802.3df D2.0 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.

Cross-Clause, Part 1

Normative reference, Clause 1, 124, 167
Comment 55

- Comment #55 pointed out the unusual wording in the IEC normative reference

Cl **1** *SC* **1.3** *P* **30** *L* **40** # 55

Dudek, Mike Marvell

Comment Type **E** *Comment Status* **D** *(bucket1)*

"One fibre rows" is strange.

SuggestedRemedy

Check the reference and correct to "One fibre row" unless the reference does have this in its title.

Proposed Response *Response Status* **W**

PROPOSED ACCEPT IN PRINCIPLE.
The referenced standard is currently in draft state. The title in the referenced draft has recently been corrected to say "One fibre row".
Change "rows" to "row".

- The current normative reference in D2.0 is

IEC 61754-7-4:2018, Fibre optic interconnecting devices and passive components—Fibre optic connector interfaces—Part 7-4: Type MPO connector family—One fibre rows 16 fibre wide.

Normative reference, Clause 1, 124, 167
Comment 55

In researching the correct wording of the normative reference, it was identified that the normative reference is still in a draft state and not publicly available.

To prevent the possibility of delaying the completion of 802.3df due to an unavailable normative reference, the editorial team decided to hijack the comment and remove the IEC reference and transition to an equivalent TIA reference.

Normative reference, Clause 1, 124, 167
Comment 55

In 124.11.3.3 replace

The MDI adapter or receptacle shall meet the dimensional specifications for interface 7-4-7: MPO adaptor interface – Opposed keyway configuration or interface 7-4-9: MPO active device receptacle, angled interface for 16 fibers, as defined in IEC 61754-7-4. The plug terminating the optical fiber cabling shall meet the dimensional specifications of interface 7-4-1: MPO female plug, down-angled interface for 16 fibers, as defined in IEC 61754-7-4. The MPO-16 female plug connector and MDI are structurally similar to those depicted in Figure 124–7, but with an angled end facet, 16 fibers, an offset keyway, and different pin diameters and locations.

With

The MDI receptacle shall meet the dimensional specifications for designation FOCIS 18 A-1-0, or designation FOCIS 18 R-1x16-1-8-1-1-2, as defined in ANSI/TIA-604-18-A:2018. The plug terminating the optical fiber cabling shall meet the dimensional specifications of designation FOCIS 18 P-1x16-1-8-2-1-1, as defined in ANSI/TIA-604-18-A:2018. The MPO-16 female plug connector and MDI are structurally similar to those depicted in Figure 124–7, but with an angled end facet, 16 fibers, an offset keyway, and different pin diameters and locations.

CC: Normative reference, Clause 1, 124, 167
Comment 55

In 167.10.3.4 replace

For option B, the MDI adapter or receptacle shall meet the dimensional specifications for interface 7-4-7: MPO adaptor interface – Opposed keyway configuration or interface 7-4-9: MPO active device receptacle, angled interface for 16 fibers, as defined in IEC 61754-7-4. The plug terminating the optical fiber cabling shall meet the dimensional specifications of interface 7-4-1: MPO female plug, down-angled interface for 16 fibers, as defined in IEC 61754-7-4. The MPO female plug connector and MDI are structurally similar to those depicted in Figure 167-9, but with 16 fibers, an offset keyway, and with different pin diameter and locations. The MDI connection shall meet the interface performance specifications of IEC 63267-1 for performance grade Bm/1m.

With

For option B, the MDI receptacle shall meet the dimensional specifications for designation FOCIS 18 A-1-0, or designation FOCIS 18 R-1x16-1-8-1-2-2, as defined in ANSI/TIA-604-18-A:2018. The plug terminating the optical fiber cabling shall meet the dimensional specifications of designation FOCIS 18 P-1x16-1-8-2-2-1, as defined in ANSI/TIA-604-18-A:2018. The MPO female plug connector and MDI are structurally similar to those depicted in Figure 167-10, but with 16 fibers, an angled end facet, an offset keyway, and with different pin diameter and locations. The MDI connection shall meet the interface performance specifications of IEC 63267-1 for performance grade Bm/1m.

CC: Normative reference, Clause 1, 124, 167
Comment 55

The new normative reference “ANSI/TIA-604-18-A:2018, FOCIS 18—Fiber Optic Connector Intermateability Standard—Type MPO-16.” is a normative reference in IEEE Std. 802.3-2022.

Delete the IEC reference in subclause 1.3 of 802.3df D2.1.

Cross-Clause, Part 2

CC: Optical PMD definitions, Clause 1
Comments 73, 74, 75, 76, 77 (part 1)

Cl 1 SC 1.4.135a P 30 L 49 # 73

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

400GBASE-DR4-2
The term "lane" is ambigous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
"IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding and 4-level pulse amplitude modulation over four lanes of single-mode fiber, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)"
to
IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding and 4-level pulse amplitude modulation over four wavelengths distributed over 4 single-mode fibes, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)

Proposed Response *Response Status* O

Cl 1 SC 1.4.184b P 31 L 6 # 74

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

800GBASE-DR8
The term "lane" is ambigous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
"IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of single-mode fiber, with reach up to at least 500 m. (See IEEE Std 802.3, Clause 124.)"
to
IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight wavelengths distibuted over 8 single-mode fibers with reaches up to at least 500 m. (See IEEE Std 802.3, Clause 124.)

Proposed Response *Response Status* O

Cl 1 SC 1.4.184c P 31 L 10 # 75

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

800GBASE-DR8-2
The term "lane" is ambigous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
"IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of single-mode fiber, with reach up to at least 500 m. (See IEEE Std 802.3, Clause 124.)"
to
IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight wavelengths distibuted over 8 single-mode fibers with reaches up to at least 2 km. (See IEEE Std 802.3, Clause 124.)

Proposed Response *Response Status* O

Cl 1 SC 1.4.184f P 31 L 20 # 76

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

800GBASE-SR8
With the introduction of WDM technology over MMF, the term "lane" is ambigous when discussing MMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
"IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of multimode fiber, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 167.)"
to
IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight wavelengths distributed over 8 multimode fibers, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 167.)

Proposed Response *Response Status* O

Cl 1 SC 1.4.184g P 31 L 24 # 77

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

800GBASE-VR8
With the introduction of WDM technology over MMF, the term "lane" is ambigous when discussing MMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
"IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of multimode fiber, with reach up to at least 50 m. (See IEEE Std 802.3, Clause 167.)"
to
IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight wavelengths distributed over 8 multimode fibers, with reach up to at least 50 m. (See IEEE Std 802.3, Clause 167.)

Proposed Response *Response Status* O

CC: Optical PMD definitions, Clause 1
Comments 73, 74, 75, 76, 77 (part 2)

1.4 Definitions

Insert the following new definition after 1.4.135 400GBASE-DR4:

1.4.135a 400GBASE-DR4-2: IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding and 4-level pulse amplitude modulation over four lanes of single-mode fiber, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)

1.4.184b 800GBASE-DR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of single-mode fiber, with reach up to at least 500 m. (See IEEE Std 802.3, Clause 124.)

1.4.184c 800GBASE-DR8-2: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of single-mode fiber, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)

1.4.184f 800GBASE-SR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of multimode fiber, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 167.)

1.4.184g 800GBASE-VR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of multimode fiber, with reach up to at least 50 m. (See IEEE Std 802.3, Clause 167.)

CC: Optical PMD definitions, Clauses 1 Comments 73, 74, 75, 76, 77 (part 3)

Change 1.4.135a from:
1.4.135a 400GBASE-DR4-2: IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding and 4-level pulse amplitude modulation over four lanes of single-mode fiber, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)
To:
1.4.135a 400GBASE-DR4-2: IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding with 4-level pulse amplitude modulation over four single-mode fibers, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)

Change 1.4.184b from:
".4.184b 800GBASE-DR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation on eight lanes over eight lanes of single-mode fiber, with reach up to at least 500 m. (See IEEE Std 802.3, Clause 124.)
To:
1.4.184b 800GBASE-DR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding with 4-level pulse amplitude modulation over eight single-mode fibers, with reach up to at least 500 m. (See IEEE Std 802.3, Clause 124.)

Change 1.4.184c from:
1.4.184c 800GBASE-DR8-2: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of single-mode fiber, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)
To:
1.4.184c 800GBASE-DR8-2: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding with 4-level pulse amplitude modulation over eight single-mode fibers, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 124.)

Change 1.4.184f from:
1.4.184f 800GBASE-SR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of multimode fiber, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 167.)
To:
1.4.184f 800GBASE-SR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding with 4-level pulse amplitude modulation over eight multimode fibers, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 167.)

Change 1.4.184g from:
1.4.184g 800GBASE-VR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding and 4-level pulse amplitude modulation over eight lanes of multimode fiber, with reach up to at least 50 m. (See IEEE Std 802.3, Clause 167.)
To:
1.4.184g 800GBASE-VR8: IEEE 802.3 Physical Layer specification for 800 Gb/s using 800GBASE-R encoding with 4-level pulse amplitude modulation over eight multimode fibers, with reach up to at least 50 m. (See IEEE Std 802.3, Clause 167.)

Make similar changes to definitions for the following:
(802.3-2022, 802.3db-2022) for 200GBASE-SR4/DR4/VR2/SR2
(802.3-2022, 802.3db-2022) for 400GBASE-SR16/SR8/SR4.2/DR4/VR4/SR4

Implement with editorial license.

CC: Optical PMD definitions, Clauses 116
Comments 63, 64 (part 1)

Cl 116 *SC* 116.1.3 *P* 95 *L* 38 # 63

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

400GBASE-DR4
The term "lane" is ambigous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change description to:
400 Gb/s PHY using 400GBASE-R encoding over 4 wavelengths distributed over 4 single-mode fibres, with reach up to at least 500 m (see Clause124)

Makes changes throughout document as appropriate with editorial license

Proposed Response *Response Status* O

Cl 116 *SC* 116.1.3 *P* 95 *L* 41 # 64

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

400GBASE-DR4-2
The term "lane" is ambigous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change description to:
400 Gb/s PHY using 400GBASE-R encoding over 4 wavelengths distributed over 4 single-mode fibres, with reach up to at least 2 km (see Clause124)

Makes changes throughout document as appropriate with editorial license

Proposed Response *Response Status* O

For 400GBASE-DR4 change:
400 Gb/s PHY using 400GBASE-R encoding over four lanes of single-mode fiber, with reach up to at least 500 m (see Clause 124)
To:
400 Gb/s PHY using 400GBASE-R encoding over four single-mode fibers, with reach up to at least 500 m (see Clause 124)

For 400GBASE-DR4-2 change:
400 Gb/s PHY using 400GBASE-R encoding over four lanes of single-mode fiber, with reach up to at least 2km (see Clause 124)
To:
400 Gb/s PHY using 400GBASE-R encoding over four single-mode fibers, with reach up to at least 2km (see Clause 124)

Make similar changes as follows:
Table 169-1 for 800GBASE-VR8/SR8/DR8/DR8-2.
Table 116-1 (802.3-2022, 802.3db-2022) for 200GBASE-SR4/DR4/VR2/SR2
Table 116-2 (802.3-2022, 802.3db-2022) for 400GBASE-SR16/SR8/SR4.2/DR4/VR4/SR4

Implement with editorial license.

CC: Optical PMD definitions, Clauses 116 Comments 63, 64 (part 2)

From 802.3-2022…

Table 116–1—200 Gb/s PHYs

Name	Description
200GBASE-KR4	200 Gb/s PHY using 200GBASE-R encoding over four lanes of an electrical backplane (see Clause 137)
200GBASE-CR4	200 Gb/s PHY using 200GBASE-R encoding over four lanes of twinaxial copper cable (see 1.4.559 and Clause 136)
200GBASE-SR4	200 Gb/s PHY using 200GBASE-R encoding over four lanes of multimode fiber (see Clause 138)
200GBASE-DR4	200 Gb/s PHY using 200GBASE-R encoding over four lanes of single-mode fiber, with reach up to at least 500 m (see Clause 121)
200GBASE-FR4	200 Gb/s PHY using 200GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 2 km (see Clause 122)
200GBASE-LR4	200 Gb/s PHY using 200GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 10 km (see Clause 122)
200GBASE-ER4	200 Gb/s PHY using 200GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 40 km (see Clause 122)

Table 116–2—400 Gb/s PHYs

Name	Description
400GBASE-SR16	400 Gb/s PHY using 400GBASE-R encoding over sixteen lanes of multimode fiber, with reach up to at least 100 m (see Clause 123)
400GBASE-SR8	400 Gb/s PHY using 400GBASE-R encoding over eight lanes of multimode fiber, with reach up to at least 100 m (see Clause 138)
400GBASE-SR4.2	400 Gb/s PHY using 400GBASE-R encoding over eight lanes on multimode fiber in a bidirectional WDM format, with reach up to at least 150 m (see Clause 150)
400GBASE-DR4	400 Gb/s PHY using 400GBASE-R encoding over four lanes of single-mode fiber, with reach up to at least 500 m (see Clause 124)
400GBASE-FR8	400 Gb/s PHY using 400GBASE-R encoding over eight WDM lanes on single-mode fiber, with reach up to at least 2 km (see Clause 122)
400GBASE-FR4	400 Gb/s PHY using 400GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 2 km (see Clause 151)
400GBASE-LR4-6	400 Gb/s PHY using 400GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 6 km (see Clause 151)
400GBASE-LR8	400 Gb/s PHY using 400GBASE-R encoding over eight WDM lanes on single-mode fiber, with reach up to at least 10 km (see Clause 122)
400GBASE-ER8	400 Gb/s PHY using 400GBASE-R encoding over eight WDM lanes on single-mode fiber, with reach up to at least 40 km (see Clause 122)

From 802.3db-2022…

Table 116–1—200 Gb/s PHYs

Name	Description
…	
200GBASE-CR4	200 Gb/s PHY using 200GBASE-R encoding over four lanes of twinaxial copper cable (see 1.4.559 and Clause 136)
200GBASE-VR2	200 Gb/s PHY using 200GBASE-R encoding over two lanes of multimode fiber, with reach up to at least 50 m (see Clause 167)
200GBASE-SR4	200 Gb/s PHY using 200GBASE-R encoding over four lanes of multimode fiber (see Clause 138)
200GBASE-SR2	200 Gb/s PHY using 200GBASE-R encoding over two lanes of multimode fiber, with reach up to at least 100 m (see Clause 167)
200GBASE-DR4	200 Gb/s PHY using 200GBASE-R encoding over four lanes of single-mode fiber, with reach up to at least 500 m (see Clause 121)
…	

Table 116–2—400 Gb/s PHYs

Name	Description
400GBASE-VR4	400 Gb/s PHY using 400GBASE-R encoding over four lanes of multimode fiber, with reach up to at least 50 m (see Clause 167)
400GBASE-SR16	400 Gb/s PHY using 400GBASE-R encoding over sixteen lanes of multimode fiber, with reach up to at least 100 m (see Clause 123)
400GBASE-SR8	400 Gb/s PHY using 400GBASE-R encoding over eight lanes of multimode fiber, with reach up to at least 100 m (see Clause 138)
400GBASE-SR4	400 Gb/s PHY using 400GBASE-R encoding over four lanes of multimode fiber, with reach up to at least 100 m (see Clause 167)
400GBASE-SR4.2	400 Gb/s PHY using 400GBASE-R encoding over eight lanes on multimode fiber in a bidirectional WDM format, with reach up to at least 150 m (see Clause 150)
…	

From 802.3df D2.0…

Table 169–1—800 Gb/s PHYs

Name	Description
800GBASE-KR8	800 Gb/s PHY using 800GBASE-R encoding over eight lanes of an electrical backplane (see Clause 163)
800GBASE-CR8	800 Gb/s PHY using 800GBASE-R encoding over eight lanes of twinaxial copper cable (see 1.4.559 and Clause 162)
800GBASE-VR8	800 Gb/s PHY using 800GBASE-R encoding over eight lanes of multimode fiber with reach up to at least 50 m (see Clause 167)
800GBASE-SR8	800 Gb/s PHY using 800GBASE-R encoding over eight lanes of multimode fiber with reach up to at least 100 m (see Clause 167)
800GBASE-DR8	800 Gb/s PHY using 800GBASE-R encoding over eight lanes of single-mode fiber with reach up to at least 500 m (see Clause 124)
800GBASE-DR8-2	800 Gb/s PHY using 800GBASE-R encoding over eight lanes of single-mode fiber with reach up to at least 2 km (see Clause 124)

CC: MAU type definitions, Clause 30 Comments 78, 79, 80, 61, 62 (part 1)

Cl 30 SC 30.5.1.1.2 P 35 L 14 # 61

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

800GBASE-SR8
With the introduction of WDM technology over MMF, the term "lane" is ambigous when discussing MMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 100 m as specified in Clause 167□
to
800GBASE-R PCS/PMA over 8 wavelengths distributed over 8 multi-mode fibres PMD with reach up to at least 100 m as specified in Clause 167

Makes changes throughout document as appropriate with editorial license

Proposed Response *Response Status* O

Cl 30 SC 30.5.1.1.2 P 35 L 16 # 62

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

800GBASE-VR8
With the introduction of WDM technology over MMF, the term "lane" is ambigous when discussing MMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 50 m as specified in Clause 167□
to
800GBASE-R PCS/PMA over 8 wavelengths distributed over 8 multi-mode fibres PMD with reach up to at least 50 m as specified in Clause 167

Makes changes throughout document as appropriate with editorial license

Proposed Response *Response Status* O

Cl 30 SC 30.5.1.1.2 P 34 L 51 # 78

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

400GBASE-DR4
The term "lane" is ambigous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124
to
400GBASE-R PCS/PMA over 4 wavelengths distributed over 4 single-mode fibres PMD with reach up to at least 500 m as specified in Clause 124

Proposed Response *Response Status* O

Cl 30 SC 30.5.1.1.2 P 35 L 8 # 79

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

800GBASE-DR8
The term "lane" is ambigous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124
to
800GBASE-R PCS/PMA over 8 wavelengths distributed over 8 single-mode fibres PMD with reach up to at least 500 m as specified in Clause 124

Makes changes throughout document as appropriate with editorial license

Proposed Response *Response Status* O

Cl 30 SC 30.5.1.1.2 P 35 L 10 # 80

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type E *Comment Status* X

800GBASE-DR8-2
The term "lane" is ambigous when discussing SMF -as a lane may be either a wavelength or a fiber.

SuggestedRemedy

Change
800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 124
to
800GBASE-R PCS/PMA over 8 wavelengths distributed over 8 single-mode fibres PDwith reach up to at least 2 km as specified in Clause 124

Makes changes throughout document as appropriate with editorial license

Proposed Response *Response Status* O

CC: MAU type definitions, Clause 30
Comments 78, 79, 80, 61, 62 (part 2)

From 802.3df D2.0…

30.5 Layer management for medium attachment units (MAUs)

30.5.1 MAU managed object class

30.5.1.1 MAU attributes

30.5.1.1.2 aMAUType

APPROPRIATE SYNTAX:

Change the 400GBASE-DR4 entry in "APPROPRIATE SYNTAX" in 30.5.1.1.2 as follows:

400GBASE-DR4	400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124

Insert the following new entry into "APPROPRIATE SYNTAX" in 30.5.1.1.2 after the entry for 400GBASE-DR4:

400GBASE-DR4-2	400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 124

Insert the following new entries into "APPROPRIATE SYNTAX" in 30.5.1.1.2 (as modified by IEEE Std 802.3cw-202x) after the entry for 400GBASE-ZR:

800GBASE-CR8	800GBASE-R PCS/PMA over 8-lane shielded balanced copper cable PMD as specified in Clause 162
800GBASE-DR8	800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124
800GBASE-DR8-2	800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 124
800GBASE-KR8	800GBASE-R PCS/PMA over an electrical backplane PMD as specified in Clause 163
800GBASE-SR8	800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 100 m as specified in Clause 167
800GBASE-VR8	800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 50 m as specified in Clause 167

From 802.3db-2022…

Insert the following new entry into "APPROPRIATE SYNTAX" in 30.5.1.1.2 after 200GBASE-R as follows:

200GBASE-SR2	200GBASE-R PCS/PMA over 2 lane multimode fiber PMD with reach up to at least 100 m as specified in Clause 167

Insert the following new entry into "APPROPRIATE SYNTAX" in 30.5.1.1.2 after 200GBASE-SR4 as follows:

200GBASE-VR2	200GBASE-R PCS/PMA over 2 lane multimode fiber PMD with reach up to at least 50 m as specified in Clause 167

Insert the following new entry into "APPROPRIATE SYNTAX" in 30.5.1.1.2 after 400GBASE-R as follows:

400GBASE-SR4	400GBASE-R PCS/PMA over 4 lane multimode fiber PMD with reach up to at least 100 m as specified in Clause 167

Insert the following new entry into "APPROPRIATE SYNTAX" in 30.5.1.1.2 after 400GBASE-SR16 as follows:

400GBASE-VR4	400GBASE-R PCS/PMA over 4 lane multimode fiber PMD with reach up to at least 50 m as specified in Clause 167

From 802.3-2022…

	up to at least 80 km as specified in Clause 154
200GBASE-CR4	200GBASE-R PCS/PMA over 4 lane shielded copper balanced cable PMD as specified in Clause 136
200GBASE-DR4	200GBASE-R PCS/PMA over 4-lane single-mode fiber PMD as specified in Clause 121
200GBASE-ER4	200GBASE-R PCS/PMA over 4 WDM lane single-mode fiber PMD with reach up to at least 40 km as specified in Clause 122
200GBASE-FR4	200GBASE-R PCS/PMA over 4 WDM lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 122
200GBASE-KR4	200GBASE-R PCS/PMA over an electrical backplane PMD as specified in Clause 137
200GBASE-LR4	200GBASE-R PCS/PMA over 4 WDM lane single-mode fiber PMD with reach up to at least 10 km as specified in Clause 122
200GBASE-R	Multi-lane PCS as specified in Clause 119 over undefined PMA/PMD
200GBASE-SR4	200GBASE-R PCS/PMA over 4 lane multimode fiber PMD as specified in Clause 138
400GBASE-DR4	400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD as specified in Clause 124
400GBASE-ER8	400GBASE-R PCS/PMA over 8 WDM lane single-mode fiber PMD with reach up to at least 40 km as specified in Clause 122
400GBASE-FR4	400GBASE-R PCS/PMA over 4 WDM lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 151
400GBASE-FR8	400GBASE-R PCS/PMA over 8 WDM lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 122
400GBASE-LR4-6	400GBASE-R PCS/PMA over 4 WDM lane single-mode fiber PMD with reach up to at least 6 km as specified in Clause 151
400GBASE-LR8	400GBASE-R PCS/PMA over 8 WDM lane single-mode fiber PMD with reach up to at least 10 km as specified in Clause 122
400GBASE-R	Multi-lane PCS as specified in Clause 119 over undefined PMA/PMD
400GBASE-SR4.2	400GBASE-R PCS/PMA over 8-lane multimode fiber PMD as specified in Clause 150
400GBASE-SR8	400GBASE-R PCS/PMA over 8-lane multimode fiber PMD as specified in Clause 138
400GBASE-SR16	400GBASE-R PCS/PMA over 16-lane multimode fiber PMD as specified in Clause 123
802.9a	Integrated services MAU as specified in IEEE Std 802.9a-1995 (withdrawn)

CC: MAU type definitions, Clause 30
Comments 78, 79, 80, 61, 62 (part 3)

For 800GBASE-SR8 change:
800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 100 m as specified in Clause 167
To:
800GBASE-R PCS/PMA over 8 multimode fibers PMD with reach up to at least 100 m as specified in Clause 167

For 800GBASE-VR8 change:
800GBASE-R PCS/PMA over 8-lane multimode fiber PMD with reach up to at least 50 m as specified in Clause 167
To:
800GBASE-R PCS/PMA over 8 multimode fibers PMD with reach up to at least 50 m as specified in Clause 167

For 800GBASE-DR8 change:
800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124
To:
800GBASE-R PCS/PMA over 8 single-mode fibers PMD with reach up to at least 500 m as specified in Clause 124

For 800GBASE-DR8-2 change:
800GBASE-R PCS/PMA over 8-lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 124
To:
800GBASE-R PCS/PMA over 8 single-mode fibers PMD with reach up to at least 2 km as specified in Clause 124

For 400GBASE-DR4 change, with appropriate editorial markups:
400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD with reach up to at least 500 m as specified in Clause 124
To:
400GBASE-R PCS/PMA over 4 single-mode fibers PMD with reach up to at least 500 m as specified in Clause 124

For 400GBASE-DR4-2 change, with appropriate editorial markups:
400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD with reach up to at least 2 km as specified in Clause 124
To:
400GBASE-R PCS/PMA over 4 single-mode fibers PMD with reach up to at least 2 km as specified in Clause 124

Make similar changes for the following:
- 200GBASE-SR4/DR4/VR2/SR2
- 400GBASE-SR16/SR8/SR4.2/DR4/VR4/SR4

Also, it was noted that the definition for the 800GBASE-R PCS is missing in 30.5.1.1.2.

In 30.5.1.1.2 add the following definition:
“800GBASE-R Multi-lane PCS as specified in Clause 172 over undefined PMA/PMD”

Implement with editorial license

Clause 45

TX EQ register, Clause 45 Comment 16

Cl **45** *SC* **45.2.1.135** *P* **45** *L* **29** # 16

Ran, Adee — Cisco

Comment Type **TR** *Comment Status* **D** *TX EQ register*

Registers 1.500 through 1.515 and 1.516 through 1.531 are mapped to variables that are used for transmitter equalization (local and remote) with AUI-C2C interfaces at 25 or 50 Gb/s per lane (defined in Annex 120B or 120D respectively). The transmit equalizer has 3 taps and specific sets of tap values (or ratios) with relatively coarse steps.

For 100 Gb/s per lane AUI-C2C, the transmitter equalization is controlled by a different set of variables, as defined in 120F.3.1.7 and 120F.3.2.6. The variables are different from and incompatible with those of Annex 120B/120D - the transmit equalizer has 5 taps and finer step size. The mapping of these variables to MDIO registers is also specified in these subclauses of 120F.

Therefore, Registers 1.500 through 1.531 should be made specific to the AUI-C2C at 25 or 50 Gb/s per lane.

This should have been done in 802.3ck, but if the subclauses of clause 45 are modified by this project, it should be done correctly.

If the suggested remedy is not within scope then, as an alternative, these subclauses of clause 45 should be deleted from 802.3df, since they are irrelevant for 800GAUI-n and thus out of scope.

SuggestedRemedy

In the title and body text of 45.2.1.135, change "50GAUI-n, 100GAUI-2, 200GAUI-n, and 400GAUI-n, and 800GAUI-n" to "50GAUI-n, 100GAUI-2, 200GAUI-8, 200GAUI-4, 400GAUI-16, and 400GAUI-8". Apply the same change in the title of Table 45-107.

Apply similarly in 45.2.1.136, 45.2.1.137 (including Table 45-108), and 45.2.1.138.

Proposed Response *Response Status* **W**

PROPOSED ACCEPT IN PRINCIPLE.
Some of the changes proposed in the suggested remedy are not within the scope of this project. However, some changes are warranted.
Delete the changes to the 45.2.1.135, 45.2.1.136, 45.2.1.137, and 45.2.1.138 subclauses from the 802.3df draft.
Other changes may be addressed through the 802.3 maintenance process.

Text subject of comment:

Change title and first sentence of 45.2.1.135 as follows:

45.2.1.135 50GAUI-n, 100GAUI-2, 200GAUI-n, ~~and~~ 400GAUI-n, and 800GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 register (Register 1.500)

The assignment of bits in the 50GAUI-n, 100GAUI-2, 200GAUI-n, ~~and~~ 400GAUI-n, and 800GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 register is shown in Table 45–107.

Change title of Table 45-107 as follows:

Table 45–107—50GAUI-n, 100GAUI-2, 200GAUI-n, ~~and~~ 400GAUI-n, and 800GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 register bit definitions

(similar changes in the per-field subclauses, and in 45.2.1.136 through 45.2.1.138)

Table 45–107—50GAUI-n, 100GAUI-2, 200GAUI-n, and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 register bit definitions

Bit(s)	Name	Description	R/W[a]
1.500.15	Request flag	1 = Change in equalization is requested 0 = No change in equalization is requested	RO
1.500.14:12	Post-cursor request	14 13 12 1 1 1 Reserved 1 1 0 Reserved 1 0 1 *Requested_eq_c1* = 5 (c(1) ratio –0.25) 1 0 0 *Requested_eq_c1* = 4 (c(1) ratio –0.2) 0 1 1 *Requested_eq_c1* = 3 (c(1) ratio –0.15) 0 1 0 *Requested_eq_c1* = 2 (c(1) ratio –0.1) 0 0 1 *Requested_eq_c1* = 1 (c(1) ratio –0.05) 0 0 0 *Requested_eq_c1* = 0 (c(1) ratio 0)	RO
1.500.11:10	Pre-cursor request	11 10 1 1 *Requested_eq_cm1* = 3 (c(–1) ratio –0.15) 1 0 *Requested_eq_cm1* = 2 (c(–1) ratio –0.1) 0 1 *Requested_eq_cm1* = 1 (c(–1) ratio –0.05) 0 0 *Requested_eq_cm1* = 0 (c(–1) ratio 0)	RO

Table in base standard

This table is relevant only for AUI-C2C with 25 and 50 Gb/s per lane; 100 Gb/s per lane requires a different equalizer (e.g., 3 precursor taps). So, it is irrelevant for 802.3df.
The "-n" in the table title should be changed to specific numbers to exclude the 100 Gb/s per lane AUIs defined in 802.3ck. This requires maintenance action.

TX EQ register, Clause 45
Comments 57, 25, 26, 27, 28, 29

Cl **45** *SC* **45.2.1.135.1** *P* **48** *L* **44** # 57

Dudek, Mike — Marvell

Comment Type **E** *Comment Status* **D** *TX EQ register*

800GAUI-16 is not being defined in this amendment and therefore 120D and 120B are not used. There is no need to make changes to these sections?

SuggestedRemedy

Remove the changes to sections 45.2.1.135.1 to 45.2.1.135.7 and other equivalent changes. (If 800GAUI-16 is to be included in this amendment then bring in Annex 120D and make appropriate changes (including Title changes)

Proposed Response *Response Status* **W**

PROPOSED ACCEPT IN PRINCIPLE.
Resolve using the response to comment #16.

Addressed by comment #16

Cl **45** *SC* **45.2.1.135.1** *P* **45** *L* **48** # 25

Slavick, Jeff — Broadcom

Comment Type **TR** *Comment Status* **D** *TX EQ register*

With the deletions the paragraph now reads a bit strangely and needs some word smithing. "The value of this bit indicates the value of the variable Request_flag in the lane 0 receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). This indicates whether the chip-to-chip device is issuing a request to change the remote transmitter equalization in the chip-to-chip lane 0 transmitter in the receive direction. If a lane 0 receiver in the receive direction is not present in the package, then the value returned for this bit should be zero."

SuggestedRemedy

Make it so the old paragraph is a full cross out text and replaced with the following paragraph:
"This bit indicates the state of the Request_flag variable of the lane 0 receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). When read as a one, the device is issuing a request to change the transmitter equalization of the transmitter driving lane 0 in the receive direction. If a lane 0 receiver in the receive direction is not present in the package, then the value returned for this bit should be zero."

Proposed Response *Response Status* **W**

PROPOSED REJECT.
Resolve using the response to comment #16.

This comment and comments 26 through 29 are overtaken by comment #16 (text subject of these comments will be removed)

Comments 17, 21 and 30

Request to delete redundant text in 45.2.3.25, 45.2.4.15 and 45.2.5.15

Change 45.2.3.25 as follows:

45.2.3.25 Multi-lane BASE-R PCS alignment status ~~3~~ registers 3 through 5 (Register 3.52, 3.53, 3.54)

The assignment of bits in the multi-lane BASE-R PCS alignment status ~~3~~ registers is shown in Table 45–254, Table 45–255, and Table 45–255a. All the bits in the multi-lane BASE-R PCS alignment status ~~3~~ registers are read only; a write to the multi-lane BASE-R PCS alignment status ~~3~~ registers shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status ~~3~~ registers. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status ~~3~~ registers that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

Clause 82 specifies 4 PCS lanes for 40GBASE-R and 50GBASE-R, and 20 PCS lanes for 100GBASE-R operation. Clause 119 specifies 8 PCS lanes for 200GBASE-R and 16 for 400GBASE-R operation. Clause 172 specifies 32 PCS lanes for 800GBASE-R operation.

Comments 22 and 23

Request to remove the word "optional" in the second sentence of 45.2.4.16a and 45.2.5.16a

Insert 45.2.4.16a and 45.2.4.16b after 45.2.4.16 as follows:

45.2.4.16a PHY XS RS-FEC codeword counter register (Register 4.300, 4.301, 4.302)

The assignment of bits in the PHY XS RS-FEC codeword counter register is shown in Table 45–326b. The PHY XS RS-FEC codeword counter register applies to the RS-FEC defined in Clause 172. See 172.3.5 for a definition of this optional counter. It is a 48-bit counter that counts once for each FEC codeword received when align status (4.50.12) is set to one. Its bits are reset to all zeros when the register is read by the management function or upon reset, and held at all ones in the case of overflow.

Clause 171

800GMII Extender FEC degrade, Clause 171 Comment 99

Cl 171	SC 171.2	P 190	L 46	# 99

Dawe, Piers — Nvidia

Comment Type TR *Comment Status* X

I don't see any the modification to the FEC degrade signaling in 171.5. It might be different to the 400GBASE-R PCS, but here we are comparing it to the 800GBASE-R PCS. I thought we sorted this out last time.

SuggestedRemedy

Delete "with the modified FEC degrade signaling defined in 171.5"

Proposed Response *Response Status* O

171.2 DTE 800GXS Sublayer

The DTE 800GXS shall be identical in function to the 800GBASE-R PCS (see Clause 172) with the ~~modified~~ exception that the FEC degrade signaling is defined in ~~171.5~~ 118.2.1. Figure 172–2 provides a functional block diagram.

171.3 PHY 800GXS Sublayer

The PHY 800GXS shall be identical in function to an 800GBASE-R PCS (see Clause 172) with the following exceptions:

— The PCS is inverted with the transmit function used for the receive direction and vice versa.

— The service interface signals are remapped as defined in 171.3.2 and 171.3.3.

— FEC degrade signaling is defined in ~~171.5~~118.2.2.

171.2 DTE 800GXS Sublayer

The DTE 800GXS shall be identical in function to the 800GBASE-R PCS (see Clause 172) with the modified FEC degrade signaling defined in 171.5. Figure 172–2 provides a functional block diagram.

171.5 FEC degrade signaling

The propagation of FEC degrade signaling across PCS and 800GXS sublayers is identical to FEC degrade signaling across PCS and 400GXS sublayers specified in 118.2. FEC degrade signaling is optional.

118.2 FEC Degrade

The propagation of FEC degrade signaling across PCS and XS sublayers is described in 116.6 and is based on the optional FEC degrade signaling described in Clause 119 with the changes described for the DTE XS in 118.2.1 and for the PHY XS in 118.2.2.

118.2.1 DTE XS FEC Degrade signaling

The variable tx_am_sf is set as follows:

tx_am_sf<2:0> = {FEC_degraded_SER + rx_local_degraded,0,0}

118.2.2 PHY XS FEC Degrade signaling

The variable tx_am_sf is set as follows:

tx_am_sf<2:0> = {PCS:rx_rm_degraded, PCS:FEC_degraded_SER + PCS:rx_local_degraded, 0}

Where PCS:rx_rm_degraded, PCS:FEC_degraded_SER, and PCS:rx_local_degraded are the rx_rm_degraded, FEC_degraded_SER, and rx_local_degraded variables from the adjacent PCS.

800GMII Extender FEC degrade, Clause 171 Comment 99 (cont'd)

119.2.4.4 Alignment marker mapping and insertion

The transmit alignment marker status field allows the local PCS to communicate the status of the FEC degraded feature to the remote PCS. If there is no extender sublayer between the PCS and the MAC, it is set as follows:

tx_am_sf<2:0> = {FEC_degraded_SER + rx_local_degraded,0,0}

If there is a Clause 118 extender sublayer between the PCS and the MAC, it is set as follows:

tx_am_sf<2:0> = {PHY_XS:rx_rm_degraded, PHY_XS:FEC_degraded_SER, 0}

Where PHY_XS:rx_rm_degraded and PHY_XS:FEC_degraded_SER are the rx_rm_degraded and FEC_degraded_SER variables from the adjacent PHY_XS sublayer.

See 119.2.5.3 for more information on the optional FEC degrade feature.

Two different definitions for degrade signalling are defined in Clause 119, depending on whether the PCS is directly below the MAC or below an MII Extender. In the case of a DTE XS, it is only the first of these that is applicable.

This is the degrade signaling (tx_am_sf<2:0>) when the PCS is directly below the MAC. This also applies to the DTE XS.

This is the degrade signaling (tx_am_sf<2:0>) when the PCS is directly below an MII Extender . This does not apply to a DTE XS.

Clause 172

Subclause organization, Clause 172 Comments 35, 37

Cl 172	*SC* 172.2.4.1.1	*P* 206	*L* 29	# 35

Huber, Tom — Nokia

Comment Type E *Comment Status* X

Per the style guide, a clause should not have a single subclause. It is however useful to have some separation between the general description and this new stateless encoder.

SuggestedRemedy

One option would be to make 172.2.4.1.1 a level-4 heading. The other would be insert a level 5 heading immediately after 172.2.4.1 with an innocuous title like 'Process description' and renumber the existing 172.2.4.1.1 to 172.2.4.1.2. In either case, the cross-reference at line 15 would also need to be updated.

Proposed Response *Response Status* O

Cl 172	*SC* 172.2.5.8.1	*P* 212	*L* 10	# 37

Huber, Tom — Nokia

Comment Type E *Comment Status* X

Per the style guide, a clause should not have a single subclause. It is however useful to have some separation between the general description and this new stateless encoder.

SuggestedRemedy

One option would be to make 172.2.5.8.1 a level-4 heading. The other would be insert a level 5 heading immediately after 172.2.5.8 with an innocuous title like 'Process description' and renumber the existing 172.2.5.8.1 to 172.2.5.8.2. In either case, the cross-reference at line 3 would also need to be updated.

Proposed Response *Response Status* O

802.3df D2.0 172.2.4.1 Comments 35, 37 (cont'd)

172.2.4 Transmit function

172.2.4.1 Encode, rate matching, and block distribution

The transmit PCS generates 66-bit blocks based on the TXD<63:0> and TXC<7:0> signals received from the 800GMII as specified by the transmit state diagram shown in Figure 119–14 or by the stateless encoder specified in 172.2.4.1.1. One 800GMII data transfer is encoded into one 66-bit block. The contents of each 66-bit block are contained in a vector tx_coded<65:0>. tx_coded<1:0> contains the sync header and the remainder of the bits contain the payload.

The transmit PCS may remove idle control characters or sequence ordered sets to compensate for the insertion of alignment markers. The transmit PCS may remove idle control characters or sequence ordered sets or may insert idle control characters to compensate for different clock domains on the 800GMII and the PMA service interface. See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules.

The 66-bit blocks are distributed to the two flows in an alternating fashion by the block distribution function such that the first 66-bit block is sent to flow 0, the second 66-bit block is sent to flow 1, the third 66-bit block is sent to flow 0, and subsequent 66-bit blocks continue the distribution procedure across the two flows.

172.2.4.1.1 PCS stateless encoder

An alternate method to that defined by the transmit state diagram shown in Figure 119–14 is specified in this subclause. This stateless encoder depends only on the current and preceding 72-bit 800GMII vectors. The encoder shall encode each 72-bit 800GMII vector (tx_raw) to a 66-bit block (tx_coded) according to the rules in Table 172–1. Constants LBLOCK_T and EBLOCK_T are defined in 119.2.6.2.1. Variables reset, tx_raw, and tx_coded are defined in 119.2.6.2.2. Functions T_TYPE and ENCODE, and the block types are defined in 119.2.6.2.3.

Table 172–1—PCS stateless encoder rules

reset	T_TYPE(tx_raw_{i-1})[a]	T_TYPE(tx_raw_i)[b]	Resulting tx_coded
1	any block type	any block type	LBLOCK_T
0	C or T	C	ENCODE(tx_raw_i)
0	C or T	S	ENCODE(tx_raw_i)
0	S or D	D	ENCODE(tx_raw_i)
0	S or D	T	ENCODE(tx_raw_i)
0	any combination not listed above		EBLOCK_T

[a] tx_raw_{i-1} is the 72-bit vector that immediately precedes tx_raw_i.
[b] tx_raw_i is the 72-bit vector that is being encoded.

Comment #35 *points out we cannot have a single subclause per the style guide.*
172.2.4.1 has a single subclause 172.2.4.1.1.

802.3df D2.0 172.2.5. Comments 35, 37 (cont'd)

172.2.5.8 Block collection, decode, and rate matching

The block collection reverses the block distribution done in the transmitter (see 172.2.4.1) by combining the 66-bit blocks from the two flows in an alternating fashion to form a single stream of 66-bit blocks. The first 66-bit block after the alignment marker group from flow 0 shall be followed by the first 66-bit block after the alignment marker group from flow 1.

The receive PCS decodes 66-bit blocks to produce RXD<63:0> and RXC<7:0> for transmission to the 800GMII as specified by the receive state diagram shown in Figure 119–15 or by the stateless decoder specified in 172.2.5.8.1. One 800GMII data transfer is decoded from each 66-bit block.

The receive PCS may insert idle control characters to compensate for the deletion of alignment markers. The receive PCS may remove idle control characters or sequence ordered sets or may insert idle control characters to compensate for different clock domains on the PMA service interface and the 800GMII. See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules.

172.2.5.8.1 PCS stateless decoder

An alternate method to that defined by the receive state diagram shown in Figure 119–15 is specified in this subclause. This stateless decoder depends only on the current and preceding 66-bit block. The decoder shall decode each 66-bit block (rx_coded) to a 72-bit 800GMII vector (rx_raw) according to the rules in Table 172–4. Constants LBLOCK_R and EBLOCK_R are defined in 119.2.6.2.1. Variables reset, rx_raw, and rx_coded are defined in 119.2.6.2.2. Functions R_TYPE and DECODE, and the block types are defined in 119.2.6.2.3.

Table 172–4—PCS stateless decoder rules

reset	R_TYPE(rx_coded_{i-1})[a]	R_TYPE(rx_coded_i)[b]	Resulting rx_raw
1	any block type	any block type	LBLOCK_R
0	any block type	E	EBLOCK_R
0	E	any block type	EBLOCK_R
0	any combination not listed above		DECODE(rx_coded_i)

[a] rx_coded_{i-1} is the 66-bit block that immediately precedes rx_coded_i.
[b] rx_coded_i is the 66-bit block that is being decoded.

Comment #37 *points out we cannot have a single subclause per the style guide.*
172.2.5.8 has a single subclause 172.2.5.8.1.

Comments 35, 37
Proposed response approach

For Comment 35

- Replace 172.2.4.1 by the following 3 subclauses :
  - 172.2.4.1 Encode
    - 172.2.4.1.1 State machine encoder
    - 172.2.4.1.2 Stateless encoder
  - 172.2.4.2 Rate matching
  - 172.2.4.3 Block distribution

For Comment 37

- Replace 172.2.5.8 by the following 3 subclauses :
  - 172.2.5.8 Block collection
  - 172.2.5.9 Decode
    - 172.2.5.9.1 State machine decoder
    - 172.2.5.9.2 Stateless decoder
  - 172.2.5.10 Rate matching

See detailed response in next two slides

background

From 802.3df D2.0…

172.2.6.2.2 Variables

The variables are the same as those specified in 119.2.6.2.2 with the following exceptions:

— 200GMII/400GMII is replaced with 800GMII.

From 802.3-2022, 119.2.6.2.2…

tx_coded<65:0>
Vector containing the output from the 64B/66B encoder. The format for this vector is shown in Figure 82–5. The leftmost bit in the figure is tx_coded<0> and the rightmost bit is tx_coded<65>.

tx_raw<71:0>
Vector containing one 200GMII/400GMII transfer. TXC<0> through TXC<7> are placed in tx_raw<0> through tx_raw<7>, respectively. TXD<0> through TXD<63> are placed in tx_raw<8> through tx_raw<71>, respectively.

rx_coded<65:0>
Vector containing the input to the 64B/66B decoder. The format for this vector is shown in Figure 82–5. The leftmost bit in the figure is rx_coded<0> and the rightmost bit is rx_coded<65>.

rx_raw<71:0>
Vector containing one 200GMII/400GMII transfer. RXC<0> through RXC<7> are from rx_raw<0> through rx_raw<7>, respectively. RXD<0> through RXD<63> are from rx_raw<8> through rx_raw<71>, respectively.

Comment 35 : Proposed replacement text

172.2.4.1 Encode

The transmit PCS generates 66-bit blocks based on the TXD<63:0> and TXC<7:0> signals received from the 800GMII. Each 800GMII transfer is encoded into one 66-bit block. The contents of each 66-bit block are contained in a vector tx_coded<65:0> with tx_coded<1:0> containing the sync header and the remainder of the bits the payload.

The transmit PCS shall use the encoding method defined in either 172.2.4.1 or 172.2.4.2.

172.2.4.1.1 State diagram encoder

The state diagram encoder generates 66-bit blocks based on the TXD<63:0> and TXC<7:0> signals received from the 800GMII as specified by the transmit state diagram shown in Figure 119–14.

172.2.4.1.2 Stateless encoder

The stateless encoder generates 66-bit blocks based only on the current and preceding 800GMII transfers. Each 800GMII transfer is mapped into a 72-bit vector tx_raw<71:0> (see 172.2.6.2.2). The encoder shall encode each tx_raw<71:0> to a 66-bit block tx_coded<65:0> according to the rules in Table 172–1. Constants LBLOCK_T and EBLOCK_T are defined in 119.2.6.2.1. Variables reset, tx_raw, and tx_coded are defined in 172.2.6.2.2. Functions T_TYPE and ENCODE, and the block types are defined in 172.2.6.2.3.

Table 172-1

172.2.4.2 Rate matching

The transmit PCS may remove idle control characters or sequence ordered sets to compensate for the insertion of alignment markers. The transmit PCS may remove idle control characters or sequence ordered sets or may insert idle control characters to compensate for different clock domains on the 800GMII and the PMA service interface. See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules.

172.2.4.3 Block distribution

The 66-bit blocks are distributed to the two flows in an alternating fashion by the block distribution function such that the first 66-bit block is sent to flow 0, the second 66-bit block is sent to flow 1, the third 66-bit block is sent to flow 0, and subsequent 66-bit blocks continue the distribution procedure across the two flows.

From 802.3df D2.0 (for comparison)…

172.2.4.1 Encode, rate matching, and block distribution

The transmit PCS generates 66-bit blocks based on the TXD<63:0> and TXC<7:0> signals received from the 800GMII as specified by the transmit state diagram shown in Figure 119–14 or by the stateless encoder specified in 172.2.4.1.1. One 800GMII data transfer is encoded into one 66-bit block. The contents of each 66-bit block are contained in a vector tx_coded<65:0>. tx_coded<1:0> contains the sync header and the remainder of the bits contain the payload.

The transmit PCS may remove idle control characters or sequence ordered sets to compensate for the insertion of alignment markers. The transmit PCS may remove idle control characters or sequence ordered sets or may insert idle control characters to compensate for different clock domains on the 800GMII and the PMA service interface. See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules.

The 66-bit blocks are distributed to the two flows in an alternating fashion by the block distribution function such that the first 66-bit block is sent to flow 0, the second 66-bit block is sent to flow 1, the third 66-bit block is sent to flow 0, and subsequent 66-bit blocks continue the distribution procedure across the two flows.

172.2.4.1.1 PCS stateless encoder

An alternate method to that defined by the transmit state diagram shown in Figure 119–14 is specified in this subclause. This stateless encoder depends only on the current and preceding 72-bit 800GMII vectors. The encoder shall encode each 72-bit 800GMII vector (tx_raw) to a 66-bit block (tx_coded) according to the rules in Table 172–1. Constants LBLOCK_T and EBLOCK_T are defined in 119.2.6.2.1. Variables reset, tx_raw, and tx_coded are defined in 119.2.6.2.2. Functions T_TYPE and ENCODE, and the block types are defined in 119.2.6.2.3.

Table 172–1—PCS stateless encoder rules

reset	T_TYPE(tx_raw_{i-1})[a]	T_TYPE(tx_raw_i)[b]	Resulting tx_coded
1	any block type	any block type	LBLOCK_T
0	C or T	C	ENCODE(tx_raw_i)
0	C or T	S	ENCODE(tx_raw_i)
0	S or D	D	ENCODE(tx_raw_i)
0	S or D	T	ENCODE(tx_raw_i)
0	any combination not listed above		EBLOCK_T

[a] tx_raw_{i-1} is the 72-bit vector that immediately precedes tx_raw_i.
[b] tx_raw_i is the 72-bit vector that is being encoded.

Comment 37 : Proposed replacement text

172.2.5.8 Block collection

The block collection reverses the block distribution done in the transmitter (see 172.2.4.3) by combining the 66-bit blocks from the two flows in an alternating fashion to form a single stream of 66-bit blocks. The first 66-bit block after the alignment marker group from flow 0 shall be followed by the first 66-bit block after the alignment marker group from flow 1.

172.2.5.9 Decode

The receive PCS decodes 66-bit blocks to produce RXD<63:0> and RXC<7:0> for transmission to the 800GMII. One 800GMII transfer is decoded from each 66-bit block.

The receive PCS shall use the decoding method defined in either 172.2.5.9.1 or 172.2.5.9.2.

172.2.5.9.1 State diagram decoder

The state diagram decoder decodes 66-bit blocks to produce RXD<63:0> and RXC<7:0> for transmission to the 800GMII as specified by the receive state diagram shown in Figure 119–15.

172.2.5.9.2 Stateless decoder

The stateless decoder generates 800GMII transfers based only on the current and preceding 66-bit blocks. The decoder shall decode each 66-bit block rx_coded<65:0> to a 72-bit vector rx_raw<71:0> (see 172.2.6.2.2) according to the rules in Table 172–4. Constants LBLOCK_R and EBLOCK_R are defined in 172.2.6.2.1. Variables reset, rx_raw, and rx_coded are defined in 172.2.6.2.2. Functions R_TYPE and DECODE, and the block types are defined in 172.2.6.2.3.

Table 172-4

172.2.5.10 Rate matching

The receive PCS may insert idle control characters to compensate for the deletion of alignment markers. The receive PCS may remove idle control characters or sequence ordered sets or may insert idle control characters to compensate for different clock domains on the PMA service interface and the 800GMII. See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules.

From 802.3df D2.0 (for comparison)…

172.2.5.8 Block collection, decode, and rate matching

The block collection reverses the block distribution done in the transmitter (see 172.2.4.1) by combining the 66-bit blocks from the two flows in an alternating fashion to form a single stream of 66-bit blocks. The first 66-bit block after the alignment marker group from flow 0 shall be followed by the first 66-bit block after the alignment marker group from flow 1.

The receive PCS decodes 66-bit blocks to produce RXD<63:0> and RXC<7:0> for transmission to the 800GMII as specified by the receive state diagram shown in Figure 119–15 or by the stateless decoder specified in 172.2.5.8.1. One 800GMII data transfer is decoded from each 66-bit block.

The receive PCS may insert idle control characters to compensate for the deletion of alignment markers. The receive PCS may remove idle control characters or sequence ordered sets or may insert idle control characters to compensate for different clock domains on the PMA service interface and the 800GMII. See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules.

172.2.5.8.1 PCS stateless decoder

An alternate method to that defined by the receive state diagram shown in Figure 119–15 is specified in this subclause. This stateless decoder depends only on the current and preceding 66-bit block. The decoder shall decode each 66-bit block (rx_coded) to a 72-bit 800GMII vector (rx_raw) according to the rules in Table 172–4. Constants LBLOCK_R and EBLOCK_R are defined in 119.2.6.2.1. Variables reset, rx_raw, and rx_coded are defined in 119.2.6.2.2. Functions R_TYPE and DECODE, and the block types are defined in 119.2.6.2.3.

Table 172–4—PCS stateless decoder rules

reset	R_TYPE(rx_coded_{j-1})[a]	R_TYPE(rx_coded_j)[b]	Resulting rx_raw
1	any block type	any block type	LBLOCK_R
0	any block type	E	EBLOCK_R
0	E	any block type	EBLOCK_R
0	any combination not listed above		DECODE(rx_coded_j)

[a] rx_coded_{j-1} is the 66-bit block that immediately precedes rx_coded_j.
[b] rx_coded_j is the 66-bit block that is being decoded.

Clause 173

PMA bit multiplexing, Clause 173 Comment 6

Cl 173	*SC* 173.4.2.1	*P* 232	*L* 7	# 6

Nicholl, Shawn AMD

Comment Type TR *Comment Status* X

In 173.4 "Functions within the PMA" the text references the undefined term "restricted bit-multiplexing" and says to "see 173.4.2.1". However, the word "restricted" does not appear in 173.4.2.1 "32:8 PMA bit-level multiplexing".

SuggestedRemedy

Propose to update the text in 173.4.2.1 "32:8 PMA bit-level multiplexing". Replace "The multiplexing function has an additional constraint ..." with "This restricted bit-multiplexing function has an additional constraint ..."

Similarly, propose to update the text in 173.4.2.2 "8:32 PMA bit-level multiplexing". Replace "The multiplexing function has an additional constraint ..." with "This restricted bit-multiplexing function has an additional constraint ..."

Likewise, propose to update the text in 173.4.2.3 "8:8 PMA bit-level multiplexing". Replace "The 4 PCSLs received on an input lane shall be mapped ..." with "This restricted bit-multiplexing function has an additional constraint that the 4 PCSLs received on an input lane shall be mapped ..."

Proposed Response *Response Status* O

Change 173.4.2.1 (page 232, line 12) from:
The bit-level multiplexing function is identical to that specified in 120.5.2, with the following exceptions:
To:
The restricted bit-level multiplexing function is identical to that specified in 120.5.2, with the following exceptions:

Change 173.4.2.2 (page 232, line 38) from:
The bit-level multiplexing function is identical to that specified in 120.5.2, with the following exceptions:
To:
The restricted bit-level multiplexing function is identical to that specified in 120.5.2, with the following exceptions:

Change 173.4.2.3 (page 233, line 4) from:
In both the transmit and receive directions, the bit-level multiplexing function is identical to that specified in 120.5.2, with the following exceptions:
To:
In both the transmit and receive directions, the restricted bit-level multiplexing function is identical to that specified in 120.5.2, with the following exceptions:

PMA location, Clause 173 comment 68

Cl 00 SC 0 P 225 L # 68

D'Ambrosia, John Futurewei, US Subsidiary of Huawei

Comment Type TR *Comment Status* X

As noted in Tables 169-2 and 169-3, 800G AUI variants are optional for both 800G copper and optical PHY types, which means you could have an 800GAUI-8 in the PHY as well as in the extender. This means you would PMA (32:8) and PMA (8:32) to support AUIs - not PMA 8:8
See Fig 173A-4 as example that a PMA (32:8) is called out for connecting to a 800GAUI-8

SuggestedRemedy

The statements regarding the 32:8 and 8:32 PMAs should reflect being present to support 800GAUIs which may not just be in the Extender as currently stated.

Proposed Response *Response Status* O

173.1.4 text is replicated here:

A 32:8 PMA sublayer is required in a PHY or an 800GMII Extender and is located immediately below either the 800GBASE-R PCS sublayer or the DTE 800GXS sublayer, respectively.

An 8:32 PMA sublayer is required in an 800GMII Extender and is located immediately above the PHY 800GXS sublayer.

An 8:8 PMA sublayer is required in a PHY with one or two physical instantiations of the PMA service interface (800GAUI-8) or in an 800GMII Extender with two physical instantiations of the PMA service interface (800GAUI-8).

Update text of 173.1.4 (above) as follows (with editorial licence):

32:8 PMA is required as follows:
-- in an 800GBASE-R PHY immediately below the 800GBASE-R PCS and immediately above either an 800GAUI-8 or 800GBASE-8 PMD
-- in an 800GMII Extender immediately below the DTE 800GXS and immediately above an 800GAUI-8

8:32 PMA is required as follows:
-- in an 800GMII Extender, located immediately below an 800GAUI-8 and immediately above the PHY 800GXS

8:8 PMA is required as follows:
-- in an 800GMII Extender with two 800GAUI-8, located immediately between the two 800GAUI-8
-- in an 800GBASE-R PHY, located immediately below an 800GAUI-8 and immediately above either an 800GAUI-8 or 800GBASE-8 PMD

Some of the extra details are not too important for this generation, but once we add the 200 Gb/s per lane AUIs there will be some ambiguity.
Splitting the details makes it clear that that you might have an 8:8 and 32:8 in both the extender and the PHY.

Add three separate subclauses under 173.4, to contain the functional overview and block diagram for the three individual PMAs, 32:8 PMA, 8:32 PMA and 8:8 PMA.