

802.3df D3.0

Comment Resolution

P802.3df editorial team

Introduction

- This slide package was assembled by the 802.3df editorial team to provide background and detailed resolutions to aid in comment resolution.

Clause 167

Clause 167, parameter labels

Comments 13 and 14

Cl 167 SC 167.7.2 P164 L26 # [redacted]

Li, Jing YOFC

Comment Type E Comment Status D (bucket1)

Receiver sensitivity (OMAouter) (max)

SuggestedRemedy
Receiver sensitivity, each lane (OMAouter) (max)

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Change "Receiver sensitivity (OMAouter) (max)" to "Receiver sensitivity (OMAouter), each lane (max)"

167.8.2 Multi-lane testing considerations

Receiver sensitivity and stressed receiver sensitivity are defined for an interface at the BER specified in 167.1.1. The interface BER is the average of the BERs of the receive lanes when they are stressed. Measurements with Pattern 3 (PRBS31Q) allow lane-by-lane BER measurements. Measurement with Pattern 5 (scrambled idle encoded by RS-FEC) gives the interface BER if all lanes are stressed at the same time.

If each lane is stressed in turn, the BER is diluted by the unstressed lanes, and the BER for that stressed lane alone is found, e.g., by multiplying by four for 400GBASE-SR4 if the unstressed lanes have low BER. In stressed receiver sensitivity measurements, unstressed lanes may be created by setting the power at the receiver under test well above its sensitivity and/or not stressing those lanes with ISI and jitter, or by other means. Each receive lane is stressed in turn while all are operated. All aggressor lanes are operated as specified. To find the interface BER, the BERs of all lanes when stressed are averaged.

Cl 167 SC 167.7.2 P164 L28 # [redacted]

Li, Jing YOFC

Comment Type E Comment Status D (bucket1)

Stressed receiver sensitivity (OMAouter)c (max)

SuggestedRemedy
Stressed receiver sensitivity, each lane (OMAouter)c (max)

Proposed Response Response Status W

PROPOSED ACCEPT.

Understanding that the interface BER is the average of the BERs, these comments should be rejected and the “each lane” label not added.

Clause 169

Clause 169, delay wording (part 1)

Comments 52, 99, 100, 101

Cl 169 SC 169.4 P182 L18 # I-52

Ran, Adee Cisco Systems, Inc.

Comment Type TR Comment Status D delay wording

The text says that bit time and pause quanta are "for 800 Gigabit Ethernet".

The title of Table 169-4 has "800GBASE", and footnotes a and b start with "For 800GBASE-R".

Although 800GBASE-R is currently the only defined PHY family, it may not be so in the future; bit time and pause quanta are independent of the PHY type, so the footnotes should not be restricted to one PHY family.

Note that the addition of such footnotes started in Clause 80 in which there were two data rates, so it was required. It isn't required in clauses that define a single data rate, such as Clause 105. If it is anticipated that Clause 169 also introduces 1.6 Terabit Ethernet, then the distinction will be required; otherwise, the data rate can be removed from the footnotes.

The table title should be consistent with the text.

SuggestedRemedy

In the table title, change "800GBASE" to "800 Gigabit Ethernet".

In footnotes a and b, either change "For 800GBASE-R" to "For 800 Gigabit Ethernet", or delete these words.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Clause 169 is specific to 800 Gigabit Ethernet or 800GBASE Physical Layer implementations, so the qualifier "(800GBASE)" in the table title is not necessary. The bit times and pause_quanta, are relevant to any 800GBASE Physical Layer implementation, so the related footnotes should not be specific to 800GBASE-R. However, it is help to be unambiguous that these numbers are specific to 800GBASE in general. In Table 169-4 title delete "(800GBASE)". In Table 169-4 footnotes 1 and 2, change "800GBASE-R" to "800GBASE".

Cl 169 SC 169.4 P182 L11 # I-99

Dawe, Piers J G NVIDIA

Comment Type T Comment Status D delay wording

This text "Predictable operation of the MAC Control PAUSE operation ... concatenation of devices." looks like it was copied from 24.6 (for 100BASE-X) when a MAC bit was about 2 m long, the largest nominal reach was 2 km (1000 bits on the line) and there were repeaters. At 800G, a MAC bit is 0.25 mm long and we expect 40 km in P802.3dj (1.6e8 bits on the line, 200,000 ns). So the medium can dominate, and one should not expect all PAUSE implementations to tolerate such long links. And, no-one talks about repeaters now.

In the proposed change, the NOTE is copied from earlier clauses.

SuggestedRemedy

Update and simplify this text, e.g. "The delay limits for each sublayer are relevant to the MAC Control PAUSE operation (Clause 31, Annex 31B).

NOTE—The physical medium interconnecting two PHYs introduces additional delay in a link.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

The introduction as written is relevant and should not be pared down.

The delays specified for the backplane (KR8) and copper cable (CR8) PMDs include allocation for the medium (14 ns or ~3 m) between Physical Layers. However, for delays specified for optical PMDs (VR8, SR8, DR8, DR8-2) include only 2 m (~10 ns) allocation for the medium between Physical Layers.

Add the following sentence after Table 169-4.

"The physical medium interconnecting two optical PHYs introduces additional delay in a link."

Clause 169, delay wording (part 2)

Comments 52, 99, 100, 101

Cl 169 SC 169.4 P182 L16 # I-100
Dawe, Piers J G NVIDIA
Comment Type T Comment Status D delay wording

Instead of "colocated", Clause 45 uses terminology like "instantiated within the same package" and "The definition of the term package is vendor specific and could be a chip, module, or other similar entity." We should use language consistent with Clause 45 if it is the same concept, as it appears to be. I suppose the key here could be whether the sublayers are the responsibilities of different parties or whether the interface between the sublayers is accessible for measurement. Also, this uses the spelling "colocated" (twice) while the base document uses "co-located" (twice in 55B). Spelling should be consistent.

SuggestedRemedy

Change the criterion to say that the delay for the sublayers within a single implementation, which might be a PCB, package, chip or module, is constrained by the sum of constraints for all of the sublayers within it.
If the word "colocated" is kept, reconcile the spelling with the base document.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

The use of the word "colocated" was (incorrectly) intended to make use of established terminology for this situation. Instead, the term "in/within the same package" is used for this purpose. See examples in 802.3ck-2022 120.1.4, 135.1.4, 162.9.4, 162.9.5.1, 163.9.2, 163.9.3.1, 120F.3.1, 120G.3.1 and 802.3-2022 45.2 (many instances).

Change: "The delay for a set of colocated sublayers may be constrained by the sum of constraints for all of the colocated sublayers."

To: "The delay for a set of sublayers within the same package may be constrained by the sum of constraints for the set of sublayers."

Cl 169 SC 169.4 P182 L28 # I-101
Dawe, Piers J G NVIDIA
Comment Type TR Comment Status D delay wording

It is not clear here whether e.g. a pair of IOs forming an AUI is one PMA sublayer or two. 173.5.4 says "up to four instances of the 800GBASE-R PMA within a Physical Layer", but the relation between instance and sublayer is not given there. 120.5.4, Delay constraints, says "...up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) but it's still ambiguous. In 173.5.4, Delay constraints, "...up to four instances of the 800GBASE-R PMA", and the numbers for the PMA in Table 173-1 (not this table 169-4) apply to an instance not a sublayer.

In 173.5.3.5 we have "group of PMAs" which is not explicitly defined: maybe it means any stack of nothing but PMA-things between PMD and PCS, which could be OK for this project but may need more careful definition if an inner FEC is put between or within PMA-things.

SuggestedRemedy

Consolidate the terminology (don't use "sublayer" and instance" for the same thing), and explicitly state somewhere whether a pair of IOs forming an AUI is one PMA sublayer or two. Add cross-references as appropriate, e.g. from the AUI annexes.
Write something like "Each instance of a PMA" in the Notes column. Change the heading of the left column to "Sublayer or instance" if appropriate.

Proposed Response Response Status W


PROPOSED ACCEPT IN PRINCIPLE.

Change "Table 169-4 contains the values of maximum sublayer delay"

To "Table 169-4 contains the values of maximum delay for each instance of a sublayer"
Implement with editorial license.

Clause 169, delay wording (part 3) “Package” examples

From IEEE Std 802.3-2022...



Each MMD contains registers 5 and 6, as defined in Table 45–2. Bits read as a one in this register indicate which MMDs are instantiated **within the same package** as the MMD being accessed. Bit 5.0 is used to indicate that Clause 22 functionality has been implemented within a Clause 45 electrical interface device. Bit 6.13 indicates that Clause 22 functionality is extended using the Clause 45 electrical interface through MMD 29. The definition of the term package is vendor specific and could be a chip, module, or other similar entity.

45.2.1 PMA/PMD registers

For devices operating at 25 Gb/s or higher speeds, the PMA may be instantiated as multiple sublayers (see 83.1.4, 109.1.4, and 120.1.4 for how MMD addresses are allocated to multiple PMA sublayers for the respective speeds). A PMA sublayer that is packaged with the PMD is addressed as MMD 1. More addressable instances of PMA sublayers, each one separated from lower addressable instances, may be implemented and addressed as MMD 8, 9, 10, and 11 where MMD 8 is the closest to the PMD and MMD 11 is the furthest from the PMD. The addresses and functions of all registers in MMD 8, 9, 10, and 11 are defined identically to MMD 1, except registers m.5 and m.6 as defined in Table 45–2.

120.1.4 PMA sublayer positioning

Change item b) in the lettered list below the fourth paragraph of 120.1.4) as follows:

- b) 200GAUI-n is a physical instantiation of the connection between two adjacent 200GBASE-R PMA sublayers with the exception of the inst.IS_SIGNAL indication which is carried outside of this physically instantiated interface. 400GAUI-n is a physical instantiation of the connection between two adjacent 400GBASE-R PMA sublayers with the exception of the inst.IS_SIGNAL indication which is carried outside of this physically instantiated interface.
 - 1) As physical instantiations, these define electrical and timing specification as well as requiring a receive re-timing function.
 - 2) 200GAUI-8 is a 26.5625 GBd by 8 lane NRZ physical instantiation of the 200 Gb/s connection. 400GAUI-16 is a 26.5625 GBd by 16 lane NRZ physical instantiation of the 400 Gb/s connection.
 - 3) 200GAUI-4 is a 26.5625 GBd by 4 lane PAM4 physical instantiation of the 200 Gb/s connection. 400GAUI-8 is a 26.5625 GBd by 8 lane PAM4 physical instantiation of the 400 Gb/s connection.
 - 4) 200GAUI-2 is a 53.125 GBd by 2 lane PAM4 physical instantiation of the 200 Gb/s connection. 400GAUI-4 is a 53.125 GBd by 4 lane PAM4 physical instantiation of the 400 Gb/s connection.
 - 5) For a PHY that includes a 200GAUI-2 interface, it is recommended that the signaling rate range for a 200GAUI-8 or 200GAUI-4 PMA output that is in the same package as the PCS be limited to 26.5625 GBd ± 50 ppm.
 - 6) For a PHY that includes a 400GAUI-4 interface, it is recommended that the signaling rate range for a 400GAUI-16 or 400GAUI-8 PMA output that is in the same package as the PCS be limited to 26.5625 GBd ± 50 ppm.

Table 163–5—Summary of transmitter specifications at TP0v (continued)

Parameter	Reference	Value	Units
Transmitter waveform			
absolute value of step size for all taps (min)	162.9.4.1.4	0.005	—
absolute value of step size for all taps (max)	162.9.4.1.4	0.025	—
value at minimum state for c(-3) (max)	162.9.4.1.5	-0.06	—
value at maximum state for c(-2) (min)	162.9.4.1.5	0.12	—
value at minimum state for c(-1) (max)	162.9.4.1.5	-0.34	—
value at minimum state for c(0) (max)	162.9.4.1.5	0.5	—
value at minimum state for c(1) (max)	162.9.4.1.5	-0.2	—
Signal-to-noise-and-distortion ratio, SNDR (min)	162.9.4.6	32.5	dB
Signal-to-residual-intersymbol-interference ratio, SNR_{ISI} (min)	162.9.4.3	28	dB
Jitter (max)^c			
J_{RMS}	162.9.4.7	0.023	UI
$J_{3\sigma}$	162.9.4.7	0.106	UI
$J_{3\sigma}$	162.9.4.7	0.115	UI
Even-odd jitter, pk-pk	162.9.4.7	0.025	UI

^a For a PMD in the same package as the PCS sublayer. In other cases, the signaling rate is derived from the input to the PMD transmit function provided by the adjacent PMA sublayer.

^b Measurement uses the method described in 93.8.1.3 with the exception that the PRBS13Q test pattern is used.

^c $J_{3\sigma}$, J_{RMS} , and even-odd jitter measurements are made with a single transmit equalizer setting selected to compensate for the loss of the transmitter package and TP0 to TP0v test fixture.

Clause 169, delay wording (part 4)

Proposed results

169.4 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementations conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. Table 169–4 contains the values of maximum ~~sublayer~~ delay [for each instance of a sublayer](#) (sum of transmit and receive delays at one end of the link) in bit times as specified in 1.4 and pause_quanta as specified in 31B.2 for 800 Gigabit Ethernet. If a PHY contains an Auto-Negotiation sublayer, the delay of the Auto-Negotiation sublayer is included within the delay of the PMD and medium. The delay for a set of ~~colocated~~ sublayers [within the same package](#) may be constrained by the sum of constraints for ~~all of the colocated the set of~~ sublayers.

[The physical medium interconnecting two optical PHYs introduces additional delay in a link.](#)

<Table 169-4>

See 80.4 for the calculation of bit time per meter of fiber or electrical cable.

See 31B.3.7 for PAUSE reaction timing constraints for stations at operating speeds of 800 Gb/s.

Table 169–4—Sublayer delay constraints (~~800GBASE~~)

Sublayer	Maximum (bit time) ^a	Maximum (pause_quanta) ^b	Maximum (ns)	Notes ^c
800G MAC, RS, and MAC Control	196 608	384	245.76	See 170.1.4.
800GBASE-R PCS or 800GXS ^d	640 000	1250	800	See 172.5.
800GBASE-R PMA	36 864	72	46.08	See 173.5.4.
800GBASE-KR8 PMD	32 768	64	40.96	Includes allocation of 14 ns for one direction through backplane medium. See 163.5.
800GBASE-CR8 PMD	32 768	64	40.96	Includes allocation of 14 ns for one direction through cable medium. See 162.5.
800GBASE-VR8 PMD	32 768	64	40.96	Includes 2 m of fiber. See 167.3.1.
800GBASE-SR8 PMD	32 768	64	40.96	Includes 2 m of fiber. See 167.3.1.
800GBASE-DR8 PMD	32 768	64	40.96	Includes 2 m of fiber. See 124.3.1.
800GBASE-DR8-2 PMD	32 768	64	40.96	Includes 2 m of fiber. See 124.3.1.

^a For 800GBASE-~~R~~, 1 bit time (BT) is equal to 1.25 ps. (See 1.4.215 for the definition of bit time.)

^b For 800GBASE-~~R~~, 1 pause_quantum is equal to 640 ps. (See 31B.2 for the definition of pause_quanta.)

^c Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.

^d If an implementation includes the 800GMII Extender, the delay associated with the 800GMII Extender includes two 800GXS sublayers.

Clause 172

Scrambler, 172.2.4.5

Comments 110, 62

Comments # 110, 62 propose changes to 172.2.4.5

172.2.4.5 Scrambler

The scrambler in each flow is identical to that specified in 119.2.4.3. Although there is no requirement on the initial value of each scrambler, if an implementation sets the scrambler state to a fixed value (e.g., when reset is asserted), the two scramblers should be set to different states.

CI 172 SC 172.2.4.5 P212 L 19 # I-62
Ran, Adeo Cisco Systems, Inc.
Comment Type T Comment Status D scrambler

The recommendation to "set to different states" deserves further explanation.

SuggestedRemedy

Add the following paragraph at the end of 172.2.4.5:

NOTE---if the two scramblers have the same state and the same input (e.g., encoded remote fault signal), their outputs will be identical. With specific choices of PMA lane muxing, this can create atypical sequences on the PMA output".

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
Resolve using the response to comment #-110.

CI 172 SC 172.2.4.5 P212 L 19 # I-110
Dawe, Piers J G NVIDIA
Comment Type TR Comment Status D scrambler

"the two scramblers should be set to different states": this is too weak, and readers do not understand the importance of this. The consequence of getting it wrong is much more than the bad spectrum or correlation issues we have seen elsewhere.

SuggestedRemedy

Change should to shall or is.
Add a sentence: This is because before the link can carry traffic, the 66-bit blocks in the two flows have the same content

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
The comment proposes to make initializing the scrambler to two different states mandatory while comment #-I-62 proposes to add a note explaining the consequences of the scrambler being initialized in the same state.
Slides will be provided to address this.

Clause 173

Clause 173, signal status (part 1)

Comment i-72

Cl 173	SC 173.5.8.1	P242	L3	# I-72
Ran, Adeo	Cisco Systems, Inc.			
Comment Type	T	Comment Status	D	signal status (bucket2)

The requirement that "data is being sent on all 32 output lanes (PMA:IS_UNITDATA_0:31.indication)" is unique to this PMA (32:8); the other two PMAs set the signal status only based on data being received on the appropriate interface.

In real implementations, an indication to the PCS that data is not being received by the PMA (which may be due to lack of a link partner) would likely be separate from an indication that data is not being transmitted (essentially a local fault). Specifying in the standard that it's the same indication is not helpful for readers.

Suggested Remedy

Delete the second item in the list.

Consider converting the list to regular paragraph text as in the other two subclauses.

Proposed Response	Response Status	W
-------------------	-----------------	---

PROPOSED ACCEPT IN PRINCIPLE.
Change the SIGNAL_OK definition to the following:
"The SIGNAL_OK parameter is set to OK when data is being received on all 8 input lanes (inst:IS_UNITDATA_0:7.indication) and the SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive is set to OK, if there is a PMD immediately below the PMA. Otherwise SIGNAL OK is set to FAIL."

The highlighted statement has some important significance, but a few extra words would make it more interpretable. Also, a similar statement should be reinstated the 8:32 PMA transmit direction. See next slide.

173.5.8 Signal status

173.5.8.1 32:8 PMA signal status

The 32:8 PMA provides signal status information to the PMA client (800GBASE-R PCS or DTE 800GXS) using the PMA:IS_SIGNAL.indication(SIGNAL_OK) service interface primitive (see 173.2 and Figure 173-3):

The SIGNAL_OK parameter is set to OK when all of the following conditions are met:

- data is being received on all 8 input lanes (inst:IS_UNITDATA_0:7.indication)
- data is being sent on all 32 output lanes (PMA:IS_UNITDATA_0:31.indication)
- the SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive is set to OK, if there is a PMD immediately below the PMA

Otherwise SIGNAL OK is set to FAIL.

173.5.8.2 8:32 PMA signal status

In the transmit direction the 8:32 PMA provides signal status information to the PHY 800GXS using the PHY_XS:IS_SIGNAL.request(SIGNAL_OK) service interface primitive (see 173.3 and Figure 173-4). The SIGNAL_OK parameter is set to OK when data is being received on all 8 input lanes (PMA:IS_UNITDATA_0:7.request). Otherwise SIGNAL OK is set to FAIL.

In the receive direction the 8:32 PMA optionally provides signal status information to the client PMA by disabling one or more output lanes (PMA:IS_UNITDATA_0:7.indication) when the PHY_XS:IS_SIGNAL.indication SIGNAL_OK parameter (see 173.3 and Figure 173-4) is FAIL.

173.5.8.3 8:8 PMA signal status

In the transmit direction the 8:8 PMA optionally provides signal status to the sublayer below by disabling one or more output lanes (inst:IS_UNITDATA_0:7.request) when data is not being received on all 8 input lanes (PMA:IS_UNITDATA_0:7.request).

In the receive direction the 8:8 PMA optionally provides signal status to the client PMA by disabling one or more output lanes (PMA:IS_UNITDATA_0:7.indication) when the PMD:IS_SIGNAL.indication SIGNAL_OK parameter (see 173.3 and Figure 173-4) is FAIL or when data is not being received on all 8 input lanes (inst:IS_UNITDATA_0:7.indication).

Clause 173, signal status (part 2)

Comment i-72

173.5.8 Signal status

173.5.8.1 32:8 PMA signal status

~~In the receive direction. The~~ the 32:8 PMA provides signal status information to the PMA client (800GBASE-R PCS or DTE 800GXS) using the PMA:IS_SIGNAL.indication(SIGNAL_OK) service interface primitive (see 173.2 and Figure 173–3):The SIGNAL_OK parameter is set to OK when all of the following conditions are met:

- data is being received on all 8 input lanes (inst:IS_UNITDATA_0:7.indication)
 - the received data is being sent on all 32 output lanes (PMA:IS_UNITDATA_0:31.indication)
 - the SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive is set to OK, if there is a PMD immediately below the PMA
- Otherwise SIGNAL OK is set to FAIL.

173.5.8.2 8:32 PMA signal status

In the transmit direction the 8:32 PMA provides signal status information to the PHY 800GXS using the PHY_XS:IS_SIGNAL.request(SIGNAL_OK) service interface primitive (see 173.3 and Figure 173–4). The SIGNAL_OK parameter is set to OK when data is being received on all 8 input lanes (PMA:IS_UNITDATA_0:7.request) and the received data is being sent on all 32 output lanes (PHY_XS:IS_UNITDATA_0:31.request). Otherwise SIGNAL OK is set to FAIL.

In the receive direction the 8:32 PMA optionally provides signal status information to the client PMA by disabling one or more output lanes (PMA:IS_UNITDATA_0:7.indication) when the PHY_XS:IS_SIGNAL.indication SIGNAL_OK parameter (see 173.3 and Figure 173–4) is FAIL.

173.5.8.3 8:8 PMA signal status

In the transmit direction the 8:8 PMA optionally provides signal status to the sublayer below by disabling one or more output lanes (inst:IS_UNITDATA_0:7.request) when data is not being received on all 8 input lanes (PMA:IS_UNITDATA_0:7.request).

In the receive direction the 8:8 PMA optionally provides signal status to the client PMA by disabling one or more output lanes (PMA:IS_UNITDATA_0:7.indication) when the PMD:IS_SIGNAL.indication SIGNAL_OK parameter (see 173.3 and Figure 173–4) is FAIL or when data is not being received on all 8 input lanes (inst:IS_UNITDATA_0:7.indication).