAM4) is 4 MHz, the low frequency ends of the jitter masks must align or be in the right order if expressed in time vs. frequency, i.e. should scale with signalling rate if in UI. If this is not done, the required depth of the LF jitter buffer in the 2:1 muxes in a 400GBASE-DR4 module is unbounded and the low frequency jitter generation requirements on the module become unreasonable. Compare 87.8.11.4 and 88.8.10: 4 MHz for 10.3125 GBd, 10 MHz for 25.78125 GBd. History:
anslow_3bs_04_0316 does not contain reasoning, refers to ghiasi_3bs_01_0316 which does not address wander and buffering. ghiasi_3bs_01a_0116.pdf#page=15 shows FIFOs but does not establish a workable spec. Slide 14 shows they can be avoided: this is what we have for 400GAUI-8 or 400GAUI-16 with 400GBASE-xR8. I have no evidence that the problems described in the [fourth] sentence have been considered or solved by the [P802.3bs] committee.

SuggestedRemedy

Add another exception for the SRS procedure, with a table like Table 121-12 replacing second row after the header row:
80 kHz < f <= 250 kHz 4e5/f
250 kHz < f <= 500 kHz 1e11/f^2
1 MHz < f <= 4 MHz 2e5/f
Or, with the UIs doubled vs. Table 121-12:
f < 40 kHz Not specified
40 kHz < f <= 4 MHz 4e5/f
4 MHz < f <= 10 LB 0.1
Increase the TDECQ limit to share the burden appropriately between transmitter and receiver.
This option means the 100G/lane receiver has to tolerate no more timing slew rate (in ps/us) than that agreed for 50G/lanes.
Or, increase jitter by 50% and corner frequency by 33%:
f < 40 kHz Not specified
40 kHz < f <= 6 MHz 4e5/f
5.333 MHz < f <= 10 LB 0.075
and add an exception in 124.8.5 that the CRU corner frequency is 5.333 MHz. Increase the TDECQ limit to share the burden between transmitter and receiver.
To do the job properly with the first option, in 124.8.5 we should add another exception to the CRU with a corner frequency of 4 MHz and a slope of 20 dB/decade (in 121.8.5.1): add a pole at 250 kHz and a zero at 500 kHz. I am advised that this can be done in hardware (in software, anything is possible).

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

One option in the suggested remedy is proposing to place an extra burden on the receiver by allowing transmitters with a higher level of TDECQ which may be due to ISI and also by requiring a higher level of jitter tolerance.
The commenter has not demonstrated that this extra burden is less onerous than putting a buffer in the PMA.
For the second option in the suggested remedy the commenter is invited to build consensus for an increase of the corner frequency to be above 4 MHz.

Cl **137** *SC* **137.9** *P* **241** *L* **1** # 20136

Dawe, Piers Mellanox

Comment Type **TR** *Comment Status* **R** *Electrical <NSR>*

We don't yet know how to write a spec for 30 dB channels that isn't bleeding edge for ICs and/or channels. This isn't Ethernet "broad market" today, it's a specialist niche.

SuggestedRemedy

Keep working on it in Working Group ballot and if things don't improve, reduce the 30 dB objective and reduce the high loss RITT loss. It might be OK to leave the channel recommended insertion loss limit if the COM spec protects the Tx and Rx.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

No specific change to the draft is suggested.

Cl **137** *SC* **137.9.2** *P* **241** *L* **22** # 20139

Dawe, Piers Mellanox

Comment Type **TR** *Comment Status* **R** *Electrical <NSR>*

Signal-to-noise-and-distortion ratio (min) 32.5 dB is too high (even worse than 120D) - probably can't measure the IC through the test fixture and cables. I suspect there is double counting of jitter in SNDR and as jitter, in COM.

SuggestedRemedy

Remove the double counting. Reduce the SNDR limit to something that can reasonably be measured, or change the measurement method.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

dawe_3cd_02_0717 was presented.

The comment highlights some issues in the current draft, but there was no consensus for adopting any of the proposed solutions.

The commenter is encourged to build consensus and bring a new proposal.

Cl **137** *SC* **137.9.2** *P* **241** *L* **21** # 20140

Dawe, Piers Mellanox

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

Output residual ISI SNR_ISI (min) 43 dB is way too high - probably can't measure the IC through the test fixture and cables, even test equipment fails this limit. The warning NOTE in 120D.3.1.7 notes the issue (for 34.8 dB), but doesn't solve it.

SuggestedRemedy

It may be necessary to move away from the SNR_ISI method.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

dawe_3cd_02_0717 was presented.

The comment highlights an issue in the current draft, but there was no consensus for adopting any of the proposed solutions.

The commenter is encourged to build consensus and bring a new proposal.

See #139.

Cl **137** *SC* **137.9.3.1** *P* **241** *L* **46** # 20141

Dawe, Piers
Mellanox

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

The low frequency RL at 14.25 dB is insignificant for signal integrity compared with the 8.7 dB at 6 GHz. This RL is much tighter than CEI-56G-LR at low (and high) frequency (although apparently looser between 4 and 9 GHz).

SuggestedRemedy

Change 14.25 - f to 12 -0.625f

Response *Response Status* **U**

ACCEPT IN PRINCIPLE.

[Editor's note: This D2.0 comment was unsatisfied.]

This issue was discussed in 802.3bs and resulted in a change to the similar specification (Comment #r02-41).

In 137.9.3.1 (Receiver input return loss), append the following text to the first paragraph: "The test fixture return loss may be de-embedded from the return loss measurements."

No need to add this in 137.9.2 (Transmitter characteristics) since it refers to 120D.3.1.1, where a similar change was applied by 802.3bs (indirectly through Table 120D-1). Update exceptions if necessary.

Implement with editorial license.

Cl **136** *SC* **136.9.3** *P* **216** *L* **11** # 20143

Dawe, Piers
Mellanox

Comment Type **TR** *Comment Status* **R** *Electrical <NSR>*

J4, now called J4u (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 J3u (1.875e-4 per bit) is a good match to the spec BER - just as J4u 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. Futher, the jitter at TP2 won't be the same as at TP0a in 137.9.2 (expected to be more).

SuggestedRemedy

Change J4 to J3u. Choose the limit at TP2 considering jitter limit at TP0a and the mated compliance board crosstalk specs, among other factors.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

The suggested remedy lacks sufficient detail required for implementation - the limits for TP2 are not included.

The commenter is encouraged to suggest and build consensus for specific limits at TP2, as well as the suggestion to change J4u to J3u.

Cl **137** *SC* **137.9.2** *P* **241** *L* **24** # 20144

Dawe, Piers — Mellanox

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

J4u in 120D (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 J3u (1.875e-4 per bit) is a good match to the spec BER - just as J4u 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 J3u, max 0.106 UI (from eq 136-6 and 7). In Eq 136-6 and 136-7 and the NOTE, change Q4=3.8906 to Q3=3.2905, Q(Q3) = 5 x10^-4.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

Note that the suggested change (J4u to J3u) seems to enable a shorter measurement while keeping the same sigma_RJ and A_DD for COM, by changing the conversion equations (136-6 and 136-7).

The task force discussed the suggested remedy. Since currently both clauses 136 and 137 use the same equations, there is preference to make changes to both clauses together. There is no consensus for changing just this clause.

See comment #143.

Cl **138** *SC* **138.7.1** *P* **262** *L* **17** # 20147

Dawe, Piers — Mellanox

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

This PMD needs more study, and knowing what TDECQ is feasible is probably the key.

SuggestedRemedy

While in WG ballot, show evidence of technical feasibility for the numbers in the spec: eyes, receiver waterfall plots, TDECQ measurements and so on. Adjust the draft as appropriate. TR because this could take a few meeting cycles.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

[Editors note: This comment is a repeat of comment 42 against draft 1.3]

No specific changes to the draft suggested.

Task force participants are encouraged to prepare consensus presentations with proposals for specific changes to the draft if necessary.

Cl **139** *SC* **139.6.1** *P* **283** *L* **36** # 20152

Dawe, Piers — Mellanox

Comment Type **TR** *Comment Status* **R** *power budget*

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen too little experimental information showing technical and economic feasibility. However, stassar_061417_3cd_adhoc-v2 shows plenty of receiver sensitivity margin (although not yet shown with SSPRQ). As more measurements with with new receiver designs and the new TDECQ method become available, it appears the optical power levels can be reduced and the spec as in this draft will be uneconomic (particularly 50GBASE-FR which should be low cost, low power, convenient for quad or octal packaging).

SuggestedRemedy

Bring more evidence for what optical power levels and TDECQ limits are right, including TDECQ measurements with SSPRQ, and correlation to actual receiver performance. Based on evidence, reduce all the optical power levels for 50GBASE-FR and 50GBASE-LR by 0.5, 1 or 1.5 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response *Response Status* **U**

REJECT.

[Editor's note: This D2.0 comment was unsatisfied.]

The suggested remedy does not propose a specific change to the draft.