Cl **0** *SC* **0** *P* *L* # r04-1

Berger, Catherine

Comment Type **G** *Comment Status* **A** *Bucket*

This draft meets all editorial requirements.

SuggestedRemedy

Response *Response Status* **C**

ACCEPT.

Cl **120D** *SC* **120D.3.1.1** *P* **352** *L* **13** # r04-10

Hidaka, Yasuo Fujitsu Laboratories of

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

The value of AC common-mode output voltage is inconsistent between Table 120D-1 (30mV) and PICS item TC5 (12mV). They must be consistent.
In the other clauses, 30mV was used for cable PHYs and 100GBASE-KP4, and 12mV was used for C2C and backplane PHYs except 100GBASE-KP4. Cable PHYs were specified as 30mV at TP2 to take account of the effects of the host board trace. 100GBASE-KP4 was specified as 30mV at TP0a, but 100GBASE-KP4 was not used in practice. In C2M spec, 17.5mV at TP1a or TP4 is used. It is also odd to have a higher value at TP0a in C2C than the value at TP1a in C2M spec. Since this C2C spec is at TP0a, I recommend to use 12mV as the consistent value, because PAM4 is more stringent than NRZ, and 30mV is the value at TP2 including the effect of host board trace.

The value in the other clauses:
100GBASE-CR4, Clause 92, Table 92-6: 30mV (TP2)
100GBASE-KR4, Clause 93, Table 93-4: 12mV (TP0a)
100GBASE-KP4, Clause 94, Table 94-13: 30mV (TP0a)
CAUI-4 (C2C), Annex 83D, Table 83D-1: 12mV (TP0a)
CAUI-4 (C2M), Annex 83E, Table 83E-1, 83E-3: 17.5mV (TP1a, TP4)
25GBASE-CR, Clause 110, Table 92-6: 30mV (TP2)
25GBASE-KR, Clause 111, Table 93-4: 12mV (TP0a)
200GAUI-4,400GAUI-8 (C2M), Annex 120E, Table 120E-1, 120E-3: 17.5mV (TP1a, TP4)

SuggestedRemedy

Change AC common-mode output voltage in Table 120D-1 from 30mV to 12mV.

Response *Response Status* **C**

ACCEPT IN PRINCIPLE.
This comment does not apply to the substantive changes between IEEE P802.3bs/D3.3 and IEEE P802.3bs/D3.4 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

However, the change to make the PICS match the normative value is an improvement to the draft that would otherwise need to be made in Maintenance.

No evidence has been provided that the normative value of 30 mV in Table 120D-1 is inadequate. In 120D.5.4.1, Item TC5, change "12 mV" to "30 mV".

Cl **120D** *SC* **120D.3.2.1** *P* **359** *L* **53** # r04-3

Hidaka, Yasuo — Fujitsu Laboratories of

Comment Type **T** *Comment Status* **R** *Bucket*

It says "higher amplitude values", but 93C.2 step 8 does not define higher amplitude values. The term "amplitude" is misleading, because it often implicates signal amplitude. We should be more specific which variable is referred here.

SuggestedRemedy

Change "If higher amplitude values are used"
to
"If a higher value of the channel noise voltage sigma_bnm is used".

Response *Response Status* **C**

REJECT.
In step 8 of 93C.2, the "channel noise voltage" is adjusted to a particular value. The text that the comment proposes to change is:
"The COM value is the target for the receiver noise level calibration defined in 93C.2 step 7. The channel noise voltage applied in 93C.2 step 8 should be as close as practical to the value needed to produce the target COM. If higher amplitude values are used, this would demonstrate margin to the specification but this is not required for compliance."
This text is not incorrect. A higher value for the channel noise voltage sigma_bnm would result in a higher amplitude value.

Cl **120D** *SC* **120D.5.4.2** *P* **366** *L* **26** # r04-11

Hidaka, Yasuo — Fujitsu Laboratories of

Comment Type **T** *Comment Status* **R** *Bucket*

An item is missing in PICS in 120D.5.4.2 corresponding to "shall" in 120D.3.2.3 Transmitter equalization feedback (optional), P360, L39.

SuggestedRemedy

Add an optional item to PICS table in 120D.5.3 Major capabilities/options as follows:

Item: *TEFB
Feature: Transmitter equalization feedback capability
Subclause: 120D.3.2.3
Value/Comment: Transmitter equalization feedback is an optional capability for a 200GAUI-4 or a 400GAUI-8 chip-to-chip receiver.
Status: O
Support: Yes [] No []

Add an item to PICS table in 120D.5.4.2 Receiver as follows:

Item: RC5
Feature: Transmitter equalization feedback
Subclause: 120D.3.2.3
Value/Comment: Operates as described in 120D.3.2.3.
Status: TEFB:M
Support: Yes [] N/A []

Response *Response Status* **C**

REJECT.
This comment does not apply to the substantive changes between IEEE P802.3bs/D3.3 and IEEE P802.3bs/D3.4 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

Since there is no requirement for every "shall" statement to be covered by a PICS item, this is a change that does not need to be made to the draft, but could be considered in maintenance.

Cl **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.

Cl **120E** *SC* **120E.3.1.7** *P* **374** *L* **5** # r04-5

Hidaka, Yasuo — Fujitsu Laboratories of

Comment Type **E** *Comment Status* **R** *Bucket*

In 2014 IEEE-SA Standards Style Manual, 13.3.2 Numerical values, it is written as "Digits should be separated into groups of three, counting from the decimal point towards the left and right. The groups should be separated by a space, and not a comma, period, or dash." and "In numbers of four digits, the space is not necessary, unless four-digit numbers are grouped in a column with number of five digits or more." The columns of G and Z1/2pi are not following this style. The column of P1/2pi is OK.

SuggestedRemedy

In columns of G and Z1/2pi, separate digits into groups of three by a space, counting from the decimal point toward the right.
For example, change "0.891251" to "0.891 251" in the cell of column G and row 1 dB peaking.

Response *Response Status* **C**

REJECT.
Despite the style manual stating that this "should" be done both to the left and right of the decimal point, the separation of digits into groups of three has only been done to the left of the decimal point in the 802.3 standard. A search of all 8 sections of the revision project Draft 2.0 found no occurrences at all of a decimal point followed by three digits, a space and another digit. Consequently, doing this for one table in this amendment would be likely to cause more confusion than clarity.

Cl **120E** *SC* **120E.3.1.7** *P* **374** *L* **14** # r04-4

Hidaka, Yasuo — Fujitsu Laboratories of

Comment Type **T** *Comment Status* **R** *Bucket*

In the row of 4dB peaking, the column of Z1/2pi has the value of "7.58765", that is one significant digit less than the other rows. The original proposal was "7.587650". Although the numerical values are same, it should have the same number of sigificiant digit as the other rows.

SuggestedRemedy

Change "7.58765" to "7.587650".

Response *Response Status* **C**

REJECT.
The use of a trailing zero in this way has meaning when the number represents a measurement result. For instance, a measurement result of 7.5 has different meaning from a value of 7.50 because the inclusion of the trailing zero is taken as indicating the precision to which the measurement has been made.

However, the numbers in Table 120E-2 are values to be inserted into Equation 120E-2. In this case trailing zeros have no meaning. For example, if the value was shown as 7.5, there is no implication that you can substitute any value between 7.45 and 7.55. Consequently, there is no difference between "7.58765" and "7.587650" being shown in this table except that showing the latter value tends to imply that there is in fact meaning to the trailing zero. This is captured in the base standard:
1.2.6 Accuracy and resolution of numerical quantities
Unless otherwise stated, numerical limits in this standard are to be taken as exact, with the number of significant digits and trailing zeros having no significance.

Cl **120E** *SC* **120E.3.4.1.1** *P* **382** *L* **15** # r04-6

Hidaka, Yasuo — Fujitsu Laboratories of

Comment Type **T** *Comment Status* **R** *Bucket*

It says "has to be greater than or equal to 7dB". IEEE-SA Standards Style Manual says that the use of the word "must" is deprecated and shall not be used when stating mandatory requirements. Although IEEE-SA Standards Style Manual does specifically state the use of the words "have to", we should deprecate it as well.

SuggestedRemedy

Change "has to be" to "should be" in line 15 and line 19.

Response *Response Status* **C**

REJECT.
The IEEE style manual does not disallow the use of "has to". There are 17 instances of the text "has to" in the base 802.3 standard. Changing this text to "should be" would soften the meaning so that it would no longer imply that this is a requirement.

Cl **120E** *SC* **120E.5.4.3** *P* **390** *L* **6** # r04-8

Hidaka, Yasuo — Fujitsu Laboratories of

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

An item is missing in PICS in 120E.5.4.3 corresponding to "shall" in 120E.3.3.2 Host stressed input test, P378, L7.

SuggestedRemedy

Add a row to PICS table in 120E.5.4.3 Host input as follows:

Item: RH2
Feature: Host stressed input test
Subclause: 120E.3.3.2
Value/Comment: Table 120E-5
Status: M
Support: Yes []

Proposed Response *Response Status* **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

This comment does not apply to the substantive changes between IEEE P802.3bs/D3.3 and IEEE P802.3bs/D3.4 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

This "shall" is already covered by PICS item RH1 in 120E.5.4.3. There is no requirement for there to be a one-to-one correspondence between PICS items and "shall" statements.

Cl **120E** *SC* **120E.5.4.4** *P* **390** *L* **14** # r04-9

Hidaka, Yasuo — Fujitsu Laboratories of

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

An item is missing in PICS in 120E.5.4.4 corresponding to "shall" in 120E.3.4.1 Module stressed input test, P380, L37.

SuggestedRemedy

Add a row to PICS table in 120E.5.4.4 Module input as follows:

Item: RM2
Feature: Module stressed input test
Subclause: 120E.3.4.1
Value/Comment: Table 120E-8
Status: M
Support: Yes []

Proposed Response *Response Status* **Z**

REJECT.

This comment was WITHDRAWN by the commenter.

This comment does not apply to the substantive changes between IEEE P802.3bs/D3.3 and IEEE P802.3bs/D3.4 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

This "shall" is already covered by PICS item RM1 in 120E.5.4.4. There is no requirement for there to be a one-to-one correspondence between PICS items and "shall" statements.

Cl **122** *SC* **122.6** *P* **248** *L* # r04-2

Burrell, Gary — Elenion Technologies

Comment Type **T** *Comment Status* **R**

IEEE 803.3.bs in Table 122-5 follows the CWDM wavelength grid defined in ITU-T G.694.2 and allows 13 nm wavelength variation within each channel. Nothing is said in the 802.3.bs about the characteristics of this variation, leading the reader to assume that all of 13 nm is random variation

As mentioned in ITU-T G.694.2 wavelength variations have two main components:
1. Laser variations around a nominal wavelength due to manufacture tolerances
2. Uncooled laser wavelength shift as a result of temperature changes

G.694.2 further states OSpecific values and allocations of this variation will be defined in individual applications.O however 802.3bs does not currently define this.

Therefore, it would be useful to specify what portion of the 13 nm is due to manufacturing random variation of the laser wavelength and which portion the variation can be expected to be similar for all of the four channel lasers.

Not specifying this could result in requiring a Reciever design that is more complex and expensive than otherwise required.

See presentation

SuggestedRemedy

Propose the random wavelength channel to channel variation shall be less than +/- 3nm and the wavelength variation due to temperature shifts shall be less than +/- 3.5 nm

This results in the same (total) Wavelength range as show in Table 122-5
No change in the Wavelength grid or (overall) tolerance from ITU G.694.2
ITU already specifies that this allocation will be defined in individual applications

Response *Response Status* **C**

REJECT.
This comment does not apply to the substantive changes between IEEE P802.3bs/D3.3 and IEEE P802.3bs/D3.4 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

The related presentation:
http://www.ieee802.org/3/bs/public/adhoc/smf/17_10_03/burrell_01_1017_smf.pdf was discussed at the IEEE P802.3bs SMF Ad Hoc call on 3 October 2017.

With the existing specification, the implementer has the possibility of trading a small wavelength range due to temperature (for instance because a cooler is used) against a wider wavelength range due to manufacturing variation. Applying the restriction proposed by this comment would remove this flexibility.

Cl **122** *SC* **122.7.1** *P* **251** *L* **48** # r04-7

Hidaka, Yasuo
Fujitsu Laboratories of

Comment Type **T** *Comment Status* **R** *Bucket*

In the footnote a of Table 122-10, it says "the total average launch power limit has to be met". IEEE-SA Standards Style Manual says that the use of the word "must" is deprecated and shall not be used when stating mandatory requirements. Although IEEE-SA Standards Style Manual does not state the use of the words "have to", we should deprecate it as well.

SuggestedRemedy

Change "As the total average launch power limit has to be met" to "As the total average launch power limit should be met".

Response *Response Status* **C**

REJECT.
This comment does not apply to the substantive changes between IEEE P802.3bs/D3.3 and IEEE P802.3bs/D3.4 or the unsatisfied negative comments from the previous ballots. Hence it is not within the scope of the recirculation ballot.

The IEEE style manual does not disallow the use of "has to". There are 17 instances of the text "has to" in the base 802.3 standard. Changing this text to "should be" would soften the meaning so that it would no longer imply that this is a requirement.