Transition time in DR transmitters

Supporting comments R2-195 and R2-200 — 8 ps to 10 ps

IEEE P802.3dj D3.2, 2nd SA recirculation ballot

Bill Simms, Piers Dawe, NVIDIA

September 2026

Supporters

Name	Affiliation
Steen Christensen	NVIDIA
Karl Muth	Broadcom
John Calvin	Keysight
Pavel Zivny	Lumilens

Proposal

Comments R2-195 and R2-200 [2], Cl 180 SC 180.7.1 P473 L9

Table	PMDs	D3.2	Proposed
180-7	**200GBASE-DR1, 400GBASE-DR2, 800GBASE-DR4, 1.6TBASE-DR8**	**8 ps**	**10 ps**
181-5	800GBASE-FR4-500	8 ps	10 ps
182-7	200GBASE-DR1-2, 400GBASE-DR2-2, 800GBASE-DR4-2, 1.6TBASE-DR8-2	8 ps	10 ps
183-6	800GBASE-FR4	8 ps	10 ps
183-6	800GBASE-LR4	13 ps	13 ps, no change

- Change the row “Transmitter transition time, each lane (max)” only. No edit to TDECQ, TECQ or SECQ limits, to the reference equalizer of Table 180-16, to the overshoot and undershoot limit, or to any receiver subclause. 180.9.15 defines the stressed receiver conformance test signal by reference to this row, so that test signal's edge widens to 10 ps with it.
- **Why 10 ps.** R2-200 proposes “10 ps or 12.7 ps for all IMDD except 800GBASE-LR4”. This contribution asks for the more conservative of the two. 10 ps is 1.06 UI: looser than today, below the 1.35 UI the draft's own reference equalizer would justify, and well inside the 1.47 UI that 800GBASE-LR4 already permits. It is also the value stated in R2-195, so one edit resolves both comments.
- R2-195 records the technical case but its suggested remedy field is empty. The edit above is offered so the comment can be implemented as written.

D3.2 sets 8 ps and 13 ps at the same signaling rate

800GBASE-FR4 and 800GBASE-LR4 are two columns of Table 183-6 [1], under one signaling rate row

PMD	Table	Rate (GBd)	Transition time	In UI	TDECQ (max)	TECQ (max)
200GBASE-DR1 and similar	180-7	106.25	8 ps	0.85	3.4 dB	3.4 dB
800GBASE-FR4-500	181-5	106.25	8 ps	0.85	3.4 dB	3.4 dB
200GBASE-DR1-2 and similar	182-7	113.4375	8 ps	0.91	3.4 dB	3.4 dB
800GBASE-FR4	183-6	113.4375	8 ps	0.91	3.4 dB	3.4 dB
800GBASE-LR4	183-6	113.4375	13 ps	1.47	3.9 dB	3.2 dB

- **FR4 and LR4 share a signaling rate and are given different edge-speed limits.** One table, one rate, 8 ps and 13 ps. Whatever edge speed is necessary at 113.4375 GBd, both of these cannot be it.
- **And the same 8 ps is reused at a second rate.** It is 0.85 UI at 106.25 GBd and 0.91 UI at 113.4375 GBd, so 8 ps is not one edge-speed requirement consistently applied either — it is one number in picoseconds reused at two signaling rates.
- **And the comparison runs the wrong way.** 800GBASE-LR4 gets the looser TDECQ, 3.9 dB against 3.4 dB, but that is a 10 km chromatic dispersion allowance and says nothing about edge speed. On **TECQ**, the back-to-back transmitter measurement that transition time belongs with, LR4 is the **tighter** of the two at 3.2 dB against 3.4 dB. The PMD given the slower edge owes the cleaner transmitter eye — which is only possible if a 13 ps edge is not itself a barrier to low eye closure.
- **And LR4 receivers are already qualified against a 13 ps edge.** 183.9.15 sets the transition time of the stressed receiver conformance test signal by reference to Table 183-6, so 800GBASE-LR4 receivers are tested with edges up to **13 ps, 1.47 UI**, while Table 183-7 requires them to work through **3.9 dB** of stressed eye closure against 3.4 dB for the DR PMDs [1]. Both clauses use the same reference equalizer. At 1.06 UI and 3.4 dB, 10 ps asks less of a DR receiver than the task force has already accepted.

A common value across Clause 180 and Clause 182 may date from when DR and DR-2 were expected to be two modes of one transmitter, for which a hardware limit in ps rather than UI is the natural form; that explains why the number is shared, not what sets it. The next slide traces where it came from. The receivers are not identical: Table 183-7 specifies a more sensitive receiver for 800GBASE-LR4, -5.5 dBm against -3.7 dBm for 800GBASE-FR4, for the 10 km loss budget. That is a noise floor rather than an equalization capability, and on the equalization axis, SECQ, LR4 is the more demanding of the two.

Where 8 ps came from, and why halving no longer fits

Scaled from a 2018 anchor while the reference receiver was strengthened twice

Generation	Source	Rate (GBd)	Transition time	In UI	TDECQ reference equalizer
50GBASE-FR/LR	802.3cd Table 139-6 [10]	26.5625	34 ps	0.903	5-tap FFE, no decision feedback
100GBASE-DR, 400GBASE-DR4	802.3cd Table 140-6, 802.3df [11]	53.125	17 ps	0.903	5-tap FFE, no decision feedback
DR, Clause 180	**P802.3dj D3.2 Table 180-7 [1]**	**106.25**	**8 ps**	**0.850**	**15-tap FFE and b(1) up to 0.33**

- **The number is arithmetic, not analysis.** 34/2 = 17 ps held the same 0.903 UI at 53.125 GBd; halving again gives 8.5 ps, and the draft carries **8 ps**, rounded down. The 34 ps anchor is a 50GBASE-FR/LR number, and king_3cd_02a_0518 [12] set it against an **ideal transmitter** — “this class of transmitter should not be ruled out” — not against any receiver limit.
- **Meanwhile the reference receiver was strengthened twice, and transition time was revisited at neither step.** 5 feed-forward taps became 15 in D1.0 [8], then D2.x added a feedback tap on straw poll O-3, **Yes 23, No 10**, on the reason that “real receivers perform much better than reference equalizer … this also leaves all the margin for RX DSP” [5]. O-3 accepted limits of 0 to 0.3; Table 180-16 now states b(1) up to 0.33.
- **So halving is the wrong scaling now.** A longer FFE and then a feedback tap both raise tolerance to a slow edge, because the feedback tap removes post-cursor ISI that an FFE can only equalize at a noise cost. Applying the same scaling to the equalizer D3.2 actually specifies gives **34/2/2/(1–0.33) = 12.7 ps**. 800GBASE-LR4 at 13 ps is consistent with that; the DR PMDs at 8 ps are consistent with a reference receiver the task force has twice agreed is weaker than a real one.

Transition time fails all while TECQ discerns

Four transmitter rolloff shapes at one bandwidth, 28 GHz, through the D3.2 reference equalizer at 113.4375 GBd, Clause 183

Transmitter rolloff shape	T20-80	vs 8 ps (max)	TECQ	vs 3.4 dB (max)
EML-like: RC limited, single pole	**8.60 ps**	**fails**	**2.42 dB**	**passes**
Gaussian	8.86 ps	fails	3.90 dB	fails
Bessel-Thomson, 4th order	9.18 ps	fails	4.30 dB	fails
VCSEL-like: damped 2nd order	9.31 ps	fails	4.23 dB	fails

- 0.71 ps of measured edge separates transmitters that TECQ spreads by 1.88 dB. Transition time is not screening an escape here; it is removing a conforming transmitter.
- **The two disagree on which is worse.** Transition time ranks the Bessel-Thomson shape ahead of the second order, 9.18 against 9.31 ps. TECQ ranks it behind, 4.30 against 4.23 dB.
- **At 10 ps nothing gets in that TECQ has not already stopped.** The three shapes TECQ fails here reach the limit between 8.38 and 8.43 ps, at about 31 GHz of transmitter bandwidth — tighter than 10 ps, so across that band TECQ is the operative limit. Slowing a low-pass transmitter always raises TECQ, so transition time has no failure mode of its own to catch.
- **TECQ is not the only backstop.** The same table bounds |TDECQ – TECQ| at 2.5 dB and overshoot and undershoot at 27% (183.9.9), with the symbol error histogram (183.9.8) and the tap coefficients (Table 180-16) [1]. Clause 180 bounds the DR PMDs the same way; backup carries the full table.

Method: MMSE taps per 180.9.6.4 with the Table 180-16 constraints, which 183.9.6 adopts by reference, SSPRQ at 64 samples per UI. Both measured columns are the no-fiber chain: 183.9.7 defines TECQ by the TDECQ method with the test fiber omitted, and 183.9.12 measures transition time at the reference receiver output. 180.9.6.4 permits alternative optimizers only if they report comparable TDECQ; a direct SER minimizer reports lower, so would not qualify; it would lower every TECQ in the table and move the 8.4 ps crossing about 1 ps later.

Summary

- **D3.2 already sets 8 ps and 13 ps at the same signaling rate**, in one table. 800GBASE-LR4 gets the looser TDECQ, 3.9 dB against 3.4 dB, as a 10 km dispersion allowance, but the tighter TECQ, 3.2 dB against 3.4 dB. On the back-to-back measurement that transition time belongs with, the PMD given the slower edge owes the cleaner eye.
- **8 ps was never derived for 200G per lane.** It is 34/2/2 = 8.5 ps rounded down, scaled from generations whose TDECQ reference equalizer was a 5-tap FFE with no decision feedback, and anchored in 2018 on an ideal transmitter rather than on any receiver limit. D3.2 now specifies a 15-tap FFE and a feedback tap up to 0.33; the same scaling corrected for it gives 12.7 ps.
- **The 8 ps limit fails transmitters that TECQ passes with a decibel to spare.** At 28 GHz, inside the transmitter bandwidth range cited in R2-200 [6], four rolloff shapes all fail 8 ps while their TECQ spans 1.88 dB and the single pole passes at 2.42 dB. Of the 36 cases studied, 6 are failed by 8 ps while conforming to TECQ, and 0 are admitted by it while failing TECQ.
- **10 ps is 1.06 UI** — looser than today, tighter than the DFE-corrected scaling, and well inside what 800GBASE-LR4 already permits. It does not expose the receiver: 180.9.15 bounds the stressed receiver conformance test signal by reference to the same row, so the qualification condition moves with it.

References

[1] IEEE P802.3dj D3.2 — Tables 180-7, 180-16, 180-18, 181-5, 182-7, 183-6, 183-8; subclauses 180.9.6.3, 180.9.8, 180.9.9, 180.9.12, 180.9.15, 183.9.6 to 183.9.9 and 183.9.12

[2] Comments R2-195 (christensen) and R2-200 (Dawe), D3.2 — www.ieee802.org/3/dj/comments/

[3] Comments I-319 to I-322 (Kimber, Semtech) and I-443, and the straw poll counts for O-5, O-6 and O-7, in 8023dj_D3p0_comments_final_id.pdf; the polls themselves are recorded as directional straw polls in minutes_3dj_2605_unapproved.pdf, IEEE 802.3 May 2026 interim, Munich

[4] Comments R2-15 and R2-204 (D3.2) — the w(0) minimum 0.9 – b(1) and the pre-post-cursor tap coefficient difference limit of Table 180-16

[5] Comments #382, #383, #384, #393 and straw poll O-3 (decision), Yes 23 No 10, in 8023dj_D2p0_comments_final_id_v2.pdf — addition of the 1-tap DFE to the TDECQ reference receiver, at coefficient limits of 0 to 0.3 as polled

[6] kimber_3dj_01a_2603 — M. Kimber, “TDECQ Limit and DFE Tap Weight Limit”, 802.3dj Vancouver March 2026; the transmitter bandwidth sweep cited in R2-200, and the recommendations to limit the DFE tap to 0.25 and TDECQ (max) to 3.0 dB

[7] kimber_3dj_adhoc_01a_260415, “Transition Time and Transmitter Bandwidth” — joint Electrical, Logic and Optics ad hoc, 15 April 2026, cited in the response to comment I-321

[8] IEEE P802.3dj D1.5 180.9.5 — 15-tap FFE with no decision feedback tap

[9] IEEE Std 802.3bs-2017 121.8.5.4 — 5 tap, T-spaced FFE TDECQ reference equalizer, with no tap coefficient constraint; Clauses 121 to 124 — no transmitter transition time specified for the optical PMDs, Table 124-6 (400GBASE-DR4) as originally published carrying no such row

[10] IEEE Std 802.3cd-2018 Tables 138-8, 139-6, 140-6 — 34 ps, 34 ps and 17 ps; 138.8.5.1, 139.7.5.4, 140.7.5.1 — sum of the equalizer tap coefficients equal to 1, and largest-magnitude tap constrained to at least 0.8; 138.8.10, 139.7.10, 140.7.10 — stressed test signal transition time by reference to the transmit characteristics table

[11] IEEE Std 802.3df-2024 — “Replace Table 124–6 with the following table”; the replacement retrofits Transmitter transition time (max) 17 ps onto 400GBASE-DR4, which IEEE Std 802.3bs-2017 had defined without one

[12] king_3cd_02a_0518 — J. King and C. Cole, “Additional PAM4 transmitter constraints”, 802.3cd interim May 2018. Origin of the transmitter transition time specification, its subclause text, and the 0.8 and 0.9 UI proposals

[13] king_3cd_02_0718 — J. King, “TDECQ map and interpretation”, 802.3cd plenary July 2018

[14] dawe_3cd_01b_0518 — P. Dawe, “TDECQ and SRS”, 802.3cd May 2018, including comment r02-54 “Bound the right (slower than expected)”; and dawe_3cd_01a_0318, in which the TDECQ map was introduced

[15] dawe_3cu_01a_031720 — P. Dawe, “TDECQ, slowness, badness and overshoot (revised)”, IEEE P802.3cu March 2020

[16] S. Echeverri-Chacón, J. J. Mohr, J. J. Vegas Olmos, P. Dawe, B. V. Pedersen, T. Franck and S. B. Christensen, “Transmitter and Dispersion Eye Closure Quaternary (TDECQ) and Its Sensitivity to Impairments in PAM4 Waveforms”, J. Lightwave Technol. 37(3) pp852-860, 1 February 2019

Backup

Supporting detail, held for discussion

Requested straw poll

- **(Direction)** I support changing Transmitter transition time, each lane (max) from 8 ps to 10 ps in Tables 180-7, 181-5, 182-7 and in the 800GBASE-FR4 column of Table 183-6, leaving 800GBASE-LR4 at 13 ps, per comments R2-195 and R2-200.

Impact of the change

Scope and risk at second recirculation

- **Transmitters:** a relaxation. Every transmitter compliant with 8 ps remains compliant. No transmitter has to be requalified.
- **Receivers:** no limit changes. The SRS conformance test signal bound in 180.9.15 follows Table 180-7 by reference, so it widens with it; the LRS signal of Table 180-18 is untouched. The receiver must tolerate transmitters with edges up to 10 ps, and 800GBASE-LR4 shows the task force has already accepted a 13 ps edge at a higher signaling rate.
- **Test and measurement:** no change to the measurement method of 180.9.12, to 180.9.15, to the reference receiver, the reference equalizer, or any test pattern.
- **Scope:** one table row in each of Clauses 180, 181, 182 and the 800GBASE-FR4 column of Clause 183. 800GBASE-LR4 is untouched.

800GBASE-LR4 is not a like-for-like precedent: its 13 ps comes with a higher TDECQ allowance, 3.9 dB against 3.4 dB, and Table 183-8 allocates it 5 dB of penalties against 3.9 dB. That allowance is for 10 km chromatic dispersion, which is why TDECQ, measured through the fiber, is looser for 800GBASE-LR4 while TECQ, measured without it, is tighter at 3.2 dB against 3.4 dB. Since transition time is a back-to-back transmitter measurement, TECQ is its counterpart, and the direction of the comparison holds.

Receivers are qualified at whatever the table says

Why relaxing the transmitter limit does not expose the receiver

Receiver test	Transition time of the conformance test signal	Specified in
Stressed receiver sensitivity (SRS)	No greater than Table 180-7, so 8 ps today	180.9.15
Low-stressed receiver sensitivity (LRS)	4.5 ps (min) to 6.5 ps (max)	Table 180-18

- **180.9.15 cites Table 180-7 rather than restating a number.** The stressed receiver conformance test signal is “no greater than the value specified in Table 180-7” [1], so changing that row to 10 ps widens the receiver conformance bound to 10 ps automatically. No edit to 180.9.15, to Table 180-18, or to any receiver limit is required, and the same by-reference sentence appears in Clauses 181, 182 and 183.
- **This is the mechanism the original authors intended.** king_3cd_02_0718 [13] found that “the transmitter most likely to be affected by receiver non-linearity, quantization, or other sampling errors, is … the Tx with maximum equalizable eye-closure (low pass filtered)”, and concluded that “the toughest receiver test condition is for a low pass filtered test source … suggests we should keep the current SRS test point”. The answer to a slow transmitter was to test the receiver there, not to forbid the transmitter.
- **Today the coupling is an upper bound only.** Nothing requires the stressed signal to be slow, so a 5 ps stressed signal is compliant, and the LRS test of Table 180-18 runs the other way entirely, requiring a fast edge in a 4.5 ps to 6.5 ps window. If the task force wants receivers demonstrably qualified at the transmitter limit, the instrument is to **require** a slow edge in the SRS calibration. That would be a stronger specification than either 8 ps or 10 ps, and this contribution would support it.

The by-reference construction is inherited, not new: IEEE Std 802.3cd-2018 [10] 138.8.10, 139.7.10 and 140.7.10 each say the transition time of the stressed receiver conformance test signal “is no greater than the value specified in” the transmit characteristics table.

The argument for keeping or tightening 8 ps

D3.0 comments I-319 to I-322 (Kimber) [3], kimber_3dj_01a_2603 [6] and kimber_3dj_adhoc_01a_260415 [7]

- **The case made.** If transmitter bandwidth is too low the BER floor rises; a bandwidth below 35 GHz is “detrimental to the BER floor”, and that “would equate to a transition time of 7.2 ps with a 53 GHz SIRC measurement filter” [3]. On that basis comments I-319 to I-322 asked to *tighten* 8 ps to 7.2 ps. All four were rejected on straw polls O-5 to O-7 as seen on the next slide.
- **What the supporting simulation varied.** kimber_3dj_01a_2603 [6] sets transmitter bandwidth with a Bessel filter, **no transmitter FIR**, and extinction ratio fixed near 4 dB. R2-200 makes the point: “other filters and transmitters would have given other results” [2]. The sweep supports a conclusion about *bandwidth*.
- **The step from bandwidth to 7.2 ps is a shape assumption, stated in the source.** kimber_3dj_adhoc_01a_260415 [7] models the low-bandwidth transmitter as a **raised cosine**, takes its transition time as 0.22/BW, and adds the 4.14 ps oscilloscope response in quadrature. That is a fair calculation for a raised cosine and no other shape. A transmitter with the same bandwidth and a different response lands on a different number of picoseconds, which is precisely why a ps limit is a poor instrument for bounding bandwidth.
- **The same contribution supports relaxing the limit.** On the feedback tap, Kimber finds that “the DFE tap weight allows very low bandwidth transmitters to be passed” [6], tabulating b(1) = 0.3 as admitting a **31.5 GHz** transmitter at 2.8 dB TDECQ. He proposed to claw that back: limit the DFE tap to 0.25 and TDECQ (max) to 3.0 dB. **D3.2 adopted neither** — Table 180-16 allows b(1) up to 0.33 and Table 180-7 allows 3.4 dB. The draft accepts the wider slow-transmitter tolerance everywhere except transition time.
- **And the draft already measures the effect directly.** If a poor transmitter does raise the error floor, the functional symbol error histogram of 180.9.8 and the ECQ limits respond to group delay as well as to amplitude response. As R2-200 puts it, a signal the ECQ equalizer “finds difficult to equalise, because of its group delay (phase) response for example, gets a bad rating anyway” [2].
- This proposal does not dispute that transmitter quality matters. It disputes that transition time is the instrument that secures it.

The straw poll record

D3.0 straw polls O-5, O-6 and O-7 [3], taken during optics comment resolution at the IEEE 802.3 May 2026 interim, Munich

Straw poll	Clauses covered	A: 8 ps (no change)	B: 7.2 ps (tighter)	C: 12 ps (looser)
O-5 (chicago)	180, 181	26	17	6
O-6 (choose 1)	180, 181	18	12	3
O-7 (choose 1)	182, 183 FR4	21	13	not offered

- **The record is presented in full because it does not favor this proposal.** In O-6, more voters wanted the limit tighter, 12, than looser, 3.
- **10 ps was never on the ballot.** O-5 and O-6 offered no change, 7.2 ps, or 12 ps. **O-7 offered no loosening option at all** — only 8 ps or 7.2 ps — so for Clauses 182 and 183 the room was never asked whether the limit should be relaxed.
- **And the comparison that matters was out of scope throughout.** 800GBASE-LR4's 13 ps was in none of the three polls, so the inconsistency this contribution raises has not yet been put to the task force.

O-5 and O-6 asked: “For the PMDs defined in clauses 180 and 181, I support setting the transition time (max) to: A: 8 ps (no change), B: 7.2 ps (per comment #I-319 and #I-320), C: 12 ps (per comment #I-443)”. O-7 asked the same question for Clause 182 and 800GBASE-FR4 with options A and B only. I-443, which sought 12 ps or removal of the item, closed with response status U, unsatisfied. The minutes of the May 2026 session record all three as directional straw polls against comments #319 and #321; the counts above are from the D3.0 comment report [3].

Transition time is a weak proxy for bandwidth

The same measurement, swept against transmitter bandwidth for each shape

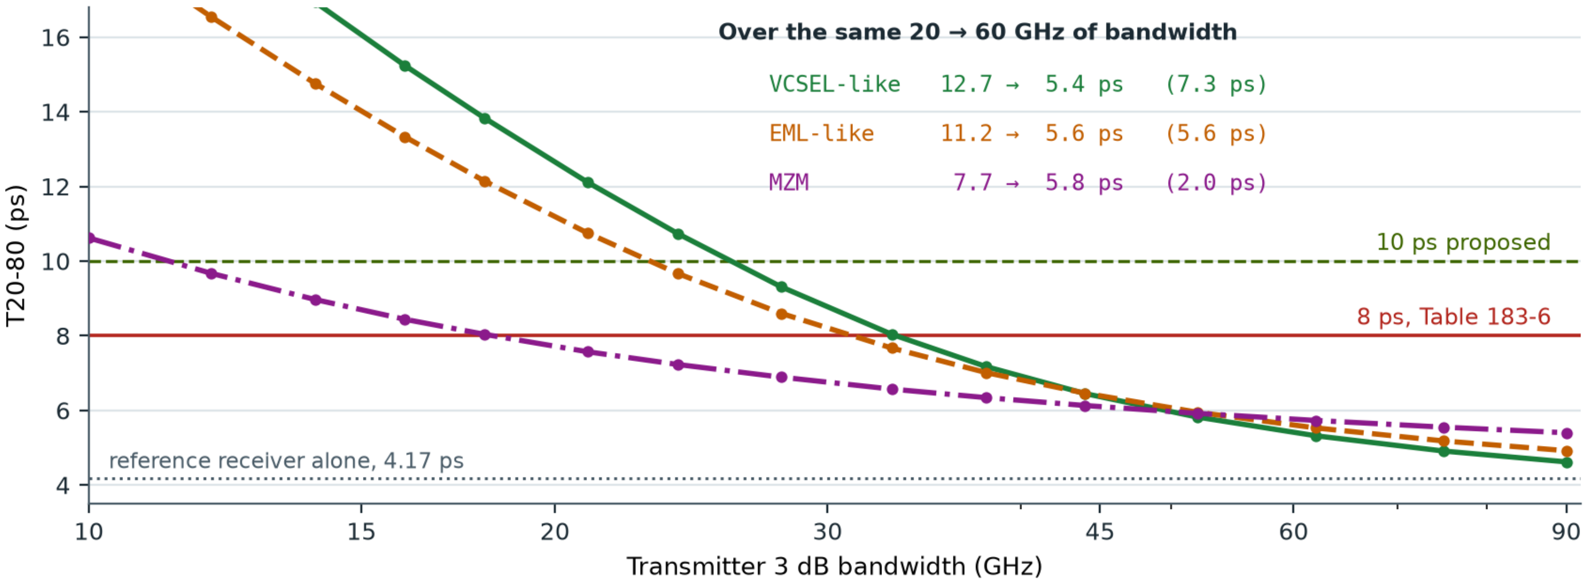

- **The proxy holds within one shape and fails across shapes.** Tripling the bandwidth of the second-order shape moves the reading 7.3 ps; doing the same to the MZM shape moves it 2.0 ps. A single transition time value therefore cannot deliver a single bandwidth, which is the objection raised in R2-200 [2].
- **Transition time is tighter on the device than it looks.** An ideal rectangular source through the reference receiver alone already reads 4.17 ps. Removing that in quadrature, 8 ps asks the transmitter's own edge to be 6.8 ps where 10 ps asks 9.1 ps, so the proposal buys the transmitter 33% more edge time, not 25%.

Same measurement and same conditions as the TECQ slide: T20-80 per 183.9.12 on the SSPRQ rising edge inside 00000333333, taken at the reference receiver output, with each shape swept over its own transmitter bandwidth. The three curves would coincide if transition time were a faithful stand-in for bandwidth.

What already constrains the transmitter

Transition time is a narrow proxy; these Clause 180 limits are not, and this proposal changes none of them

- **Transition time is a narrow measurement.** It is one number from one edge of one symbol sequence: with SSPRQ, the pattern Table 180-14 [1] specifies, the rising edge measured is the one inside 00000333333 and the falling edge the one inside 33333000000, with the 0% and 100% levels taken from the OMAouter procedure on other parts of the pattern. ECQ is computed over the whole pattern.
- **And it does not map cleanly to bandwidth.** The relation between bandwidth and transition time depends on the filter type, so the limit does not reliably deliver the bandwidth it is standing in for.

Limit	Reference	Value
TDECQ, each lane (max)	180.9.6	3.4 dB
TECQ, each lane (max)	180.9.7	3.4 dB
| TDECQ – TECQ |, each lane (max)	Table 180-7	2.5 dB
Main tap coefficient limit w(0)	Table 180-16	0.9 – b(1) to 2.5
Pre-post-cursor tap coefficient difference limit	Table 180-16	0.25 max
Transmitter overshoot and undershoot (max)	180.9.9	27%
Transmitter functional test (TFT) symbol error histogram	180.9.8	Mask in Table 180-17

- These catch undesirable signals by their effect rather than by proxy. A transmitter that is genuinely hard to equalize fails one of them, whatever its measured edge speed. A transition time limit substantially tighter than 800GBASE-LR4's is not needed as well.
- **One caution on reading the table.** The w(0) floor is not a constant: it is coupled to the feedback tap, so the real constraint is **w(0) + b(1)** ≥ **0.9** [4]. Main-tap gain and feedback gain trade against each other and are not independently available.

Where 8 ps came from

Scaled from a 2018 anchor, king_3cd_02a_0518 [12], not derived from 200G/lane analysis

PMD	Source	Rate (GBd)	Transition time	In UI	Reference equalizer
50GBASE-FR/LR	802.3cd Table 139-6	26.5625	34 ps	0.903	5-tap FFE
100GBASE-DR	802.3cd Table 140-6	53.125	17 ps	0.903	5-tap FFE
400GBASE-DR4	802.3df Table 124-6	53.125	17 ps	0.903	5-tap FFE
DR, Clause 180	P802.3dj D3.2 Table 180-7	106.25	8 ps	0.850	15-tap FFE from D1.0

- **The purpose on the record is equalizer cost and interoperability.** King and Cole gave one reason for creating the spec: to “exclude very slow transmitters to ensure that the range of permitted transmitters is compatible with **low power equalizers** and to ease interoperability” [12]. The subclause text they drafted became 138.8.7, 139.7.7 and 140.7.7, and is the ancestor of D3.2 180.9.12.
- **Transition time was chosen as a stand-in for tap-weight limits, on measurement grounds.** It “avoids specifying an indirect characteristic such as EQ tap weights” and “allows simple, practical measurements at sub-component level” [12].
- **The value was anchored on an ideal transmitter, not on a receiver limit.** “For a Nyquist bandwidth Tx, with ideal raised cosine response, T10-90 ~ 0.75 UI out of the transmitter … through a Nyquist bandwidth receiver this becomes ~1 UI. **This class of transmitter should not be ruled out.**” [12] Nothing in the 2018 record derives a ceiling.
- **Everything since has been arithmetic.** 34/2 = 17 held the same 0.903 UI at 53.125 GBd. Continuing the halving gives 34/2/2 = **8.5 ps** at 106.25 GBd; the draft carries **8 ps**, rounded down. The 34 ps anchor is a 50GBASE-FR/LR number, and it was set against a **5-tap** FFE.

king_3cd_02a_0518 proposed max T20-80 of 0.8 UI for SMF and 0.9 UI for 50G MMF; 802.3cd [10] adopted the looser 0.903 UI for both, and 34 ps was itself looser than the 27 ps to 30 ps the supporting simulations proposed. ***The DR family had no such limit for seven years.*** *IEEE Std 802.3bs-2017 [9], which defined 400GBASE-DR4, specifies no transmitter transition time for its optical PMDs: the term does not appear in Clauses 121 to 124, and Table 124-6 as published carries no such row. IEEE Std 802.3df-2024 [11] then directs “Replace Table 124–6 with the following table”, and the replacement carries Transmitter transition time (max) 17 ps. The limit was retrofitted onto a shipping PMD when the table was rewritten to add 800GBASE-DR8 and the -2 variants.*

Transition time in UI, which removes the line rate

Published standards, IEEE P802.3dj D3.2, and this proposal

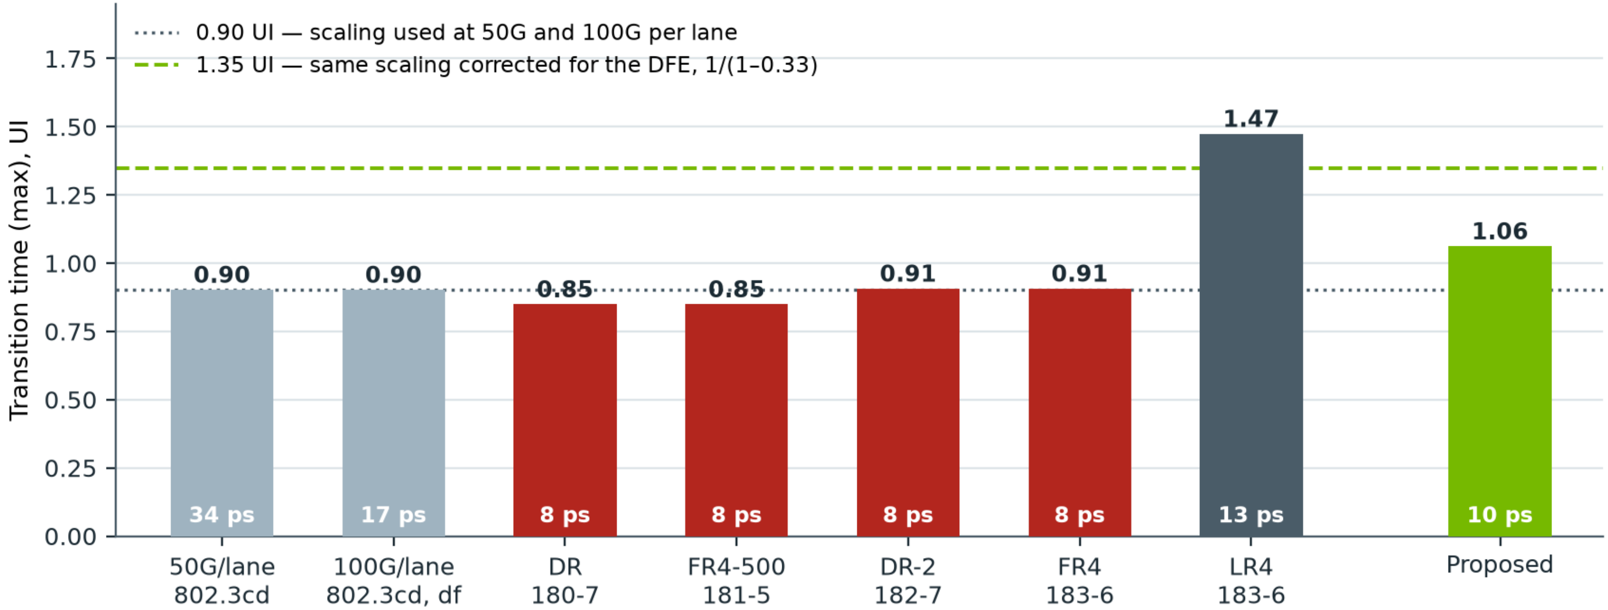

- **Scaled on UI, the draft's 8 ps is tighter than any shipped generation.** 0.85 UI at 106.25 GBd and 0.91 UI at 113.4375 GBd, against the 0.90 UI used at 50G and 100G per lane.
- **Corrected for the reference equalizer D3.2 actually specifies, that scaling gives 1.35 UI.** Only 800GBASE-LR4, at 1.47 UI, is above it.
- **10 ps is 1.06 UI and 1.13 UI at the two rates** — above the old scaling, still below the DFE-corrected line, and at least 23% tighter than 800GBASE-LR4.

The reference receiver got stronger twice

5 taps to 15 taps, then a feedback tap on top, with no change to the limit

Draft or standard	TDECQ reference equalizer	Defined in
802.3bs, cd, cu, df	5 tap, T-spaced, feed-forward equalizer (FFE). No decision feedback.	802.3bs-2017 121.8.5.4, 802.3cd-2018 139.7.5.4
P802.3dj D1.0 to D2.0	15-tap, T-spaced FFE. Three times the taps, still no decision feedback.	D1.5 180.9.5 [8]
P802.3dj D2.x to D3.2	15 feed-forward taps and one feedback tap, b(1) up to 0.33	D3.2 180.9.6.3, Table 180-16

- A longer FFE, and then a feedback tap, both raise tolerance to a slow edge. The feedback tap removes post-cursor intersymbol interference that an FFE otherwise equalizes at a noise cost.
- The feedback tap was requested against D2.0 by comments #382 to #384 and #393, and accepted in principle on straw poll O-3, **Yes 23, No 10**. The reason given and accepted was that “real receivers perform much better than reference equalizer … this also leaves all the margin for RX DSP” [5].
- Applying the same scaling with the DFE the draft itself specifies: **34/2/2/(1–0.33) = 12.7 ps**.
- The DR transition time was revisited at neither step. 800GBASE-LR4, at 13 ps, is consistent with the corrected scaling. The DR PMDs, at 8 ps, are consistent with a reference receiver the task force has twice agreed is weaker than a real one.

802.3df adds no decision feedback tap and inherits the 5-tap Clause 121 equalizer. The only DFE parameters in 802.3cd are COM parameters for the electrical PHYs of Clauses 136 and 137. Straw poll O-3 itself reads “limits of 0 to 0.3”, and the comment it resolved had asked for –0.4 to 0, so the sign convention as well as the value has moved since. The figure used here is the limit as D3.2 Table 180-16 now states it, b(1) from 0 to 0.33 referenced to OMATDECQ/2 [1], and it is the value R2-200 uses to obtain 12.7 ps.

The 2018 anchor expressed in UI

king_3cd_02a_0518, and every limit since, against King's ideal transmitter

Limit	In UI	Over King's ideal transmitter
King's ideal Nyquist raised-cosine Tx, in T20-80	0.657	—
king_3cd_02a_0518 proposal, SMF	0.800	+22%
D3.2 8 ps at 106.25 GBd	0.850	+29%
802.3cd adopted, 34 ps and 17 ps	0.903	+38%
D3.2 8 ps at 113.4375 GBd	0.907	+38%
Proposed 10 ps at 106.25 GBd	**1.062**	**+62%**
D3.2 13 ps, 800GBASE-LR4	1.475	+125%

- **King's anchor is a floor, not a ceiling.** It says the limit must be loose enough not to exclude an ideal Nyquist-bandwidth transmitter. It was never offered as the maximum a receiver can tolerate, and nothing in the 2018 record derives a maximum.
- The ceiling is set by what a low power equalizer can afford. That is the question the reference equalizer has answered twice since 2018, first at 15 taps and then with the feedback tap, and it is why the limit should have moved and did not.
- Dawe had asked for tighter still in comment r02-54, 28 ps for Clause 139 and 15 ps for Clause 140, 0.74 and 0.80 UI, revised on the same slide to 27 to 28 ps. 802.3cd adopted neither his values nor King's, settling on 0.903 UI.

The T20-80 to T10-90 conversion used here, 0.657, is that of a Gaussian edge. king_3cd_02a_0518 gives the anchor in T10-90 and the proposal in T20-80 without stating a conversion of its own.

The TDECQ map, and what the right-hand bound was for

dawe_3cd_01b_0518 [14], king_3cd_02_0718 [13], and Echeverri-Chacon et al., JLT 37(3) 2019 [16]

- **The map.** TDECQ decomposes into C = 10log10(Ceq), the equalizable or “slowness” penalty, and K = TDECQ – 10log10(Ceq), the non-equalizable or “badness” penalty. Slow, low-pass-filtered transmitters sit to the **right**. The decomposition and the bounded map both come out of the 2018 802.3cd work by the same group — dawe_3cd_01a_0318 and dawe_3cd_01b_0518 [14] — and the peer-reviewed account is the JLT paper [16]. **Neither the paper nor the map sets a transition time limit.**
- **What the right-hand bound was justified by.** Comment r02-54, “Bound the right (slower than expected)”, gives one reason: receiving such a signal “would place an extra burden on the receive equalizer e.g. **better linearity and/or finer AtoD or tap resolution**”. The exclusion box drawn on the right of the map itself carries the same reason and no other — it “should be excluded because it **requires strong tap weights not useful in practice**, for SMF would have failed T/2-spaced TDECQ” [14]. The concern on the record is tap weights, equalizer and converter cost. D3.2 Table 180-16 [1] now limits tap weights directly.
- **Clock recovery is cited in that document, but for the opposite side of the map.** The single mention of the CDR in dawe_3cd_01b_0518 is on the **left** bound, against over-emphasis: going too far left “requires significant tap weights of the opposite sign to normal, could confuse the CDR, no benefit to Tx” [14]. When Dawe revisited the same four bounds for 802.3cu in dawe_3cu_01a_031720 [15], clock recovery does not appear at all. A timing-extraction rationale for the right-hand bound is a new argument rather than an inherited one, and rests on its own evidence.
- **The author of the transition time spec cautioned against corner-trimming.** On the map, king_3cd_02_0718: it is “usually surrounded by warnings of what is bad for a receiver, but **without supporting data from real receivers, or analysis to back up the claims**”, and because C and K sum to TDECQ the compliant region “is a triangle – which always looks like it could do with trimming at the corners” [13].

802.3cd [10] carried a “TDECQ – 10log10(Ceq) (max)” row in Tables 138-8, 139-6 and 140-6, and Table 180-7 has no equivalent. Two things about that row. In each of the three tables it was set equal to the TDECQ limit — 4.5 against 4.5 in Table 138-8, 3.4 against 3.4 in Table 140-6 — so it bit on pre-emphasis, not on slowness. And D3.2 does not simply drop it: Table 180-7 carries TECQ (max) 3.4 dB and | TDECQ – TECQ | (max) 2.5 dB, which bound the transmitter's own eye closure back to back, alongside the Table 180-16 tap constraints that 802.3cd lacked. This is why the proposal corrects the transition time value rather than removing the item.

In 2018 no tap limit bounded a slow signal; in D3.2 one does

802.3cd constrained the taps from below only. D3.2 constrains them from above.

Constraint on the TDECQ reference equalizer	802.3cd-2018	P802.3dj D3.2
Number of feed-forward taps	5	15
Decision feedback tap	none	1, b(1) up to 0.33
Main tap coefficient (min) — bounds pre-emphasis	0.8	w(0) ≥ 0.9 – b(1)
Main tap coefficient (max) — bounds slowness	**none**	**w(0) ≤ 2.5**
Per-tap normalized coefficient limits	none	all 15, Table 180-16
Feed-forward tap DC gain, the sum of the w(i)	= 1 exactly	≤ 1, no minimum
Pre-post-cursor difference limit	none	0.25
Transition time (max)	34 ps, 17 ps	8 ps, 13 ps

- **802.3cd did constrain the taps, but only from below.** 138.8.5.1, 139.7.5.4 and 140.7.5.1 each require the largest-magnitude tap to be “at least 0.8” [10]. That is a floor, and a floor limits *pre-emphasis*. Dawe's own account in dawe_3cu_01a_031720: “Equalizable overshoot is constrained by minimum of largest tap coefficient (0.8) — towards the **left** on the TDECQ map” [15].
- **Nothing in 802.3cd bounded a slow signal through the taps.** There was no maximum. Dawe's r02-54 had offered one as the alternative to a transition time limit — a “maximum cursor strength limit” [14] — and it was not adopted, so transition time went in as the only bound on the slow corner.
- **D3.2 has the maximum that 802.3cd lacked.** Table 180-16 [1] caps the main tap at **w(0) ≤ 2.5** with the feed-forward DC gain capped at 1. A cap on the main tap with the DC gain capped is a cap on the high-frequency boost the reference equalizer may apply, and so a bound on how slow a signal can pass. The slow corner is now held by the mechanism 802.3cd declined, and the transition time limit is a second, cruder bound on the same axis.

If 10 ps is not acceptable: the consistency-only alternative

A narrower change that requires no judgment about receivers

Clauses	Rate (GBd)	D3.2	In UI	Consistency-only	In UI
182, 183	113.4375	8 ps	0.907	8 ps, no change	0.907
180, 181	**106.25**	**8 ps**	**0.850**	**8.5 ps**	**0.907**

- The same 8 ps cannot be correct at both signaling rates. Matching Clause 180 and 181 to the UI of Clause 182 and 183 gives **8.5 ps** at 106.25 GBd — a 7% relaxation rather than 25%.
- This asks the task force only to apply its own scaling consistently. It takes no position on the reference equalizer, on the 800GBASE-LR4 comparison, or on what a receiver can tolerate.
- It is offered as a fallback, not as the proposal. It leaves the DR PMDs at a limit still derived from a 5-tap reference equalizer, and leaves the 1.6-to-1 gap against 800GBASE-LR4 in place.