Style	Commenter First Name	Commenter Last Name	Commenter Email	Affiliations	Category	Page	Subclause	Line #	Comment	Proposed Change
Non-Ballot Comment	Adee	Ran	aran@cisco.com	Cisco systems, inc.	Technical	898	178B.8.3.1	9	The current text can be read as requiring deassertion of local_rx_ready only at 100 ms or later. It should be clarified that implementations may deassert local_rx_ready earlier, but deassertion is mandatory under the stated conditions.	Replace the text from "The specific conditions for setting this variable to true are implementation dependent" through the end of the local_rx_ready definition with the following: "The specific conditions for setting this variable to true or false are implementation dependent; however, this variable shall be false when either of the following applies: (a) an input signal corresponding to tx_disable = true in the peer is received for 100 ms (the minimum duration of the quiet_timer) or longer; (b) the training control state diagram (Figure 178B-10a) is in RECOVERY state."
Non-Ballot Comment	Adee	Ran	aran@cisco.com	Cisco systems, inc.	Editorial	841	176D.8.12.1	9	The representation of the MCB and HCB in figures 176D-8a and 176D-8b is backwards. It should be similar to that of figure 176D-7b, where the MCB has a receptacle and the HCB plugs into it.	Flip the receptable/plug shapes horizontally in **figures 176D-8a and 176D-8b.**
Non-Ballot Comment	Adee	Ran	aran@cisco.com	Cisco systems, inc.	Technical	894	178B.8.2.1	38	The variable name training_status causes a lot of consternation because it is also used when mr_training_enable is false. In fact it represents the status of the interface, and indicates also the presence of loss of peer signal, which is not necessarily related to training. A more appropriate name is interface_status. Similarly for lane_training_status, for which a more appropriate name is lane_status. This variable is enumerated and includes TRAINED as one of the possible values, but this value is also used when training is disabled. The value TRAINED corresponds to READY in the interface-level variable. These variables are new in 802.3dj so can be safely renamed. We should consider doing that to reduce confusion for readers who are not interested in the training aspects of Annex 178B.	Rename the variable training_status to interface_status or ilt_status. Rename the variable lane_training_status to lane_status, and the enumerated value TRAINED to READY. Implement across Annex 178B including stat diagrams, and in clause 45.
Non-Ballot Comment	Adee	Ran	aran@cisco.com	Cisco systems, inc.	Editorial	896	178B.8.2.4	1	Following up on unsatisfied comment I-206. The RTS update state diagram (Figure 179B-9) has two states with no assignments. These can be merged into other states. Previous proposals in <https://www.ieee802.org/3/dj/public/26_05/mascitto_3dj_01_2605.pdf> and <https://www.ieee802.org/3/dj/public/26_03/slavick_3dj_03_2603.pdf> have been considered, but were not fully equivalent. This suggested remedy is a slight modification to previous proposal mascitto_01, fixing an issue to make it fully equivalent, in case the CRG wants to go in this direction. Since it is fully equivalent, the comment is editorial. The suggested remedy includes the textual diagram representation. A graphical representation will be provided in a separate contribution.	[*] --> START: (reset + mr_restart + !adjacent_intf_rx_ready) START: local_rts false; USE_TX_CLOCK(local) START --> SWITCH_CLOCK: ((isl_ready + !mr_training_enable) * uses_recovered_clock) START --> FORWARD_RTS: ((isl_ready + !mr_training_enable) * !uses_recovered_clock) SWITCH_CLOCK: USE_TX_CLOCK(recovered); start forward_rts_timer SWITCH_CLOCK --> FORWARD_RTS: ((isl_ready + !mr_training_enable) * forward_rts_timer_done) FORWARD_RTS: local_rts true

Non-Ballot Comment	Adee	Ran	aran@cisco.com	Cisco systems, inc.	Technical	889	178B.7.8	51	The 2nd through 4th paragraphs define when an interface can send requests and when it needs to respond to requests. The 4th paragraph states that in states other than TRAIN_LOCAL and TRAIN_REMOTE requests are not sent and not serviced; It is stated that requests may be made in TRAIN_LOCAL, but it is not stated explicitly that sending requests is not allowed in TRAIN_REMOTE. Additionally, for training pattern settings (in 178B.7.9), and for modulation and coding requests, there are not restrictions at all. It makes sense that these requests are also allowed only in TRAIN_LOCAL and responses to them are required in TRAIN_LOCAL and TRAIN_REMOTE. The suggested remedy is to repurpose subclause 178B.7.10 "Handshake timing" which already deals with requests and responses. Alternatively, a new subclause can be created.	Change the title of 178B.7.10 from "Handshake timing" to "Requests and responses", and insert the following paragraphs at the beginning: "Requests for the peer interface (changes of the initial condition request, modulation and precoding request, training pattern request, coefficient select, and coefficient request bits) shall not be sent when the training control state diagram (Figure 178B-10) is in any state other than TRAIN_LOCAL. When the training control state diagram (Figure 178B–10a) is in any state other than TRAIN_LOCAL or TRAIN_REMOTE, the interface shall ignore any request received from the peer interface." Delete the 3rd and 4th paragraphs in 178B.7.8 (this subclause is specific to equalization control).
Non-Ballot Comment	Adee	Ran	aran@cisco.com	Cisco systems, inc.	Technical	488	180.9.8.1	26	Following up on unsatisfied comment #I-25. As the response indicates, there is an issue to be considered. "The VOA level is set to the value given by Equation (180–17), in which the first and second terms normalize differences in test fibers and optical receivers, respectively, so that symbol error counts are repeatable across different conditions" The response claims that receiver differences with respect to jitter tracking are addressed by setting the attenuation based on the functional receiver's jitter tolerance. However, the calibration procedure uses the RS, not SRS, and RS does not include a jitter stress. Even if SRS was used, jitter tolerance is a function of jitter profile. not a single number, and cannot be converted into a single attenuation value; the jitter applied in SRS test is a single frequency plus specific aggressors, and does not represent the jitter profile of the transmitter under test. One functional receiver might have no margin at one point of the stressed test (and thus have no extra sensitivity) but still have margin with respect to the jitter profile of the transmitter under test; another functional receiver may behave differently and have no margin. The problem of using a specific receiver as a test equipment is fundamental and cannot be circumvented. Users of the standard would expect to get the same results with any functional receiver (as the quoted sentence explicitly suggests), but it is not a realistic expectation. It is possible that the same problem exists with different test fibers as well.	Change the quoted sentence to read: "The VOA level is set to the value given by Equation (180–17)". Add the following NOTE at the end 180.9.8.1: "NOTE---The intent of the first and second terms of Equation (180–17) is to calibrate differences in test fibers and optical receivers, respectively, to reduce variability across different test conditions. Test results can still vary when significantly different fibers and receivers are used".

Non-Ballot Comment	Adee	Ran	aran@cisco.com	Cisco systems, inc.	Technical	838	176D.8.11	25	Following up on unsatisfied comment I-106. The comment addressed the impossibility of measuring the histogram to the low probabilities required by the tables in the "Error ratio requirements" subclauses. Extrapolation is required in practice. An informative note about extrapolation is provided in all PMD clauses, but not in the AUI annexes. In the AUIs the problem is even worse because the required probabilities are lower, and cannot be proven in billion years of test time. See slides 10 and 11 of <https://www.ieee802.org/3/dj/public/adhoc/electrical/26_0421/ran_3dj_adhoc_01_260421.pdf>. As shown in these slides, meeting the required probabilities for bins up to 4 can be demonstrated in about a minute for all cases. Therefore, extrapolation from these bins is practical. A case that should be addressed is that there are no hits in some bins during the test time (which can happen in a good receiver). A common practice in BER tests is to require 0 errors in a 3/BER bits (for a confidence level of 95%). Thus, the highest bin used for the linear fit should be the first one with 0 count (and at least bin 4) and its value should be taken as 3. Any 0 in lower bins should be excluded from the linear fit. This requires a NOTE in the AUI annexes which is slightly modified compared to the NOTE currently used in the PMD clauses. Also applies to 176C.6.4.4.	Add the following NOTE after the text in 176C.6.4.4 and in 176D.8.11: NOTE—In order to predict whether a receiver meets the BLER requirement in a short test time, extrapolation of the measured histogram (see 174A.9.3) to Hm(16) can be performed. One way of doing this is using a linear fit of log10(Hm(k)) for k = 0 to n, where n is at least 4 and Hm(k)=0 for all k>=n. Hm(n) is set to 3 for the purpose of the linear fit, and if Hm(k)=0 for any k<n, it is removed from the linear fit.
Non-Ballot Comment	Adee	Ran	aran@cisco.com	Cisco systems, inc.	Technical	918	179B.2.1	1	Following up on unsatisfied comment I-238. The large variation allowed between test fixture is a concern for both PMD/AUI component measurements and for cable assemblies. The IL of actual test fixtures seems to be systematically lower than the reference IL, which makes measured results of transmitters and cable assemblies too optimistic unless the difference is accounted for. But this accounting is not obvious, and the draft does not include any guidance.	Specify a procedure for calibrating the measurement using S-parameters of an actual HCB. A presentation is planned.

Non-Ballot Comment	Adee	Ran	adee@ieee.org	Cisco systems, inc.	Technical	484	180.9.6.4	41	Following up on unsatisfied comment I-473. The purpose of using two separated histograms is to account for varying sampling phase in a real receiver. The width of each histogram does not model any real behavior, and is only widening each histogram and increases variability of measurement results. Ideally a zero width histogram could be used to capture the vertical opening at each specific phase. Some maximum should be allowed for implementation purposes, which should match a reasonable oversampling ratio. But there should be no minimum width - specifying a minimum requires artificially increasing the oversampling ratio to capture more phases, which also increases measurement complexity. The distance between the two histograms should also match a reasonable oversampling ratio. It should be at least two phases apart to allow for sampling the middle for MMSE equalizer optimization. With an oversampling ratio of 32, as noted in comment I-473, a distance of ±1/32=0.03125 would be significantly smaller than the current ±0.05 UI and cause a degradation. Doubling would make it too large. With a ratio of 40, a distance of ±2/40=0.05 matches the current specification. This enables specifying the maximum width of each histogram as 1/40=0.025.	Change from "Each histogram window has a width of 0.04 UI" to "The maximum width of each histogram is 0.025 UI".
Non-Ballot Comment	Adee	Ran	adee@ieee.org	Cisco systems, inc.	Technical	484	180.9.6.4	38	Real receivers optimize their equalizer based on an average sampling phase (which is represented by phi_0), and the deviations due to jitter are symmetrical around the average. The concept of moving the histograms away from being symmetrically spaced around a sampling phase is a remnant from the old TDECQ in which there was no phi_0 concept. This additional degree of freedom increases the space on which TDECQ should be optimized, and is not justified; expecting a real receiver to optimize at one phase but sample at another phase is not realistic (at least not for a minimally compliant receiver, which is what TDECQ should assume). If transmitters create signals that have significant asymmetry after equalization with a strong equalizer as used in TDECQ, they should be required to have margin at the worst phase, because it will dominate the performance during the periods that the receiver samples around that phase.	Delete the sentence "The precise time position of the pair of histogram windows is adjusted to minimize TDECQ while keeping the histogram windows spaced 0.1 UI apart". Change "nominally centered" to "centered".

Non-Ballot Comment	Adee	Ran	adee@ieee.org	Cisco systems, inc.	Technical	484	180.9.6.4	3	"Alternative optimization methods may be used to determine equalizer tap weights if they report equal or lower mean values of TDECQ" This sentence makes any measurement result dubious - one cannot prove that it is the correct result without considering all alternative optimization methods. Which cannot be proven without an exhaustive search on a huge space. Transmitter specifications should use a reference receiver that represents a minimally capable receiver. The currently described method for setting the equalizer coefficients using MMSE is a reasonable assumption of real receiver capability. But assuming a real receiver can achieve a global minimum using other (unspecified) optimization methods is not realistic. Therefore, even if there is a set of coefficients that results in lower TDECQ, it should not be considered the "true" TDECQ because it is not achievable by a minimally capable receiver. Allowing a transmitter to pass the TDECQ requirement using "alternative optimization methods" can lead to lack of interoperability with receivers that meet the sensitivity requirements.	Delete the quoted sentence.