Cl **162** *SC* **162.9.3.3** *P* **170** *L* **31** # I-53

Ran, Adee Cisco Systems, Inc.

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

The definition of SNDR refers back to 120D which does not state what the Tx equalization should be in this measurement. Based on a previous specification in clause 92, it may be understood that the limit in Table 162–10 applies to any valid equalization setting.

Since transmitters typically have noise sources that are independent of equalization, and applying equalization reduces the pulse peak, it is expected that increasing the "strength" of Tx equalization would degrade the measured SNDR. We can assume equalization settings with c(0) close to 0.5, which would reduce the measured pulse peak by 5-6 dB; this makes the SNDR spec more difficult than it seems.

A related concern is that the noise injected in the receiver ITT is also after Tx equalization (like realistic transmitters), and it is calibrated by measuring SNDR and using the results as TX_SNR. However, TX_SNR in COM represents a white noise source _before_ the Tx equalization, since it should have the same spectrum as the victim signal.

There seems to be a mismatch between the effect of TX_SNR in COM and the effect of SNDR in real links.

This may also affect SNDR and/or SNR_TX in clause 163 and annex 120F, although the receiver test signal is calibrated differently.

SuggestedRemedy

The definition of SNDR and/or the calculation of the effect of SNR_Tx in COM may need to be changed.

A detailed presentation is planned.

Response *Response Status* **U**

REJECT.

The following presentations were reviewed by the task force:
https://www.ieee802.org/3/ck/public/22_01/ran_3ck_01_0122.pdf
https://www.ieee802.org/3/ck/public/22_01/ran_3ck_03a_0122.pdf

Per straw poll #24, there is no consensus to make the proposed changes. Further analysis and consensus building is required.

Straw poll #24 (direction)
I support adopting SNDR and sigma_tx calculation as proposed on slide 3 of ran_3ck_03a_0122.
Yes: 7
No: 2
Need more information: 20

Cl **162** *SC* **162.11.5** *P* **184** *L* **33** # I-57

Ran, Adee Cisco Systems, Inc.

Comment Type **TR** *Comment Status* **R** *CA ILcd*

Equation 162-19 lets the difference between ILcd and ILdd be 10 dB up to half of (an old Nyquist frequency) and then linearly lower at higher frequencies. This does not make sense physically, and open the door to poor cables. The Tx output common mode noise problem is exacerbated by strong conversion from common mode to differential signal.

Note that COM does not cover the conversion loss term, so we should strive to make it negligible, rather than allowing it to be large.

At low frequencies we expect low ILdd and high ILcd, and the difference is much larger than 10 dB. Even at high frequencies up to 40 GHz, channels submitted to 802.3ck do not exceed 10 dB. We should not allow less than 10 dB difference across the upper half of the spectrum.

Based on samples of submitted channels and some measured channels it is suggested to tighten this specification to be 24 dB at the lowest frequency, linear slope to 10 dB at Nyquist/2, and constant 10 dB at maximum frequency.

This also holds for the specification in clause 163 (channel construction may be different but the arguments above still hold and the effect on the link budget is the same).

A presentation of some contributed data compared to the proposed limit is planned. Any contradictory data would be welcome.

SuggestedRemedy

Change equation 162-19 limit to be
24 - 13.56/f *14 | 0.05 <= f <= 13.56
10 | 13.56 <= f <= 40

Change Figure 162–9 accordingly.

Response *Response Status* **U**

REJECT.

Commenter has requested to update suggested remedy to:
Change equation 162-19 limit to be
30 – 8f |0.05 </= f </= 2.5
10 |2.5 </= f </= 25
10 – (f-25)/3 |25 </= f </= 40

This proposed responses is shown plotted along with the current limit line and responses of posted channels on slide 28 of the following presentation:
https://www.ieee802.org/3/ck/public/22_01/heck_3ck_01a_0122.pdf

There is interest in aligning the limit line with recently adopted test methodology for TX common-mode AC noise. However, a complete proposal with consensus is required.

There is no consensus to adopt the proposed changes at this time.

Cl **120G** *SC* **120G.3.1** *P* **258** *L* **21** # I-107

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *HO eye width*

ESMW/EW were removed in draft 1.4 with the introduction of the +/- 50 mUI rectangular window with VEO and VEC limits not passing the task force introduced Gaussian window which in effect reduces implicit minimum receiver eye opening. With current Gaussian window for typical high loss channel EW can be as little as 120 mUI, in comparisons CL120E min ESMW=220 mU. The 120 mUI can be further degraded for lower loss channel with pathological reflections/jitter may result in EW <100 mUI. Eye width opening is as critical as VEC/VEO, without explicit EW specifications and with current Gaussian window there is significant interoperability risk.

SuggestedRemedy

An explicit ESMW>=175 mUI specifications which is available in the scope might be the simplest, other alternative would be to go back to rectangular mask with +/- 50 mUI or introduce 10 sides mask as demonstrated in https://www.ieee802.org/3/ck/public/21_01/dawe_3ck_01_0121.pdf

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed changes.

For details, see the reponse to comment i-211.

Cl **120G** *SC* **120G.3.2** *P* **261** *L* **12** # I-108

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *MO eye width*

ESMW/EW were removed in draft 1.4 with the introduction of the +/- 50 mUI rectangular window with VEO and VEC limits not passing the task force introduced Gaussian window which in effect reduces implicit minimum receiver eye opening. With current Gaussian window for typical high loss channel EW can be as little as 120 mUI, in comparisons CL120E min farend ESMW=200 mU. The 120 mUI can be further degraded for lower loss channel with pathological reflections/jitter may result in EW <100 mUI. Eye width opening is as critical as VEC/VEO, without explicit EW specifications and with current Gaussian window there is significant interoperability risk.

SuggestedRemedy

An explicit ESMW>=150 mUI specifications which is available in the scope might be the simplest, other alternative would be to go back to rectangular mask with +/- 50 mUI or introduce 10 sides mask as demonstrated in https://www.ieee802.org/3/ck/public/21_01/dawe_3ck_01_0121.pdf

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed changes.

For details, see the reponse to comment i-211.

Cl **120G** *SC* **120G.3.2.2** *P* **262** *L* **27** # I-109

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *Test configuration*

Fig 120G-7 shows the most trivial component in the the capacitors, why not other components such as CDR, TX/RX optics?

SuggestedRemedy

Suggest removing what is inside the module just show a box for module under test

Response *Response Status* **U**

REJECT.
The capacitors are intended to show that the module input and output are AC-coupled.

Cl **120G** *SC* **120G.3.3.5.2** *P* **267** *L* **39** # I-115

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *HI eye width*

ESMW/EW were removed in draft 1.4 with the introduction of the +/- 50 mUI rectangular window with VEO and VEC limits not passing the task force introduced Gaussian window which in effect reduces implicit minimum receiver eye opening. With current Gaussian window for typical high loss channel EW can be as little as 120 mUI, in comparisons CL120E min farend ESMW=200 mU. The 120 mUI can be further degraded for lower loss channel with pathological reflections/jitter may result in EW <100 mUI. Eye width opening is as critical as VEC/VEO, without explicit EW specifications and with current Gaussian window there is significant interoperability risk.

SuggestedRemedy

An explicit ESMW>=150 mUI specifications which is available in the scope might be the simplest, other alternative would be to go back to rectangular mask with +/- 50 mUI or introduce 10 sides mask as demonstrated in https://www.ieee802.org/3/ck/public/21_01/dawe_3ck_01_0121.pdf

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed changes.

For details, see the reponse to comment i-211.

Cl **120G** *SC* **120G.3.4** *P* **269** *L* **19** # I-116

Ghiasi, Ali — Ghiasi Quantum LLC,Marvell Semiconductor, Inc.

Comment Type **TR** *Comment Status* **R** *MI eye width*

ESMW/EW were removed in draft 1.4 with the introduction of the +/- 50 mUI rectangular window with VEO and VEC limits not passing the task force introduced Gaussian window which in effect reduces implicit minimum receiver eye opening. With current Gaussian window for typical high loss channel EW can be as little as 120 mUI, in comparisons CL120E min ESMW=220 mU. The 120 mUI can be further degraded for lower loss channel with pathological reflections/jitter may result in EW <100 mUI. Eye width opening is as critical as VEC/VEO, without explicit EW specifications and with current Gaussian window there is significant interoperability risk.

SuggestedRemedy

An explicit ESMW>=175 mUI specifications which is available in the scope might be the simplest, other alternative would be to go back to rectangular mask with +/- 50 mUI or introduce 10 sides mask as demonstrated in https://www.ieee802.org/3/ck/public/21_01/dawe_3ck_01_0121.pdf

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed changes.

For details, see the reponse to comment i-211.

Cl **162** *SC* **162.9.3** *P* **166** *L* **32** # I-170

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *CR loss budget*

The draft CR loss budget wastes 3 dB in nearly every case. The relative range of host losses, 6.875/2.3 = 3:1, is too small for switch layout yet not needed for NICs.
The recommendation for the host traces plus BGA footprint and host connector footprint, 6.875 dB, compares very poorly with C2M's host insertion loss up to 11.9 dB, making passive copper to this draft expensive and unattractive for a switch, yet a full range of NICs can be made with only 3.75 dB.
C2M already has short and long ports.
Server-switch links are asymmetric in form factor (e.g. QSFP-DD to 2 x QSFP) and will get made with an asymmetric loss budget, so it would be better for the standard to regularise what will happen anyway with industry-standard registers.
This change would also benefit CR switch-switch links because the low loss of the shortest ports would be recognised, so more of the ports in a switch (with higher loss) could be used for CR switch-switch links.
The symmetric budget is used for some designs under way and may be useful in future for LOM, so it is kept here as "B", and the better way (A and C) added.

SuggestedRemedy

As in dawe_3ck_01a_0721.pdf:
3 classes of CR ports, host loss allocations of A 9.5, B 6.875, C 3.75 dB. B is as D2.1.
A connects to C, B to B or C, C to A, B or C.
Use 2 bits in the training control field to advertise A, B or C to the other end.
In Table 162-10, add limits A and C for linear fit pulse peak ratio (min). Change text in 162.9.3.1.2 to refer to the table.
In Table 162-14, add columns for Test 2 (high loss), A and C, with test channel insertion loss: A: 6.875-3.75 = 3.125 dB lower (20.5 dB to 21.5 dB), and C: 9.5-6.875 = 2.625 dB higher (26.25 dB to 27.25 dB). No change needed for Test 1.
In 162A.4, add equations for IL_PCBmax and ILHostMax A and B and show them in Fig 162A-1 and 2. In 162A.5, add Value columns A, C in Table 162A-1 (ILChmin and ILMaxHost differ). Adjust figures 162A-3 and 4.
Add MDIO registers to report local and remote host ability to station management, for inventory and diagnostics.

Response *Response Status* **U**

REJECT.

Per straw poll # 14 there is no consensus to make the proposed change.

Straw poll #6 (direction) -- taken on 2022/1/25
I support P802.3ck specifying multiple CR host types as proposed in comment i-170.
A. Yes
B. No
C. Abstain
Results: A: 11 B: 11 C: 5

Straw poll #14 (decision) -- taken on 2022/2/16
I support P802.3ck specifying multiple CR host types as proposed in comment i-170.
A. Yes
B. No
C. Abstain
Results: A: 8 B: 31 C: 2

Cl **162** *SC* **162.9.3.6** *P* **172** *L* **27** # I-178

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *TX RLcc*

As for the mated test fixtures and the cable, this common mode return loss spec RLcc becomes useless at the frequency when the MCB loss is 2/2 dB, which is only 10 GHz. The spec should trend down with the MCB trace loss at 0.1 dB/GHz.

SuggestedRemedy

Use a frequency-dependent mask 2 dB 0.2 <= f <= 4, 1.6+0.1*f dB 4 < f <= 30, 8.5-0.13f 30 < f <= 40. f is in GHz. See another comment for cable RLcc, 162.11.6.

Response *Response Status* **U**

REJECT.

Per straw poll #21, there is not consensus to make the proposed change.

Straw poll #21 (decision)
I support changing the CR TX RLcc as proposed in the suggested remedy in comment i-178.
Yes: 9
No: 10

Cl **162** *SC* **162.11** *P* **181** *L* **31** # I-180

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *CR loss budget*

The poor max cable loss makes CR unattractive, while all NICs and some ports on any switch have host loss budget going to waste. Enabling longer cables on a minority of links is needed.
In the remedy, each host knows the other host's loss class through the training protocol and the cable's loss class from its I2C compliance code, so no extra management features needed in the spec for the long cable class.

SuggestedRemedy

2 classes of cable, which could be called "short" (19.75 dB, as today) and "long", 19.75+2*(6.875-3.75) - 0.5 = 19.75+6.25 - 0.5 = 25.5 dB max (achievable cable length 3 m). Long cables connect port types C (see another comment) at both ends, short cables connect a valid combination of A, B, C.
In 162.11.2, cable assembly insertion loss, change text "less than or equal to 19.75 dB" to refer to Table 162-17 instead.
In 162.11.7.1.1, add zp = 30.7 mm for the "short" cable.
In Table 162A-1, add a column for the A-short-A scenario (ILCamax is 25.5 dB).
Illustrate in figures 162A-3 and 162A-4.

Response *Response Status* **U**

REJECT.

The suggested remedy is predicated on the adoption of comment i-170.

Resolve using the response to comment i-170.

Cl **162** *SC* **162.11.6** *P* **185** *L* **28** # I-181

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *CA RLcc*

We need a common mode return loss spec RLcc to stop large common-mode voltages building up through multiple low-loss reflections. As we know, this common mode return loss spec RLcc becomes useless at the frequency when the MCB loss is 1.8/2 dB, which is only 8.5 GHz. The impedance the cable presents is mostly related to the connector, so it's much like the mated test fixtures' RLcc, except at the very lowest frequencies where the cable loss is very small and both connectors can be seen by the measurement. This proposal allows for that.

SuggestedRemedy

Use a frequency-dependent mask 1.2 dB 0.05 <= f <= 4, 0.76+0.11*f dB 4 < f <= 30 GHz. f is in GHz. See another comment for Tx, Table 162-11, 162.9.3.6.

Response *Response Status* **U**

REJECT.

Per straw poll #22, there is not consensus to make the proposed change.

Straw poll #22 (decision)
I support changing the CA RLcc as proposed in the suggested remedy in comment i-181.
Yes: 10
No: 10

Cl **120G** *SC* **120G.3.2** *P* **261** *L* **11** # I-187

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MO EH*

On one hand: the eye height measurement method is very inaccurate, host receivers that implement CR can cope with much smaller eye height than this, VEC is much more important. Receiver noise is already in the measurement, C2M drivers are traditionally 900/1200 as strong as CR/KR drivers, and the end-to-end loss is lower by a much larger ratio. So a small EH is acceptable.
On the other hand: if the eye height limit is the same at near end as at far end, there is huge margin at near end and the implementer can optimise beyond far end, only limited by the NE VEC spec, while we want modules to be set up consistently, for the full range from near to far. NE and FE EH naturally differ, and the spec should reflect that. Also, host designers know their own loss and low-loss hosts (NICs) can take advantage of a naturally larger signal that cost the module nothing. This applies to both the short and long modes.

SuggestedRemedy

Change the far end eye height so that it is 2 dB below near end: if near can remain at 15 mV, far becomes 12 mV. Far end remains the one with less margin, that the implementer should tune the module for.

Response *Response Status* **U**

REJECT.

The comment makes reference to the capabilities of a CR SERDES. Annex 120G is specifying C2M recievers and transmitters. Although it is true that the host might have a CR-capable SERDES that may not be universally the case. Note that there are different host channel budgets for CR and C2M.

The comment does not provide sufficient justification for the proposed changes. Analysis is required to demonstrate the need.

There is no consensus to make the proposed changes.

Cl **120G** *SC* **120G.3.2** *P* **261** *L* **11** # I-188

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MO EH/VEC*

The module output eye height and VEC have to comply at both near end and far end, and depending on the cleanliness of its signal, a module can be tuned to either end or somewhere in the middle, or even somewhere outside the range. The host stressed input signal is tuned to far end, only, so the host isn't required to receive those other tuning choices. This is inconsistent and a serious flaw in the spec. Yet we would rather not have multiple host stress tests, nor require the host to receive unnecessary and sub-optimal signal tunings, so we need to make sure that modules are tuned correctly.

SuggestedRemedy

Tighten the equaliser limits for module output so that modules are tuned consistently across the industry. Because the channel losses in short and long mode testing are significantly different, in Table 20G-11 use separate gDC limits for short and long mode (see other comments). To discourage module implementers from mis-tuning modules so they are optimised significantly beyond the far end, in Table 120G-3, ensure that each near end VEC is 0.5 dB less (better) than its corresponding far end VEC, and the far end EHs are 2 dB less than the corresponding near end EHs. Note other comments that address what these values should be.

Response *Response Status* **U**

REJECT.

The comment provides insufficient evidence evidence that the proposed changes are necessary or improve the interoperability.

Cl **120G** *SC* **120G.3.2.2.1** *P* **263** *L* **14** # I-189

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MO test channel*

If we include an allowance for host transmitter package loss for the host stressed input test, it would make sense to include the same allowance for far-end module output specs. As the change is to the reference host channel which is in software, it's convenient to do, rather than rely on extrapolation.

SuggestedRemedy

Increase the two far-end lengths by 2.2 dB (taking 16 dB to 18.2 dB, aligning with 120G.3.4.3.2). In Table 120G-11, increase bbmax(1) from 0.4 to 0.55. Reduce module output eye height by 2.2 dB.

Response *Response Status* **U**

REJECT.

The total host side insertion loss prescribed is 9.6 dB for the synthetic transmission line and 2.3 dB for the module compliance board for a total of 11.9 dB, which matches with the maximum host insertion loss recommendation in Figure 120G-2.

The comment proposes that the module output should be measured with the maximum host insertion loss plus an allocation similar to that used in the frequency-dependant attenuator in 120G.3.4.3.2 then scale the eye height proportionally and increase the DFE tap range.

The reasoning for making the changes seems sounds, but insufficient analysis has been provided to show that the changes to the DFE tap range and the eye height value are appropriate.

There is some interest in increasing the channel loss as proposed, but there is insufficient analysis provided to support the proposed new values for bbmax and eye height. Further analysis and consensus is encouraged.

Cl **120G** *SC* **120G.3.3.5.1** *P* **265** *L* **50** # I-193

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *HI SIT PG*

For module output, the optimum for postcursor with a module package and channel is zero. We want consistent stressed signals across the industry so we should give guidance where we can.
The same point applies to module stressed input signal generator, but 120G.3.4.3.1 refers back to here.

SuggestedRemedy

Say that in practice, a postcursor may be used to make the PG output like that from a module with zero postcursor. Modify "The tap coefficients are not specified with the exception that".

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient evidence to support proposed changes. The suggested remedy does not provide sufficient detail to implement.

Cl **120G** *SC* **120G.3.4.3.2** *P* **271** *L* **33** # I-204

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *HI SIT calibration*

We have a gDC + gDC2 max limit for the high loss module stressed input case to ensure that the module can equalise a very slow signal. Likewise, there should be max/min limits for gDC + gDC2 for the low loss case to set the contract for faster signals. In Table 120G-11, gDC+gDC2 can be -2 for TP1a and -1 for TP4 near-end. dudek_3ck_01_0921 slide 5 indicates that a range of -3 to -1 dB would be suitable.

SuggestedRemedy

Add an exception that for the low-loss case, the reference receiver CTLE setting that minimizes VEC has gDC + gDC2 in the range -3 dB to -1 dB.
It may be preferable to put the exceptions for both low- and high-loss cases in Table 120G-10.

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient justification to make the proposed changes.

Further analysis to justify the suggested remedy is required.

Cl **120G** *SC* **120G.5.2** *P* **275** *L* **27** # I-206

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MO gDC values*

The limits for TP4 gDC, gDC2 should not be the same for short and long output modes.
The range of losses in a module is much less than the range of losses of the four reference host channels. So, obviously, different channels will need different CTLE settings.
Obviously, CTLE settings that represent signals outside what the spec makes a host capable of receiving in a particular mode, should be excluded, to make module implementers set up their product correctly.

SuggestedRemedy

Create separate limits for TP4 short and long output modes, so 4 sets for TP4+, in the style of TP1a. See other comments.

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient justification for the proposed changes nor does the suggested remedy provide sufficient detail to implement.

Cl **120G** *SC* **120G.5.2** *P* **275** *L* **28** # I-207

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MO gDC values*

The maximum gDC is -2 for TP1a and -1 for TP4 near-end. As the MCB loss and HCB loss are within 0.2 dB of each other, these specs are inconsistent by 0.8 dB.
dudek_3ck_01_0921 slide 5 shows that -1 is reasonable for a 12 mm package trace, and shorter traces are possible, e.g. an on-board repeater. Hosts and modules with less loss than the MCB and HCB respectively may have to receive a signal less filtered at the point of use than in the module or host output measurement.
ghiasi_3ck_adhoc_01a_042121 slide 9 says that -1 is needed for 5 dB ball to ball, 1.6 dB less than the mated compliance boards' loss.
On the other hand, things go bad rapidly with too much emphasis. It would be safer to set both at -2, which would require retuning the short setting in ghiasi_3ck_adhoc_01a_042121 with reduced output emphasis - which should be OK.
See other comments that give specific ranges for the stressed signals to ensure that inputs are tested with representative low-loss signals.

SuggestedRemedy

For TP4 gDC, change -1 to -2.

Response *Response Status* **U**

REJECT.

The comment does not provide sufficient justification for the proposed changes.

Analysis is required to determine the need and impact of the proposed change.

Cl **120G** *SC* **120G.5.2** *P* **275** *L* **34** # I-208

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MI gDC values*

The weakest (max, least -ve) gDC + gDC2 is -2 for TP1a, -2 for TP4 near end, -3 for TP4 far end and -10.5 for module stressed input high loss. There is about 10 dB loss difference between short near end and long far end, but 1 dB difference in max gDC + gDC2 which is far too little. It looks like TP4 far end (-9 to -2 in the draft) is out of step, with a much wider range than TP4 near end. TP4 LONG far end should never use this wide range as most of the channel loss is fixed. We should not be encouraging modules to try to do a job the host receiver does better, and we want modules to be set up consistently so that the short/long mode choice means something.
Also, if we include an allowance for host transmitter package loss for the host stressed input test, it would make sense to include the same allowance for far-end module output specs.

SuggestedRemedy

Impose a max gDC + gDC2 limit of -5 for TP4 long far end, e.g. with gDC, gDC2 ranges in the same style as TP1a:
Range for gDC2 = 0 -9 to -5
Range for -1 <= gDC2 < 0 -9 to -4
Range for -2 <= gDC2 < -1 -9 to -3
Range for -3 <= gDC2 < -2 -9 to -2

Response *Response Status* **U**

REJECT.

There is some agreement with the direction of the proposal but further analysis is required to determine appropriate values.

Cl **120G** *SC* **120G.5.2** *P* **275** *L* **34** # I-209

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *MO gDC values*

As a most of the channel for TP4 far-end is known exactly and the max loss to TP4 far end is less than to TP1a, the range of gDC, gDC2 combinations should be a subset of the TP1a ones.

SuggestedRemedy

For continuous time filter, DC gain for TP4 short far-end (gDC), change to sets of limits that depend on gDC2 in the same style as for TP1a. The allowed values should be subsets of those for TP1a.
See another comment for TP4 long far end.
For TP4 short far end, change from -9 to -2, to:
Range for gDC2 = 0 -7 to -3
Range for -1 <= gDC2 < 0 -7 to -2
Range for -2 <= gDC2 < -1 -7 to -2
Range for -3 <= gDC2 < -2 -7 to -2

Response *Response Status* **U**

REJECT.

There is some agreement with the direction of the proposal but further analysis is required to determine appropriate values.

Cl **120G** *SC* **120G.5.2** *P* **277** *L* **6** # I-211

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *EH/VEC method mask*

This draft has a (de-)weighted rectangular eye mask spec with mask height = max(EHmin, EA/VECmax) and effective mask width ~2x0.03 to 2x0.035 UI, although it is described as a histogram 2x0.05 UI wide. This is too narrow; compare 120E with ESMW of 0.2 or 0.22 UI. It's half as wide as TDECQ with histograms extending to +/-0.07 UI.

This de-weighted histogram might have worked if there had been a guarantee that no host or module would ever produce a fast, highly jittered eye, but we don't have that guarantee. Work needs to be done to repair the hole in the spec.

See healey_3ck_01a_1020 slide 6, orange dots for +/-0.025 UI which is the closest to the current draft. For VEC of 10 dB, EW can be anywhere in the range 160 to 290 mUI: an almost 2:1 range. Driver risetime is not reported; if it is always the COM default slowest-reasonable 7.5 ps, then even worse EW is possible with faster or peaked drivers. This is too much worse than 120E. As the plot shows, a wide range of eye widths are possible, so we don't need to allow the worst ones by an oversight.

De-weighting the sides of the histogram with flat top and bottom, rather than chamfering the corners, means that infringing the corners by a mile is counted the same as infringing by an inch, which is bad.
Most of the weight of samples is in the middle of the eye which is a waste of measurement time; we know the corners will fail first so we should measure them, not the middle Hence the 2-offsets approach of TDEC and healey_3ck_01a_1020.
The effective BER criterion of the (de-)weighted mask seems to be around 1e-4, not 1e-5 as before.

The distribution of repeated measurements is very skewed.

We need an eye mask that's more eye shaped, so that a higher proportion of the samples near the boundary are measured at full weight and contribute properly to the measurement. Eye mask measurement with a 10-sided mask has been pre-programmed into scopes for about 20 years, we should use established tools and methods where they work well.

The 10-sided mask controls the eye on the diagonal more strongly than the rectangular uniform histogram/mask because hits are collected over the time of the chamfer, rather than just in corners. The de-weighted rectangular histogram controls the eye on the diagonal more weakly than the rectangular uniform histogram/mask because hits are collected just in corners, and de-weighted.

SuggestedRemedy

Change from a 4-cornered weighted mask with corners at t = ts+/-0.05, V = y +/-H/2 to a 10-cornered unweighted mask with corners at t = ts+/-1/16, ts+/-0.05, ts+/-3/32, V = y +/-H/2, y +/-H*0.4, y. y is near VCmid, VCupp or VClow (vertically floating, as in D3.0).
H is max(EHmin, Eye Amplitude * 10^(-VECmax/20)). Eye Amplitude is AVupp, AVmid or AVlow, as today.

This simple scalable method gives VEC results 0.5 to 1 dB more optimistic than the unweighted rectangular mask. It can remain as the EH and VEC limits are revised in the light of experience.

Response *Response Status* **U**

REJECT.

Straw polls #8 and #9 indicate strong consensus to continue with a weighted window approach. Straw polls #10 and #11 indicate strong consensus to continue with the currently specified weighting function.

There is no consensus to make the proposed changes to the draft.

Straw poll #8 (chicago rules)
Straw poll #9 (choose one)
I support the following direction of the eye opening specification method:
A. weighted window per Draft 3.0 (as is or with some improvements)
B. revert to uniform weighted window per D2.1 (D3.0 comment #212)
C. 10pt mask per D3.0 comment #211
#8 A: 31 B: 12 C: 6
#9 A: 27 B: 5 C: 1

Note: Straw poll #8 and #9 are the same question and answers except #8 is chicago rules (pick any) and #9 is choose one.

Straw poll #10 (chicago rules)
Straw poll #11 (choose one)
To address eye width issues expressed, I support the following method to modify the weighted window:
A. no change
B. “wider” weighting mask (e.g., larger sigma, alternate distribution shape)
C. add jitter specification
D. add eye width specification (i.e., per D3.0 comments 107, 108, 115, 116)
#10 A: 26 B: 15 C: 9 D:9
#11 A: 19 B: 5 C: 3 D: 4

Note: Straw poll #10 and #11 are the same question and answers except #10 is chicago rules (pick any) and #11 is choose one.

Cl **120G** *SC* **120G.5.2** *P* **277** *L* **6** # I-212

Dawe, Piers J G NVIDIA

Comment Type **TR** *Comment Status* **R** *EH/VEC method mask*

The Gaussian weighting has the effect of destroying the histogram width, allowing bad fast eyes to pass, while failing less bad slow eyes. It gives the false impression that the histogram width still applies. With a weighting standard deviation of 0.02 UI, the eye height is measured at around +/-0.035 UI rather than the +/-0.05 UI with the unweighted histogram - depending on eye shape. Compare 120E with ESMW of 0.2 or 0.22 UI, and TDECQ with histograms extending twice as wide, to +/-0.07 UI.
This weighting is equivalent to relaxing the VEC spec by 1.5 to 2 dB - but it depends on the eye shape, it weakens the spec most for the worst-shaped eyes, which is bad. It applies a worse BER criterion than the 1e-5 intended.

SuggestedRemedy

Remove the Gaussian weighting and set the eye height and VEC limits (which need revision anyway) appropriately. ghiasi_3ck_01_0721, which was not given the presentation time it deserved, says that the minimum eye height in particular needs to be reduced for TP1 and TP4 far end.

Response *Response Status* **U**

REJECT.

There is no consensus to make the proposed changes.

For details, see the reponse to comment i-211.