

IEEE P802.3bs D3.4 200 Gb/s & 400 Gb/s Ethernet 4th Sponsor recirculation ballot comments

Cl 120E SC 120E.3.1 P 369 L 19 # i-119
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

The host is allowed to output a signal with large peak-to-peak amplitude but very small EH - in other words, a very bad signal. If the module is exactly like the reference receiver, that would work - but that's not a reasonable "if".

SuggestedRemedy

We may need some other spec to protect the module from unexpected signals.

Response Response Status U

REJECT.
 No remedy provided. The commenter is encouraged to provide a presentation on this subject.

Cl 121 SC 121.8.5.3 P 228 L 9 # i-140
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

It may be 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.

SuggestedRemedy

Define TDECQrms = $10 \cdot \log_{10}(C_{dc} \cdot A_{RMS} / (s^3 \cdot Q_t \cdot R))$ where A_RMS is the standard deviation of the measured signal after the 19.34 GHz filter response and s is the standard deviation of a fast clean signal with OMA=0.5 and without emphasis, observed through the 19.34 GHz filter response (from memory I believe s is about 0.82). Require that TDECQrms shall not exceed the limit for TDECQ. If we think it's justified, we could allow a slightly higher limit for TDECQrms.

Response Response Status U

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.

Cl 120D SC 120D.3.1.1 P 353 L 24 # r01-36
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status A

Transmitter Output residual ISI SNR_ISI (max) 38 dB is too high - probably can't measure the IC through the test fixture and cables.

SuggestedRemedy

Start by checking whether Gaussian assumptions are tripping us up.

Response Response Status U

ACCEPT IN PRINCIPLE.

See response to comment #r01-22

[Editor's note added after comment resolution completed.

The response to comment r01-22 is:

In Table 120D-1:

Change the minimum SNR_ISI value from 38 to 34.8 dB.

Change the minimum SNDR from 31 to 31.5 dB.

Change Linear fit pulse peak (min) from $0.736 \cdot V_f$ to $0.76 \cdot V_f$

In Table 120D-8:

Change Av and Afe values from 0.45 to 0.44

Add another NOTE at the end of 120D.3.1.7:

NOTE 2--The observed SNR_ISI can be significantly influenced by the measurement setup, e.g. reflections in cables and connectors. High-precision measurement and careful calibration of the setup are recommended.

]

Cl 120D SC 120D.3.1.8 P 358 L 46 # r01-41
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

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.

SuggestedRemedy

Change $14.25 - f$ to $12 - 0.625f$

Response Response Status U

REJECT.
 No consensus to make a change at this time, but further investigation is encouraged.

[Editor's note added after comment resolution completed. The consensus view was that further investigation of the effect of Return Loss at low frequencies should take place, but no change to the equation can be justified at this time.]

IEEE P802.3bs D3.4 200 Gb/s & 400 Gb/s Ethernet 4th Sponsor recirculation ballot comments

Cl 124 SC 124.8.9 P 302 L 31 # r01-55
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Following up on D3.0 comment 153: if the jitter corner frequency for 26.5625 GBd (NRZ and PAM4) is 4 MHz, the low frequency (sloping) part of the jitter mask should scale with signalling rate, i.e. align if expressed in time vs. frequency, to avoid a need for a poorly specified wander buffer in the 2:1 muxes in a 400GBASE-DR4 module. 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.

SuggestedRemedy

Add another exception for the SRS procedure, with a table like Table 121-12 but with the frequencies doubled.

Or, replacing second row after the header row:

80 kHz < f <= 500 kHz	4e5/f
500 kHz < f <= 1 MHz	2e11/f^2
1 MHz < f <= 4 MHz	2e5/f

Response Response Status U

REJECT.

This issue was already discussed in response to comment i-153 to D3.0 which was: "The jitter corner frequency was extensively discussed within the Task Force with multiple presentations on the topic. The CRU corner frequency was chosen to be 4 MHz for all interfaces (including 400GBASE-DR4) in the March 2016 TF meeting as recorded in: http://www.ieee802.org/3/bs/public/16_03/anslow_3bs_04_0316.pdf."

The possible need for a buffer was discussed in presentations made leading up to this decision. For example, see: http://www.ieee802.org/3/bs/public/16_01/ghiasi_3bs_01a_0116.pdf#page=15

There was no consensus to make a change to the draft.

Cl 121 SC 121.7.1 P 221 L 25 # r02-28
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen far too little experimental information showing technical and economic feasibility. It looks like this PMD can be made to work but as measurements with the new TDECQ method and with new receiver designs become available, we expect the optical power levels can be reduced and the spec as in this draft will 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 200GBASE-DR4 by 0.5, 1 or 1.5 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response Response Status U

REJECT.

This comment does not apply to the substantive changes between IEEE P802.3bs/D3.2 and IEEE P802.3bs/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

The suggested remedy does not propose any changes to the draft.

Cl 121 SC 121.8.5.1 P 226 L 49 # r02-31
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Using the same pattern on the aggressor lanes (correlated crosstalk) is very unusual. Does what we gain in correctly handling the spectrum of the deterministic part of the crosstalk outweigh what we lose in inconsistency vs. UI- and sub-UI phasing? As D3.1 comment 13 points out, using the conventional uncorrelated crosstalk can simplify the PMA. It should be possible to calculate the relative measurement accuracy of the two approaches.

SuggestedRemedy

Work out which is better; change the crosstalk patterns here and the related pattern generator options in Clause 120 as appropriate.

Response Response Status U

REJECT.

The suggested remedy does not propose any changes to the draft.

The commenter is invited to perform the calculation suggested in the comment and prepare a consensus presentation with proposed changes to the draft.

IEEE P802.3bs D3.4 200 Gb/s & 400 Gb/s Ethernet 4th Sponsor recirculation ballot comments

CI 121 SC 121.8.5.3 P 229 L 42 # r02-35
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Updating D3.0 comment 140:
 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 2.5/2 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.

SuggestedRemedy

Define TDECQrms = $10 \cdot \log_{10}(A_RMS / (s^3 \cdot Qt \cdot R))$ where A_RMS is the standard deviation of the measured signal after the 13.28125 GHz filter response. s is close to the standard deviation of a fast clean signal with OMA=0.5 and without emphasis, observed through the 13.28125 GHz filter response, according to what level of dirty-but-emphasised signal we decide is acceptable. Require that TDECQrms shall not exceed the limit for TDECQ.

Response Response Status U

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.

CI 122 SC 122.7.1 P 252 L 14 # r02-36
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

PAM4 optics is still new and raw, we are still debugging the specification methodology, and we have seen far 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, consider reducing all the optical power levels in this clause except the -30 dBm signal detect limit by 0.5 or 1 dB (with other adjustments for other reasons). Review the TDECQ limits.

Response Response Status U

REJECT.
 This comment does not apply to the substantive changes between IEEE P802.3bs/D3.2 and IEEE P802.3bs/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

The suggested remedy does not propose any changes to the draft.

CI 124 SC 124.7.1 P 298 L 4 # r02-37
 Dawe, Piers J G Mellanox Technologie

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 400GBASE-DR4 by 0.5 or 1 dB (with other adjustments for other reasons). Review the TDECQ limit.

Response Response Status U

REJECT.
 This comment does not apply to the substantive changes between IEEE P802.3bs/D3.2 and IEEE P802.3bs/D3.1 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

The suggested remedy does not propose any changes to the draft.

IEEE P802.3bs D3.4 200 Gb/s & 400 Gb/s Ethernet 4th Sponsor recirculation ballot comments

CI 124 SC 124.8.9 P 302 L 46 # r02-40
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Following up on 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 second sentence have been considered or solved by the 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.

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.

CI 120D SC 120D.3.1.1 P 353 L 24 # r02-42
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Signal-to-noise-and-distortion ratio (min) 31.5 dB is too high (increased by D3.1 comment 22, so even worse than before) - 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.

The presentation:

http://www.ieee802.org/3/bs/public/17_07/dawe_3bs_04_0717.pdf was reviewed.

Changing the SNDR limit to 28.5 dB is considered to be placing too great a burden on the receiver and it has not been demonstrated that implementations cannot meet the current specification.

CI 120D SC 120D.3.1.1 P 353 L 26 # r02-43
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Following D3.1 comments 22 and 36: transmitter Output residual ISI SNR_ISI (min) 34.8 dB is still 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 shows the issue, but doesn't solve it.

SuggestedRemedy

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

Response Response Status U

REJECT.

No remedy provided

IEEE P802.3bs D3.4 200 Gb/s & 400 Gb/s Ethernet 4th Sponsor recirculation ballot comments

Cl **120D** SC **120D.3.1.1** P **354** L **36** # **r02-44**
 Dawe, Piers J G Mellanox Technologie

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

Following D3.1 comment 41: 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-MR 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**

REJECT.
 Re-statement of comment r01-41 which was rejected with the response:
 No consensus to make a change at this time, but further investigation is encouraged.
 [Editor's note added after comment resolution completed. The consensus view was that further investigation of the effect of Return Loss at low frequencies should take place, but no change to the equation can be justified at this time.]

While additional work has been done on this topic, there is still no consensus to make a change.

Cl **120E** SC **120E.3.1** P **371** L **20** # **r02-46**
 Dawe, Piers J G Mellanox Technologie

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

Building on D3.0 comment 119: The host is allowed to output a signal with 900 mV peak-to-peak amplitude but only 32 mV eye height - a very bad signal. If the module is exactly like the reference receiver, that would work, but with a good but slightly different receiver the eye will collapse.

SuggestedRemedy

We need some other spec to protect the module from such unexpected signals. A vertical eye closure spec will probably work. I'll try to bring a presentaitaon.

Response Response Status **U**

REJECT.
 No presentation providing a suggested remedy for this comment was submitted.
 While a vertical eye closure specification was considered worth further investigation, no consensus was reached to make a change to the draft.

Cl **121** SC **121.8.5.3** P **228** L **43** # **r03-27**
 Dawe, Piers J G Mellanox Technologie

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 2.5/2 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.
 D3.0 comment 140, D3.2 r02-35

SuggestedRemedy

Define $TDECQ_{rms} = 10 \cdot \log_{10}(A_{RMS}/(s^3 \cdot Q_t \cdot R))$ where A_{RMS} is the standard deviation of the measured signal after the 13.28125 GHz filter response. We choose s , which is close to the standard deviation of a fast clean signal with $OMA=0.5$ and without emphasis, observed through the 13.28125 GHz filter response, according to what level of dirty-but-emphasised signal we decide is acceptable. Q_t and R are as in Eq 121-12. Require that $TDECQ_{rms}$ shall not exceed the limit for TDECQ.

Response Response Status **U**

REJECT.
 This is related to unsatisfied comments i-140 and r02-35.
 The resolution to comment r02-35 was: 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.

The proposed remedy is almost identical to the one proposed in r02-35.
 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, has not been provided.

IEEE P802.3bs D3.4 200 Gb/s & 400 Gb/s Ethernet 4th Sponsor recirculation ballot comments

Cl 120D SC 120D.3.1.1 P 353 L 24 # r03-30
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Signal-to-noise-and-distortion ratio (min), increased to 31.5 dB for all Tx emphasis settings, is too high: see daw_3bs_04_0717 and daw_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.
 D3.2 r02-42

SuggestedRemedy

Either apply the SNDR spec for no emphasis only, and adjust eq 93A-30 for the way sigma_e varies with emphasis (not much, the equation might get simpler), or apply a SNDR limit that accounts for the way sigma_e varies with emphasis:
 $SNDR_{0+20\log_{10}(P_{max_equalized}/P_{max_unequalized})}$

Response Response Status U

REJECT.
 This is an extension of comment r02-42, which was rejected after review of presentation: http://www.ieee802.org/3/bs/public/17_07/daw_3bs_04_0717.pdf at the July meeting, with this justification:

Changing the SNDR limit to 28.5 dB is considered to be placing too great a burden on the receiver and it has not been demonstrated that implementations cannot meet the current specification.

Noise is treated in the COM calculation as independent of the Tx equalization, just as in this test.
 There was no consensus to apply either change in the suggested remedy.

Cl 120D SC 120D.3.1.1 P 353 L 26 # r03-31
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Transmitter output residual ISI SNR_ISI (min) 34.8 dB is still too high see daw_3bs_04_0717 and daw_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.
 D3.1 comments 22 and 36, D3.2 comment 43

SuggestedRemedy

In 120D.3.1.7, change "The SNR_ISI specification shall be met for all transmit equalization settings" to "The SNR_ISI is measured with Local_eq_cm1 and Local_eq_c1 set to zero".

Response Response Status U

REJECT.
 Re-statement of comment r02-43 which was rejected with the response: "No remedy provided."
 A remedy is now provided, however there was no consensus for the suggested remedy to be adopted since it is not expected that SNR_ISI will change significantly with transmit equalization setting and poor SNR_ISI with transmit equalization turned on would cause poor performance.

Cl 120D SC 120D.3.1.1 P 353 L 36 # r03-32
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

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-MR at low (and high) frequency (although apparently looser between 4 and 9 GHz). Also it is tighter at low frequencies than the new channel return loss limit, which seems wrong.
 Following D3.1 comment 41, D3.2 r02-44

SuggestedRemedy

Particularly now we have a channel return loss limit, we can change 14.25 - f to 12 -0.625f

Response Response Status U

REJECT.
 Re-statement of comment r02-44 which was rejected with the response: "While additional work has been done on this topic, there is still no consensus to make a change."

There is still no consensus to make the suggested change since the effect that this relaxation would have on system performance due to the interaction between the channel and the Tx and Rx devices has not been shown.

IEEE P802.3bs D3.4 200 Gb/s & 400 Gb/s Ethernet 4th Sponsor recirculation ballot comments

Cl 120D SC 120D.3.2 P 359 L 36 # r03-34
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Changing the return loss spec for the receiver was a mistake, because the effects of receiver reflections to a nominal-impedance channel and transmitter are in the receiver interference tolerance test, and the extra reflections to a channel and transmitter with different impedances are controlled/accounted for by the channel COM, now based on nominal impedances, the new channel return loss spec and the transmitter return loss spec. From the simple formula for reflection at an impedance mismatch, one can see that these effects are close to additive, so controlling/accounting for them separately is OK. In other words, the receiver pays for its own reflections in the interference tolerance test, so we don't have to tell the receiver designer how to do his job in this regard.

SuggestedRemedy

Revert 120D.3.1.1, Equation (120D-2) to 93.8.1.4, Equation (93-3).

Response Response Status U

REJECT.
 The change in definition of receiver return loss was the direct result of the resolution of comment r02-60. There was consensus for this change.

The commenter made a revised proposal in regard of this comment as shown in http://www.ieee802.org/3/bs/public/17_09/dawe_3bs_02a_0917.pdf

There was no consensus to make the suggested change in this presentation since the effect that this relaxation would have on system performance due to the interaction between the channel and the Rx device has not been shown.

Cl 120E SC 120E.3.1 P 372 L 20 # r03-40
 Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

The host is allowed to output a signal with 900 mV peak-to-peak amplitude but only 32 mV eye height - a very bad signal. If the module is exactly like the reference receiver, that would work, but with a good but slightly different receiver the eye will collapse with not enough margin for e.g. temperature changes causing mistuning. The module can't inconvenience the host in the same way because its peak-to-peak output voltage is measured before most of the loss. D3.0 comment 119, D3.2 r02-46.

SuggestedRemedy

Add a vertical eye closure spec to protect the module from such unexpected signals. VEC defined as largest of three ratios for the three sub-eyes, limit in the low teens of dB.

Response Response Status U

REJECT.
 Re-statement of comment r02-46 which was rejected with the response:
 "No presentation providing a suggested remedy for this comment was submitted. While a vertical eye closure specification was considered worth further investigation, no consensus was reached to make a change to the draft."

No consensus was reached for the suggested change as there is evidence that signals with large amplitude and small eyes will be seen in practice and evidence for what the limiting ratio for these should be has not been provided.

IEEE P802.3bs D3.4 200 Gb/s & 400 Gb/s Ethernet 4th Sponsor recirculation ballot comments

CI 120E SC 120E.3.1 P 371 L 20 # r04-12
Dawe, Piers J G Mellanox Technologie

Comment Type TR Comment Status R

Following up on previous comments: The host is allowed to output a signal with 900 mV peak-to-peak amplitude but only 32 mV eye height - a very bad signal. If the module is exactly like the reference receiver, that would work, but with a good but slightly different receiver the eye will collapse with not enough margin for e.g. temperature changes causing mistuning. The module can't inconvenience the host in the same way because its peak-to-peak output voltage is measured before most of the loss.
D3.0 comment 119, D3.2 r02-46, D3.3 r03-40.

SuggestedRemedy

Add a vertical eye closure spec to protect the module from such unexpected signals. VEC defined as largest of three ratios for the three sub-eyes. A reference bad signal (the module stressed input signal) could have VEC ~8 dB, a very bad low loss host to the D3.4 spec could have 16 dB, so set a limit e.g. max 12 dB. See presentation.

Response Response Status U

REJECT.

The associated presentation:

http://www.ieee802.org/3/bs/public/adhoc/elect/05Oct_17/dawe_01b_100517_elect.pdf
was discussed at the IEEE P802.3bs Electrical Interface Ad Hoc call on 5 October 2017.

There is no agreement that this issue will be seen in practical systems and there has been no validation that the proposed VEC limit of 12 dB would solve the problem. Also, there may be unforeseen consequences for introducing this limit. Consequently, there was no consensus to make this change to the draft.