

IEEE P802.3cw D2.3 400 Gb/s over DWDM systems 3rd Working Group recirculation ballot comments

CI 00 SC 0 P 0 L 0 # 1

Brown, Matt Alphawave

Comment Type E Comment Status X

802.3cw is now preceded by 802.3df and will be amendment 10. 802.3df has been added to cover page (page 1) and the amendment lines (page 13) but references elsewhere have not been updated.

SuggestedRemedy

In clauses being amended by 802.3cw (1, 30, 45, 116, 118)...
Change any amendments to include references to 802.3df and changes made in 802.3df, as appropriate.
Implement with editorial license.

Proposed Response Response Status O

CI 155 SC 155.2.2 P 46 L 7 # 2

Brown, Matt Alphawave

Comment Type E Comment Status X

"When communicating" phrase is deceiving since it implies that sometimes it does not communicate with the other layer. I think the intent was to provide a reference to each of the two interfaces. Also, the PCS does not communicate *with* the 400GMII, it communicates *via* the 400GMII with the RS or PHY 400GXS above. Similar for communication with the PMA.

SuggestedRemedy

Change "When communicating with the 400GMII, the 400GBASE-ZR PCS uses an eight octet-wide, synchronous data path, with packet delineation being provided by transmit control signals (TXC) and receive control signals (RXC) (see 81.3). When communicating with the 400GBASE-ZR PMA in the transmit direction, the 400GBASE-ZR PCS provides codewords (see 155.3.2.1) of a systematic (128, 119) double-extended Hamming code (denoted SD-FEC within this clause) to the 400GBASE-ZR PMA. When communicating with the 400GBASE-ZR PMA in the receive direction, the 400GBASE-ZR PCS receives 128 x m bit SD-FEC codewords (see 155.3.2.2.1) from the 400GBASE-ZR PMA, where m is the implementation dependent sampling resolution of each component of the DP-16QAM symbol in bits."

To: "For communication via the 400GMII, the 400GBASE-ZR PCS uses an eight octet-wide, synchronous data path, with packet delineation being provided by transmit control signals (TXC) and receive control signals (RXC) (see 81.3). For communication with the 400GBASE-ZR PMA in the transmit direction, the 400GBASE-ZR PCS provides codewords (see 155.3.2.1) of a systematic (128, 119) double-extended Hamming code (denoted SD-FEC within this clause) to the 400GBASE-ZR PMA. For communication with the 400GBASE-ZR PMA in the receive direction, the 400GBASE-ZR PCS receives 128 x m bit SD-FEC codewords (see 155.3.2.2.1) from the 400GBASE-ZR PMA, where m is the implementation dependent sampling resolution of each component of the DP-16QAM symbol in bits."

Proposed Response Response Status O

CI 155 SC 155.2.5.3 P 48 L 13 # 3

Bruckman, Leon Huawei

Comment Type T Comment Status X

The text "Idle blocks are removed from the 257-bit encoded data at a rate of 163 832/163 840" is not clear

SuggestedRemedy

Change: "Idle blocks are removed from the 257-bit encoded data at a rate of 163 832/163 840" to: "Idle blocks are removed from the 257-bit encoded data to reduce the rate by 163 832/163 840"

Proposed Response Response Status O

IEEE P802.3cw D2.3 400 Gb/s over DWDM systems 3rd Working Group recirculation ballot comments

CI 155 SC 155.2.6.8 P 58 L 12 # 4

Bruckman, Leon Huawei

Comment Type T Comment Status X

The text "Idle blocks are added to the stream of 257-bit data blocks at a rate of 163 832 / 163 840." is not clear

SuggestedRemedy

Change: "Idle blocks are added to the stream of 257-bit data blocks at a rate of 163 832 / 163 840." to "Idle blocks are added to the stream of 257-bit data blocks to increase the rate by 163 832 / 163 840."

Proposed Response Response Status O

CI 155 SC 155.3.1.3 P 60 L 35 # 5

Bruckman, Leon Huawei

Comment Type E Comment Status X

Wrong line label in Figure 155-10. Twice PMD_IS_UNITDATA.request

SuggestedRemedy

Replace right arrow PMD_IS_UNITDATA.request with PMD_IS_UNITDATA.indication

Proposed Response Response Status O

CI 45 SC 45.2.1.6 P 24 L 38 # 6

Marris, Arthur Cadence Design Systems

Comment Type E Comment Status X

802.3df is also modifying bits 1.7.6:0

SuggestedRemedy

Add as modified by IEEE Std 802.3df-202x
add extra bit 7 to make it bits 1.7.7.0
Change to 0 1 1 1 1 1 1 = 400GBASE-ZR PMA/PMD

Proposed Response Response Status O

CI 45 SC 45.2.3 P 31 L 22 # 7

Marris, Arthur Cadence Design Systems

Comment Type E Comment Status X

Table 45-233—PCS registers has been modified by 802.3df

SuggestedRemedy

Add as modified by IEEE Std 802.3df-202x
Change 3.632 to 3.664

Proposed Response Response Status O

CI 117 SC 117.1 P 38 L 29 # 8

Marris, Arthur Cadence Design Systems

Comment Type T Comment Status X

Missing -R

SuggestedRemedy

Change 200GBASE to 200GBASE-R
Change 400GBASE to 400GBASE-R

Proposed Response Response Status O

CI 45 SC 45.2.1.6 P 24 L 27 # 9

Huber, Thomas Nokia

Comment Type E Comment Status X

Table 45-7 is modified by 802.3df. Since 802.3cw is now after 802.3df, the editing instruction should include 802.3df.

SuggestedRemedy

Change "as modified by IEEE Std 802.3db-2022" to "as modified by IEEE Std 802.3db-2022 and IEEE Std 802.3df-202x"

Proposed Response Response Status O

IEEE P802.3cw D2.3 400 Gb/s over DWDM systems 3rd Working Group recirculation ballot comments

CI 45 SC 45.2.1.6 P 24 L 36 # 10

Huber, Thomas

Nokia

Comment Type E Comment Status X

Since 802.3df also modifies Table 45-7, and 802.3cw is now after 802.3df, the changes need to be based on the table as it exists in 802.3df rather than in 802.3db.

SuggestedRemedy

In the table, change the value in the Bits column to 1.7.7:0. Change the Description column to show the value 0 1 1 1 1 1 1 1 = reserved being changed to 0 1 1 1 1 1 1 1 = 400GBASE-ZR PMA/PMD

Proposed Response Response Status O

CI 45 SC 45.2.1.22 P 26 L 3 # 11

Huber, Thomas

Nokia

Comment Type E Comment Status X

Since 802.3cw is now after 802.3df, the editing instruction should include 802.3df.

SuggestedRemedy

Change "as modified by IEEE Std 802.3db-2022" to "as modified by IEEE Std 802.3db-2022 and IEEE Std 802.3df-202x"

Proposed Response Response Status O

CI 155 SC 155.3.1 P 60 L 29 # 12

Zimmerman, George

CME Consulting/APL Gp, Cisco, Marvell, OnSemi, Se

Comment Type T Comment Status X

(Figure 155-10) This is a comment related to unsatisfied comment 345 (on d2p0). I appreciate much of the clean up that the Task Force and editorial team have done to remove implementation. Most of the instances of the ADC and DAC are removed; however, the ADC and DAC are still present in Figure 155-10, which is supposed to be a functional block diagram, not an implementation diagram. If, for example, I had an analog chromatic dispersion equalizer, the functional diagram might still be met, but there would be no ADC at the location shown. If this comment is accepted, comment 345 will be satisfied.

SuggestedRemedy

Suggest remove blocks labeled DAC & ADC from Figure 155-10, leaving 2 pairs of output arrows from PS field insertion (X) and PS field insertion (Y), labeled X_I, X_Q and Y_I, Y_Q on the transmit side, and 4 input arrows to Chromatic dispersion equalizer (labeled X_I, X_Q, Y_I, and Y_Q, if comment labeled FIG3 is accepted).

Proposed Response Response Status O

CI 155 SC 155.3.1 P 60 L 35 # 13

Zimmerman, George

CME Consulting/APL Gp, Cisco, Marvell, OnSemi, Se

Comment Type T Comment Status X

(Figure 155-10) Also related to unsatisfied comment 345 (dp20). In redrawing the figure, it appears an error was created in the primitive interface at the bottom of the figure. Both directions to & from the PMD are labeled. PMD_IS_UNITDATA.request . I believe (confirmed by figures 156-2 and 156-3), the receive side (right hand side) should be "indication".

SuggestedRemedy

Suggest change right hand side "PMD_IS_UNITDATA.request" to "PMD_IS_UNITDATA.indication"

Proposed Response Response Status O

CI 155 SC 155.3.1 P 60 L 31 # 14

Zimmerman, George

CME Consulting/APL Gp, Cisco, Marvell, OnSemi, Se

Comment Type E Comment Status X

(Figure 155-10) Also related to unsatisfied comment 345 (dp20). The diagram doesn't show how the signals labeled X_I, X_Q, Y_I, and Y_Q relate to the PMD_IS_UNITDATA.request at the output. I believe these are the four components of the request / indications (specified in 156.2.1.1 and 156.2.1.2), but they are not called out as such. It also doesn't show any such label for the receiver (should be indication) side, although the primitive is labeled with the four components. (note the text, e.g., 155.3.3 suggests these are also X_I, X_Q, Y_I, and Y_Q). (Comment labeled FIG3)

SuggestedRemedy

Suggest: Label X_I, X_Q, Y_I, and Y_Q on the receive side; Add a Note to the figure: "NOTE - X_I, X_Q, Y_I, and Y_Q are the four (two complex) components of the inputs and outputs to the PMD, which are the parameters of the primitives PMD_IS_UNITDATA.request and PMD_IS_UNITDATA.indication.

Proposed Response Response Status O

IEEE P802.3cw D2.3 400 Gb/s over DWDM systems 3rd Working Group recirculation ballot comments

Cl 155 SC 155.2.5.5.2 P 49 L 42 # 15

Zimmerman, George CME Consulting/APL Gp, Cisco, Marvell, OnSemi, Se

Comment Type E Comment Status X

Style - the style guide says you spell out single digit numbers - "It is set to one" vs. "It is set to 1". We misuse this a LITTLE in IEEE Std 802.3 (29 "is set to 1" instances vs. over 300 "is set to one". Also, we usually try to avoid pronouns (It) and instead say specifically what we mean - helps out editing when things are moved around.

SuggestedRemedy

Suggest changing "It is set to 1" to "The remote PHY fault indication bit is set to one", and changing "otherwise it is set to 0" to "otherwise it is set to zero".

Proposed Response Response Status O

Cl 155 SC 155.7.4.1 P 82 L 37 # 16

Zimmerman, George CME Consulting/APL Gp, Cisco, Marvell, OnSemi, Se

Comment Type T Comment Status X

This is related to unsatisfied comment 346. The requirements in d2p3 are much improved, but the PICS, which are also part of comment 346 appear to be simply a list of the section headers. From comment 346, "The style of IEEE SA standards (and IEEE Std 802.3) is that requirements use the term "shall". Each PICS item should have an associated "shall" and each "shall" should have a PICS." In many cases this is now OK, as there is only a single shall per subclause. Not ideal, but OK. But in some cases (155.2.5.9, 155.3.3.1.3, 155.4.3, 155.4.5, and 155.5) there are multiple shalls in teh subclause, and each should have its own PICS item.

Because the hard part (putting the shalls in the text) has been done, I plan to mark 346 satisfied. I realize this is a lot of work, and would be OK with a commitment to do the work of collating PICs to shalls in initial SA ballot. (note, I have tagged this in clause 155, it doesn't look like a problem in the other clauses).

SuggestedRemedy

Suggest PICS be rewritten to reflect shalls. This can be done by searching (using advanced search in Adobe) for all instances of "shall" and then collating each to a PIC. For example, PIC TF9 should be broken into multiple PICS (one for each shall) - this may cause you to write some of the "shall's out of text, where they may not be appropriate, e.g., "operation shall be functionally equivalent... and... polynomial shall be..." should probably just be "with the generating polynomial of " (or simply, "shall be functionally equivalent to the frame-synchronous scrambler in Figure 153-5). Again, this is a lot of work, and willing to work with editors to do this later - rather than fill up the comments.

Proposed Response Response Status O

Cl 155 SC 155 P 42 L 4 # 17

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comment 278: this project is too slow, and has descended to only 25 comments from only four commenters when there is a lot to fix still. The moment for doing this spec in 802.3 has passed, it doesn't add significantly to 400ZR, it lacks momentum and there are not enough willing participants in P802.3cw to justify it.

SuggestedRemedy

Cancel this project.
Encourage those interested to feed their learnings into OIF's "400ZR" maintenance.
Re-use relevant parts of the draft in P802.3dj when the time comes.

Proposed Response Response Status O

Cl 155 SC 155 P 42 L 4 # 18

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comment 281: this PCS/PMA is way too complicated for just a "directive" specification. We need examples, as in Annex 91A, RS-FEC codeword examples, or Annex 76A, FEC Encoding example, or the OIF test vectors for 400ZR.

SuggestedRemedy

Publish examples of e.g. FEC and other blocks before and after coding. Smallish ones can go in the document, all can be uploaded to the directory that IEEE provides for these things.
If no-one does the work needed, cancel the project.

Proposed Response Response Status O

Cl 155 SC 155.3.3 P 62 L 37 # 19

Dawe, Piers Nvidia

Comment Type ER Comment Status X

Avoid inconsistent terminology, use the usual 802.3 terminology

SuggestedRemedy

Change "symbol rate" to "signaling rate", several places.

Proposed Response Response Status O

IEEE P802.3cw D2.3 400 Gb/s over DWDM systems 3rd Working Group recirculation ballot comments

CI 156 SC 156.9 P 102 L 13 # 20

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comment 285, optical parameters are inadequately defined.

SuggestedRemedy

Review the 400ZR maintenance projects' activities for corrections and improvements and changes that would apply to this draft, including to EVM.

Proposed Response Response Status O

CI 156 SC 156.9.1 P 102 L 42 # 21

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise, not clear how it would be measured if the transmitter is transmitting Pattern 5. I don't believe that laser frequency noise can be defined with Pattern 5 (scrambled idle). It would have to be a static pattern such as PRBS7Q, PRBS9Q or PRBS11Q in each dimension, or (undesirable) without modulation.

SuggestedRemedy

Set a suitable pattern for laser frequency noise.

Proposed Response Response Status O

CI 156 SC 156.7.1 P 98 L 11 # 22

Dawe, Piers Nvidia

Comment Type E Comment Status X

20ppm

SuggestedRemedy

Insert space. Also in the next table.

Proposed Response Response Status O

CI 156 SC 156.8 P 101 L 31 # 23

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comment 284: It is hard to grasp what this table is meant to say, based on what is in this section, and one cannot see what shape the mask is without plotting it out. The spec should do that job, once, so that every reader doesn't have to.

SuggestedRemedy

1. Insert a sentence: The limit for adjacent channel spectral isolation is given in Table 156-10 and illustrated in Figure 156-xx. Adjacent channel spectral isolation is defined in 156.9.31.
2. Provide the graph to illustrate it. x axis Frequency offset from -75 GHz to 75 GHz, linear scale. y axis Adjacent channel spectral isolation, linear scale in dB.

Proposed Response Response Status O

CI 156 SC 156.9.4 P 104 L 2 # 24

Dawe, Piers Nvidia

Comment Type E Comment Status X

Figures 156-6 and 7 are in a serif font, unlike the others.

SuggestedRemedy

Change to Arial

Proposed Response Response Status O

IEEE P802.3cw D2.3 400 Gb/s over DWDM systems 3rd Working Group recirculation ballot comments

CI 156 SC 156.9.6 P 105 L 8 # 25

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise and write down how it may be measured. For example, it is not stated what is measured in Hz². It is not stated adequately what to do with the two sidebands. The table column header says one-sided, but that's the wrong place to attempt a definition, and does it mean one folds both sidebands together, explicitly or as in a self-homodyne measurement, or takes the worst of the two, or what? It is not stated whether +ve and -ve frequencies are taken into account or just +ve. It seems that this extremely arcane term is more of a concept, or at most a laser modeller's input parameter, than an observable output, so it is not clear that it is the right thing to be specifying, as it may not be measurable.

SuggestedRemedy

Define and specify something relevant and measurable, clearly and completely, with an explanation of how it may be measured and what instrument may be used, and references as necessary. Probably an example is needed. Phase noise is a better-known parameter with some literature, although it needs careful definition to avoid ambiguity. See e.g. IEC 61280-1-3, Fibre optic communication subsystem test procedures--Part 1-3: General communication subsystems--Central wavelength and spectral width measurement for an example of a measurement spec that can be referred to in a definition.

Proposed Response Response Status O

CI 156 SC 156.9.6 P 105 L 9 # 26

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. This text says "The mask frequencies are relative to the laser center frequency from *less than* 100 Hz to half the signaling rate", Table 156-13 has 10² to 10⁹ Hz, and Figure 156-7 shows 10² to something indeterminate above 10¹⁰.

SuggestedRemedy

Reconcile the frequency range for this spec, with clear and consistent lower and upper frequencies. For example, 100 Hz to 59.84375/2 = 29.921875 GHz, or 100 Hz to 30 GHz, or 100 Hz to 30.8 GHz to match the transmit spectrum.

Proposed Response Response Status O

CI 156 SC 156.9.1 P 102 L 45 # 27

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. The header for this column is "Parameter" but "Laser frequency noise mask" is not an observable property of a signal, not even hypothetically. It's a mask, a property of the spec.

SuggestedRemedy

Change "Laser frequency noise mask" here, in Table 156-7 and in the title of 156.9.6. In 156.9.6, start by saying what frequency noise is before discussing the mask.

Proposed Response Response Status O

CI 156 SC 156.9.6 P 105 L 9 # 28

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise and write down how it may be measured. The laser frequency noise is supposed to be controlled down to less than 100 Hz. That's too vague for a spec. No indication is given of how it might be measured, but instruments that can measure GHz often don't measure kHz and below.

SuggestedRemedy

Either don't say anything about frequencies lower than the spec range, or use a separate recommendation (not expected to be testable). Review whether 100 Hz is feasible or necessary, change the limit if appropriate.

Proposed Response Response Status O

CI 156 SC 156.10.1.2.4 P 112 L 47 # 29

Dawe, Piers Nvidia

Comment Type E Comment Status X

"using a RRC filter with a beta = 0.2" is too terse, as "RRC" doesn't appear in the 7000 pages of the base standard, nor elsewhere in 156.10. "a beta" reads oddly. Unnecessary use of a symbol in a sentence, unlike the way it's done in 156.9.4.

SuggestedRemedy

Change to "using a RRC filter (see 156.9.4) with a roll-off factor beta of 0.2"

Proposed Response Response Status O

IEEE P802.3cw D2.3 400 Gb/s over DWDM systems 3rd Working Group recirculation ballot comments

CI 156 SC 156.9.6 P 105 L 21 # 30

Dawe, Piers

Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise and write down how it may be measured. This says "One-sided frequency noise power spectral density (Hz²/Hz)". I can see that a spectral density can be per hertz. Power has dimensions of energy per time, while Hz² is time⁻². These are incompatible.

SuggestedRemedy

If the units are not changed, delete "power" in the table row header and caption, and Figure 156-7, both y axis and caption.

Proposed Response Response Status O

CI 156 SC 156.9.1 P 102 L 45 # 31

Dawe, Piers

Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. The header for this column is "Parameter" but "Laser frequency noise mask" is not an observable property of a signal, not even hypothetically. It's a mask, a property of the spec.

SuggestedRemedy

Change "Laser frequency noise mask" here, in Table 156-7 and in the title of 156.9.6. In 156.9.6, start by saying what frequency noise is before discussing the mask.

Proposed Response Response Status O

CI 156 SC 156.9.4 P 104 L 49 # 32

Dawe, Piers

Nvidia

Comment Type E Comment Status X

T and f should be italic, as in 156A.3

SuggestedRemedy

per comment

Proposed Response Response Status O

CI 156 SC 156.9.6 P 105 L 10 # 33

Dawe, Piers

Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. You can't have a "should" in a definition, it has to be decisive.

SuggestedRemedy

Change "should" to "is" (not "shall" to avoid a trivial PICS). Similarly in 156.10.1.1, "coherent receiver should have", "ENOB and sampling rate of the digitizers should be".

Proposed Response Response Status O

CI 156 SC 156.9.5 P 106 L 1 # 34

Dawe, Piers

Nvidia

Comment Type TR Comment Status X

D2.1 comment 285, optical parameters are inadequately defined. This says "The spectral floor is the limit of the upper mask as defined in 156.9.4 and shall be within the limits given in Table 156-7." There is nothing an implementer can do to affect the limit of the upper mask as defined in 156.9.4, that's a property of the spec. Also causing an upper limit a "floor" is weird; the transmitted spectrum might have a floor, not the mask. The -20 dB limit is given in 156.9.4 anyway. This term is not needed.

SuggestedRemedy

Delete the subclause, and the row for "Spectral floor" in Table 156-7.

Proposed Response Response Status O

CI 156 SC 156.9.3 P 109 L 35 # 35

Dawe, Piers

Nvidia

Comment Type TR Comment Status X

D2.1 comment 285, optical parameters are inadequately defined. It is not clear if the adjacent channel spectral isolation spec applies to frequencies between the ones given in Table 156-10, and if so, whether the intermediate limits are interpolated linearly, lin-log (linear in dB) or stepwise as in in Table 52-8 and Figure 52-3.

SuggestedRemedy

Define the limit fully. If possible, refer to a document that indicates how this can be measured.

Proposed Response Response Status O

IEEE P802.3cw D2.3 400 Gb/s over DWDM systems 3rd Working Group recirculation ballot comments

CI 156 SC 156.9.6 P 105 L 8 # 36

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. The method of interpolation for the laser frequency noise mask is not specified. Figure 156-7 implies log-log interpolation but that is illustrative not normative.

SuggestedRemedy

State that log-log interpolation is used to build the mask is not specified.

Proposed Response Response Status O

CI 156 SC 156.9.6 P 105 L 15 # 37

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.1 comments 285, optical parameters are inadequately defined, and 286, define frequency noise. This says "The definition of maximum laser linewidth is provided in ITU-T G.698.2." G.698.2, 7.2.8 Maximum laser linewidth, says "The laser linewidth is defined as: The level of the white noise component of the power spectrum density of the instantaneous laser frequency multiplied by pi." We need a definition of linewidth, not maximum laser linewidth. A power spectrum density would be in the dimensions of power per frequency, which is not inverse time, so this definition is not satisfactory as it stands.

SuggestedRemedy

Use another reference with a dimensionally correct definition, or write one for laser linewidth (not "maximum laser linewidth" here).

Proposed Response Response Status O

CI 155 SC 155.2.5.11 P 54 L 30 # 38

Dawe, Piers Nvidia

Comment Type TR Comment Status X

D2.0 comment 463: generic operation ... in ITU-T G.709.3 Annex D: but that contains undefined symbols and terms. As it seems it is not very long, write it out cleanly here This is supposed to be a spec, we need a specific definition, not "generic". G.709.3 Annex D describes GMP (as referenced in 155.2.5.3), not the Hamming SD-FEC scheme. Also, G.709.3 is in revision. 400ZR 10.5, Inner Hamming Code, which is about one page long, specifically addresses a systematic (128, 119) double-extended Hamming code.

SuggestedRemedy

Copy the material from 400ZR 10.5, changing some of the b to m if appropriate to match the usual FEC notation in 802.3, and replacing the undefined symbols that look like ^ and V with the ones usually used in 802.3. Whatever symbols are used, say what they mean.

Proposed Response Response Status O