
**Information technology —
Telecommunications and information
exchange between systems — Local
and metropolitan area networks —
Specific requirements —**

**Part 3:
Standard for Ethernet**

**AMENDMENT 10: Media access control
parameters, physical layers, and
management parameters for 200 Gb/s
and 400 Gb/s operation**

*Technologies de l'information — Télécommunications et échange
d'information entre systèmes — Réseaux locaux et métropolitains —
Prescriptions spécifiques —*

Partie 3: Norme pour Ethernet

*AMENDEMENT 10: Paramètres de commande d'accès aux supports,
couches physiques et paramètres de gestion pour le fonctionnement
en 200 Gb/s et 400 Gb/s*





COPYRIGHT PROTECTED DOCUMENT

© IEEE 2017

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO or IEEE at the respective address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Fax: +41 22 749 09 47
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Institute of Electrical and Electronics Engineers, Inc
3 Park Avenue, New York
NY 10016-5997, USA

Email: stds.ipr@ieee.org
Website: www.ieee.org

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of ISO documents should be noted (see www.iso.org/directives).

IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. The IEEE develops its standards through a consensus development process, approved by the American National Standards Institute, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Institute and serve without compensation. While the IEEE administers the process and establishes rules to promote fairness in the consensus development process, the IEEE does not independently evaluate, test, or verify the accuracy of any of the information contained in its standards.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT), see www.iso.org/iso/foreword.html.

ISO/IEC/IEEE 8802-3:2017/Amd.10 was prepared by the LAN/MAN of the IEEE Computer Society (as IEEE Std 802.3bs-2017) and drafted in accordance with its editorial rules. It was adopted, under the "fast-track procedure" defined in the Partner Standards Development Organization cooperation agreement between ISO and IEEE, by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC6, *Telecommunications and information exchange between systems*.

A list of all parts in the ISO/IEC/IEEE 8802 series can be found on the ISO website.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/AMD10:2019

IEEE Std 802.3bs™-2017

(Amendment to
IEEE Std 802.3™-2015
as amended by
IEEE Std 802.3bw™-2015,
IEEE Std 802.3by™-2016,
IEEE Std 802.3bq™-2016,
IEEE Std 802.3bp™-2016,
IEEE Std 802.3br™-2016,
IEEE Std 802.3bn™-2016,
IEEE Std 802.3bz™-2016,
IEEE Std 802.3bu™-2016,
IEEE Std 802.3bv™-2017, and
IEEE Std 802.3-2015/Cor 1-2017)

IEEE Standard for Ethernet

Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters for 200 Gb/s and 400 Gb/s Operation

**LAN/MAN Standards Committee
of the
IEEE Computer Society**

**Approved 6 December 2017
of the
IEEE-SA Standards Board**

Abstract: Clause 116 through Clause 124 and Annex 119A through Annex 120E are added by this amendment to IEEE Std 802.3-2015. This amendment includes IEEE 802.3 Media Access Control (MAC) parameters, Physical Layer specifications, and management parameters for the transfer of IEEE 802.3 format frames at 200 Gb/s and 400 Gb/s.

Keywords: 200 Gb/s Ethernet, 200GAUI-4, 200GAUI-8, 200GBASE-DR4, 200GBASE-FR4, 200GBASE-LR4, 200GBASE-R, 200GMII, 200GXS, 400 Gb/s Ethernet, 400GAUI-8, 400GAUI-16, 400GBASE-DR4, 400GBASE-FR8, 400GBASE-LR8, 400GBASE-SR16, 400GBASE-R, 400GMII, 400GXS, EEE, Energy Efficient Ethernet, Ethernet, FEC, forward error correction, IEEE 802.3™, IEEE 802.3bs™, MMF, PAM4, Physical Medium Dependent sublayer, PMD, SMF

The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

Copyright © 2017 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 12 December 2017. Printed in the United States of America.

IEEE and 802 are registered trademarks in the U.S. Patent & Trademark Office, owned by The Institute of Electrical and Electronics Engineers, Incorporated.

PDF: ISBN 978-1-5044-4450-7 STD22871
Print: ISBN 978-1-5044-4451-4 STDPD22871

IEEE prohibits discrimination, harassment and bullying. For more information, visit <http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html>.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher.

Important Notices and Disclaimers Concerning IEEE Standards Documents

IEEE documents are made available for use subject to important notices and legal disclaimers. These notices and disclaimers, or a reference to this page, appear in all standards and may be found under the heading “Important Notices and Disclaimers Concerning IEEE Standards Documents.” They can also be obtained on request from IEEE or viewed at <http://standards.ieee.org/IPR/disclaimers.html>.

Notice and Disclaimer of Liability Concerning the Use of IEEE Standards Documents

IEEE Standards documents (standards, recommended practices, and guides), both full-use and trial-use, are developed within IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (“IEEE-SA”) Standards Board. IEEE (“the Institute”) develops its standards through a consensus development process, approved by the American National Standards Institute (“ANSI”), which brings together volunteers representing varied viewpoints and interests to achieve the final product. IEEE Standards documents are developed through scientific, academic, and industry-based technical working groups. Volunteers in IEEE working groups are not necessarily members of the Institute and participate without compensation from IEEE. While IEEE administers the process and establishes rules to promote fairness in the consensus development process, IEEE does not independently evaluate, test, or verify the accuracy of any of the information or the soundness of any judgments contained in its standards.

IEEE Standards do not guarantee or ensure safety, security, health, or environmental protection, or ensure against interference with or from other devices or networks. Implementers and users of IEEE Standards documents are responsible for determining and complying with all appropriate safety, security, environmental, health, and interference protection practices and all applicable laws and regulations.

IEEE does not warrant or represent the accuracy or content of the material contained in its standards, and expressly disclaims all warranties (express, implied and statutory) not included in this or any other document relating to the standard, including, but not limited to, the warranties of: merchantability; fitness for a particular purpose; non-infringement; and quality, accuracy, effectiveness, currency, or completeness of material. In addition, IEEE disclaims any and all conditions relating to: results; and workmanlike effort. IEEE standards documents are supplied “AS IS” and “WITH ALL FAULTS.”

Use of an IEEE standard is wholly voluntary. The existence of an IEEE standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard.

In publishing and making its standards available, IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity nor is IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing any IEEE Standards document, should rely upon his or her own independent judgment in the exercise of reasonable care in any given circumstances or, as appropriate, seek the advice of a competent professional in determining the appropriateness of a given IEEE standard.

IN NO EVENT SHALL IEEE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO: PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE PUBLICATION, USE OF, OR RELIANCE UPON ANY STANDARD, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE AND REGARDLESS OF WHETHER SUCH DAMAGE WAS FORESEEABLE.

Translations

The IEEE consensus development process involves the review of documents in English only. In the event that an IEEE standard is translated, only the English version published by IEEE should be considered the approved IEEE standard.

Official statements

A statement, written or oral, that is not processed in accordance with the IEEE-SA Standards Board Operations Manual shall not be considered or inferred to be the official position of IEEE or any of its committees and shall not be considered to be, or be relied upon as, a formal position of IEEE. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position of IEEE.

Comments on standards

Comments for revision of IEEE Standards documents are welcome from any interested party, regardless of membership affiliation with IEEE. However, IEEE does not provide consulting information or advice pertaining to IEEE Standards documents. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Since IEEE standards represent a consensus of concerned interests, it is important that any responses to comments and questions also receive the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to comments or questions except in those cases where the matter has previously been addressed. For the same reason, IEEE does not respond to interpretation requests. Any person who would like to participate in revisions to an IEEE standard is welcome to join the relevant IEEE working group.

Comments on standards should be submitted to the following address:

Secretary, IEEE-SA Standards Board
445 Hoes Lane
Piscataway, NJ 08854 USA

Laws and regulations

Users of IEEE Standards documents should consult all applicable laws and regulations. Compliance with the provisions of any IEEE Standards document does not imply compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

Copyrights

IEEE draft and approved standards are copyrighted by IEEE under U.S. and international copyright laws. They are made available by IEEE and are adopted for a wide variety of both public and private uses. These include both use, by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making these documents available for use and adoption by public authorities and private users, IEEE does not waive any rights in copyright to the documents.

Photocopies

Subject to payment of the appropriate fee, IEEE will grant users a limited, non-exclusive license to photocopy portions of any individual standard for company or organizational internal use or individual, non-commercial use only. To arrange for payment of licensing fees, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; +1 978 750 8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Updating of IEEE Standards documents

Users of IEEE Standards documents should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect.

Every IEEE standard is subjected to review at least every ten years. When a document is more than ten years old and has not undergone a revision process, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE standard.

In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit the IEEE Xplore at <http://ieeexplore.ieee.org/> or contact IEEE at the address listed previously. For more information about the IEEE-SA or IEEE's standards development process, visit the IEEE-SA Website at <http://standards.ieee.org>.

Errata

Errata, if any, for all IEEE standards can be accessed on the IEEE-SA Website at the following URL: <http://standards.ieee.org/findstds/errata/index.html>. Users are encouraged to check this URL for errata periodically.

Patents

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken by the IEEE with respect to the existence or validity of any patent rights in connection therewith. If a patent holder or patent applicant has filed a statement of assurance via an Accepted Letter of Assurance, then the statement is listed on the IEEE-SA Website at <http://standards.ieee.org/about/sasb/patcom/patents.html>. Letters of Assurance may indicate whether the Submitter is willing or unwilling to grant licenses under patent rights without compensation or under reasonable rates, with reasonable terms and conditions that are demonstrably free of any unfair discrimination to applicants desiring to obtain such licenses.

Essential Patent Claims may exist for which a Letter of Assurance has not been received. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patents Claims, or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this standard are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

Participants

The following individuals were officers and members of the IEEE 802.3 working group at the beginning of the IEEE P802.3bs working group ballot. Individuals may have not voted or may have voted for approval, disapproval, or abstention on this amendment.

David J. Law, *IEEE 802.3 Working Group Chair*
Adam Healey, *IEEE 802.3 Working Group Vice-Chair*
Pete Anslow, *IEEE 802.3 Working Group Secretary*
Steven B. Carlson, *IEEE 802.3 Working Group Executive Secretary*
Valerie Maguire, *IEEE 802.3 Working Group Treasurer*

John D'Ambrosia, *IEEE P802.3bs 200 Gb/s and 400 Gb/s Ethernet Task Force Chair*
Pete Anslow, *IEEE P802.3bs 200 Gb/s and 400 Gb/s Ethernet Task Force Editor-in-Chief*
Mark Gustlin, *IEEE P802.3bs 200 Gb/s and 400 Gb/s Ethernet Task Force Editor*
for Clauses 117, 118, 119, Annex 119A
Steve Trowbridge, *IEEE P802.3bs 200 Gb/s and 400 Gb/s Ethernet Task Force Editor*
for Clause 120, Annex 120A
Peter Stassar, *IEEE P802.3bs 200 Gb/s and 400 Gb/s Ethernet Task Force Editor*
for Clauses 121, 122, 124
Jonathan King,¹ *IEEE P802.3bs 200 Gb/s and 400 Gb/s Ethernet Task Force Editor*
for Clause 123
Andre Szczepanek, *IEEE P802.3bs 200 Gb/s and 400 Gb/s Ethernet Task Force Editor*
for Annexes 120D, 120E

Justin Abbott
David Abramson
Shadi Abughazaleh
Mohammad Ahmed
Eric Baden
Amrik Bains
Thananya Baldwin
Denis Beaudoin
Christian Beia
Michael Bennett
Vipul Bhatt
William Bliss
Brad Booth
Martin Bouda
Ralf-Peter Braun
Theodore Brillhart
Paul Brooks
Alan Brown
Matthew Brown
Chris Bullock
Jairo Bustos Heredia
Adrian Butter
Francesco Caggioni
Anthony Calbone
Clark Carty
Craig Chabot
Geoffrey Chacon Simon
Mandeep Chadha

David Chalupsky
Jacky Chang
Xin Chang
Ahmad Chini
Keng Hua Chuang
Christopher R. Cole
Yair Darshan
Piers Dawe
Fred Dawson
Wael Diab
Eric DiBiaso
John Dillard
Daniel Dillow
Thuyen Dinh
Curtis Donahue
Dan Dove
Mike Dudek
David Dwelley
Frank Effenberger
Hesham Elbakoury
David Estes
John Ewen
Ramin Farjad
Shahar Feldman
James Fife
Alan Flatman
Matthias Fritsche
Richard Frosch

Andrew Gardner
Claude Gauthier
Ali Ghiasi
Joel Goergen
Volker Goetzfried
Zhigang Gong
Steven Gorshe
Robert Grow
Marek Hajduczenia
Takehiro Hayashi
Yasuo Hidaka
Rita Horner
Bernd Hormmeyer
Victor Hou
Yasuhiro Hyakutake
Hideki Isono
Tom Issenhuth
Kenneth Jackson
Andrew Jimenez
Chad Jones
Peter Jones
Manabu Kagami
Upen Kareti
Keisuke Kawahara
Yasuaki Kawatsu
Michael Kelsen
Scott Kipp
Michael Klempa

¹Not a member of the IEEE 802.3 working group at the beginning of the working group ballot.

Curtis Knittle
Shigeru Kobayashi
Daniel Koehler
Paul Kolesar
Tom Kolze
Glen Kramer
Hans Lackner
Jeffrey Lapak
Mark Laubach
Han Hyub Lee
David Lewis
Jon Lewis
Mike Peng Li
Jane Lim
Dekun Liu
Hai-Feng Liu
William Lo
Miklos Lukacs
Kent Lusted
Jeffery Maki
David Malicoat
Yonatan Malkiman
Arthur Marris
Takeo Masuda
Erdem Matoglu
Naoki Matsuda
Mick McCarthy
Brett McClellan
Thomas McDermott
John McDonough
Larry McMillan
Richard Mei
Richard Mellitz
Bryan Moffitt
Ardeshir Mohammadian
Paul Mooney
Dale Murray
Henry Muyschondt
James Nadolny

Edward Nakamoto
Gary Nicholl
Kevin Noll
Mark Nowell
David Ofelt
Tom Palkert
Hui Pan
Sesha Panguluri
Vasu Parthasarathy
Petar Pepeljugoski
Gerald Pepper
Ruben Perez De Aranda Alonso
Michael Peters
Phong Pham
Jean Picard
William Powell
Rick Rabinovich
Adee Ran
Alon Regev
Duane Remein
Victor Renteria
Christopher Roth
Salvatore Rotolo
Toshiaki Sakai
Jorge Salinger
Sam Sambasivan
Edward Sayre
Dieter Schickelanz
Fred Schindler
Hossein Sedarat
Naoshi Serizawa
Masood Shariff
Ramin Shirani
Tom Skaar
Jeff Slavick
Daniel Smith
Scott Sommers
Yoshiaki Sone
Tom Souvignier
Heath Stewart

Robert Stone
David Stover
Junqing Sun
Ken-Ichi Suzuki
Steve Swanson
William Szeto
Bharat Tailor
Takayuki Tajima
Satoshi Takahashi
Kohichi Tamura
Brian Teipen
Geoffrey Thompson
Pirooz Tooyserkani
Albert Tretter
Yoshihiro Tsukamoto
Ed Ulrichs
Alexander Umnov
Sterling A. Vaden
Stefano Valle
Paul Vanderlaan
Robert Wagner
Dylan Walker
Haifei Wang
Roy Wang
Tongtong Wang
Xinyuan Wang
Matthias Wendt
Oded Wertheim
Natalie Wienckowski
Ludwig Winkel
Peter Wu
Dayin Xu
Yu Xu
Jun Yi
Lennart Yseboodt
Hayato Yuki
Andrew Zambell
Yan Zhuang
George Zimmerman

The following members of the individual balloting committee voted on this amendment. Balloters may have voted for approval, disapproval, or abstention.

| | | |
|----------------------|-------------------|---------------------|
| Mohammad Ahmed | Werner Hoelzl | Glenn Parsons |
| Thomas Alexander | Rita Horner | Bansi Patel |
| Pete Anslow | Noriyuki Ikeuchi | Arumugam Paventhan |
| Butch Anton | Sergiu Iordanescu | Michael Peters |
| Stefan Aust | Osamu Ishida | David Piehler |
| Eric Baden | Atsushi Ito | Rick Pimpinella |
| Saman Behtash | Raj Jain | Adee Ran |
| Ralf-Peter Braun | SangKwon Jeong | Alon Regev |
| Nancy Bravin | Piotr Karocki | Maximilian Riegel |
| Theodore Brillhart | Stuart Kerry | Robert Robinson |
| Matthew Brown | Yongbum Kim | Toshiaki Sakai |
| Jairo Bustos Heredia | Jonathan King | Osman Sakr |
| William Byrd | Paul Kolesar | Dieter Schickelanz |
| Steven B. Carlson | Mark Laubach | Takeshi Shimizu |
| Juan Carreon | David J. Law | Kapil Shrikhande |
| David Chalupsky | June Hee Lee | Jeff Slavick |
| Boung Wook Cho | David Lewis | Thomas Starai |
| Keith Chow | Jon Lewis | Peter Stassar |
| Keng Hua Chuang | Mike-Peng Li | Rene Struik |
| Charles Cook | Arthur H. Light | Walter Struppler |
| Eugene Dai | Kent Lusted | Mitsutoshi Sugawara |
| John D'Ambrosia | Elvis Maculuba | Patrik Sundstrom |
| Piers J. G. Dawe | Valerie Maguire | James Theodoras |
| Patrick Diamond | Jeffery Maki | David Thompson |
| Michael Dudek | Arthur Marris | Geoffrey Thompson |
| John French | Mick McCarthy | Michael Thompson |
| Matthias Fritsche | Brett McClellan | Steven Tilden |
| Yukihiro Fujimoto | Thomas McDermott | Steve Trowbridge |
| Ali Ghiasi | Michael McInnis | Mark-Rene Uchida |
| Joel Goergen | Richard Mellitz | Alexander Umnov |
| Zhigang Gong | Tremont Miao | Paul Vanderlaan |
| James Graba | Jeffrey Moore | Dmitri Varsanofiev |
| Randall Groves | Charles Moorwood | George Vlantis |
| Robert Grow | Jose Morales | Khurram Waheed |
| Mark Gustlin | Michael Newman | Oded Wertheim |
| Adam Healey | Nick S. A. Nikjoo | Andreas Wolf |
| Marco Hernandez | Paul Nikolich | Peter Wu |
| David Hess | Mark Nowell | Jun Xu |
| Yasuo Hidaka | Satoshi Obara | Oren Yuen |
| Guido Hiertz | Thomas Palkert | Zhen Zhou |

When the IEEE-SA Standards Board approved this amendment on 6 December 2017, it had the following membership:

Jean-Philippe Faure, *Chair*
Gary Hoffman, *Vice Chair*
John D. Kulick, *Past Chair*
Konstantinos Karachalios, *Secretary*

| | | |
|--------------------|-----------------------|-----------------|
| Chuck Adams | Thomas Kochy | Robby Robson |
| Masayuki Ariyoshi | Joseph L. Koepfinger* | Dorothy Stanley |
| Ted Burse | Kevin Lu | Adrian Stephens |
| Stephen Dukes | Daleep Mohla | Mehmet Ulema |
| Doug Edwards | Damir Novosel | Phil Wennblom |
| J. Travis Griffith | Ronald C. Petersen | Howard Wolfman |
| Michael Janezic | Annette D. Reilly | Yu Yuan |

*Member Emeritus

Introduction

This introduction is not part of IEEE Std 802.3bs-2017, IEEE Standard for Ethernet—Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters for 200 Gb/s and 400 Gb/s Operation.

IEEE Std 802.3™ was first published in 1985. Since the initial publication, many projects have added functionality or provided maintenance updates to the specifications and text included in the standard. Each IEEE 802.3 project/amendment is identified with a suffix (e.g., IEEE Std 802.3ba™-2010).

The half duplex Media Access Control (MAC) protocol specified in IEEE Std 802.3-1985 is Carrier Sense Multiple Access with Collision Detection (CSMA/CD). This MAC protocol was key to the experimental Ethernet developed at Xerox Palo Alto Research Center, which had a 2.94 Mb/s data rate. Ethernet at 10 Mb/s was jointly released as a public specification by Digital Equipment Corporation (DEC), Intel and Xerox in 1980. Ethernet at 10 Mb/s was approved as an IEEE standard by the IEEE Standards Board in 1983 and subsequently published in 1985 as IEEE Std 802.3-1985. Since 1985, new media options, new speeds of operation, and new capabilities have been added to IEEE Std 802.3. A full duplex MAC protocol was added in 1997.

Some of the major additions to IEEE Std 802.3 are identified in the marketplace with their project number. This is most common for projects adding higher speeds of operation or new protocols. For example, IEEE Std 802.3u™ added 100 Mb/s operation (also called Fast Ethernet), IEEE Std 802.3z added 1000 Mb/s operation (also called Gigabit Ethernet), IEEE Std 802.3ae added 10 Gb/s operation (also called 10 Gigabit Ethernet), IEEE Std 802.3ah™ specified access network Ethernet (also called Ethernet in the First Mile) and IEEE Std 802.3ba added 40 Gb/s operation (also called 40 Gigabit Ethernet) and 100 Gb/s operation (also called 100 Gigabit Ethernet). These major additions are all now included in and are superseded by IEEE Std 802.3-2015 and are not maintained as separate documents.

At the publication date of IEEE Std 802.3bs-2017, IEEE Std 802.3 is composed of the following documents:

IEEE Std 802.3-2015

Section One—Includes Clause 1 through Clause 20 and Annex A through Annex H and Annex 4A. Section One includes the specifications for 10 Mb/s operation and the MAC, frame formats and service interfaces used for all speeds of operation.

Section Two—Includes Clause 21 through Clause 33 and Annex 22A through Annex 33E. Section Two includes management attributes for multiple protocols and speed of operation as well as specifications for providing power over twisted pair cabling for multiple operational speeds. It also includes general information on 100 Mb/s operation as well as most of the 100 Mb/s Physical Layer specifications.

Section Three—Includes Clause 34 through Clause 43 and Annex 36A through Annex 43C. Section Three includes general information on 1000 Mb/s operation as well as most of the 1000 Mb/s Physical Layer specifications.

Section Four—Includes Clause 44 through Clause 55 and Annex 44A through Annex 55B. Section Four includes general information on 10 Gb/s operation as well as most of the 10 Gb/s Physical Layer specifications.

Section Five—Includes Clause 56 through Clause 77 and Annex 57A through Annex 76A. Clause 56 through Clause 67 and Clause 75 through Clause 77, as well as associated annexes, specify subscriber access and other Physical Layers and sublayers for operation from 512 kb/s to 10 Gb/s, and defines services and protocol elements that enable the exchange of IEEE Std 802.3 format frames between

stations in a subscriber access network. Clause 68 specifies a 10 Gb/s Physical Layer specification. Clause 69 through Clause 74 and associated annexes specify Ethernet operation over electrical backplanes at speeds of 1000 Mb/s and 10 Gb/s.

Section Six—Includes Clause 78 through Clause 95 and Annex 83A through Annex 93C. Clause 78 specifies Energy-Efficient Ethernet. Clause 79 specifies IEEE 802.3 Organizationally Specific Link Layer Discovery Protocol (LLDP) type, length, and value (TLV) information elements. Clause 80 through Clause 95 and associated annexes includes general information on 40 Gb/s and 100 Gb/s operation as well the 40 Gb/s and 100 Gb/s Physical Layer specifications. Clause 90 specifies Ethernet support for time synchronization protocols.

IEEE Std 802.3bw™-2015

Amendment 1—This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 96. This amendment adds 100 Mb/s Physical Layer (PHY) specifications and management parameters for operation on a single balanced twisted-pair copper cable.

IEEE Std 802.3by™-2016

Amendment 2—This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 105 through Clause 112, Annex 109A, Annex 109B, Annex 109C, Annex 110A, Annex 110B, and Annex 110C. This amendment adds MAC parameters, Physical Layers, and management parameters for the transfer of IEEE 802.3 format frames at 25 Gb/s.

IEEE Std 802.3bq™-2016

Amendment 3—This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 113 and Annex 113A. This amendment adds new Physical Layers for 25 Gb/s and 40 Gb/s operation over balanced twisted-pair structured cabling systems.

IEEE Std 802.3bp™-2016

Amendment 4—This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 97 and Clause 98. This amendment adds point-to-point 1 Gb/s Physical Layer (PHY) specifications and management parameters for operation on a single balanced twisted-pair copper cable in automotive and other applications not utilizing the structured wiring plant.

IEEE Std 802.3br™-2016

Amendment 5—This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 99. This amendment adds a MAC Merge sublayer and a MAC Merge Service Interface to support for Interspersing Express Traffic over a single link.

IEEE Std 802.3bn™-2016

Amendment 6—This amendment adds the Physical Layer specifications and management parameters for symmetric and/or asymmetric operation of up to 10 Gb/s on point-to-multipoint Radio Frequency (RF) distribution plants comprising either amplified or passive coaxial media. It also extends the operation of Ethernet Passive Optical Networks (EPON) protocols, such as Multipoint Control Protocol (MPCP) and Operation Administration and Management (OAM).

IEEE Std 802.3bz™-2016

Amendment 7—This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 125 and Clause 126. This amendment adds new rates of 2.5 Gb/s and 5 Gb/s and new Physical Layers for operation at 2.5 Gb/s and 5 Gb/s over balanced twisted-pair structured cabling systems.

IEEE Std 802.3bu™-2016

Amendment 8—This amendment includes changes to IEEE Std 802.3-2015 to define a methodology for the provision of power via a single twisted pair to connected Data Terminal Equipment (DTE) with IEEE 802.3 single twisted-pair interfaces.

IEEE Std 802.3bv™-2017

Amendment 9—This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 115 and Annex 115A. This amendment adds point-to-point 1000 Mb/s Physical Layer (PHY) specifications and management parameters for operation on duplex plastic optical fiber (POF) targeting use in automotive, industrial, home-network, and other applications.

IEEE Std 802.3™-2015/Cor 1-2017

This corrigendum clarifies which lane of the media dependent interface (MDI) of a multi-lane Physical Layer entity (PHY) is used as the timestamping reference point.

IEEE Std 802.3bs™-2017

Amendment 10—This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 116 through Clause 124 and Annex 119A through Annex 120E. This amendment adds MAC parameters, Physical Layers, and management parameters for the transfer of IEEE 802.3 format frames at 200 Gb/s and 400 Gb/s.

A companion document IEEE Std 802.3.1 describes Ethernet management information base (MIB) modules for use with the Simple Network Management Protocol (SNMP). IEEE Std 802.3.1 is updated to add management capability for enhancements to IEEE Std 802.3 after approval of the enhancements.

IEEE Std 802.3 will continue to evolve. New Ethernet capabilities are anticipated to be added within the next few years as amendments to this standard.

Contents

| | |
|---|----|
| 1. Introduction..... | 31 |
| 1.1 Overview..... | 31 |
| 1.1.3 Architectural perspectives..... | 31 |
| 1.1.3.2 Compatibility interfaces..... | 31 |
| 1.3 Normative references..... | 32 |
| 1.4 Definitions..... | 32 |
| 1.5 Abbreviations..... | 34 |
| 4. Media Access Control..... | 35 |
| 4.4 Specific implementations..... | 35 |
| 4.4.2 MAC parameters..... | 35 |
| 30. Management..... | 36 |
| 30.2 Managed objects..... | 36 |
| 30.2.5 Capabilities..... | 36 |
| 30.3 Layer management for DTEs..... | 36 |
| 30.3.2 PHY device managed object class..... | 36 |
| 30.3.2.1 PHY device attributes..... | 36 |
| 30.3.2.1.2 aPhyType..... | 36 |
| 30.3.2.1.3 aPhyTypeList..... | 37 |
| 30.3.2.1.5 aSymbolErrorDuringCarrier..... | 37 |
| 30.5 Layer management for medium attachment units (MAUs)..... | 37 |
| 30.5.1 MAU managed object class..... | 37 |
| 30.5.1.1 MAU attributes..... | 37 |
| 30.5.1.1.2 aMAUType..... | 37 |
| 30.5.1.1.4 aMediaAvailable..... | 38 |
| 30.5.1.1.12 aLaneMapping..... | 38 |
| 30.5.1.1.15 aFECAbility..... | 38 |
| 30.5.1.1.17 aFECCorrectedBlocks..... | 39 |
| 30.5.1.1.18 aFECUncorrectableBlocks..... | 39 |
| 30.5.1.1.32 aPCSFECIndicationAbility..... | 40 |
| 30.5.1.1.33 aPCSFECIndicationEnable..... | 40 |
| 45. Management Data Input/Output (MDIO) Interface..... | 41 |
| 45.2 MDIO Interface Registers..... | 41 |
| 45.2.1 PMA/PMD registers..... | 41 |
| 45.2.1.1 PMA/PMD control 1 register (Register 1.0)..... | 43 |
| 45.2.1.1.3 Speed selection (1.0.13, 1.0.6, 1.0.5:2)..... | 44 |
| 45.2.1.1.4 PMA remote loopback (1.0.1)..... | 44 |
| 45.2.1.1.5 PMA local loopback (1.0.0)..... | 44 |
| 45.2.1.2 PMA/PMD status 1 register (Register 1.1)..... | 44 |
| 45.2.1.2.3 Fault (1.1.7)..... | 44 |
| 45.2.1.4 PMA/PMD speed ability (Register 1.4)..... | 45 |
| 45.2.1.4.aaa 400G capable (1.4.15)..... | 45 |
| 45.2.1.4.ac 200G capable (1.4.12)..... | 45 |
| 45.2.1.6 PMA/PMD control 2 register (Register 1.7)..... | 45 |
| 45.2.1.6.3 PMA/PMD type selection (1.7.65:0)..... | 47 |
| 45.2.1.7 PMA/PMD status 2 register (Register 1.8)..... | 48 |

| | | |
|---------------|--|----|
| 45.2.1.7.4 | Transmit fault (1.8.11) | 48 |
| 45.2.1.7.5 | Receive fault (1.8.10) | 48 |
| 45.2.1.8 | PMD transmit disable register (Register 1.9) | 48 |
| 45.2.1.8.1 | PMD transmit disable 914 (1.9.1015) | 49 |
| 45.2.1.8.2 | PMD transmit disable 4, 5, 6, 7, 8 through 13 (1.9.5, 1.9.6, 1.9.7, 1.9.8, 1.9.9 through 1.9.14) | 50 |
| 45.2.1.9 | PMD receive signal detect register (Register 1.10) | 50 |
| 45.2.1.9.1 | PMD receive signal detect 914 (1.10.1015) | 50 |
| 45.2.1.9.2 | PMD receive signal detect 4, 5, 6, 7, 8 through 13 (1.10.5, 1.10.6, 1.10.7, 1.10.8, 1.10.9 through 1.10.14) | 51 |
| 45.2.1.10 | PMA/PMD extended ability register (Register 1.11) | 51 |
| 45.2.1.10.aab | 200G/400G extended abilities (1.11.13) | 51 |
| 45.2.1.14e | 200G PMA/PMD extended ability register (Register 1.23) | 51 |
| 45.2.1.14e.1 | 200G PMA remote loopback ability (1.23.15) | 52 |
| 45.2.1.14e.2 | 200GBASE-LR4 ability (1.23.5) | 52 |
| 45.2.1.14e.3 | 200GBASE-FR4 ability (1.23.4) | 52 |
| 45.2.1.14e.4 | 200GBASE-DR4 ability (1.23.3) | 52 |
| 45.2.1.14f | 400G PMA/PMD extended ability register (Register 1.24) | 52 |
| 45.2.1.14f.1 | 400G PMA remote loopback ability (1.24.15) | 53 |
| 45.2.1.14f.2 | 400GBASE-LR8 ability (1.24.5) | 53 |
| 45.2.1.14f.3 | 400GBASE-FR8 ability (1.24.4) | 53 |
| 45.2.1.14f.4 | 400GBASE-DR4 ability (1.24.3) | 53 |
| 45.2.1.14f.5 | 400GBASE-SR16 ability (1.24.2) | 53 |
| 45.2.1.14g | PMD transmit disable extension register (Register 1.27) | 53 |
| 45.2.1.14g.1 | PMD transmit disable 15 (1.27.0) | 54 |
| 45.2.1.14h | PMD receive signal detect extension register (Register 1.28) | 54 |
| 45.2.1.14h.1 | PMD receive signal detect 15 (1.28.0) | 54 |
| 45.2.1.116a | 200GAUI-8 and 400GAUI-16 chip-to-module recommended CTLE, lane 0 register (Register 1.400) | 55 |
| 45.2.1.116a.1 | Recommended CTLE peaking (1.400.4:1) | 55 |
| 45.2.1.116b | 200GAUI-8 and 400GAUI-16 chip-to-module recommended CTLE, lane 1 through lane 7 registers (Registers 1.401 through 1.407) | 55 |
| 45.2.1.116c | 400GAUI-16 chip-to-module recommended CTLE, lane 8 through lane 15 registers (Registers 1.408 through 1.415) | 55 |
| 45.2.1.116d | 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 register (Register 1.500) | 56 |
| 45.2.1.116d.1 | Request flag (1.500.15) | 57 |
| 45.2.1.116d.2 | Post-cursor request (1.500.14:12) | 57 |
| 45.2.1.116d.3 | Pre-cursor request (1.500.11:10) | 57 |
| 45.2.1.116d.4 | Post-cursor remote setting (1.500.9:7) | 57 |
| 45.2.1.116d.5 | Pre-cursor remote setting (1.500.6:5) | 57 |
| 45.2.1.116d.6 | Post-cursor local setting (1.500.4:2) | 57 |
| 45.2.1.116d.7 | Pre-cursor local setting (1.500.1:0) | 57 |
| 45.2.1.116e | 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 1 through lane 15 registers (Registers 1.501 through 1.515) | 58 |
| 45.2.1.116f | 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 0 register (Register 1.516) | 58 |
| 45.2.1.116f.1 | Request flag (1.516.15) | 59 |
| 45.2.1.116f.2 | Post-cursor request (1.516.14:12) | 59 |
| 45.2.1.116f.3 | Pre-cursor request (1.516.11:10) | 59 |
| 45.2.1.116f.4 | Post-cursor remote setting (1.516.9:7) | 59 |
| 45.2.1.116f.5 | Pre-cursor remote setting (1.516.6:5) | 60 |
| 45.2.1.116f.6 | Post-cursor local setting (1.516.4:2) | 60 |
| 45.2.1.116f.7 | Pre-cursor local setting (1.516.1:0) | 60 |

| | | |
|--------------|---|----|
| 45.2.1.116g | 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 1 through lane 15 registers (Registers 1.517 through 1.531) | 60 |
| 45.2.1.123 | Test-pattern ability (Register 1.1500) | 60 |
| 45.2.1.124 | PRBS pattern testing control (Register 1.1501) | 62 |
| 45.2.1.125 | Square wave testing control (Register 1.1510) | 63 |
| 45.2.1.125a | PRBS13Q testing control (Register 1.1512) | 64 |
| 45.2.1.126 | PRBS Tx pattern testing error counter (Register 1.1600 through 1.1615, 1.1601, 1.1602, 1.1603, 1.1604, 1.1605, 1.1606, 1.1607, 1.1608, 1.1609) | 65 |
| 45.2.1.127 | PRBS Rx pattern testing error counter (Register 1.1700 through 1.1715, 1.1701, 1.1702, 1.1703, 1.1704, 1.1705, 1.1706, 1.1707, 1.1708, 1.1709) | 66 |
| 45.2.3 | PCS registers | 66 |
| 45.2.3.1 | PCS control 1 register (Register 3.0) | 67 |
| 45.2.3.2 | PCS status 1 register (Register 3.1) | 67 |
| 45.2.3.2.7 | PCS receive link status (3.1.2) | 67 |
| 45.2.3.4 | PCS speed ability (Register 3.4) | 67 |
| 45.2.3.4.8 | 200G capable (3.4.8) | 68 |
| 45.2.3.4.9 | 400G capable (3.4.9) | 68 |
| 45.2.3.6 | PCS control 2 register (Register 3.7) | 68 |
| 45.2.3.6.1 | PCS type selection (3.7.3:0) | 69 |
| 45.2.3.7a | PCS status 3 register (Register 3.9) | 69 |
| 45.2.3.7a.1 | 400GBASE-R capable (3.9.1) | 69 |
| 45.2.3.7a.2 | 200GBASE-R capable (3.9.0) | 69 |
| 45.2.3.9a | EEE control and capability 2 (Register 3.21) | 69 |
| 45.2.3.9a.a | 400GBASE-R EEE fast wake supported (3.21.5) | 70 |
| 45.2.3.9a.b | 200GBASE-R EEE fast wake supported (3.21.3) | 70 |
| 45.2.3.13 | BASE-R and MultiGBASE-T PCS status 1 register (Register 3.32) | 70 |
| 45.2.3.13.1 | BASE-R and MultiGBASE-T receive link status (3.32.12) | 70 |
| 45.2.3.17 | BASE-R PCS test-pattern control register (Register 3.42) | 70 |
| 45.2.3.46 | Lane 0 mapping register (Register 3.400) | 70 |
| 45.2.3.47h | PCS FEC symbol error counter lane 0 (Register 3.600, 3.601) | 71 |
| 45.2.3.47i | PCS FEC symbol error counter lane 1 through 15 (Registers 3.602 through 3.631) | 71 |
| 45.2.3.47j | PCS FEC control register (Register 3.800) | 71 |
| 45.2.3.47j.1 | PCS FEC degraded SER enable (3.800.2) | 72 |
| 45.2.3.47j.2 | PCS FEC bypass indication enable (3.800.1) | 72 |
| 45.2.3.47k | PCS FEC status register (Register 3.801) | 72 |
| 45.2.3.47k.1 | Local degraded SER received (3.801.6) | 72 |
| 45.2.3.47k.2 | Remote degraded SER received (3.801.5) | 73 |
| 45.2.3.47k.3 | PCS FEC degraded SER (3.801.4) | 73 |
| 45.2.3.47k.4 | PCS FEC degraded SER ability (3.801.3) | 73 |
| 45.2.3.47k.5 | PCS FEC high SER (3.801.2) | 73 |
| 45.2.3.47k.6 | PCS FEC bypass indication ability (3.801.1) | 73 |
| 45.2.3.47l | PCS FEC corrected codewords counter (Register 3.802, 3.803) | 73 |
| 45.2.3.47m | PCS FEC uncorrected codewords counter (Register 3.804, 3.805) | 74 |
| 45.2.3.47n | PCS FEC degraded SER activate threshold register (Register 3.806, 3.807) | 74 |
| 45.2.3.47o | PCS FEC degraded SER deactivate threshold register (Register 3.808, 3.809) | 74 |
| 45.2.3.47p | PCS FEC degraded SER interval register (Register 3.810, 3.811) | 75 |
| 45.2.4 | PHY XS registers | 75 |
| 45.2.4.1 | PHY XS control 1 register (Register 4.0) | 76 |
| 45.2.4.4 | PHY XS speed ability (Register 4.4) | 76 |
| 45.2.4.4.a | 400G capable (4.4.9) | 77 |
| 45.2.4.4.b | 200G capable (4.4.8) | 77 |
| 45.2.4.11a | BASE-R PHY XS status 1 register (Register 4.32) | 77 |
| 45.2.4.11a.1 | BASE-R PHY XS receive link status (4.32.12) | 77 |

| | | |
|--------------|--|----|
| 45.2.4.11b | BASE-R PHY XS test-pattern control register (Register 4.42)..... | 78 |
| 45.2.4.11b.1 | Transmit test-pattern enable (4.42.3)..... | 78 |
| 45.2.4.11c | Multi-lane BASE-R PHY XS alignment status 1 register (Register 4.50)..... | 78 |
| 45.2.4.11c.1 | PHY XS lane alignment status (4.50.12)..... | 78 |
| 45.2.4.11d | Multi-lane BASE-R PHY XS alignment status 3 register (Register 4.52)..... | 79 |
| 45.2.4.11d.1 | Lane 7 aligned (4.52.7)..... | 79 |
| 45.2.4.11d.2 | Lane 6 aligned (4.52.6)..... | 79 |
| 45.2.4.11d.3 | Lane 5 aligned (4.52.5)..... | 79 |
| 45.2.4.11d.4 | Lane 4 aligned (4.52.4)..... | 80 |
| 45.2.4.11d.5 | Lane 3 aligned (4.52.3)..... | 80 |
| 45.2.4.11d.6 | Lane 2 aligned (4.52.2)..... | 80 |
| 45.2.4.11d.7 | Lane 1 aligned (4.52.1)..... | 80 |
| 45.2.4.11d.8 | Lane 0 aligned (4.52.0)..... | 80 |
| 45.2.4.11e | Multi-lane BASE-R PHY XS alignment status 4 register (Register 4.53)..... | 80 |
| 45.2.4.11e.1 | Lane 15 aligned (4.53.7)..... | 81 |
| 45.2.4.11e.2 | Lane 14 aligned (4.53.6)..... | 81 |
| 45.2.4.11e.3 | Lane 13 aligned (4.53.5)..... | 81 |
| 45.2.4.11e.4 | Lane 12 aligned (4.53.4)..... | 81 |
| 45.2.4.11e.5 | Lane 11 aligned (4.53.3)..... | 81 |
| 45.2.4.11e.6 | Lane 10 aligned (4.53.2)..... | 82 |
| 45.2.4.11e.7 | Lane 9 aligned (4.53.1)..... | 82 |
| 45.2.4.11e.8 | Lane 8 aligned (4.53.0)..... | 82 |
| 45.2.4.11f | PHY XS lane mapping, lane 0 register (Register 4.400)..... | 82 |
| 45.2.4.11g | PHY XS lane mapping, lane 1 through lane 15 registers (Registers 4.401 through 4.415)..... | 82 |
| 45.2.4.11h | PHY XS FEC symbol error counter lane 0 (Register 4.600, 4.601)..... | 82 |
| 45.2.4.11i | PHY XS FEC symbol error counter lane 1 through 15 (Registers 4.602 through 4.631)..... | 83 |
| 45.2.4.11j | PHY XS FEC control register (Register 4.800)..... | 83 |
| 45.2.4.11j.1 | PHY XS FEC degraded SER enable (4.800.2)..... | 83 |
| 45.2.4.11j.2 | PHY XS FEC bypass indication enable (4.800.1)..... | 83 |
| 45.2.4.11k | PHY XS FEC status register (Register 4.801)..... | 84 |
| 45.2.4.11k.1 | Remote degraded SER received (4.801.5)..... | 84 |
| 45.2.4.11k.2 | PHY XS FEC degraded SER (4.801.4)..... | 84 |
| 45.2.4.11k.3 | PHY XS FEC degraded SER ability (4.801.3)..... | 84 |
| 45.2.4.11k.4 | PHY XS FEC high SER (4.801.2)..... | 85 |
| 45.2.4.11k.5 | PHY XS FEC bypass indication ability (4.801.1)..... | 85 |
| 45.2.4.11l | PHY XS FEC corrected codewords counter (Register 4.802, 4.803)..... | 85 |
| 45.2.4.11m | PHY XS FEC uncorrected codewords counter (Register 4.804, 4.805)..... | 85 |
| 45.2.4.11n | PHY XS FEC degraded SER activate threshold register (Register 4.806, 4.807)..... | 86 |
| 45.2.4.11o | PHY XS FEC degraded SER deactivate threshold register (Register 4.808, 4.809)..... | 86 |
| 45.2.4.11p | PHY XS FEC degraded SER interval register (Register 4.810, 4.811)..... | 86 |
| 45.2.5 | DTE XS registers..... | 87 |
| 45.2.5.1 | DTE XS control 1 register (Register 5.0)..... | 88 |
| 45.2.5.4 | DTE XS speed ability (Register 5.4)..... | 88 |
| 45.2.5.4.a | 400G capable (5.4.9)..... | 88 |
| 45.2.5.4.b | 200G capable (5.4.8)..... | 89 |
| 45.2.5.11a | BASE-R DTE XS status 1 register (Register 5.32)..... | 89 |
| 45.2.5.11a.1 | BASE-R DTE XS receive link status (5.32.12)..... | 89 |
| 45.2.5.11b | BASE-R DTE XS test-pattern control register (Register 5.42)..... | 89 |
| 45.2.5.11b.1 | Transmit test-pattern enable (5.42.3)..... | 90 |
| 45.2.5.11c | Multi-lane BASE-R DTE XS alignment status 1 register (Register 5.50)..... | 90 |

| | | |
|--------------|---|-----|
| 45.2.5.11c.1 | DTE XS lane alignment status (5.50.12) | 90 |
| 45.2.5.11d | Multi-lane BASE-R DTE XS alignment status 3 register (Register 5.52) | 90 |
| 45.2.5.11d.1 | Lane 7 aligned (5.52.7) | 91 |
| 45.2.5.11d.2 | Lane 6 aligned (5.52.6) | 91 |
| 45.2.5.11d.3 | Lane 5 aligned (5.52.5) | 91 |
| 45.2.5.11d.4 | Lane 4 aligned (5.52.4) | 91 |
| 45.2.5.11d.5 | Lane 3 aligned (5.52.3) | 91 |
| 45.2.5.11d.6 | Lane 2 aligned (5.52.2) | 92 |
| 45.2.5.11d.7 | Lane 1 aligned (5.52.1) | 92 |
| 45.2.5.11d.8 | Lane 0 aligned (5.52.0) | 92 |
| 45.2.5.11e | Multi-lane BASE-R DTE XS alignment status 4 register (Register 5.53) | 92 |
| 45.2.5.11e.1 | Lane 15 aligned (5.53.7) | 93 |
| 45.2.5.11e.2 | Lane 14 aligned (5.53.6) | 93 |
| 45.2.5.11e.3 | Lane 13 aligned (5.53.5) | 93 |
| 45.2.5.11e.4 | Lane 12 aligned (5.53.4) | 93 |
| 45.2.5.11e.5 | Lane 11 aligned (5.53.3) | 93 |
| 45.2.5.11e.6 | Lane 10 aligned (5.53.2) | 93 |
| 45.2.5.11e.7 | Lane 9 aligned (5.53.1) | 93 |
| 45.2.5.11e.8 | Lane 8 aligned (5.53.0) | 93 |
| 45.2.5.11f | DTE XS lane mapping, lane 0 register (Register 5.400) | 93 |
| 45.2.5.11g | DTE XS lane mapping, lane 1 through lane 15 registers (Registers 5.401 through 5.415) | 94 |
| 45.2.5.11h | DTE XS FEC symbol error counter lane 0 (Register 5.600, 5.601) | 94 |
| 45.2.5.11i | DTE XS FEC symbol error counter lane 1 through 15 (Registers 5.602 through 5.631) | 94 |
| 45.2.5.11j | DTE XS FEC control register (Register 5.800) | 95 |
| 45.2.5.11j.1 | DTE XS FEC degraded SER enable (5.800.2) | 95 |
| 45.2.5.11j.2 | DTE XS FEC bypass indication enable (5.800.1) | 95 |
| 45.2.5.11k | DTE XS FEC status register (Register 5.801) | 95 |
| 45.2.5.11k.1 | Local degraded SER received (5.801.6) | 96 |
| 45.2.5.11k.2 | Remote degraded SER received (5.801.5) | 96 |
| 45.2.5.11k.3 | DTE XS FEC degraded SER (5.801.4) | 96 |
| 45.2.5.11k.4 | DTE XS FEC degraded SER ability (5.801.3) | 96 |
| 45.2.5.11k.5 | DTE XS FEC high SER (5.801.2) | 96 |
| 45.2.5.11k.6 | DTE XS FEC bypass indication ability (5.801.1) | 97 |
| 45.2.5.11l | DTE XS FEC corrected codewords counter (Register 5.802, 5.803) | 97 |
| 45.2.5.11m | DTE XS FEC uncorrected codewords counter (Register 5.804, 5.805) | 97 |
| 45.2.5.11n | DTE XS FEC degraded SER activate threshold register (Register 5.806, 5.807) | 97 |
| 45.2.5.11o | DTE XS FEC degraded SER deactivate threshold register (Register 5.808, 5.809) | 98 |
| 45.2.5.11p | DTE XS FEC degraded SER interval register (Register 5.810, 5.811) | 98 |
| 78 | Energy-Efficient Ethernet (EEE) | 99 |
| 78.1 | Overview | 99 |
| 78.1.4 | PHY types optionally supporting EEE | 99 |
| 78.5 | Communication link access latency | 100 |
| 78.5.1 | 10 Gb/s PHY extension using extender sublayers XGXS | 100 |
| 90 | Ethernet support for time synchronization protocols | 101 |
| 90.1 | Introduction | 101 |

| | |
|--|-----|
| 116. Introduction to 200 Gb/s and 400 Gb/s networks | 102 |
| 116.1 Overview..... | 102 |
| 116.1.1 Scope..... | 102 |
| 116.1.2 Relationship of 200 Gigabit and 400 Gigabit Ethernet to the ISO OSI reference model | 102 |
| 116.1.3 Nomenclature..... | 103 |
| 116.1.4 Physical Layer signaling systems | 104 |
| 116.2 Summary of 200 Gigabit and 400 Gigabit Ethernet sublayers | 105 |
| 116.2.1 Reconciliation Sublayer (RS) and Media Independent Interface | 105 |
| 116.2.2 200GMII and 400GMII Extender Sublayers (200GXS and 400GXS)..... | 105 |
| 116.2.3 Physical Coding Sublayer (PCS)..... | 106 |
| 116.2.4 Physical Medium Attachment (PMA) sublayer..... | 106 |
| 116.2.5 Physical Medium Dependent (PMD) sublayer | 106 |
| 116.2.6 Management interface (MDIO/MDC)..... | 106 |
| 116.2.7 Management..... | 106 |
| 116.3 Service interface specification method and notation | 106 |
| 116.3.1 Inter-sublayer service interface..... | 107 |
| 116.3.2 Instances of the Inter-sublayer service interface..... | 107 |
| 116.3.3 Semantics of inter-sublayer service interface primitives..... | 107 |
| 116.3.3.1 IS_UNITDATA_i.request..... | 107 |
| 116.3.3.1.1 Semantics of the service primitive..... | 108 |
| 116.3.3.1.2 When generated | 109 |
| 116.3.3.1.3 Effect of receipt | 109 |
| 116.3.3.2 IS_UNITDATA_i.indication | 110 |
| 116.3.3.2.1 Semantics of the service primitive..... | 110 |
| 116.3.3.2.2 When generated | 110 |
| 116.3.3.2.3 Effect of receipt | 110 |
| 116.3.3.3 IS_SIGNAL.indication | 110 |
| 116.3.3.3.1 Semantics of the service primitive..... | 110 |
| 116.3.3.3.2 When generated | 110 |
| 116.3.3.3.3 Effect of receipt | 110 |
| 116.4 Delay constraints..... | 111 |
| 116.5 Skew constraints | 112 |
| 116.6 FEC Degrade..... | 115 |
| 116.7 State diagrams..... | 116 |
| 116.8 Protocol implementation conformance statement (PICS) proforma..... | 117 |
| 117. Reconciliation Sublayer (RS) and Media Independent Interface for 200 Gb/s and 400 Gb/s operation (200GMII and 400GMII)..... | 118 |
| 117.1 Overview..... | 118 |
| 117.1.1 Summary of major concepts | 119 |
| 117.1.2 Application..... | 119 |
| 117.1.3 Rate of operation..... | 119 |
| 117.1.4 Delay constraints..... | 119 |
| 117.1.5 Allocation of functions | 120 |
| 117.1.6 200GMII/400GMII structure | 120 |
| 117.1.7 Mapping of 200GMII/400GMII signals to PLS service primitives..... | 120 |
| 117.2 200GMII/400GMII data stream..... | 120 |
| 117.3 200GMII/400GMII functional specifications..... | 120 |
| 117.4 LPI Assertion and Detection..... | 120 |

| | | |
|-----------|---|-----|
| 117.5 | Protocol implementation conformance statement (PICS) proforma for Clause 117, Reconciliation Sublayer (RS) and Media Independent Interface for 200 Gb/s and 400 Gb/s operation (200GMII and 400GMII)..... | 121 |
| 117.5.1 | Introduction..... | 121 |
| 117.5.2 | Identification..... | 121 |
| 117.5.2.1 | Implementation identification..... | 121 |
| 117.5.2.2 | Protocol summary..... | 121 |
| 117.5.3 | Major capabilities/options..... | 122 |
| 117.5.4 | PICS proforma tables for Reconciliation Sublayer (RS) and Media Independent Interface for 200 Gb/s and 400 Gb/s operation (200GMII and 400GMII)..... | 122 |
| 117.5.4.1 | General..... | 122 |
| 117.5.4.2 | Mapping of PLS service primitives..... | 122 |
| 117.5.4.3 | Data stream structure..... | 123 |
| 117.5.4.4 | 200GMII/400GMII signal functional specifications..... | 123 |
| 117.5.4.5 | Link fault signaling state diagram..... | 124 |
| 117.5.4.6 | LPI functions..... | 124 |
| 118. | 200GMII Extender, 400GMII Extender, 200GMII Extender Sublayer (200GXS), and 400GMII Extender Sublayer (400GXS)..... | 125 |
| 118.1 | Overview..... | 125 |
| 118.1.1 | Summary of major concepts..... | 126 |
| 118.1.2 | 200GXS/400GXS Sublayer..... | 126 |
| 118.1.3 | 200GAUI-n/400GAUI-n..... | 126 |
| 118.2 | FEC Degrade..... | 126 |
| 118.2.1 | DTE XS FEC Degrade signaling..... | 126 |
| 118.2.2 | PHY XS FEC Degrade signaling..... | 127 |
| 118.3 | 200GXS and 400GXS partitioning example..... | 127 |
| 118.4 | 200GXS and 400GXS MDIO function mapping..... | 127 |
| 118.5 | Protocol implementation conformance statement (PICS) proforma for Clause 118, 200GMII Extender, 400GMII Extender, 200GMII Extender Sublayer (200GXS), and 400GMII Extender Sublayer (400GXS)..... | 131 |
| 118.5.1 | Introduction..... | 131 |
| 118.5.2 | Identification..... | 131 |
| 118.5.2.1 | Implementation identification..... | 131 |
| 118.5.2.2 | Protocol summary..... | 131 |
| 118.5.3 | Major capabilities/options..... | 132 |
| 118.5.4 | PICS proforma tables for 200GMII Extender, 400GMII Extender, 200GMII Extender Sublayer (200GXS), and 400GMII Extender Sublayer (400GXS)..... | 132 |
| 118.5.4.1 | Transmit function..... | 132 |
| 118.5.4.2 | Receive function..... | 133 |
| 118.5.4.3 | 64B/66B coding rules..... | 133 |
| 118.5.4.4 | Scrambler and descrambler..... | 134 |
| 118.5.4.5 | Alignment markers..... | 134 |
| 118.5.5 | Test-pattern modes..... | 134 |
| 118.5.6 | Bit order..... | 134 |
| 118.5.7 | Management..... | 135 |
| 118.5.7.1 | State diagrams..... | 135 |
| 118.5.7.2 | Loopback..... | 135 |
| 118.5.7.3 | Delay constraints..... | 136 |
| 119. | Physical Coding Sublayer (PCS) for 64B/66B, type 200GBASE-R and 400GBASE-R..... | 137 |
| 119.1 | Overview..... | 137 |

| | | |
|-------------|--|-----|
| 119.1.1 | Scope..... | 137 |
| 119.1.2 | Relationship of 200GBASE-R and 400GBASE-R to other standards | 137 |
| 119.1.3 | Physical Coding Sublayer (PCS) | 137 |
| 119.1.4 | Inter-sublayer interfaces | 138 |
| 119.1.4.1 | PCS service interface (200GMII/400GMII) | 138 |
| 119.1.4.2 | Physical Medium Attachment (PMA) service interface | 138 |
| 119.1.5 | Functional block diagram | 139 |
| 119.2 | Physical Coding Sublayer (PCS) | 140 |
| 119.2.1 | Functions within the PCS | 140 |
| 119.2.2 | Use of blocks | 140 |
| 119.2.3 | 64B/66B code | 141 |
| 119.2.3.1 | Notation conventions | 141 |
| 119.2.3.2 | 64B/66B block structure | 141 |
| 119.2.3.3 | Control codes | 141 |
| 119.2.3.4 | Valid and invalid blocks | 142 |
| 119.2.3.5 | Idle (/I/) | 142 |
| 119.2.3.6 | Start (/S/) | 142 |
| 119.2.3.7 | Terminate (/T/) | 142 |
| 119.2.3.8 | Ordered set (/O/) | 142 |
| 119.2.3.9 | Error (/E/) | 142 |
| 119.2.4 | Transmit | 142 |
| 119.2.4.1 | Encode and rate matching | 142 |
| 119.2.4.2 | 64B/66B to 256B/257B transcoder | 143 |
| 119.2.4.3 | Scrambler | 145 |
| 119.2.4.4 | Alignment marker mapping and insertion | 145 |
| 119.2.4.4.1 | AM creation for the 200GBASE-R PCS | 146 |
| 119.2.4.4.2 | AM creation for the 400GBASE-R PCS | 148 |
| 119.2.4.5 | Pre-FEC distribution | 150 |
| 119.2.4.6 | Reed-Solomon encoder | 150 |
| 119.2.4.7 | Symbol distribution | 152 |
| 119.2.4.8 | Transmit bit ordering and distribution | 153 |
| 119.2.4.9 | Test-pattern generators | 155 |
| 119.2.5 | Receive function | 155 |
| 119.2.5.1 | Alignment lock and deskew | 155 |
| 119.2.5.2 | Lane reorder and de-interleave | 155 |
| 119.2.5.3 | Reed-Solomon decoder | 155 |
| 119.2.5.4 | Post-FEC interleave | 156 |
| 119.2.5.5 | Alignment marker removal | 156 |
| 119.2.5.6 | Descrambler | 156 |
| 119.2.5.7 | 256B/257B to 64B/66B transcoder | 157 |
| 119.2.5.8 | Decode and rate matching | 157 |
| 119.2.6 | Detailed functions and state diagrams | 158 |
| 119.2.6.1 | State diagram conventions | 158 |
| 119.2.6.2 | State variables | 158 |
| 119.2.6.2.1 | Constants | 158 |
| 119.2.6.2.2 | Variables | 158 |
| 119.2.6.2.3 | Functions | 160 |
| 119.2.6.2.4 | Counters | 162 |
| 119.2.6.3 | State diagrams | 162 |
| 119.3 | PCS management | 167 |
| 119.3.1 | PCS MDIO function mapping | 167 |
| 119.4 | Loopback | 168 |
| 119.5 | Delay constraints | 168 |

| | | |
|--------------|--|-----|
| 119.6 | Protocol implementation conformance statement (PICS) proforma for Clause 119, Physical Coding Sublayer (PCS) for 64B/66B, type 200GBASE-R and 400GBASE-R | 169 |
| 119.6.1 | Introduction | 169 |
| 119.6.2 | Identification | 169 |
| 119.6.2.1 | Implementation identification | 169 |
| 119.6.2.2 | Protocol summary | 169 |
| 119.6.3 | Major capabilities/options | 170 |
| 119.6.4 | PICS proforma tables for Physical Coding Sublayer (PCS) 64B/66B, type 200GBASE-R and 400GBASE-R | 170 |
| 119.6.4.1 | Transmit function | 170 |
| 119.6.4.2 | Receive function | 171 |
| 119.6.4.3 | 64B/66B coding rules | 171 |
| 119.6.4.4 | Scrambler and descrambler | 172 |
| 119.6.4.5 | Alignment markers | 172 |
| 119.6.4.6 | Test-pattern modes | 172 |
| 119.6.4.7 | Bit order | 173 |
| 119.6.4.8 | Management | 173 |
| 119.6.4.9 | State diagrams | 173 |
| 119.6.4.10 | Loopback | 173 |
| 119.6.4.11 | Delay constraints | 174 |
| 120. | Physical Medium Attachment (PMA) sublayer, type 200GBASE-R and 400GBASE-R | 175 |
| 120.1 | Overview | 175 |
| 120.1.1 | Scope | 175 |
| 120.1.2 | Position of the PMA in the 200GBASE-R and 400GBASE-R sublayers | 175 |
| 120.1.3 | Summary of functions | 175 |
| 120.1.4 | PMA sublayer positioning | 176 |
| 120.2 | PMA interfaces | 178 |
| 120.3 | PMA service interface | 178 |
| 120.4 | Service interface below PMA | 181 |
| 120.5 | Functions within the PMA | 182 |
| 120.5.1 | Per input-lane clock and data recovery | 182 |
| 120.5.2 | Bit-level multiplexing | 182 |
| 120.5.3 | Skew and Skew Variation | 183 |
| 120.5.3.1 | Skew generation toward SP1 | 183 |
| 120.5.3.2 | Skew tolerance at SP1 | 183 |
| 120.5.3.3 | Skew generation toward SP2 | 183 |
| 120.5.3.4 | Skew tolerance at SP5 | 185 |
| 120.5.3.5 | Skew generation at SP6 | 185 |
| 120.5.3.6 | Skew tolerance at SP6 | 185 |
| 120.5.4 | Delay constraints | 185 |
| 120.5.5 | Clocking architecture | 185 |
| 120.5.6 | Signal drivers | 186 |
| 120.5.7 | Gray mapping for PAM4 encoded lanes | 186 |
| 120.5.8 | Link status | 186 |
| 120.5.9 | PMA local loopback mode (optional) | 187 |
| 120.5.10 | PMA remote loopback mode (optional) | 187 |
| 120.5.11 | PMA test patterns (optional) | 187 |
| 120.5.11.1 | Test patterns for NRZ encoded signals | 188 |
| 120.5.11.1.1 | PRBS31 test pattern | 188 |
| 120.5.11.1.2 | PRBS9 test pattern | 189 |
| 120.5.11.1.3 | Square wave test pattern | 189 |
| 120.5.11.2 | Test patterns for PAM4 encoded signals | 189 |

| | | |
|--------------|---|-----|
| 120.5.11.2.1 | PRBS13Q test pattern | 190 |
| 120.5.11.2.2 | PRBS31Q test pattern | 190 |
| 120.5.11.2.3 | SSPRQ test pattern | 192 |
| 120.5.11.2.4 | Square wave (quaternary) test pattern | 193 |
| 120.6 | PMA MDIO function mapping | 193 |
| 120.7 | Protocol implementation conformance statement (PICS) proforma for Clause 120, Physical Medium Attachment (PMA) sublayer, type 200GBASE-R and 400GBASE-R | 198 |
| 120.7.1 | Introduction | 198 |
| 120.7.2 | Identification | 198 |
| 120.7.2.1 | Implementation identification | 198 |
| 120.7.2.2 | Protocol summary | 198 |
| 120.7.3 | Major capabilities/options | 199 |
| 120.7.4 | Skew generation and tolerance | 201 |
| 120.7.5 | Test patterns | 201 |
| 120.7.6 | Loopback modes | 202 |
| 121. | Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-DR4 | 203 |
| 121.1 | Overview | 203 |
| 121.1.1 | Bit error ratio | 203 |
| 121.2 | Physical Medium Dependent (PMD) service interface | 204 |
| 121.3 | Delay and Skew | 205 |
| 121.3.1 | Delay constraints | 205 |
| 121.3.2 | Skew constraints | 205 |
| 121.4 | PMD MDIO function mapping | 206 |
| 121.5 | PMD functional specifications | 206 |
| 121.5.1 | PMD block diagram | 206 |
| 121.5.2 | PMD transmit function | 207 |
| 121.5.3 | PMD receive function | 207 |
| 121.5.4 | PMD global signal detect function | 207 |
| 121.5.5 | PMD lane-by-lane signal detect function | 208 |
| 121.5.6 | PMD reset function | 208 |
| 121.5.7 | PMD global transmit disable function (optional) | 208 |
| 121.5.8 | PMD lane-by-lane transmit disable function (optional) | 209 |
| 121.5.9 | PMD fault function (optional) | 209 |
| 121.5.10 | PMD transmit fault function (optional) | 209 |
| 121.5.11 | PMD receive fault function (optional) | 209 |
| 121.6 | Lane assignments | 209 |
| 121.7 | PMD to MBI optical specifications for 200GBASE-DR4 | 209 |
| 121.7.1 | 200GBASE-DR4 transmitter optical specifications | 210 |
| 121.7.2 | 200GBASE-DR4 receive optical specifications | 210 |
| 121.7.3 | 200GBASE-DR4 illustrative link power budget | 211 |
| 121.8 | Definition of optical parameters and measurement methods | 212 |
| 121.8.1 | Test patterns for optical parameters | 212 |
| 121.8.2 | Wavelength | 212 |
| 121.8.3 | Average optical power | 213 |
| 121.8.4 | Outer Optical Modulation Amplitude (OMA _{outer}) | 213 |
| 121.8.5 | Transmitter and dispersion eye closure for PAM4 (TDECQ) | 213 |
| 121.8.5.1 | TDECQ conformance test setup | 213 |
| 121.8.5.2 | Channel requirements | 214 |
| 121.8.5.3 | TDECQ measurement method | 215 |
| 121.8.5.4 | TDECQ reference equalizer | 218 |
| 121.8.6 | Extinction ratio | 218 |
| 121.8.7 | Relative intensity noise (RIN _{21.4OMA}) | 218 |

| | | |
|------------|--|-----|
| 121.8.8 | Receiver sensitivity | 218 |
| 121.8.9 | Stressed receiver sensitivity | 218 |
| 121.8.9.1 | Stressed receiver conformance test block diagram | 219 |
| 121.8.9.2 | Stressed receiver conformance test signal characteristics and calibration | 220 |
| 121.8.9.3 | Stressed receiver conformance test signal verification | 220 |
| 121.8.9.4 | Sinusoidal jitter for receiver conformance test | 221 |
| 121.9 | Safety, installation, environment, and labeling | 221 |
| 121.9.1 | General safety | 221 |
| 121.9.2 | Laser safety | 222 |
| 121.9.3 | Installation | 222 |
| 121.9.4 | Environment | 222 |
| 121.9.5 | Electromagnetic emission | 222 |
| 121.9.6 | Temperature, humidity, and handling | 222 |
| 121.9.7 | PMD labeling requirements | 222 |
| 121.10 | Fiber optic cabling model | 223 |
| 121.11 | Characteristics of the fiber optic cabling (channel) | 223 |
| 121.11.1 | Optical fiber cable | 224 |
| 121.11.2 | Optical fiber connection | 224 |
| 121.11.2.1 | Connection insertion loss | 224 |
| 121.11.2.2 | Maximum discrete reflectance | 224 |
| 121.11.3 | Medium Dependent Interface (MDI) | 225 |
| 121.11.3.1 | Optical lane assignments | 225 |
| 121.11.3.2 | Medium Dependent Interface (MDI) requirements | 225 |
| 121.12 | Protocol implementation conformance statement (PICS) proforma for Clause 121, Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-DR4 | 226 |
| 121.12.1 | Introduction | 226 |
| 121.12.2 | Identification | 226 |
| 121.12.2.1 | Implementation identification | 226 |
| 121.12.2.2 | Protocol summary | 226 |
| 121.12.3 | Major capabilities/options | 227 |
| 121.12.4 | PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-DR4 | 227 |
| 121.12.4.1 | PMD functional specifications | 227 |
| 121.12.4.2 | Management functions | 228 |
| 121.12.4.3 | PMD to MDI optical specifications for 200GBASE-DR4 | 228 |
| 121.12.4.4 | Optical measurement methods | 229 |
| 121.12.4.5 | Environmental specifications | 229 |
| 121.12.4.6 | Characteristics of the fiber optic cabling and MDI | 229 |
| 122 | Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 | 230 |
| 122.1 | Overview | 230 |
| 122.1.1 | Bit error ratio | 231 |
| 122.2 | Physical Medium Dependent (PMD) service interface | 232 |
| 122.3 | Delay and Skew | 232 |
| 122.3.1 | Delay constraints | 232 |
| 122.3.2 | Skew constraints | 232 |
| 122.4 | PMD MDIO function mapping | 233 |
| 122.5 | PMD functional specifications | 234 |
| 122.5.1 | PMD block diagram | 234 |
| 122.5.2 | PMD transmit function | 234 |
| 122.5.3 | PMD receive function | 235 |
| 122.5.4 | PMD global signal detect function | 235 |

| | | |
|------------|---|-----|
| 122.5.5 | PMD lane-by-lane signal detect function | 236 |
| 122.5.6 | PMD reset function | 236 |
| 122.5.7 | PMD global transmit disable function (optional) | 236 |
| 122.5.8 | PMD lane-by-lane transmit disable function | 236 |
| 122.5.9 | PMD fault function (optional) | 236 |
| 122.5.10 | PMD transmit fault function (optional) | 236 |
| 122.5.11 | PMD receive fault function (optional) | 237 |
| 122.6 | Wavelength-division-multiplexed lane assignments | 237 |
| 122.7 | PMD to MDI optical specifications for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 | 238 |
| 122.7.1 | 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 transmitter optical specifications | 238 |
| 122.7.2 | 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 receive optical specifications | 241 |
| 122.7.3 | 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 illustrative link power budgets | 243 |
| 122.8 | Definition of optical parameters and measurement methods | 243 |
| 122.8.1 | Test patterns for optical parameters | 243 |
| 122.8.2 | Wavelength | 244 |
| 122.8.3 | Average optical power | 244 |
| 122.8.4 | Outer Optical Modulation Amplitude (OMA _{outer}) | 244 |
| 122.8.5 | Transmitter and dispersion eye closure for PAM4 (TDECQ) | 245 |
| 122.8.5.1 | TDECQ conformance test setup | 245 |
| 122.8.5.2 | Channel requirements | 246 |
| 122.8.5.3 | TDECQ measurement method | 247 |
| 122.8.5.4 | TDECQ reference equalizer | 247 |
| 122.8.6 | Extinction ratio | 247 |
| 122.8.7 | Relative intensity noise (RIN _{16.5OMA} and RIN _{15.1OMA}) | 247 |
| 122.8.8 | Receiver sensitivity | 247 |
| 122.8.9 | Stressed receiver sensitivity | 247 |
| 122.8.9.1 | Stressed receiver conformance test block diagram | 248 |
| 122.8.9.2 | Stressed receiver conformance test signal characteristics and calibration | 248 |
| 122.8.9.3 | Stressed receiver conformance test signal verification | 248 |
| 122.9 | Safety, installation, environment, and labeling | 249 |
| 122.9.1 | General safety | 249 |
| 122.9.2 | Laser safety | 249 |
| 122.9.3 | Installation | 250 |
| 122.9.4 | Environment | 250 |
| 122.9.5 | Electromagnetic emission | 250 |
| 122.9.6 | Temperature, humidity, and handling | 250 |
| 122.9.7 | PMD labeling requirements | 250 |
| 122.10 | Fiber optic cabling model | 251 |
| 122.11 | Characteristics of the fiber optic cabling (channel) | 252 |
| 122.11.1 | Optical fiber cable | 252 |
| 122.11.2 | Optical fiber connection | 252 |
| 122.11.2.1 | Connection insertion loss | 252 |
| 122.11.2.2 | Maximum discrete reflectance | 252 |
| 122.11.3 | Medium Dependent Interface (MDI) requirements | 253 |
| 122.12 | Protocol implementation conformance statement (PICS) proforma for Clause 122, Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-FR4, 200GBASE- LR4, 400GBASE-FR8, and 400GBASE-LR8 | 254 |
| 122.12.1 | Introduction | 254 |
| 122.12.2 | Identification | 254 |
| 122.12.2.1 | Implementation identification | 254 |

| | | |
|------------|---|-----|
| 122.12.2.2 | Protocol summary | 254 |
| 122.12.3 | Major capabilities/options | 255 |
| 122.12.4 | PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 | 255 |
| 122.12.4.1 | PMD functional specifications | 255 |
| 122.12.4.2 | Management functions | 256 |
| 122.12.4.3 | PMD to MDI optical specifications for 200GBASE-FR4 | 257 |
| 122.12.4.4 | PMD to MDI optical specifications for 200GBASE-LR4 | 257 |
| 122.12.4.5 | PMD to MDI optical specifications for 400GBASE-FR8 | 257 |
| 122.12.4.6 | PMD to MDI optical specifications for 400GBASE-LR8 | 257 |
| 122.12.4.7 | Optical measurement methods | 258 |
| 122.12.4.8 | Environmental specifications | 258 |
| 122.12.4.9 | Characteristics of the fiber optic cabling and MDI | 258 |
| 123. | Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-SR16 | 259 |
| 123.1 | Overview | 259 |
| 123.1.1 | Bit error ratio | 260 |
| 123.2 | Physical Medium Dependent (PMD) service interface | 260 |
| 123.3 | Delay and Skew | 261 |
| 123.3.1 | Delay constraints | 261 |
| 123.3.2 | Skew constraints | 261 |
| 123.4 | PMD MDIO function mapping | 262 |
| 123.5 | PMD functional specifications | 262 |
| 123.5.1 | PMD block diagram | 262 |
| 123.5.2 | PMD transmit function | 263 |
| 123.5.3 | PMD receive function | 263 |
| 123.5.4 | PMD global signal detect function | 264 |
| 123.5.5 | PMD lane-by-lane signal detect function | 264 |
| 123.5.6 | PMD reset function | 264 |
| 123.5.7 | PMD global transmit disable function (optional) | 265 |
| 123.5.8 | PMD lane-by-lane transmit disable function (optional) | 265 |
| 123.5.9 | PMD fault function (optional) | 265 |
| 123.5.10 | PMD transmit fault function (optional) | 265 |
| 123.5.11 | PMD receive fault function (optional) | 265 |
| 123.6 | Lane assignments | 265 |
| 123.7 | PMD to MDI optical specifications for 400GBASE-SR16 | 266 |
| 123.7.1 | 400GBASE-SR16 transmitter optical specifications | 266 |
| 123.7.2 | 400GBASE-SR16 receive optical specifications | 266 |
| 123.7.3 | 400GBASE-SR16 illustrative link power budget | 266 |
| 123.8 | Definition of optical parameters and measurement methods | 266 |
| 123.8.1 | Test patterns for optical parameters | 266 |
| 123.8.2 | Center wavelength and spectral width | 267 |
| 123.8.3 | Average optical power | 267 |
| 123.8.4 | Optical Modulation Amplitude (OMA) | 267 |
| 123.8.5 | Transmitter and dispersion eye closure (TDEC) | 267 |
| 123.8.6 | Extinction ratio | 267 |
| 123.8.7 | Transmitter optical waveform (transmit eye) | 267 |
| 123.8.8 | Stressed receiver sensitivity | 267 |
| 123.9 | Safety, installation, environment, and labeling | 268 |
| 123.9.1 | General safety | 268 |
| 123.9.2 | Laser safety | 268 |
| 123.9.3 | Installation | 268 |

| | | |
|------------|--|-----|
| 123.9.4 | Environment..... | 268 |
| 123.9.5 | Electromagnetic emission | 268 |
| 123.9.6 | Temperature, humidity, and handling..... | 268 |
| 123.9.7 | PMD labeling requirements | 269 |
| 123.10 | Fiber optic cabling model | 269 |
| 123.11 | Characteristics of the fiber optic cabling (channel) | 269 |
| 123.11.1 | Optical fiber cable..... | 270 |
| 123.11.2 | Optical fiber connection..... | 270 |
| 123.11.2.1 | Connection insertion loss..... | 270 |
| 123.11.2.2 | Maximum discrete reflectance..... | 270 |
| 123.11.3 | Medium Dependent Interface (MDI) | 270 |
| 123.11.3.1 | Optical lane assignments | 271 |
| 123.11.3.2 | Medium Dependent Interface (MDI) requirements | 271 |
| 123.12 | Protocol implementation conformance statement (PICS) proforma for Clause 123, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-SR16..... | 272 |
| 123.12.1 | Introduction..... | 272 |
| 123.12.2 | Identification..... | 272 |
| 123.12.2.1 | Implementation identification..... | 272 |
| 123.12.2.2 | Protocol summary | 272 |
| 123.12.3 | Major capabilities/options..... | 273 |
| 123.12.4 | PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-SR16 | 273 |
| 123.12.4.1 | PMD functional specifications..... | 273 |
| 123.12.4.2 | Management functions..... | 274 |
| 123.12.4.3 | PMD to MDI optical specifications for 40GBASE-SR16 | 274 |
| 123.12.4.4 | Optical measurement methods..... | 275 |
| 123.12.4.5 | Environmental specifications..... | 275 |
| 123.12.4.6 | Characteristics of the fiber optic cabling and MDI..... | 275 |
| 124. | Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-DR4 | 277 |
| 124.1 | Overview..... | 277 |
| 124.1.1 | Bit error ratio | 277 |
| 124.2 | Physical Medium Dependent (PMD) service interface | 278 |
| 124.3 | Delay and Skew | 279 |
| 124.3.1 | Delay constraints..... | 279 |
| 124.3.2 | Skew constraints | 279 |
| 124.4 | PMD MDIO function mapping..... | 279 |
| 124.5 | PMD functional specifications..... | 280 |
| 124.5.1 | PMD block diagram..... | 280 |
| 124.5.2 | PMD transmit function | 281 |
| 124.5.3 | PMD receive function..... | 281 |
| 124.5.4 | PMD global signal detect function | 281 |
| 124.5.5 | PMD lane-by-lane signal detect function | 282 |
| 124.5.6 | PMD reset function..... | 282 |
| 124.5.7 | PMD global transmit disable function (optional) | 282 |
| 124.5.8 | PMD lane-by-lane transmit disable function (optional) | 283 |
| 124.5.9 | PMD fault function (optional) | 283 |
| 124.5.10 | PMD transmit fault function (optional) | 283 |
| 124.5.11 | PMD receive fault function (optional)..... | 283 |
| 124.6 | Lane assignments..... | 283 |
| 124.7 | PMD to MDI optical specifications for 40GBASE-DR4 | 283 |
| 124.7.1 | 40GBASE-DR4 transmitter optical specifications | 284 |
| 124.7.2 | 40GBASE-DR4 receive optical specifications | 284 |

| | | |
|-----------------------|--|-----|
| 124.7.3 | 400GBASE-DR4 illustrative link power budget | 285 |
| 124.8 | Definition of optical parameters and measurement methods | 286 |
| 124.8.1 | Test patterns for optical parameters | 286 |
| 124.8.2 | Wavelength | 286 |
| 124.8.3 | Average optical power | 287 |
| 124.8.4 | Outer Optical Modulation Amplitude (OMA _{outer}) | 287 |
| 124.8.5 | Transmitter and dispersion eye closure for PAM4 (TDECQ) | 287 |
| 124.8.6 | Extinction ratio | 287 |
| 124.8.7 | Relative intensity noise (RIN _{21.4OMA}) | 288 |
| 124.8.8 | Receiver sensitivity | 288 |
| 124.8.9 | Stressed receiver sensitivity | 288 |
| 124.9 | Safety, installation, environment, and labeling | 288 |
| 124.9.1 | General safety | 288 |
| 124.9.2 | Laser safety | 288 |
| 124.9.3 | Installation | 289 |
| 124.9.4 | Environment | 289 |
| 124.9.5 | Electromagnetic emission | 289 |
| 124.9.6 | Temperature, humidity, and handling | 289 |
| 124.9.7 | PMD labeling requirements | 289 |
| 124.10 | Fiber optic cabling model | 289 |
| 124.11 | Characteristics of the fiber optic cabling (channel) | 290 |
| 124.11.1 | Optical fiber cable | 290 |
| 124.11.2 | Optical fiber connection | 291 |
| 124.11.2.1 | Connection insertion loss | 291 |
| 124.11.2.2 | Maximum discrete reflectance | 291 |
| 124.11.3 | Medium Dependent Interface (MDI) | 291 |
| 124.11.3.1 | Optical lane assignments | 291 |
| 124.11.3.2 | Medium Dependent Interface (MDI) requirements | 292 |
| 124.12 | Protocol implementation conformance statement (PICS) proforma for Clause 124, Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-DR4 | 293 |
| 124.12.1 | Introduction | 293 |
| 124.12.2 | Identification | 293 |
| 124.12.2.1 | Implementation identification | 293 |
| 124.12.2.2 | Protocol summary | 293 |
| 124.12.3 | Major capabilities/options | 294 |
| 124.12.4 | PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-DR4 | 294 |
| 124.12.4.1 | PMD functional specifications | 294 |
| 124.12.4.2 | Management functions | 295 |
| 124.12.4.3 | PMD to MDI optical specifications for 400GBASE-DR4 | 295 |
| 124.12.4.4 | Optical measurement methods | 296 |
| 124.12.4.5 | Environmental specifications | 296 |
| 124.12.4.6 | Characteristics of the fiber optic cabling and MDI | 296 |
| Annex A (informative) | Bibliography | 297 |
| Annex 4A (normative) | Simplified full duplex media access control | 298 |
| 4A.4 | Specific implementations | 298 |
| 4A.4.2 | MAC parameters | 298 |
| Annex 31B (normative) | MAC Control PAUSE operation | 299 |
| 31B.3 | Detailed specification of PAUSE operation | 299 |

| | |
|--|-----|
| 31B.3.7 Timing considerations for PAUSE operation | 299 |
| 31B.4 Protocol implementation conformance statement (PICS) proforma for MAC Control PAUSE operation | 299 |
| 31B.4.3 Major capabilities/options | 299 |
| 31B.4.6 PAUSE command MAC timing considerations | 300 |
| Annex 93A (normative) Specification methods for electrical channels | 301 |
| 93A.1 Channel Operating Margin | 301 |
| 93A.1.2 Transmitter and receiver device package models | 302 |
| 93A.1.2.3 Two-port network for the package transmission line | 302 |
| 93A.1.4 Filters | 302 |
| 93A.1.4.3 Receiver equalizer | 302 |
| 93A.1.6 Determination of variable equalizer parameters | 303 |
| 93A.1.7 Interference and noise amplitude | 303 |
| Annex 119A (informative) 200GBASE-R and 400GBASE-R PCS FEC codeword examples | 304 |
| Annex 120A (informative) 200 Gb/s and 400 Gb/s PMA sublayer partitioning examples | 310 |
| 120A.1 Partitioning example supporting 400GBASE-SR16 | 310 |
| 120A.2 Partitioning examples supporting 200GBASE-DR4/FR4/LR4 and 400GBASE-FR8/LR8 | 311 |
| 120A.3 Partitioning examples supporting 400GBASE-DR4 | 313 |
| 120A.4 Partitioning example using 200GXS and 400GXS | 314 |
| Annex 120B (normative) Chip-to-chip 200 Gb/s eight-lane Attachment Unit Interface (200GAUI-8 C2C) and 400 Gb/s sixteen-lane Attachment Unit Interface (400GAUI-16 C2C) | 315 |
| 120B.1 Overview | 315 |
| 120B.2 200GAUI-8 and 400GAUI-16 chip-to-chip compliance point definition | 317 |
| 120B.3 200GAUI-8 and 400GAUI-16 chip-to-chip electrical characteristics | 318 |
| 120B.3.1 200GAUI-8 and 400GAUI-16 C2C transmitter characteristics | 318 |
| 120B.3.2 200GAUI-8 and 400GAUI-16 C2C receiver characteristics | 318 |
| 120B.4 200GAUI-8 and 400GAUI-16 chip-to-chip channel characteristics | 319 |
| 120B.5 Protocol implementation conformance statement (PICS) proforma for Annex 120B, Chip- to-chip 200 Gb/s eight-lane Attachment Unit Interface (200GAUI-8 C2C) and 400 Gb/s sixteen-lane Attachment Unit Interface (400GAUI-16 C2C) | 320 |
| 120B.5.1 Introduction | 320 |
| 120B.5.2 Identification | 320 |
| 120B.5.2.1 Implementation identification | 320 |
| 120B.5.2.2 Protocol summary | 320 |
| 120B.5.3 Major capabilities/options | 321 |
| 120B.5.4 PICS proforma tables for Chip-to-chip 200 Gb/s eight-lane Attachment Unit Interface (200GAUI-8 C2C) and 400 Gb/s sixteen-lane Attachment Unit Interface (400GAUI-16 C2C) | 321 |
| 120B.5.4.1 Transmitter | 321 |
| 120B.5.4.2 Receiver | 322 |
| 120B.5.4.3 Channel | 322 |
| Annex 120C (normative) Chip-to-module 200 Gb/s eight-lane Attachment Unit Interface (200GAUI-8 C2M) and 400 Gb/s sixteen-lane Attachment Unit Interface (400GAUI-16 C2M) | 323 |
| 120C.1 Overview | 323 |

| | |
|--|-----|
| 120C.1.1 Bit error ratio | 325 |
| 120C.2 200GAUI-8 and 400GAUI-16 chip-to-module compliance point definitions | 325 |
| 120C.3 200GAUI-8 and 400GAUI-16 chip-to-module electrical characteristics | 325 |
| 120C.3.1 200GAUI-8 and 400GAUI-16 C2M host output characteristics | 325 |
| 120C.3.2 200GAUI-8 and 400GAUI-16 C2M module output characteristics | 325 |
| 120C.3.3 200GAUI-8 and 400GAUI-16 C2M host input characteristics | 325 |
| 120C.3.4 200GAUI-8 and 400GAUI-16 C2M module input characteristics | 326 |
| 120C.4 200GAUI-8 and 400GAUI-16 C2M measurement methodology | 326 |
| 120C.5 Protocol implementation conformance statement (PICS) proforma for Annex 120C, Chip-to-module 200 Gb/s eight-lane Attachment Unit Interface (200GAUI-8 C2M) and 400 Gb/s sixteen-lane Attachment Unit Interface (400GAUI-16 C2M) | 327 |
| 120C.5.1 Introduction | 327 |
| 120C.5.2 Identification | 327 |
| 120C.5.2.1 Implementation identification | 327 |
| 120C.5.2.2 Protocol summary | 327 |
| 120C.5.3 Major capabilities/options | 328 |
| 120C.5.4 PICS proforma tables for Chip-to-module 200 Gb/s eight-lane Attachment Unit Interface (200GAUI-8 C2M) and 400 Gb/s sixteen-lane Attachment Unit Interface (400GAUI-16 C2M) | 328 |
| 120C.5.4.1 Host output | 328 |
| 120C.5.4.2 Module output | 329 |
| 120C.5.4.3 Host input | 329 |
| 120C.5.4.4 Module input | 329 |
| Annex 120D (normative) Chip-to-chip 200 Gb/s four-lane Attachment Unit Interface (200GAUI-4 C2C) and 400 Gb/s eight-lane Attachment Unit Interface (400GAUI-8 C2C) | 330 |
| 120D.1 Overview | 330 |
| 120D.2 200GAUI-4 and 400GAUI-8 chip-to-chip compliance point definition | 333 |
| 120D.3 200GAUI-4 and 400GAUI-8 chip-to-chip electrical characteristics | 333 |
| 120D.3.1 200GAUI-4 and 400GAUI-8 C2C transmitter characteristics | 333 |
| 120D.3.1.1 Transmitter differential output return loss | 334 |
| 120D.3.1.2 Transmitter linearity | 335 |
| 120D.3.1.2.1 Measurement of mean signal levels | 335 |
| 120D.3.1.3 Linear fit to the measured waveform | 336 |
| 120D.3.1.4 Steady-state voltage and linear fit pulse peak | 336 |
| 120D.3.1.5 Transmitter equalization settings | 336 |
| 120D.3.1.6 Transmitter output noise and distortion | 338 |
| 120D.3.1.7 Transmitter output residual ISI | 338 |
| 120D.3.1.8 Output jitter | 338 |
| 120D.3.1.8.1 J4u and JRMS jitter | 339 |
| 120D.3.1.8.2 Even-odd Jitter | 340 |
| 120D.3.2 200GAUI-4 and 400GAUI-8 C2C receiver characteristics | 340 |
| 120D.3.2.1 Receiver interference tolerance | 340 |
| 120D.3.2.2 Receiver jitter tolerance | 342 |
| 120D.3.2.3 Transmitter equalization feedback (optional) | 342 |
| 120D.4 200GAUI-4 and 400GAUI-8 chip-to-chip channel characteristics | 343 |
| 120D.4.1 Channel Operating Margin | 343 |
| 120D.4.2 Channel return loss | 344 |
| 120D.5 Protocol implementation conformance statement (PICS) proforma for Annex 120D, Chip-to-chip 200 Gb/s four-lane Attachment Unit Interface (200GAUI-4 C2C) and 400 Gb/s eight-lane Attachment Unit Interface (400GAUI-8 C2C) | 346 |
| 120D.5.1 Introduction | 346 |

| | |
|--|-----|
| 120D.5.2 Identification | 346 |
| 120D.5.2.1 Implementation identification | 346 |
| 120D.5.2.2 Protocol summary | 346 |
| 120D.5.3 Major capabilities/options | 347 |
| 120D.5.4 PICS proforma tables for Chip-to-chip 200 Gb/s four-lane Attachment Unit Interface (200GAUI-4 C2C) and 400 Gb/s eight-lane Attachment Unit Interface (400GAUI-8 C2C) | 347 |
| 120D.5.4.1 Transmitter | 347 |
| 120D.5.4.2 Receiver | 348 |
| 120D.5.4.3 Channel | 348 |
| Annex 120E (normative) Chip-to-module 200 Gb/s four-lane Attachment Unit Interface (200GAUI-4 C2M) and 400 Gb/s eight-lane Attachment Unit Interface (400GAUI-8 C2M) | 349 |
| 120E.1 Overview | 349 |
| 120E.1.1 Bit error ratio | 351 |
| 120E.2 200GAUI-4 and 400GAUI-8 chip-to-module compliance point definitions | 351 |
| 120E.3 200GAUI-4 and 400GAUI-8 chip-to-module electrical characteristics | 352 |
| 120E.3.1 200GAUI-4 and 400GAUI-8 C2M host output characteristics | 352 |
| 120E.3.1.1 Signaling rate and range | 353 |
| 120E.3.1.2 Signal levels | 353 |
| 120E.3.1.3 Output return loss | 354 |
| 120E.3.1.4 Differential termination mismatch | 354 |
| 120E.3.1.5 Transition time | 354 |
| 120E.3.1.6 Host output eye width and eye height | 354 |
| 120E.3.1.7 Reference receiver for eye width and eye height evaluation | 355 |
| 120E.3.2 200GAUI-4 and 400GAUI-8 C2M module output characteristics | 357 |
| 120E.3.2.1 Module output eye width, eye height, and pre-cursor ISI ratio | 358 |
| 120E.3.2.1.1 Reference receiver for module output evaluation | 359 |
| 120E.3.2.1.2 Far-end pre-cursor ISI ratio | 359 |
| 120E.3.3 200GAUI-4 and 400GAUI-8 C2M host input characteristics | 359 |
| 120E.3.3.1 Input return loss | 359 |
| 120E.3.3.2 Host stressed input test | 360 |
| 120E.3.3.2.1 Host stressed input test procedure | 360 |
| 120E.3.4 200GAUI-4 and 400GAUI-8 C2M module input characteristics | 362 |
| 120E.3.4.1 Module stressed input test | 362 |
| 120E.3.4.1.1 Module stressed input test procedure | 362 |
| 120E.4 200GAUI-4 and 400GAUI-8 C2M measurement methodology | 364 |
| 120E.4.1 HCB/MCB characteristics | 364 |
| 120E.4.2 Eye width and eye height measurement method | 365 |
| 120E.5 Protocol implementation conformance statement (PICS) proforma for Annex 120E, Chip-to-module 200 Gb/s four-lane Attachment Unit Interface (200GAUI-4 C2M) and 400 Gb/s eight-lane Attachment Unit Interface (400GAUI-8 C2M) | 368 |
| 120E.5.1 Introduction | 368 |
| 120E.5.2 Identification | 368 |
| 120E.5.2.1 Implementation identification | 368 |
| 120E.5.2.2 Protocol summary | 368 |
| 120E.5.3 Major capabilities/options | 369 |
| 120E.5.4 PICS proforma tables for Chip-to-module 200 Gb/s four-lane Attachment Unit Interface (200GAUI-4 C2M) and 400 Gb/s eight-lane Attachment Unit Interface (400GAUI-8 C2M) | 369 |
| 120E.5.4.1 Host output | 369 |
| 120E.5.4.2 Module output | 370 |

| | |
|------------------------------|-----|
| 120E.5.4.3 Host input..... | 370 |
| 120E.5.4.4 Module input..... | 370 |

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/AMD10:2019

IEEE Standard for Ethernet

Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters for 200 Gb/s and 400 Gb/s Operation

NOTE—The editing instructions contained in this amendment define how to merge the material contained therein into the existing base standard and its amendments to form the comprehensive standard.

The editing instructions are shown in ***bold italic***. Four editing instructions are used: change, delete, insert, and replace. ***Change*** is used to make corrections in existing text or tables. The editing instruction specifies the location of the change and describes what is being changed by using ~~striking through~~ (to remove old material) and underline (to add new material). ***Delete*** removes existing material. ***Insert*** adds new material without disturbing the existing material. Deletions and insertions may require renumbering. If so, renumbering instructions are given in the editing instruction. ***Replace*** is used to make changes in figures or equations by removing the existing figure or equation and replacing it with a new one. Editing instructions, change markings, and this NOTE will not be carried over into future editions because the changes will be incorporated into the base standard.

Cross references that refer to clauses, tables, equations, or figures not covered by this amendment are highlighted in green.

1. Introduction

1.1 Overview

1.1.3 Architectural perspectives

1.1.3.2 Compatibility interfaces

Insert the following compatibility interfaces after “100 Gb/s Parallel Physical Interface (CPPI)”:

- q) *200 Gb/s Media Independent Interface (200GMII).* The 200GMII is designed to connect a 200 Gb/s capable MAC to a 200 Gb/s PHY. While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 200 Gb/s speeds. The 200GMII is a logical interconnection intended for use as an intra-chip interface. No mechanical connector is specified for use with the 200GMII. The 200GMII is optional.

- r) *200 Gb/s Attachment Unit Interface (200GAUI-n)*. The 200GAUI-n is a physical instantiation of the PMA service interface to extend the connection between 200 Gb/s capable PMAs. While conformance with implementation of this interface is not necessary to ensure communication, it is recommended, since it allows maximum flexibility in intermixing PHYs and DTEs at 200 Gb/s speeds. The 200GAUI-n is intended for use as a chip-to-chip or a chip-to-module interface. Two widths of 200GAUI-n are defined: an eight-lane version (200GAUI-8) in Annex 120B and Annex 120C, and a four-lane version (200GAUI-4) in Annex 120D and Annex 120E. No mechanical connector is specified for use with the 200GAUI-n. The 200GAUI-n is optional.
- s) *400 Gb/s Media Independent Interface (400GMII)*. The 400GMII is designed to connect a 400 Gb/s capable MAC to a 400 Gb/s PHY. While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 400 Gb/s speeds. The 400GMII is a logical interconnection intended for use as an intra-chip interface. No mechanical connector is specified for use with the 400GMII. The 400GMII is optional.
- t) *400 Gb/s Attachment Unit Interface (400GAUI-n)*. The 400GAUI-n is a physical instantiation of the PMA service interface to extend the connection between 400 Gb/s capable PMAs. While conformance with implementation of this interface is not necessary to ensure communication, it is recommended, since it allows maximum flexibility in intermixing PHYs and DTEs at 400 Gb/s speeds. The 400GAUI-n is intended for use as a chip-to-chip or a chip-to-module interface. Two widths of 400GAUI-n are defined: a sixteen-lane version (400GAUI-16) in Annex 120B and Annex 120C, and an eight-lane version (400GAUI-8) in Annex 120D and Annex 120E. No mechanical connector is specified for use with the 400GAUI-n. The 400GAUI-n is optional.

1.3 Normative references

Insert the following references in alphanumeric order:

ANSI/TIA-604-18:2015, FOCIS 18—Fiber Optic Connector Intermateability Standard—Type MPO-16

IEC 61753-021-2:2007, Fibre optic interconnecting devices and passive components performance standard—Part 021-2: Grade C/3 single-mode fibre optic connectors for category C—Controlled environment.

Change the reference for IEC 60793-2-10 as follows:

IEC 60793-2-10:2014, Optical fibres—Part 2-10: Product specifications—Sectional specification for category A1 multimode fibres.

1.4 Definitions

Insert the following 17 new definitions into the list after 1.4.72a 40GBASE-T (as inserted by IEEE Std 802.3bq-2016):

1.4.72b 200GBASE-DR4: IEEE 802.3 Physical Layer specification for 200 Gb/s using 200GBASE-R encoding and 4-level pulse amplitude modulation over four lanes of single-mode fiber, with reach up to at least 500 m. (See IEEE Std 802.3, Clause 121.)

1.4.72c 200GBASE-FR4: IEEE 802.3 Physical Layer specification for 200 Gb/s using 200GBASE-R encoding and 4-level pulse amplitude modulation over four WDM lanes on single-mode fiber, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 122.)

1.4.72d 200GBASE-LR4: IEEE 802.3 Physical Layer specification for 200 Gb/s using 200GBASE-R encoding and 4-level pulse amplitude modulation over four WDM lanes on single-mode fiber, with reach up to at least 10 km. (See IEEE Std 802.3, Clause 122.)

1.4.72e 200GBASE-R: An IEEE 802.3 family of Physical Layer devices using the physical coding sublayer defined in Clause 119 for 200 Gb/s operation. (See IEEE Std 802.3, Clause 119.)

1.4.72f 200 Gb/s Attachment Unit Interface (200GAUI-n): A physical instantiation of the PMA service interface to extend the connection between 200 Gb/s capable PMAs over n lanes, used for chip-to-chip or chip-to-module interconnections. Two widths of 200GAUI-n are defined: an eight-lane version (200GAUI-8), and a four-lane version (200GAUI-4). (See IEEE Std 802.3, Annex 120B and Annex 120C for 200GAUI-8, or Annex 120D and Annex 120E for 200GAUI-4.)

1.4.72g 200 Gb/s Media Independent Interface (200GMII): The interface between the Reconciliation Sublayer (RS) and the Physical Coding Sublayer (PCS) for 200 Gb/s operation. (See IEEE Std 802.3, Clause 117.)

1.4.72h 200GMII Extender: The 200 Gb/s Media Independent Interface Extender extends the reach of the 200GMII and consists of two 200GXS sublayers with a 200GAUI-n between them. (See IEEE Std 802.3, Clause 118.)

1.4.72i 200GXS: The 200 Gb/s Extender Sublayer (200GXS) is part of the 200GMII Extender. In functionality, it is almost identical to the 200GBASE-R PCS Sublayer defined in Clause 119. Two types of 200GXS are defined: the DTE 200GXS adjacent to the RS sublayer and the PHY 200GXS adjacent to the PHY. (See IEEE Std 802.3, Clause 118.)

1.4.72j 400GBASE-DR4: 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 500 m. (See IEEE Std 802.3, Clause 124.)

1.4.72k 400GBASE-FR8: IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding and 4-level pulse amplitude modulation over eight WDM lanes on single-mode fiber, with reach up to at least 2 km. (See IEEE Std 802.3, Clause 122.)

1.4.72l 400GBASE-LR8: IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding and 4-level pulse amplitude modulation over eight WDM lanes on single-mode fiber, with reach up to at least 10 km. (See IEEE Std 802.3, Clause 122.)

1.4.72m 400GBASE-R: An IEEE 802.3 family of Physical Layer devices using the physical coding sublayer defined in Clause 119 for 400 Gb/s operation. (See IEEE Std 802.3, Clause 119.)

1.4.72n 400GBASE-SR16: IEEE 802.3 Physical Layer specification for 400 Gb/s using 400GBASE-R encoding over sixteen lanes of multimode fiber, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 123.)

1.4.72o 400 Gb/s Attachment Unit Interface (400GAUI-n): A physical instantiation of the PMA service interface to extend the connection between 400 Gb/s capable PMAs over n lanes, used for chip-to-chip or chip-to-module interconnections. Two widths of 400GAUI-n are defined: a sixteen-lane version (400GAUI-16), and an eight-lane version (400GAUI-8). (See IEEE Std 802.3, Annex 120B and Annex 120C for 400GAUI-16, or Annex 120D and Annex 120E for 400GAUI-8.)

1.4.72p 400 Gb/s Media Independent Interface (400GMII): The interface between the Reconciliation Sublayer (RS) and the Physical Coding Sublayer (PCS) for 400 Gb/s operation. (See IEEE Std 802.3, Clause 117.)

1.4.72q 400GMII Extender: The 400 Gb/s Media Independent Interface Extender extends the reach of the 400GMII and consists of two 400GXS sublayers with a 400GAUI-n between them. (See IEEE Std 802.3, Clause 118.)

1.4.72r 400GXS: The 400 Gb/s Extender Sublayer (400GXS) is part of the 400GMII Extender. In functionality, it is almost identical to the 400GBASE-R PCS Sublayer defined in Clause 119. Two types of 400GXS are defined: the DTE 400GXS adjacent to the RS sublayer and the PHY 400GXS adjacent to the PHY. (See IEEE Std 802.3, Clause 118.)

Change 1.4.107 (as modified by IEEE Std 802.3by-2016) as follows:

1.4.107 BASE-R: An IEEE 802.3 family of Physical Layer devices using the 64B/66B encoding defined in Clause 49, Clause 82, ~~or~~ Clause 107, or Clause 119. (See IEEE Std 802.3, Clause 49, Clause 82, ~~or~~ Clause 107, or Clause 119.)

Change 1.4.325 as follows:

1.4.325 PCS lane (PCSL): In 40GBASE-R, ~~and~~ 100GBASE-R, 200GBASE-R, and 400GBASE-R, the PCS distributes encoded data to multiple logical lanes, these logical lanes are called PCS lanes. One or more PCS lanes can be multiplexed and carried on a physical lane together at the PMA service interface. (See IEEE Std 802.3, Clause 83 and Clause 120.)

1.5 Abbreviations

Insert the following new abbreviations into the list, in alphanumeric order:

| | |
|-----------|---|
| 200GAUI-n | 200 Gb/s Attachment Unit Interface over n lanes |
| 200GMII | 200 Gb/s Media Independent Interface |
| 200GXS | 200GMII Extender Sublayer |
| 400GAUI-n | 400 Gb/s Attachment Unit Interface over n lanes |
| 400GMII | 400 Gb/s Media Independent Interface |
| 400GXS | 400GMII Extender Sublayer |
| SER | symbol error ratio |

4. Media Access Control

4.4 Specific implementations

4.4.2 MAC parameters

Change Table 4-2 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bz-2016) as follows:

Table 4–2—MAC parameters

| Parameters | MAC data rate | | | |
|-----------------------------|------------------------------|----------------------|---|----------------------|
| | Up to and including 100 Mb/s | 1 Gb/s | 2.5 Gb/s, 5 Gb/s, 25 Gb/s, 40 Gb/s, and 100 Gb/s, <u>200 Gb/s, and 400 Gb/s</u> | 10 Gb/s |
| slotTime | 512 bit times | 4096 bit times | not applicable | not applicable |
| interPacketGap ^a | 96 bits | 96 bits | 96 bits | 96 bits |
| attemptLimit | 16 | 16 | not applicable | not applicable |
| backoffLimit | 10 | 10 | not applicable | not applicable |
| jamSize | 32 bits | 32 bits | not applicable | not applicable |
| maxBasicFrameSize | 1518 octets | 1518 octets | 1518 octets | 1518 octets |
| maxEnvelopeFrameSize | 2000 octets | 2000 octets | 2000 octets | 2000 octets |
| minFrameSize | 512 bits (64 octets) | 512 bits (64 octets) | 512 bits (64 octets) | 512 bits (64 octets) |
| burstLimit | not applicable | 65 536 bits | not applicable | not applicable |
| ipgStretchRatio | not applicable | not applicable | not applicable | 104 bits |

^a References to interFrameGap or interFrameSpacing in other clauses (e.g., 13, 35, and 42) shall be interpreted as interPacketGap.

Change Note 7 below Table 4-2 as follows:

NOTE 7—For 40 Gb/s, ~~and 100 Gb/s, 200 Gb/s, and 400 Gb/s~~ operation, the received interpacket gap (the spacing between two packets, from the last bit of the FCS field of the first packet to the first bit of the Preamble of the second packet) can have a minimum value of 8 BT (bit times), as measured at the XLGMII, ~~or~~ CGMII, 200GMII, or 400GMII receive signals at the DTE due to clock tolerance and lane alignment requirements.

30. Management

30.2 Managed objects

30.2.5 Capabilities

Insert two new rows between the rows for *aRSFECIndicationEnable* and *aSNRopMarginChn1A* in Table 30-1e (as modified by IEEE Std 802.3bq-2016) as follows (unchanged rows not shown):

Table 30-1e—Capabilities

| | | | DTE | Repeater | MAU | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|--|--|--|---------------------------|----------|-----|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|
| | | | Basic Package (mandatory) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

30.3 Layer management for DTEs

30.3.2 PHY device managed object class

30.3.2.1 PHY device attributes

30.3.2.1.2 aPhyType

Insert 200GBASE-R and 400GBASE-R PHY types at the end of “APPROPRIATE SYNTAX” section of 30.3.2.1.2 as follows:

APPROPRIATE SYNTAX:

| | |
|------------|--|
| ... | |
| 200GBASE-R | Clause 119 200 Gb/s multi-PCS lane 64B/66B |
| 400GBASE-R | Clause 119 400 Gb/s multi-PCS lane 64B/66B |

30.3.2.1.3 aPhyTypeList

Insert 200GBASE-R and 400GBASE-R PHY types at the end of “APPROPRIATE SYNTAX” section of 30.3.2.1.3 as follows:

APPROPRIATE SYNTAX:

| | |
|------------|--|
| ... | |
| 200GBASE-R | Clause 119 200 Gb/s multi-PCS lane 64B/66B |
| 400GBASE-R | Clause 119 400 Gb/s multi-PCS lane 64B/66B |

30.3.2.1.5 aSymbolErrorDuringCarrier

Change the fourth paragraph of the “BEHAVIOUR DEFINED AS” section of 30.3.2.1.5 (as modified by IEEE Std 802.3by-2016) as follows:

BEHAVIOUR DEFINED AS:

For operation at 10 Gb/s, 25 Gb/s, 40 Gb/s, ~~and 100 Gb/s~~, 200 Gb/s, and 400 Gb/s, it is a count of the number of times the receiving media is non-idle (the time between the Start of Packet Delimiter and the End of Packet Delimiter as defined by 46.2.5 and 81.2.5) for a period of time equal to or greater than minFrameSize, and during which there was at least one occurrence of an event that causes the PHY to indicate “Receive Error” on the media independent interface (see Table 46-4 and Table 81-3).

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

Insert new MAU types immediately above 802.9a in “APPROPRIATE SYNTAX” section of 30.5.1.1.2 as follows:

APPROPRIATE SYNTAX:

| | |
|---------------|---|
| ... | |
| 200GBASE-R | Multi-lane PCS as specified in Clause 119 over undefined PMA/PMD |
| 200GBASE-DR4 | 200GBASE-R PCS/PMA over 4-lane single-mode fiber PMD as specified in Clause 121 |
| 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-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 |
| 400GBASE-R | Multi-lane PCS as specified in Clause 119 over undefined PMA/PMD |
| 400GBASE-SR16 | 400GBASE-R PCS/PMA over 16-lane multimode fiber PMD as specified in Clause 123 |
| 400GBASE-DR4 | 400GBASE-R PCS/PMA over 4-lane single-mode fiber PMD as specified in Clause 124 |
| 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-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 |
| ... | |

Change the last paragraph of the “BEHAVIOUR DEFINED AS” section of 30.5.1.1.2 (as modified by IEEE Std 802.3by-2016) as follows:

The enumerations 1000BASE-X, 1000BASE-XHD, 1000BASE-XFD, 10GBASE-X, 10GBASE-R, 10GBASE-W, 25GBASE-R, 40GBASE-R, ~~and 100GBASE-R, 200GBASE-R, and 400GBASE-R~~ shall only be returned if the underlying PMD type is unknown.;

30.5.1.1.4 aMediaAvailable

Change the sixth paragraph of the “BEHAVIOUR DEFINED AS” section of 30.5.1.1.4 (as modified by IEEE Std 802.3bq-2016) as follows:

For 40 Gb/s, ~~and 100 Gb/s~~, 200 Gb/s, and 400 Gb/s, the enumerations map to value of the link_fault variable (see 81.3.4) within the Link Fault Signaling state diagram (see 81.3.4.1 and Figure 46–11) as follows: the values OK and Link Interruption map to the enumeration “available” the value Local Fault maps to the enumeration “not available” and the value Remote Fault maps to the enumeration “remote fault.”

30.5.1.1.12 aLaneMapping

Change 30.5.1.1.12 as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of INTEGERS.

BEHAVIOUR DEFINED AS:

For 40/100/~~200/400~~GBASE-R PHYs and 100GBASE-P PHYs, an array of PCS lane identifiers. The indices of this array (0 to n – 1) denote the service interface lane number where n is the number of PCS lanes in use. Each element of this array contains the PCS lane number for the PCS lane that has been detected in the corresponding service interface lane.

If a Clause 45 MDIO Interface to the PCS is present, then this attribute will map to the Lane mapping registers (see 45.2.3.46 and 45.2.3.47).;

30.5.1.1.15 aFECAbility

Change 30.5.1.1.15 (as modified by IEEE Std 802.3by-2016) as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that meets the requirement of the description below

| | |
|---------------|--|
| unknown | initializing, true state not yet known |
| supported | FEC supported |
| not supported | FEC not supported |

BEHAVIOUR DEFINED AS:

A read-only value that indicates if the PHY supports a FEC sublayer for forward error correction (see 65.2, Clause 74, Clause 91, ~~and Clause 108~~, and Clause 119).

If a Clause 45 MDIO Interface is present and support for FEC is optional, then this attribute maps to the FEC capability register (see 45.2.8.2 or 45.2.1.92).;

30.5.1.1.17 aFECCorrectedBlocks

Change 30.5.1.1.17 (as modified by IEEE Std 802.3by-2016) as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of generalized nonresettable counters. Each counter has a maximum increment rate of 1 200 000 counts per second for 1000 Mb/s implementations, 5 000 000 counts per second for 10 Gb/s, 25 Gb/s, and 40 Gb/s implementations, ~~and 2 500 000 counts per second for 100 Gb/s implementations, 40 000 000 counts per second for 200 Gb/s implementations, and 80 000 000 counts per second for 400 Gb/s implementations.~~

BEHAVIOUR DEFINED AS:

For 1000BASE-PX, 10/25/40/100/200/400GBASE-R, 100GBASE-P, 10GBASE-PR, or 10/1GBASE-PRX PHYs, an array of corrected FEC block counters. The counters do not increment for other PHY types. The indices of this array (0 to N – 1) denote the FEC sublayer instance number where N is the number of FEC sublayer instances in use. The number of FEC sublayer instances in use is set to one for PHYs that do not use PCS lanes or use a single FEC instance for ~~all multiple FEC lanes~~. Each element of this array contains a count of corrected FEC blocks for that FEC sublayer instance.

Increment the counter by one for each received block that is corrected by the FEC function in the PHY for the corresponding lane or FEC sublayer instance.

If a Clause 45 MDIO Interface is present, then this attribute maps to the FEC corrected blocks counter(s) (see 45.2.8.5 and 45.2.1.94 for 10GBASE-R, 45.2.3.39 for 10GBASE-PR and 10/1GBASE-PRX, 45.2.1.116 for BASE-R, ~~and 45.2.1.103 for RS-FEC, and 45.2.3.471 for PCS FEC~~);

30.5.1.1.18 aFECUncorrectableBlocks

Change 30.5.1.1.18 (as modified by IEEE Std 802.3by-2016) as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of generalized nonresettable counters. Each counter has a maximum increment rate of 1 200 000 counts per second for 1000 Mb/s implementations, and 5 000 000 counts per second for 10 Gb/s, 25 Gb/s, and 40 Gb/s implementations, ~~and 2 500 000 counts per second for 100 Gb/s implementations, 40 000 000 counts per second for 200 Gb/s implementations, and 80 000 000 counts per second for 400 Gb/s implementations.~~

BEHAVIOUR DEFINED AS:

For 1000BASE-PX, 10/25/40/100/200/400GBASE-R, 100GBASE-P, 10GBASE-PR, or 10/1GBASE-PRX PHYs, an array of uncorrectable FEC block counters. The counters do not increment for other PHY types. The indices of this array (0 to N – 1) denote the FEC sublayer instance number where N is the number of FEC sublayer instances in use. The number of FEC sublayer instances in use is set to one for PHYs that do not use PCS lanes or use a single FEC instance for ~~all multiple FEC lanes~~. Each element of this array contains a count of uncorrectable ~~corrected~~ FEC blocks for that FEC sublayer instance.

Increment the counter by one for each FEC block that is determined to be uncorrectable by the FEC function in the PHY for the corresponding lane or FEC sublayer instance.

If a Clause 45 MDIO Interface is present, then this attribute maps to the FEC uncorrectable blocks counter(s) (see 45.2.8.6 and 45.2.1.95 for 10GBASE-R, 45.2.3.40 for 10GBASE-PR and 10/1GBASE-PRX, 45.2.1.117 for BASE-R, and 45.2.1.104 for RS-FEC, and 45.2.3.47m for PCS FEC).;

Insert 30.5.1.1.32 and 30.5.1.1.33 after 30.5.1.1.31 as follows:

30.5.1.1.32 aPCSFECIndicationAbility

ATTRIBUTE

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that meets the requirement of the following description:

| | |
|---------------|---|
| unknown | initializing, true state not yet known |
| supported | PCS FEC error indication bypass ability supported |
| not supported | PCS FEC error indication bypass ability not supported |

BEHAVIOUR DEFINED AS:

A read-only value that indicates if the PHY supports an optional PCS FEC error indication bypass ability (see 119.2.5.3).

If a Clause 45 MDIO Interface is present, then this attribute maps to the PCS FEC status register (see 45.2.3.47k).;

30.5.1.1.33 aPCSFECIndicationEnable

ATTRIBUTE

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that meets the requirement of the following description:

| | |
|----------|--|
| unknown | initializing, true state not yet known |
| disabled | PCS FEC error indication bypass disabled |
| enabled | PCS FEC error indication bypass enabled |

BEHAVIOUR DEFINED AS:

A read-write value that indicates the mode of operation of the PCS FEC error indication bypass function (see 119.2.5.3).

A GET operation returns the current mode of operation of the PCS FEC. A SET operation changes the mode of operation of the PCS FEC to the indicated value.

If a Clause 45 MDIO Interface is present, then this attribute maps to the RS-FEC control register (see 45.2.3.47i).;

45. Management Data Input/Output (MDIO) Interface

45.2 MDIO Interface Registers

45.2.1 PMA/PMD registers

Change the first sentence of 45.2.1 (as modified by IEEE Std 802.3by-2016) as follows:

For devices operating at 25 Gb/s or higher speeds, the PMA may be instantiated as multiple sublayers (see 83.1.4, and 109.1.4, and 120.1.4 for how MMD addresses are allocated to multiple PMA sublayers for the respective speeds).

Change the reserved row for 1.23 through 1.29 in Table 45-3 (as modified by IEEE Std 802.3bv-2017) as follows (unchanged rows not shown):

Table 45-3—PMA/PMD registers

| Register address | Register name | Subclause |
|--------------------------|--|-------------------|
| ... | | |
| <u>1.23</u> | <u>200G PMA/PMD extended ability</u> | <u>45.2.1.14e</u> |
| <u>1.24</u> | <u>400G PMA/PMD extended ability</u> | <u>45.2.1.14f</u> |
| <u>1.25, 1.26</u> | <u>Reserved</u> | |
| <u>1.27</u> | <u>PMD transmit disable extension</u> | <u>45.2.1.14g</u> |
| <u>1.28</u> | <u>PMD receive signal detect extension</u> | <u>45.2.1.14h</u> |
| <u>1.23 through 1.29</u> | <u>Reserved</u> | |
| ... | | |

Change the reserved row for 1.340 through 1.699 in Table 45-3 as follows (unchanged rows not shown):

Table 45-3—PMA/PMD registers

| Register address | Register name | Subclause |
|---------------------|---|-----------------------------|
| ... | | |
| 1.340 through 1.399 | Reserved | |
| 1.400 through 1.407 | 200GAUI-8 and 400GAUI-16 chip-to-module recommended CTLE, lane 0 through lane 7 | 45.2.1.116a, 45.2.1.116b |
| 1.408 through 1.415 | 400GAUI-16 chip-to-module recommended CTLE, lane 8 through lane 15 | 45.2.1.116c |
| 1.416 through 1.499 | Reserved | |
| 1.500 through 1.515 | 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 through lane 15 | 45.2.1.116d, 45.2.1.116e |
| 1.516 through 1.531 | 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 0 through lane 15 | 45.2.1.116f, 45.2.1.116g |
| 1.532 through 1.699 | Reserved | |
| ... | | |

Change the reserved row for 1.1511 through 1.1599 in Table 45-3 as follows (unchanged rows not shown):

Table 45-3—PMA/PMD registers

| Register address | Register name | Subclause |
|-----------------------|-------------------------|-------------|
| ... | | |
| 1.1511 through 1.1599 | Reserved | |
| 1.1512 | PRBS13Q testing control | 45.2.1.125a |
| 1.1513 through 1.1599 | Reserved | |
| ... | | |

Change the rows for 1.1600 through 1.1799 in Table 45-3 as follows (unchanged rows not shown):

Table 45-3—PMA/PMD registers

| Register address | Register name | Subclause |
|-------------------------------------|--|------------|
| ... | | |
| 1.1600 through 1.161 599 | PRBS Tx error counters, lane 0 through lane 159 | 45.2.1.126 |
| 1.161 60 through 1.1699 | Reserved | |
| 1.1700 through 1.171 599 | PRBS Rx error counters, lane 0 through lane 159 | 45.2.1.127 |
| 1.171 60 through 1.1799 | Reserved | |
| ... | | |

45.2.1.1 PMA/PMD control 1 register (Register 1.0)

Change the row for bits 1.0.5:2 in Table 45-4 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bz-2016) as follows (unchanged rows not shown):

Table 45-4—PMA/PMD control 1 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|---------|-----------------|--|------------------|
| ... | | | |
| 1.0.5:2 | Speed selection | 5 4 3 2 1 * * * = Reserved 1 1 x x = Reserved 1 0 1 x = Reserved 1 0 0 1 = 400 Gb/s 1 0 0 0 = 200 Gb/s 0 1 1 1 = 5 Gb/s 0 1 1 0 = 2.5 Gb/s 0 1 0 1 = Reserved 0 1 0 0 = 25 Gb/s 0 0 1 1 = 100 Gb/s 0 0 1 0 = 40 Gb/s 0 0 0 1 = 10PASS-TS/2BASE-TL 0 0 0 0 = 10 Gb/s | R/W |
| ... | | | |

^a R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.1.1.3 Speed selection (1.0.13, 1.0.6, 1.0.5:2)

Change the last paragraph of 45.2.1.1.3 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bz-2016) as follows:

When bits 5 through 2 are set to 0010 the use of a 40G PMA/PMD is selected; when set to 0011 the use of a 100G PMA/PMD is selected; when set to 0100 the use of a 25G PMA/PMD is selected; when set to 0110 the use of a 2.5G PMA/PMD is selected; when set to 0111 the use of a 5G PMA/PMD is selected; when set to 1000 the use of a 200G PMA/PMD is selected; when set to 1001 the use of a 400G PMA/PMD is selected. More specific selection is performed using the PMA/PMD control 2 register (Register 1.7) (see 45.2.1.6.3).

45.2.1.1.4 PMA remote loopback (1.0.1)

Change the last two sentences of the second paragraph of 45.2.1.1.4 as follows:

~~For 40/100 Gb/s operation, the remote loopback functionality is detailed in 83.5.9. For 25/40/100 Gb/s operation, the remote loopback ability bit is specified in register 1.13 the 40G/100G PMA/PMD extended ability register. For 200 Gb/s and 400 Gb/s operation, the remote loopback ability bit is specified in registers 1.23 and 1.24, respectively.~~

45.2.1.1.5 PMA local loopback (1.0.0)

Change the second paragraph of 45.2.1.1.5 as follows:

The local loopback function is mandatory for the 100GBASE-KX, 10GBASE-KR, 10GBASE-X, 40GBASE-KR4, 40GBASE-CR4, and 100GBASE-CR10 port type and optional for all other port types, except 2BASE-TL, 10PASS-TS, and 10/1GBASE-PRX, which do not support loopback. ~~A device's ability to perform the local loopback function is advertised in the local loopback ability bit of the related speed-dependent status register. A PMA that is unable to perform the local loopback function shall ignore writes to this bit and shall return a value of zero when read. For 10 Gb/s operation, the~~ The local loopback functionality is detailed in the relevant PMA clause ~~48.3.3 and 51.8. For 40/100 Gb/s operation, the local loopback functionality is detailed in 83.5.8. For 10/25/40/100/200/400 Gb/s operation, the local loopback ability bit is specified in the PMA/PMD status 2 register.~~

45.2.1.2 PMA/PMD status 1 register (Register 1.1)

45.2.1.2.3 Fault (1.1.7)

Change the fourth sentence of 45.2.1.2.3 (as modified by IEEE Std 802.3by-2016) as follows:

For 10/25/40/100/200/400 Gb/s operation, bit 1.1.7 is set to a one when either of the fault bits (1.8.11, 1.8.10) located in register 1.8 are set to a one.

45.2.1.4 PMA/PMD speed ability (Register 1.4)

Change the reserved rows of Table 45–6 (as modified by IEEE Std 802.3bz-2016) as follows (unchanged rows not shown):

Table 45–6—PMA/PMD speed ability register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|--------|--|--|------------------|
| 1.4.15 | Reserved for future speeds- 400G capable | Value always 0 1 = PMA/PMD is capable of operating at 400 Gb/s 0 = PMA/PMD is not capable of operating at 400 Gb/s | RO |
| ... | | | |
| 1.4.12 | Reserved for future speeds- 200G capable | Value always 0 1 = PMA/PMD is capable of operating at 200 Gb/s 0 = PMA/PMD is not capable of operating at 200 Gb/s | RO |
| ... | | | |

^a RO = Read only

Insert 45.2.1.4.aaa before 45.2.1.4.aa (as inserted by IEEE Std 802.3bz-2016) as follows:

45.2.1.4.aaa 400G capable (1.4.15)

When read as a one, bit 1.4.15 indicates that the PMA/PMD is able to operate at a data rate of 400 Gb/s. When read as a zero, bit 1.4.15 indicates that the PMA/PMD is not able to operate at a data rate of 400 Gb/s.

Insert 45.2.1.4.ac after 45.2.1.4.ab (as inserted by IEEE Std 802.3bz-2016) and before 45.2.1.4.a (as inserted by IEEE Std 802.3by-2016) as follows:

45.2.1.4.ac 200G capable (1.4.12)

When read as a one, bit 1.4.12 indicates that the PMA/PMD is able to operate at a data rate of 200 Gb/s. When read as a zero, bit 1.4.12 indicates that the PMA/PMD is not able to operate at a data rate of 200 Gb/s.

45.2.1.6 PMA/PMD control 2 register (Register 1.7)

Change Table 45–7 (as modified by IEEE Std 802.3bw-2015, IEEE Std 802.3by-2016, IEEE Std 802.3bq-2016, IEEE Std 802.3bp-2016, IEEE Std 802.3bn-2016, IEEE Std 802.3bz-2016, and IEEE Std 802.3bv-2017) as follows:

Table 45–7—PMA/PMD control 2 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|----------|-----------------------------|------------------|
| 1.7.15:10 | Reserved | Value always 0 | RO |
| 1.7.9 | PIASE | PMA ingress AUI stop enable | R/W |
| 1.7.8 | PEASE | PMA egress AUI stop enable | R/W |

Table 45–7—PMA/PMD control 2 register bit definitions (*continued*)

| Bit(s) | Name | Description | R/W ^a |
|--------------------------------------|------------------------|--|------------------|
| 1.7.7:6 <u>1.7.7</u> | Reserved | Value always 0 | RO |
| 1.7.5:0 <u>1.7.6:0</u> | PMA/PMD type selection | <u>6 5 4 3 2 1 0</u> <u>1 1 x x x x x = reserved</u> <u>1 0 1 1 1 1 x = reserved</u> <u>1 0 1 1 1 0 1 = reserved</u> <u>1 0 1 1 1 0 0 = 400GBASE-LR8 PMA/PMD</u> <u>1 0 1 1 0 1 1 = 400GBASE-FR8 PMA/PMD</u> <u>1 0 1 1 0 1 0 = 400GBASE-DR4 PMA/PMD</u> <u>1 0 1 1 0 0 1 = 400GBASE-SR16 PMA/PMD</u> <u>1 0 1 1 0 0 0 = reserved</u> <u>1 0 1 0 1 1 x = reserved</u> <u>1 0 1 0 1 0 1 = 200GBASE-LR4 PMA/PMD</u> <u>1 0 1 0 1 0 0 = 200GBASE-FR4 PMA/PMD</u> <u>1 0 1 0 0 1 1 = 200GBASE-DR4 PMA/PMD</u> <u>1 0 1 0 0 1 0 = reserved</u> <u>1 0 1 0 0 0 x = reserved</u> <u>1 0 0 x x x x = reserved</u> <u>0 1 1 1 1 1 x = reserved</u> <u>0 1 1 1 1 0 1 = BASE-T1 PMA/PMD^b</u> <u>0 1 1 1 1 0 0 = reserved</u> <u>0 1 1 1 0 1 1 = reserved</u> <u>0 1 1 1 0 1 0 = 25GBASE-SR PMA/PMD</u> <u>0 1 1 1 0 0 1 = 25GBASE-KR or 25GBASE-KR-S PMA/PMD</u> <u>0 1 1 1 0 0 0 = 25GBASE-CR or 25GBASE-CR-S PMA/PMD</u> <u>0 1 1 0 1 1 1 = 25GBASE-T PMA</u> <u>0 1 1 0 1 1 0 = reserved</u> <u>0 1 1 0 1 0 1 = reserved</u> <u>0 1 1 0 1 0 0 = BASE-H PMA/PMD^c</u> <u>0 1 1 0 0 1 1 = 10GPASS-XR-U PMA/PMD</u> <u>0 1 1 0 0 1 0 = 10GPASS-XR-D PMA/PMD</u> <u>0 1 1 0 0 0 1 = 5GBASE-T PMA</u> <u>0 1 1 0 0 0 0 = 2.5GBASE-T PMA</u> <u>0 1 0 1 1 1 1 = 100GBASE-SR4 PMA/PMD</u> <u>0 1 0 1 1 1 0 = 100GBASE-CR4 PMA/PMD</u> <u>0 1 0 1 1 0 1 = 100GBASE-KR4 PMA/PMD</u> <u>0 1 0 1 1 0 0 = 100GBASE-KP4 PMA/PMD</u> <u>0 1 0 1 0 1 1 = 100GBASE-ER4 PMA/PMD</u> <u>0 1 0 1 0 1 0 = 100GBASE-LR4 PMA/PMD</u> <u>0 1 0 1 0 0 1 = 100GBASE-SR10 PMA/PMD</u> <u>0 1 0 1 0 0 0 = 100GBASE-CR10 PMA/PMD</u> <u>0 1 0 0 1 1 1 = reserved</u> | R/W |

Table 45–7—PMA/PMD control 2 register bit definitions (*continued*)

| Bit(s) | Name | Description | R/W ^a |
|---------------------------------|--|---|------------------|
| 1.7.6:0 (<i>continued</i>) | PMA/PMD type selection (<i>continued</i>) | 0100110 = 40GBASE-T PMA 0100101 = 40GBASE-ER4 PMA/PMD 0100100 = 40GBASE-FR PMA/PMD 0100011 = 40GBASE-LR4 PMA/PMD 0100010 = 40GBASE-SR4 PMA/PMD 0100001 = 40GBASE-CR4 PMA/PMD 0100000 = 40GBASE-KR4 PMA/PMD 0011111 = 10/1GBASE-PRX-U4 0011110 = 10GBASE-PR-U4 0011101 = 10/1GBASE-PRX-D4 0011100 = 10GBASE-PR-D4 0011011 = reserved 0011010 = 10GBASE-PR-U3 0011001 = 10GBASE-PR-U1 0011000 = 10/1GBASE-PRX-U3 0010111 = 10/1GBASE-PRX-U2 0010110 = 10/1GBASE-PRX-U1 0010101 = 10GBASE-PR-D3 0010100 = 10GBASE-PR-D2 0010011 = 10GBASE-PR-D1 0010010 = 10/1GBASE-PRX-D3 0010001 = 10/1GBASE-PRX-D2 0010000 = 10/1GBASE-PRX-D1 0001111 = 10BASE-T PMA/PMD 0001110 = 100BASE-TX PMA/PMD 0001101 = 100BASE-KX PMA/PMD 0001100 = 100BASE-T PMA/PMD 0001011 = 10GBASE-KR PMA/PMD 0001010 = 10GBASE-KX4 PMA/PMD 0001001 = 10GBASE-T PMA 0001000 = 10GBASE-LRM PMA/PMD 0000111 = 10GBASE-SR PMA/PMD 0000110 = 10GBASE-LR PMA/PMD 0000101 = 10GBASE-ER PMA/PMD 0000100 = 10GBASE-LX4 PMA/PMD 0000011 = 10GBASE-SW PMA/PMD 0000010 = 10GBASE-LW PMA/PMD 0000001 = 10GBASE-EW PMA/PMD 0000000 = 10GBASE-CX4 PMA/PMD | |

^a R/W = Read/Write, RO = Read only

^b If BASE-T1 is selected, bits 1.2100.3:0 are used to differentiate which BASE-T1 PMA/PMD is selected.

^c If BASE-H PMA/PMD is selected, register 1.900 is used to differentiate which BASE-H PMA/PMD is selected.

Change the title and text of 45.2.1.6.3 as follows:

45.2.1.6.3 PMA/PMD type selection (1.7.6:0)

The PMA/PMD type of the PMA/PMD shall be selected using bits 56 to 0. The PMA/PMD type abilities of the PMA/PMD are advertised in bits 9 and 7 through 0 of the PMA/PMD status 2 register; the PMA/PMD extended ability register; and the 40G/100G PMA/PMD extended ability register; the 200G PMA/PMD extended ability register; and the 400G PMA/PMD extended ability register. A PMA/PMD shall ignore writes to the PMA/PMD type selection bits that select PMA/PMD types it has not advertised. It is the responsibility of the STA entity to ensure that mutually acceptable MMD types are applied consistently across all the MMDs on a particular PHY.

The PMA/PMD type selection defaults to a supported ability.

45.2.1.7 PMA/PMD status 2 register (Register 1.8)**45.2.1.7.4 Transmit fault (1.8.11)**

Insert the following rows at the bottom of Table 45-9 as follows (unchanged rows not shown):

Table 45-9—Transmit fault description location

| PMA/PMD | Description location |
|---|----------------------|
| ... | |
| 200GBASE-DR4 | 121.5.10 |
| 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, 400GBASE-LR8 | 122.5.10 |
| 400GBASE-SR16 | 123.5.10 |
| 400GBASE-DR4 | 124.5.10 |

45.2.1.7.5 Receive fault (1.8.10)

Insert the following rows at the bottom of Table 45-10 as follows (unchanged rows not shown):

Table 45-10—Receive fault description location

| PMA/PMD | Description location |
|---|----------------------|
| ... | |
| 200GBASE-DR4 | 121.5.11 |
| 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, 400GBASE-LR8 | 122.5.11 |
| 400GBASE-SR16 | 123.5.11 |
| 400GBASE-DR4 | 124.5.11 |

45.2.1.8 PMD transmit disable register (Register 1.9)

Change the first paragraph of 45.2.1.8 as follows:

The assignment of bits in the PMD transmit disable register is shown in Table 45-11. The transmit disable functionality is optional and a PMD's ability to perform the transmit disable functionality is advertised in the PMD transmit disable ability bit 1.8.8. A PMD that does not implement the transmit disable functionality shall ignore writes to the PMD transmit disable register and may return a value of zero for all bits. A PMD device that operates using a single lane and has implemented the transmit disable function shall use bit 1.9.0 to control the function. Such devices shall ignore writes to bits 1.9.1 to 1.9.15 and return a value of zero for those bits when they are read. The description of the transmit disable function for the various PMA/PMDs is given in Table 45-12.

Change the reserved row and footnote a of Table 45–11 and insert new rows at the top of the table as follows (unchanged rows not shown):

Table 45–11—PMD transmit disable register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|----------------------|-------------------------|---|------------------|
| 1.9.15:11 | Reserved | Value always 0 | RO |
| 1.9.15 | PMD transmit disable 14 | 1 = Disable output on transmit lane 14 0 = Enable output on transmit lane 14 | R/W |
| 1.9.14 | PMD transmit disable 13 | 1 = Disable output on transmit lane 13 0 = Enable output on transmit lane 13 | R/W |
| 1.9.13 | PMD transmit disable 12 | 1 = Disable output on transmit lane 12 0 = Enable output on transmit lane 12 | R/W |
| 1.9.12 | PMD transmit disable 11 | 1 = Disable output on transmit lane 11 0 = Enable output on transmit lane 11 | R/W |
| 1.9.11 | PMD transmit disable 10 | 1 = Disable output on transmit lane 10 0 = Enable output on transmit lane 10 | R/W |
| ... | | | |

^a RO = Read only, R/W = Read/Write

Insert the following rows at the bottom of Table 45-12 as follows (unchanged rows not shown):

Table 45–12—Transmit disable description location

| PMA/PMD | Description location |
|---|----------------------|
| ... | |
| 200GBASE-DR4 | 121.5.7 |
| 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 | 122.5.7 |
| 400GBASE-SR16 | 123.5.7 |
| 400GBASE-DR4 | 124.5.7 |

Change the title and text of 45.2.1.8.1 as follows:

45.2.1.8.1 PMD transmit disable 914 (1.9.4015)

When bit 1.9.4015 is set to a one, the PMD shall disable output on lane 914 of the transmit path. When bit 1.9.4015 is set to zero, the PMD shall enable output on lane 914 of the transmit path.

The default value for bit 1.9.4015 is zero.

NOTE—Transmission will not be enabled when this bit is set to zero unless the global PMD transmit disable bit is also zero.

Change the title and text of 45.2.1.8.2 as follows:

45.2.1.8.2 PMD transmit disable ~~4, 5, 6, 7, 8~~ through 13 (1.9.5, ~~1.9.6, 1.9.7, 1.9.8, 1.9.9~~ through 1.9.14)

~~These bits 1.9.5 through 1.9.14 are defined similarly to bit 1.9.4015 for lanes 4 through 13, 5, 6, 7, and 8, respectively.~~

45.2.1.9 PMD receive signal detect register (Register 1.10)

Change the text of 45.2.1.9 as follows:

The assignment of bits in the PMD receive signal detect register is shown in Table 45–13. The 10G-PMD receive signal detect register is mandatory. PMD types that use only a single lane indicate the status of the receive signal detect using bit 1.10.0 and return a value of zero for bits 1.10.4015:1. PMD types that use multiple wavelengths or lanes indicate the status of each lane in bits 1.10.4015:1 and the logical AND of those bits in bit 1.10.0.

Change the reserved row of Table 45–13 and insert new rows at the top of the table as follows (unchanged rows not shown)

Table 45–13—PMD receive signal detect register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------------------|------------------------------|--|------------------|
| 1.10.15:11 | Reserved | Value always 0 | RO |
| 1.10.15 | PMD receive signal detect 14 | 1 = Signal detected on receive lane 14 0 = Signal not detected on receive lane 14 | RO |
| 1.10.14 | PMD receive signal detect 13 | 1 = Signal detected on receive lane 13 0 = Signal not detected on receive lane 13 | RO |
| 1.10.13 | PMD receive signal detect 12 | 1 = Signal detected on receive lane 12 0 = Signal not detected on receive lane 12 | RO |
| 1.10.12 | PMD receive signal detect 11 | 1 = Signal detected on receive lane 11 0 = Signal not detected on receive lane 11 | RO |
| 1.10.11 | PMD receive signal detect 10 | 1 = Signal detected on receive lane 10 0 = Signal not detected on receive lane 10 | RO |
| ... | | | |

^a RO = Read only

Change the title and text of 45.2.1.9.1 as follows:

45.2.1.9.1 PMD receive signal detect ~~914~~ (1.10.4015)

When bit 1.10.4015 is read as a one, a signal has been detected on lane ~~914~~ of the PMD receive path. When bit 1.10.4015 is read as a zero, a signal has not been detected on lane ~~914~~ of the PMD receive path.

Change the title and text of 45.2.1.9.2 as follows:

45.2.1.9.2 PMD receive signal detect ~~4, 5, 6, 7, 8~~ through 13 (1.10.5, ~~1.10.6, 1.10.7, 1.10.8, 1.10.9~~ through 1.10.14)

~~These bits 1.10.5 through 1.10.14~~ are defined similarly to bit 1.10.4015 for lanes ~~4 through 13, 5, 6, 7, and 8~~, respectively.

45.2.1.10 PMA/PMD extended ability register (Register 1.11)

Change the row for 1.11.13 in Table 45-14 (as modified by IEEE Std 802.3bz-2016) as follows (unchanged rows not shown):

Table 45-14—PMA/PMD Extended Ability register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|---------|---|---|------------------|
| ... | | | |
| 1.11.13 | Reserved 200G/400G extended abilities | Value always zero 1 = PMA/PMD has 200G/400G extended abilities listed in register 1.23 or register 1.24 0 = PMA/PMD does not have 200G/400G extended abilities | RO |
| ... | | | |

^a RO = Read only

Insert 45.2.1.10.aab after 45.2.1.10.aaa (as inserted by IEEE Std 802.3bz-2016) and before 45.2.1.10.aa (as inserted by IEEE Std 802.3bz-2016) as follows:

45.2.1.10.aab 200G/400G extended abilities (1.11.13)

When read as a one, bit 1.11.13 indicates that the PMA/PMD has 200G extended abilities listed in register 1.23 or 400G extended abilities listed in register 1.24. When read as a zero, bit 1.11.13 indicates that the PMA/PMD does not have 200G or 400G extended abilities.

Insert 45.2.1.14e through 45.2.1.14h after 45.2.1.14d (as inserted by IEEE Std 802.3bv-2017) as follows:

45.2.1.14e 200G PMA/PMD extended ability register (Register 1.23)

The assignment of bits in the 200G PMA/PMD extended ability register is shown in Table 45-17e.

Table 45-17e—200G PMA/PMD extended ability register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|----------------------------------|--|------------------|
| 1.23.15 | 200G PMA remote loopback ability | 1 = 200G PMA has the ability to perform a remote loopback function 0 = 200G PMA does not have the ability to perform a remote loopback function | RO |
| 1.23.14:6 | Reserved | Value always 0 | RO |

Table 45–17e—200G PMA/PMD extended ability register bit definitions (continued)

| Bit(s) | Name | Description | R/W ^a |
|----------|----------------------|--|------------------|
| 1.23.5 | 200GBASE-LR4 ability | 1 = PMA/PMD is able to perform 200GBASE-LR4 0 = PMA/PMD is not able to perform 200GBASE-LR4 | RO |
| 1.23.4 | 200GBASE-FR4 ability | 1 = PMA/PMD is able to perform 200GBASE-FR4 0 = PMA/PMD is not able to perform 200GBASE-FR4 | RO |
| 1.23.3 | 200GBASE-DR4 ability | 1 = PMA/PMD is able to perform 200GBASE-DR4 0 = PMA/PMD is not able to perform 200GBASE-DR4 | RO |
| 1.23.2:0 | Reserved | Value always 0 | RO |

^a RO = Read only**45.2.1.14e.1 200G PMA remote loopback ability (1.23.15)**

When read as a one, bit 1.23.15 indicates that the 200G PMA is able to perform the remote loopback function. When read as a zero, bit 1.23.15 indicates that the 200G PMA is not able to perform the remote loopback function. If a PMA is able to perform the remote loopback function, then it is controlled using the PMA remote loopback bit 1.0.1 (see 45.2.1.1.4).

45.2.1.14e.2 200GBASE-LR4 ability (1.23.5)

When read as a one, bit 1.23.5 indicates that the PMA/PMD is able to operate as a 200GBASE-LR4 PMA/PMD type. When read as a zero, bit 1.23.5 indicates that the PMA/PMD is not able to operate as a 200GBASE-LR4 PMA/PMD type.

45.2.1.14e.3 200GBASE-FR4 ability (1.23.4)

When read as a one, bit 1.23.4 indicates that the PMA/PMD is able to operate as a 200GBASE-FR4 PMA/PMD type. When read as a zero, bit 1.23.4 indicates that the PMA/PMD is not able to operate as a 200GBASE-FR4 PMA/PMD type.

45.2.1.14e.4 200GBASE-DR4 ability (1.23.3)

When read as a one, bit 1.23.3 indicates that the PMA/PMD is able to operate as a 200GBASE-DR4 PMA/PMD type. When read as a zero, bit 1.23.3 indicates that the PMA/PMD is not able to operate as a 200GBASE-DR4 PMA/PMD type.

45.2.1.14f 400G PMA/PMD extended ability register (Register 1.24)

The assignment of bits in the 400G PMA/PMD extended ability register is shown in Table 45–17f.

Table 45–17f—400G PMA/PMD extended ability register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|----------------------------------|--|------------------|
| 1.24.15 | 400G PMA remote loopback ability | 1 = 400G PMA has the ability to perform a remote loopback function 0 = 400G PMA does not have the ability to perform a remote loopback function | RO |
| 1.24.14:6 | Reserved | Value always 0 | RO |

Table 45–17f—400G PMA/PMD extended ability register bit definitions (continued)

| Bit(s) | Name | Description | R/W ^a |
|----------|-----------------------|--|------------------|
| 1.24.5 | 400GBASE-LR8 ability | 1 = PMA/PMD is able to perform 400GBASE-LR8 0 = PMA/PMD is not able to perform 400GBASE-LR8 | RO |
| 1.24.4 | 400GBASE-FR8 ability | 1 = PMA/PMD is able to perform 400GBASE-FR8 0 = PMA/PMD is not able to perform 400GBASE-FR8 | RO |
| 1.24.3 | 400GBASE-DR4 ability | 1 = PMA/PMD is able to perform 400GBASE-DR4 0 = PMA/PMD is not able to perform 400GBASE-DR4 | RO |
| 1.24.2 | 400GBASE-SR16 ability | 1 = PMA/PMD is able to perform 400GBASE-SR16 0 = PMA/PMD is not able to perform 400GBASE-SR16 | RO |
| 1.24.1:0 | Reserved | Value always 0 | RO |

^a RO = Read only

45.2.1.14f.1 400G PMA remote loopback ability (1.24.15)

When read as a one, bit 1.24.15 indicates that the 400G PMA is able to perform the remote loopback function. When read as a zero, bit 1.24.15 indicates that the 400G PMA is not able to perform the remote loopback function. If a PMA is able to perform the remote loopback function, then it is controlled using the PMA remote loopback bit 1.0.1 (see 45.2.1.1.4).

45.2.1.14f.2 400GBASE-LR8 ability (1.24.5)

When read as a one, bit 1.24.5 indicates that the PMA/PMD is able to operate as a 400GBASE-LR8 PMA/PMD type. When read as a zero, bit 1.24.5 indicates that the PMA/PMD is not able to operate as a 400GBASE-LR8 PMA/PMD type.

45.2.1.14f.3 400GBASE-FR8 ability (1.24.4)

When read as a one, bit 1.24.4 indicates that the PMA/PMD is able to operate as a 400GBASE-FR8 PMA/PMD type. When read as a zero, bit 1.24.4 indicates that the PMA/PMD is not able to operate as a 400GBASE-FR8 PMA/PMD type.

45.2.1.14f.4 400GBASE-DR4 ability (1.24.3)

When read as a one, bit 1.24.3 indicates that the PMA/PMD is able to operate as a 400GBASE-DR4 PMA/PMD type. When read as a zero, bit 1.24.3 indicates that the PMA/PMD is not able to operate as a 400GBASE-DR4 PMA/PMD type.

45.2.1.14f.5 400GBASE-SR16 ability (1.24.2)

When read as a one, bit 1.24.2 indicates that the PMA/PMD is able to operate as a 400GBASE-SR16 PMA/PMD type. When read as a zero, bit 1.24.5 indicates that the PMA/PMD is not able to operate as a 400GBASE-SR16 PMA/PMD type.

45.2.1.14g PMD transmit disable extension register (Register 1.27)

The assignment of bits in the PMD transmit disable extension register is shown in Table 45–17g. The transmit disable functionality is optional and a PMD's ability to perform the transmit disable functionality is advertised in the PMD transmit disable ability bit 1.8.8. A PMD that does not implement the transmit disable functionality shall ignore writes to the PMD transmit disable extension register and may return a value of

zero for all bits. A PMD device that operates using a single lane and has implemented the transmit disable function shall use bit 1.9.0 to control the function. Such devices shall ignore writes to bit 1.27.0 and return a value of zero for that bit when read.

NOTE—Disabling the transmitter on one or more lanes stops the entire link from carrying data.

Table 45–17g—PMD transmit disable extension register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|-------------------------|---|------------------|
| 1.27.15:1 | Reserved | Value always 0 | RO |
| 1.27.0 | PMD transmit disable 15 | 1 = Disable output on transmit lane 15 0 = Enable output on transmit lane 15 | R/W |

^a RO = Read only, R/W = Read/Write

45.2.1.14g.1 PMD transmit disable 15 (1.27.0)

When bit 1.27.0 is set to a one, the PMD shall disable output on lane 15 of the transmit path. When bit 1.27.0 is set to a zero, the PMD shall enable output on lane 15 of the transmit path.

The default value for bit 1.27.0 is zero.

NOTE—Transmission will not be enabled when this bit is set to a zero unless the global PMD transmit disable bit is also zero.

45.2.1.14h PMD receive signal detect extension register (Register 1.28)

The assignment of bits in the PMD receive signal detect extension register is shown in Table 45–17h. PMD types that use only a single lane indicate the status of the receive signal detect using bit 1.10.0 and return a value of zero for bit 1.28.0. PMD types that use multiple wavelengths or lanes indicate the status of each lane in bits 1.10.15:1 and 1.28.0 (if needed) and the logical AND of those bits in bit 1.10.0.

Table 45–17h—PMD receive signal detect extension register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|------------------------------|--|------------------|
| 1.28.15:1 | Reserved | Value always 0 | RO |
| 1.28.0 | PMD receive signal detect 15 | 1 = Signal detected on receive lane 15 0 = Signal not detected on receive lane 15 | RO |

^a RO = Read only

45.2.1.14h.1 PMD receive signal detect 15 (1.28.0)

When bit 1.28.0 is read as a one, a signal has been detected on lane 15 of the PMD receive path. When bit 1.28.0 is read as a zero, a signal has not been detected on lane 15 of the PMD receive path.

Insert 45.2.1.116a through 45.2.1.116g after 45.2.1.116 as follows:

45.2.1.116a 200GAUI-8 and 400GAUI-16 chip-to-module recommended CTLE, lane 0 register (Register 1.400)

The assignment of bits in the 200GAUI-8 and 400GAUI-16 chip-to-module recommended CTLE, lane 0 register is shown in Table 45–90aa.

Table 45–90aa—200GAUI-8 and 400GAUI-16 chip-to-module recommended CTLE, lane 0 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--------------------------|---|------------------|
| 1.400.15:5 | Reserved | Value always 0 | RO |
| 1.400.4:1 | Recommended CTLE peaking | 4 3 2 1 1 1 x x = Reserved 1 0 1 x = Reserved 1 0 0 1 = 9 dB 1 0 0 0 = 8 dB 0 1 1 1 = 7 dB 0 1 1 0 = 6 dB 0 1 0 1 = 5 dB 0 1 0 0 = 4 dB 0 0 1 1 = 3 dB 0 0 1 0 = 2 dB 0 0 0 1 = 1 dB 0 0 0 0 = Reserved | R/W |
| 1.400.0 | Reserved | Value always 0 | RO |

^a R/W = Read/Write, RO = Read only

45.2.1.116a.1 Recommended CTLE peaking (1.400.4:1)

The value of these bits sets the CTLE peaking value recommended for lane 0 by a host that implements the optional 200GAUI-8 or 400GAUI-16 chip-to-module interface defined in Annex 120C (see 120C.3.1). The module may optionally use this information to adjust its CTLE setting.

45.2.1.116b 200GAUI-8 and 400GAUI-16 chip-to-module recommended CTLE, lane 1 through lane 7 registers (Registers 1.401 through 1.407)

The 200GAUI-8 and 400GAUI-16 chip-to-module recommended CTLE, lane 1 through lane 7 registers are defined similarly to register 1.400 (which is used for lane 0, see 45.2.1.116a) but for lanes 1 through 7, respectively.

45.2.1.116c 400GAUI-16 chip-to-module recommended CTLE, lane 8 through lane 15 registers (Registers 1.408 through 1.415)

The 400GAUI-16 chip-to-module recommended CTLE, lane 8 through lane 15 registers are defined similarly to register 1.400 (which is used for lane 0, see 45.2.1.116a) but for lanes 8 through 15, respectively.

45.2.1.116d 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 register (Register 1.500)

The assignment of bits in the 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 0 register is shown in Table 45–90ab. The transmitter, receive direction, is the transmitter that sends data towards the MAC.

Table 45–90ab—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 <i>Requested_eq_c1</i> = 5 (c(1) ratio –0.25) 1 0 0 <i>Requested_eq_c1</i> = 4 (c(1) ratio –0.2) 0 1 1 <i>Requested_eq_c1</i> = 3 (c(1) ratio –0.15) 0 1 0 <i>Requested_eq_c1</i> = 2 (c(1) ratio –0.1) 0 0 1 <i>Requested_eq_c1</i> = 1 (c(1) ratio –0.05) 0 0 0 <i>Requested_eq_c1</i> = 0 (c(1) ratio 0) | RO |
| 1.500.11:10 | Pre-cursor request | 11 10 1 1 <i>Requested_eq_cm1</i> = 3 (c(–1) ratio –0.15) 1 0 <i>Requested_eq_cm1</i> = 2 (c(–1) ratio –0.1) 0 1 <i>Requested_eq_cm1</i> = 1 (c(–1) ratio –0.05) 0 0 <i>Requested_eq_cm1</i> = 0 (c(–1) ratio 0) | RO |
| 1.500.9:7 | Post-cursor remote setting | 9 8 7 1 1 1 Reserved 1 1 0 Reserved 1 0 1 <i>Remote_eq_c1</i> = 5 (c(1) ratio –0.25) 1 0 0 <i>Remote_eq_c1</i> = 4 (c(1) ratio –0.2) 0 1 1 <i>Remote_eq_c1</i> = 3 (c(1) ratio –0.15) 0 1 0 <i>Remote_eq_c1</i> = 2 (c(1) ratio –0.1) 0 0 1 <i>Remote_eq_c1</i> = 1 (c(1) ratio –0.05) 0 0 0 <i>Remote_eq_c1</i> = 0 (c(1) ratio 0) | R/W |
| 1.500.6:5 | Pre-cursor remote setting | 6 5 1 1 <i>Remote_eq_cm1</i> = 3 (c(–1) ratio –0.15) 1 0 <i>Remote_eq_cm1</i> = 2 (c(–1) ratio –0.1) 0 1 <i>Remote_eq_cm1</i> = 1 (c(–1) ratio –0.05) 0 0 <i>Remote_eq_cm1</i> = 0 (c(–1) ratio 0) | R/W |
| 1.500.4:2 | Post-cursor local setting | 4 3 2 1 1 1 Reserved 1 1 0 Reserved 1 0 1 <i>Local_eq_c1</i> = 5 (c(1) ratio –0.25) 1 0 0 <i>Local_eq_c1</i> = 4 (c(1) ratio –0.2) 0 1 1 <i>Local_eq_c1</i> = 3 (c(1) ratio –0.15) 0 1 0 <i>Local_eq_c1</i> = 2 (c(1) ratio –0.1) 0 0 1 <i>Local_eq_c1</i> = 1 (c(1) ratio –0.05) 0 0 0 <i>Local_eq_c1</i> = 0 (c(1) ratio 0) | R/W |
| 1.500.1:0 | Pre-cursor local setting | 1 0 1 1 <i>Local_eq_cm1</i> = 3 (c(–1) ratio –0.15) 1 0 <i>Local_eq_cm1</i> = 2 (c(–1) ratio –0.1) 0 1 <i>Local_eq_cm1</i> = 1 (c(–1) ratio –0.05) 0 0 <i>Local_eq_cm1</i> = 0 (c(–1) ratio 0) | R/W |

^a R/W = Read/Write, RO = Read only

45.2.1.116d.1 Request flag (1.500.15)

The value of this bit indicates the value of the variable *Request_flag* in the lane 0 200GAUI-n or 400GAUI-n receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). This indicates whether the 200GAUI-n or 400GAUI-n chip-to-chip device is issuing a request to change the remote transmitter equalization in the 200GAUI-n or 400GAUI-n chip-to-chip lane 0 transmitter in the receive direction. If a lane 0 200GAUI-n or 400GAUI-n receiver in the receive direction is not present in the package, then the value returned for this bit should be zero.

45.2.1.116d.2 Post-cursor request (1.500.14:12)

The value of these bits indicates the value of the variable *Requested_eq_c1* in the lane 0 200GAUI-n or 400GAUI-n receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). When *Request_flag* is equal to 1, this value indicates the ratio of the post-cursor coefficient $c(1)$, which is requested for the transmitter equalization in the 200GAUI-n or 400GAUI-n chip-to-chip lane 0 transmitter in the receive direction.

45.2.1.116d.3 Pre-cursor request (1.500.11:10)

The value of these bits indicates the value of the variable *Requested_eq_cml* in the lane 0 200GAUI-n or 400GAUI-n receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). When *Request_flag* is equal to 1, this value indicates the ratio of the pre-cursor coefficient $c(-1)$, which is requested for the transmitter equalization in the 200GAUI-n or 400GAUI-n chip-to-chip lane 0 transmitter in the receive direction.

45.2.1.116d.4 Post-cursor remote setting (1.500.9:7)

The value of these bits sets the variable *Remote_eq_c1* for the lane 0 200GAUI-n or 400GAUI-n receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). This is used by a 200GAUI-n or 400GAUI-n receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the post-cursor coefficient $c(1)$ being used in lane 0 of the 200GAUI-n or 400GAUI-n transmitter in the receive direction (see 120B.3.1 and 120D.3.1.5). It may be used to generate values for the request flag and the request bits. If a lane 0 200GAUI-n or 400GAUI-n receiver in the receive direction is not present in the package, then these bits have no effect.

45.2.1.116d.5 Pre-cursor remote setting (1.500.6:5)

The value of these bits sets the variable *Remote_eq_cml* for the lane 0 200GAUI-n or 400GAUI-n receiver in the receive direction (see 120B.3.2 and 120D.3.2.3). This is used by a 200GAUI-n or 400GAUI-n receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the pre-cursor coefficient $c(-1)$ being used in lane 0 of the 200GAUI-n or 400GAUI-n transmitter in the receive direction (see 120B.3.1 and 120D.3.1.5). It may be used to generate values for the request flag and the request bits. If a lane 0 200GAUI-n or 400GAUI-n receiver in the receive direction is not present in the package, then these bits have no effect.

45.2.1.116d.6 Post-cursor local setting (1.500.4:2)

The value of these bits sets the variable *Local_eq_c1* for the lane 0 200GAUI-n or 400GAUI-n transmitter in the receive direction (see 120B.3.1 and 120D.3.1.5), which controls the weight of the transmitter equalization post-cursor coefficient $c(1)$. If a lane 0 200GAUI-n or 400GAUI-n transmitter in the receive direction is not present in the package, then these bits have no effect.

45.2.1.116d.7 Pre-cursor local setting (1.500.1:0)

The value of these bits sets the variable *Local_eq_cml* for the lane 0 200GAUI-n or 400GAUI-n transmitter in the receive direction (see 120B.3.1 and 120D.3.1.5), which controls the weight of the transmitter

equalization pre-cursor coefficient $c(-1)$. If a lane 0 200GAUI-n or 400GAUI-n transmitter in the receive direction is not present in the package, then these bits have no effect.

45.2.1.116e 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 1 through lane 15 registers (Registers 1.501 through 1.515)

The 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, receive direction, lane 1 through lane 15 registers are defined similarly to register 1.500 (which is used for lane 0, see 45.2.1.116d) but for lanes 1 through 15, respectively. The transmitter, receive direction, is the transmitter that sends data towards the MAC.

45.2.1.116f 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 0 register (Register 1.516)

The assignment of bits in the 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 0 register is shown in Table 45–90ac. The transmitter, transmit direction, is the transmitter that sends data towards the PMD.

Table 45–90ac—200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 0 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-------------|----------------------------|--|------------------|
| 1.516.15 | Request flag | 1 = Change in equalization is requested 0 = No change in equalization is requested | RO |
| 1.516.14:12 | Post-cursor request | 14 13 12 1 1 1 Reserved 1 1 0 Reserved 1 0 1 <i>Requested_eq_c1</i> = 5 (c(1) ratio –0.25) 1 0 0 <i>Requested_eq_c1</i> = 4 (c(1) ratio –0.2) 0 1 1 <i>Requested_eq_c1</i> = 3 (c(1) ratio –0.15) 0 1 0 <i>Requested_eq_c1</i> = 2 (c(1) ratio –0.1) 0 0 1 <i>Requested_eq_c1</i> = 1 (c(1) ratio –0.05) 0 0 0 <i>Requested_eq_c1</i> = 0 (c(1) ratio 0) | RO |
| 1.516.11:10 | Pre-cursor request | 11 10 1 1 <i>Requested_eq_cm1</i> = 3 (c(–1) ratio –0.15) 1 0 <i>Requested_eq_cm1</i> = 2 (c(–1) ratio –0.1) 0 1 <i>Requested_eq_cm1</i> = 1 (c(–1) ratio –0.05) 0 0 <i>Requested_eq_cm1</i> = 0 (c(–1) ratio 0) | RO |
| 1.516.9:7 | Post-cursor remote setting | 9 8 7 1 1 1 Reserved 1 1 0 Reserved 1 0 1 <i>Remote_eq_c1</i> = 5 (c(1) ratio –0.25) 1 0 0 <i>Remote_eq_c1</i> = 4 (c(1) ratio –0.2) 0 1 1 <i>Remote_eq_c1</i> = 3 (c(1) ratio –0.15) 0 1 0 <i>Remote_eq_c1</i> = 2 (c(1) ratio –0.1) 0 0 1 <i>Remote_eq_c1</i> = 1 (c(1) ratio –0.05) 0 0 0 <i>Remote_eq_c1</i> = 0 (c(1) ratio 0) | R/W |
| 1.516.6:5 | Pre-cursor remote setting | 6 5 1 1 <i>Remote_eq_cm1</i> = 3 (c(–1) ratio –0.15) 1 0 <i>Remote_eq_cm1</i> = 2 (c(–1) ratio –0.1) 0 1 <i>Remote_eq_cm1</i> = 1 (c(–1) ratio –0.05) 0 0 <i>Remote_eq_cm1</i> = 0 (c(–1) ratio 0) | R/W |

Table 45–90ac—200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 0 register bit definitions (continued)

| Bit(s) | Name | Description | R/W ^a |
|-----------|---------------------------|---|------------------|
| 1.516.4:2 | Post-cursor local setting | 4 3 2 1 1 1 Reserved 1 1 0 Reserved 1 0 1 <i>Local_eq_c1</i> = 5 (c(1) ratio –0.25) 1 0 0 <i>Local_eq_c1</i> = 4 (c(1) ratio –0.2) 0 1 1 <i>Local_eq_c1</i> = 3 (c(1) ratio –0.15) 0 1 0 <i>Local_eq_c1</i> = 2 (c(1) ratio –0.1) 0 0 1 <i>Local_eq_c1</i> = 1 (c(1) ratio –0.05) 0 0 0 <i>Local_eq_c1</i> = 0 (c(1) ratio 0) | R/W |
| 1.516.1:0 | Pre-cursor local setting | 1 0 1 1 <i>Local_eq_cm1</i> = 3 (c(–1) ratio –0.15) 1 0 <i>Local_eq_cm1</i> = 2 (c(–1) ratio –0.1) 0 1 <i>Local_eq_cm1</i> = 1 (c(–1) ratio –0.05) 0 0 <i>Local_eq_cm1</i> = 0 (c(–1) ratio 0) | R/W |

^a R/W = Read/Write, RO = Read only

45.2.1.116f.1 Request flag (1.516.15)

The value of this bit indicates the value of the variable *Request_flag* in the lane 0 200GAUI-n or 400GAUI-n receiver in the transmit direction (see 120B.3.2 and 120D.3.2.3). This indicates whether the 200GAUI-n or 400GAUI-n chip-to-chip device is issuing a request to change the remote transmitter equalization in the 200GAUI-n or 400GAUI-n chip-to-chip lane 0 transmitter in the transmit direction. If a lane 0 200GAUI-n or 400GAUI-n receiver in the transmit direction is not present in the package, then the value returned for this bit should be zero.

45.2.1.116f.2 Post-cursor request (1.516.14:12)

The value of these bits indicates the value of the variable *Requested_eq_c1* in the lane 0 200GAUI-n or 400GAUI-n receiver in the transmit direction (see 120B.3.2 and 120D.3.2.3). When *Request_flag* is equal to 1, this value indicates the ratio of the post-cursor coefficient c(1), which is requested for the transmitter equalization in the 200GAUI-n or 400GAUI-n chip-to-chip lane 0 transmitter in the transmit direction.

45.2.1.116f.3 Pre-cursor request (1.516.11:10)

The value of these bits indicates the value of the variable *Requested_eq_cm1* in the lane 0 200GAUI-n or 400GAUI-n receiver in the transmit direction (see 120B.3.2 and 120D.3.2.3). When *Request_flag* is equal to 1, this value indicates the ratio of the pre-cursor coefficient c(–1), which is requested for the transmitter equalization in the 200GAUI-n or 400GAUI-n chip-to-chip lane 0 transmitter in the transmit direction.

45.2.1.116f.4 Post-cursor remote setting (1.516.9:7)

The value of these bits sets the variable *Remote_eq_c1* for the lane 0 200GAUI-n or 400GAUI-n receiver in the transmit direction (see 120B.3.2 and 120D.3.2.3). This is used by a 200GAUI-n or 400GAUI-n receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the post-cursor coefficient c(1) being used in lane 0 of the 200GAUI-n or 400GAUI-n transmitter in the transmit direction (see 120B.3.1 and 120D.3.1.5). It may be used to generate values for the request flag and the request bits. If a lane 0 200GAUI-n or 400GAUI-n receiver in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.116f.5 Pre-cursor remote setting (1.516.6:5)

The value of these bits sets the variable *Remote_eq_cml* for the lane 0 200GAUI-n or 400GAUI-n receiver in the transmit direction (see 120B.3.2 and 120D.3.2.3). This is used by a 200GAUI-n or 400GAUI-n receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the pre-cursor coefficient $c(-1)$ being used in lane 0 of the 200GAUI-n or 400GAUI-n transmitter in the transmit direction (see 120B.3.1 and 120D.3.1.5). It may be used to generate values for the request flag and the request bits. If a lane 0 200GAUI-n or 400GAUI-n receiver in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.116f.6 Post-cursor local setting (1.516.4:2)

The value of these bits sets the variable *Local_eq_c1* for the lane 0 200GAUI-n or 400GAUI-n transmitter in the transmit direction (see 120B.3.1 and 120D.3.1.5), which controls the weight of the transmitter equalization post-cursor coefficient $c(1)$. If a lane 0 200GAUI-n or 400GAUI-n transmitter in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.116f.7 Pre-cursor local setting (1.516.1:0)

The value of these bits sets the variable *Local_eq_cml* for the lane 0 200GAUI-n or 400GAUI-n transmitter in the transmit direction (see 120B.3.1 and 120D.3.1.5), which controls the weight of the transmitter equalization pre-cursor coefficient $c(-1)$. If a lane 0 200GAUI-n or 400GAUI-n transmitter in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.116g 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 1 through lane 15 registers (Registers 1.517 through 1.531)

The 200GAUI-n and 400GAUI-n chip-to-chip transmitter equalization, transmit direction, lane 1 through lane 15 registers are defined similarly to register 1.516 (which is used for lane 0, see 45.2.1.116f) but for lanes 1 through 15, respectively. The transmitter, transmit direction, is the transmitter that sends data towards the PMD.

45.2.1.123 Test-pattern ability (Register 1.1500)

Change 45.2.1.123 as follows:

The test-pattern ability register is used for PHY types that implement SSPRO, square wave, testing and PRBS testing in the PMA. These functions are described in 83.5.10 and 120.5.11. The assignment of bits in the test-pattern ability register is shown in Table 45–92.

Table 45–92—Test-pattern ability register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|---|-----------------------------------|--|------------------|
| 1.1500.15:13 1.1500.15:14 | Reserved | Value always 0 | RO |
| <u>1.1500.13</u> | <u>SSPRO Tx generator ability</u> | <u>1 = SSPRO transmit direction pattern generator supported</u> <u>0 = SSPRO transmit direction pattern generator not supported</u> | <u>RO</u> |
| 1.1500.12 | Square wave test ability | 1 = Square wave testing supported 0 = Square wave testing not supported | RO |
| 1.1500.11:6 | Reserved | Value always 0 | RO |

Table 45–92—Test-pattern ability register bit definitions (*continued*)

| Bit(s) | Name | Description | R/W ^a |
|-----------|------------------------------|--|------------------|
| 1.1500.11 | PRBS13Q Tx generator ability | 1 = PRBS13Q transmit direction pattern generator supported 0 = PRBS13Q transmit direction pattern generator not supported | RO |
| 1.1500.10 | PRBS13Q Rx generator ability | 1 = PRBS13Q receive direction pattern generator supported 0 = PRBS13Q receive direction pattern generator not supported | RO |
| 1.1500.9 | PRBS31Q Tx generator ability | 1 = PRBS31Q transmit direction pattern generator supported 0 = PRBS31Q transmit direction pattern generator not supported | RO |
| 1.1500.8 | PRBS31Q Tx checker ability | 1 = PRBS31Q transmit direction pattern checker supported 0 = PRBS31Q transmit direction pattern checker not supported | RO |
| 1.1500.7 | PRBS31Q Rx generator ability | 1 = PRBS31Q receive direction pattern generator supported 0 = PRBS31Q receive direction pattern generator not supported | RO |
| 1.1500.6 | PRBS31Q Rx checker ability | 1 = PRBS31Q receive direction pattern checker supported 0 = PRBS31Q receive direction pattern checker not supported | RO |
| 1.1500.5 | PRBS9 Tx generator ability | 1 = PRBS9 transmit direction pattern generator supported 0 = PRBS9 transmit direction pattern generator not supported | RO |
| 1.1500.4 | PRBS9 Rx generator ability | 1 = PRBS9 receive direction pattern generator supported 0 = PRBS9 receive direction pattern generator not supported | RO |
| 1.1500.3 | PRBS31 Tx generator ability | 1 = PRBS31 transmit direction pattern generator supported 0 = PRBS31 transmit direction pattern generator not supported | RO |
| 1.1500.2 | PRBS31 Tx checker ability | 1 = PRBS31 transmit direction pattern checker supported 0 = PRBS31 transmit direction pattern checker not supported | RO |
| 1.1500.1 | PRBS31 Rx generator ability | 1 = PRBS31 receive direction pattern generator supported 0 = PRBS31 receive direction pattern generator not supported | RO |
| 1.1500.0 | PRBS31 Rx checker ability | 1 = PRBS31 receive direction pattern checker supported 0 = PRBS31 receive direction pattern checker not supported | RO |

^a RO = Read only

If SSFQ pattern testing is supported and register 1.1500 is implemented then bit 1.1500.13 shall indicate the generation ability in the transmit direction. The SSFQ pattern test is controlled by register 1.1501 (see 45.2.1.124).

If square wave testing is supported and this register 1.1500 is implemented then bit 1.1500.12 shall be set to 1. The square wave test is controlled by register 1.1510 (see 45.2.1.125).

If PRBS13Q pattern testing is supported and register 1.1500 is implemented then bit 1.1500.11 shall indicate the generation ability in the transmit direction and bit 1.1500.10 shall indicate the generation ability in the receive direction. The PRBS13Q pattern test is controlled by register 1.1512 (see 45.2.1.125a).

If PRBS31Q pattern testing is supported and register 1.1500 is implemented then bit 1.1500.9 shall indicate the generation ability in the transmit direction and bit 1.1500.7 shall indicate the generation ability in the receive direction. Bit 1.1500.8 shall indicate the checker ability in the transmit direction and bit 1.1500.6 shall indicate the checker ability in the receive direction. The PRBS31Q pattern test is controlled by register 1.1501 (see 45.2.1.124).

If PRBS9 pattern testing is supported and this register 1.1500 is implemented then bit 1.1500.5 shall indicate the generation ability in the transmit direction and bit 1.1500.4 shall indicate the generation ability in the receive direction. The PRBS pattern test is controlled by register 1.1501 (see 45.2.1.124).

If PRBS31 pattern testing is supported and this register 1.1500 is implemented then bit 1.1500.3 shall indicate the generation ability in the transmit direction and bit 1.1500.1 shall indicate the generation ability in the receive direction. Bit 1.1500.2 shall indicate the checker ability in the transmit direction and bit 1.1500.0 shall indicate the checker ability in the receive direction. The PRBS pattern test is controlled by register 1.1501 (see 45.2.1.124).

45.2.1.124 PRBS pattern testing control (Register 1.1501)

Change the first paragraph of 45.2.1.124 and Table 45–93 as follows:

The PRBS pattern testing control register is used for PHY types that implement PRBS pattern testing in the PMA. This function is described in 83.5.10 and 120.5.11. The assignment of bits in the PRBS pattern testing control register is shown in Table 45–93.

Table 45–93—PRBS pattern testing control register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-------------|---|---|------------------|
| 1.1501.15:4 | Reserved | Value always 0 | RO |
| 1.1501.14 | SSPRO pattern enable | 1 = Enable SSPRO test-pattern 0 = Disable SSPRO test-pattern | R/W |
| 1.1501.13 | PRBS31Q pattern enable | 1 = Enable PRBS31Q test-pattern 0 = Disable PRBS31Q test-pattern | R/W |
| 1.1501.12 | Reserved | Value always 0 | RO |
| 1.1501.11 | Transmitter linearity test pattern enable | 1 = Enable transmitter linearity test-pattern 0 = Disable transmitter linearity test-pattern | R/W |
| 1.1501.10 | QPRBS13 pattern enable | 1 = Enable QPRBS13 test-pattern 0 = Disable QPRBS13 test-pattern | R/W |
| 1.1501.9 | JP03B pattern enable | 1 = Enable JP03B test-pattern 0 = Disable JP03B test-pattern | R/W |
| 1.1501.8 | JP03A pattern enable | 1 = Enable JP03A test-pattern 0 = Disable JP03A test-pattern | R/W |
| 1.1501.7 | PRBS31 pattern enable | 1 = Enable PRBS31 test-pattern 0 = Disable PRBS31 test-pattern | R/W |
| 1.1501.6 | PRBS9 pattern enable | 1 = Enable PRBS9 test-pattern 0 = Disable PRBS9 test-pattern | R/W |
| 1.1501.5:4 | Reserved | Value always 0 | RO |
| 1.1501.3 | Tx generator enable | 1 = Enable transmit direction test-pattern generator 0 = Disable transmit direction test-pattern generator | R/W |
| 1.1501.2 | Tx checker enable | 1 = Enable transmit direction test-pattern checker 0 = Disable transmit direction test-pattern checker | R/W |
| 1.1501.1 | Rx generator enable | 1 = Enable receive direction test-pattern generator 0 = Disable receive direction test-pattern generator | R/W |
| 1.1501.0 | Rx checker enable | 1 = Enable receive direction test-pattern checker 0 = Disable receive direction test-pattern checker | R/W |

^a R/W = Read/Write, RO = Read only

Change the last two paragraphs of 45.2.1.124 as follows:

Register 1.1501, bit 3 enables pattern PRBS-generation in the transmit direction. Register 1.1501, bit 2 enables pattern PRBS-checking in the transmit direction. Register 1.1501, bit 1 enables pattern PRBS-generation in the receive direction. Register 1.1501, bit 0 enables pattern PRBS-checking in the receive direction. If none ~~neither~~ of the bits 3:0 ~~7 and 6~~ are asserted then bits 7 and 6 ~~3:0~~ have no effect.

Register 1.1501 bit 8 enables testing with the JP03A pattern defined in 94.2.9.1 for 100GBASE-KP4 PMA/PMD. Register 1.1501 bit 9 enables testing with the JP03B pattern defined in 94.2.9.2 for 100GBASE-KP4 PMA/PMD. Register ~~field~~ 1.1501 bit 10 enables testing with the QPRBS13 pattern defined in 94.2.9.3 for 100GBASE-KP4 PMA/PMD. Register ~~field~~ 1.1501 bit 11 enables the transmitter linearity test pattern defined in 94.2.9.4 for 100GBASE-KP4 PMA/PMD. Register 1.1501 bit 13 enables the PRBS31Q test pattern defined in 120.5.11.2.2. Register 1.1501 bit 14 enables the SSPRO test pattern defined in 120.5.11.2.3. The assertion of bits 1.1501.8, 1.1501.9, 1.1501.10, 1.1501.11, 1.1501.13, and 1.1501.14 are mutually exclusive. If more than one bit is asserted, the behavior is undefined. The assertion of 1.1501.8, 1.1501.9, 1.1501.10, and 1.1501.11 operates in conjunction with register 1.1501 bit 3 for 100GBASE-KP4 PMA/PMD. For other PMA/PMD types or if bit 1.1501.3 is not asserted, then 1.1501.8, 1.1501.9, 1.1501.10, and 1.1501.11 have no effect. For the 200GBASE-R and 400GBASE-R PMAs, the assertion of register 1.1501 bit 13 operates in conjunction with register 1.1501 bits 3 and 1. If neither bit is asserted, then register 1.1501 bit 13 has no effect. Also for the 200GBASE-R and 400GBASE-R PMAs, the assertion of register 1.1501 bit 14 operates in conjunction with register 1.1501 bit 3. If bit 1.1501.3 is not asserted, then register 1.1501 bit 14 has no effect. For other PMA/PMD types register 1.1501 bits 13 and 14 have no effect.

45.2.1.125 Square wave testing control (Register 1.1510)

Change 45.2.1.125 as follows:

The square wave testing control register is used for PHY types that implement transmit square wave testing in the PMA. This function is described in 83.5.10 and 120.5.11. The assignment of bits in the square wave testing control register is shown in Table 45–94.

Table 45–94—Square wave testing control register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-------------------------|--------------------------|---|------------------|
| 1.1510.15:10 | Reserved | Value always 0 | R0 |
| <u>1.1510.15</u> | <u>Lane 15 SW enable</u> | <u>1 = Enable square wave on lane 15</u> <u>0 = Disable square wave on lane 15</u> | <u>R/W</u> |
| <u>1.1510.14</u> | <u>Lane 14 SW enable</u> | <u>1 = Enable square wave on lane 14</u> <u>0 = Disable square wave on lane 14</u> | <u>R/W</u> |
| <u>1.1510.13</u> | <u>Lane 13 SW enable</u> | <u>1 = Enable square wave on lane 13</u> <u>0 = Disable square wave on lane 13</u> | <u>R/W</u> |
| <u>1.1510.12</u> | <u>Lane 12 SW enable</u> | <u>1 = Enable square wave on lane 12</u> <u>0 = Disable square wave on lane 12</u> | <u>R/W</u> |
| <u>1.1510.11</u> | <u>Lane 11 SW enable</u> | <u>1 = Enable square wave on lane 11</u> <u>0 = Disable square wave on lane 11</u> | <u>R/W</u> |
| <u>1.1510.10</u> | <u>Lane 10 SW enable</u> | <u>1 = Enable square wave on lane 10</u> <u>0 = Disable square wave on lane 10</u> | <u>R/W</u> |

Table 45–94—Square wave testing control register bit definitions (continued)

| Bit(s) | Name | Description | R/W ^a |
|----------|------------------|---|------------------|
| 1.1510.9 | Lane 9 SW enable | 1 = Enable square wave on lane 9 0 = Disable square wave on lane 9 | R/W |
| 1.1510.8 | Lane 8 SW enable | 1 = Enable square wave on lane 8 0 = Disable square wave on lane 8 | R/W |
| 1.1510.7 | Lane 7 SW enable | 1 = Enable square wave on lane 7 0 = Disable square wave on lane 7 | R/W |
| 1.1510.6 | Lane 6 SW enable | 1 = Enable square wave on lane 6 0 = Disable square wave on lane 6 | R/W |
| 1.1510.5 | Lane 5 SW enable | 1 = Enable square wave on lane 5 0 = Disable square wave on lane 5 | R/W |
| 1.1510.4 | Lane 4 SW enable | 1 = Enable square wave on lane 4 0 = Disable square wave on lane 4 | R/W |
| 1.1510.3 | Lane 3 SW enable | 1 = Enable square wave on lane 3 0 = Disable square wave on lane 3 | R/W |
| 1.1510.2 | Lane 2 SW enable | 1 = Enable square wave on lane 2 0 = Disable square wave on lane 2 | R/W |
| 1.1510.1 | Lane 1 SW enable | 1 = Enable square wave on lane 1 0 = Disable square wave on lane 1 | R/W |
| 1.1510.0 | Lane 0 SW enable | 1 = Enable square wave on lane 0 0 = Disable square wave on lane 0 | R/W |

^a R/W = Read/Write, RO = Read-only

Register 1.1510, bits 0 through 159 enable square wave output on PMA lanes 0 through 159 respectively. Lanes for which a square wave pattern is not enabled act as determined by other registers.

Insert 45.2.1.125a after 45.2.1.125 as follows:

45.2.1.125a PRBS13Q testing control (Register 1.1512)

The PRBS13Q testing control register is used for PHY types that implement PRBS13Q testing in the PMA. This function is described in 120.5.11.2.1. The assignment of bits in the PRBS13Q testing control register is shown in Table 45–94a.

Table 45–94a—PRBS13Q testing control register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-------------|-----------------------|---|------------------|
| 1.1512.15:8 | Reserved | Value always 0 | RO |
| 1.1512.7 | Lane 7 PRBS13Q enable | 1 = Enable PRBS13Q on lane 7 0 = Disable PRBS13Q on lane 7 | R/W |
| 1.1512.6 | Lane 6 PRBS13Q enable | 1 = Enable PRBS13Q on lane 6 0 = Disable PRBS13Q on lane 6 | R/W |

Table 45–94a—PRBS13Q testing control register bit definitions (continued)

| Bit(s) | Name | Description | R/W ^a |
|----------|-----------------------|---|------------------|
| 1.1512.5 | Lane 5 PRBS13Q enable | 1 = Enable PRBS13Q on lane 5 0 = Disable PRBS13Q on lane 5 | R/W |
| 1.1512.4 | Lane 4 PRBS13Q enable | 1 = Enable PRBS13Q on lane 4 0 = Disable PRBS13Q on lane 4 | R/W |
| 1.1512.3 | Lane 3 PRBS13Q enable | 1 = Enable PRBS13Q on lane 3 0 = Disable PRBS13Q on lane 3 | R/W |
| 1.1512.2 | Lane 2 PRBS13Q enable | 1 = Enable PRBS13Q on lane 2 0 = Disable PRBS13Q on lane 2 | R/W |
| 1.1512.1 | Lane 1 PRBS13Q enable | 1 = Enable PRBS13Q on lane 1 0 = Disable PRBS13Q on lane 1 | R/W |
| 1.1512.0 | Lane 0 PRBS13Q enable | 1 = Enable PRBS13Q on lane 0 0 = Disable PRBS13Q on lane 0 | R/W |

^a R/W = Read/Write, RO = Read only

Register 1.1512, bits 0 through 7 enable PRBS13Q output on PMA lanes 0 through 7 respectively. The assertion of register 1.1512 bits 0 through 7 operates in conjunction with register 1.1501 bits 3 and 1. If neither bit is asserted, then register 1.1512 bits 0 through 7 have no effect. Lanes for which a PRBS13Q pattern is not enabled act as determined by other registers.

Change the title and text of 45.2.1.126 as follows:

45.2.1.126 PRBS Tx pattern testing error counter (Register 1.1600 through 1.1615, ~~1.1601, 1.1602, 1.1603, 1.1604, 1.1605, 1.1606, 1.1607, 1.1608, 1.1609~~)

The PRBS Tx pattern testing error counter registers are used for PHY types that implement PRBS Tx pattern testing in the PMA. This function is described in 83.5.10 and 120.5.11. The assignment of bits in the PRBS Tx pattern testing error counter registers are shown in Table 45–95. Register 1.1600 contains the PRBS Tx pattern testing error counter for lane 0, register 1.1601 contains the PRBS Tx pattern testing error counter for lane 1, and registers 1.1602 through 1.1615 contain the PRBS Tx pattern testing error counters for lanes 2 through 159 respectively. Counters corresponding to lanes that are not implemented in a PMA shall read all zeros.

Table 45–95—PRBS Tx pattern testing error counter

| Bit(s) | Name | Description | R/W ^a |
|---------------------------|---------------|-------------|------------------|
| 1.1600 ^b .15:0 | Error counter | | RO, NR |

^a RO = Read only, NR = Non Roll-over.

^b All instances of address 1.1600 also apply to addresses 1.1601 through 1.1615.

The PRBS Tx pattern testing error counter is a 16-bit counter as defined in 83.5.10 and 120.5.11. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

Change the title and content of 45.2.1.127 as follows:

45.2.1.127 PRBS Rx pattern testing error counter (Register 1.1700 through 1.1715, ~~1.1701, 1.1702, 1.1703, 1.1704, 1.1705, 1.1706, 1.1707, 1.1708, 1.1709~~)

The PRBS Rx pattern testing error counter registers are used for PHY types that implement PRBS Rx pattern testing in the PMA. This function is described in 83.5.10 and 120.5.11. The assignment of bits in the PRBS Rx pattern testing error counter registers is identical to the PRBS Tx pattern testing error counter as shown in Table 45–95. Register 1.1700 contains the PRBS Rx pattern testing error counter for lane 0, register 1.1701 contains the PRBS Rx pattern testing error counter for lane 1, and registers 1.1702 through 1.1715 contain the PRBS Rx pattern testing error counters for lanes 2 through 15 respectively. Counters corresponding to lanes that are not implemented in a PMA shall read all zeros.

The PRBS Rx pattern testing error counter is a 16-bit counter as defined in 83.5.10 and 120.5.11. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.3 PCS registers

Change the reserved row for 3.9 through 3.13 in Table 45-119 as follows (unchanged rows not shown):

Table 45–119—PCS registers

| Register address | Register name | Subclause |
|---------------------------|---------------------|------------------|
| ... | | |
| <u>3.9</u> | <u>PCS status 3</u> | <u>45.2.3.7a</u> |
| <u>3.109</u> through 3.13 | Reserved | |
| ... | | |

Change the reserved row for 3.523 through 3.1799 in Table 45-119 (as modified by IEEE Std 802.3bv-2017) as follows (unchanged rows not shown):

Table 45–119—PCS registers

| Register address | Register name | Subclause |
|--|---|--|
| ... | | |
| 3.523 through 3.1799 <u>3.599</u> | Reserved | |
| <u>3.600</u> through 3.631 | <u>PCS FEC symbol error counter, lane 0 to 15</u> | <u>45.2.3.47h</u> <u>45.2.3.47i</u> |
| <u>3.632</u> through 3.799 | Reserved | |
| <u>3.800</u> | <u>PCS FEC control</u> | <u>45.2.3.47j</u> |
| <u>3.801</u> | <u>PCS FEC status</u> | <u>45.2.3.47k</u> |
| <u>3.802, 3.803</u> | <u>PCS FEC corrected codewords counter</u> | <u>45.2.3.47l</u> |
| <u>3.804, 3.805</u> | <u>PCS FEC uncorrected codewords counter</u> | <u>45.2.3.47m</u> |
| <u>3.806, 3.807</u> | <u>PCS FEC degraded SER activate threshold</u> | <u>45.2.3.47n</u> |

Table 45–119—PCS registers (*continued*)

| Register address | Register name | Subclause |
|-----------------------------|--|-------------------|
| <u>3.808, 3.809</u> | <u>PCS FEC degraded SER deactivate threshold</u> | <u>45.2.3.47o</u> |
| <u>3.810, 3.811</u> | <u>PCS FEC degraded SER interval</u> | <u>45.2.3.47p</u> |
| <u>3.812 through 3.1799</u> | <u>Reserved</u> | |
| ... | | |

45.2.3.1 PCS control 1 register (Register 3.0)

Change the row for bits 3.0.5:2 in Table 45-120 (as modified by IEEE Std 802.3bz-2016) as follows (unchanged rows not shown):

Table 45–120—PCS control 1 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|---------|-----------------|--|------------------|
| ... | | | |
| 3.0.5:2 | Speed selection | 5 4 3 2 1 1 x x = Reserved 1 0 1 1 = Reserved 1 0 1 1 = Reserved 1 0 1 0 = 400 Gb/s 1 0 0 1 = 200 Gb/s-Reserved 1 0 0 0 = 5 Gb/s 0 1 1 1 = 2.5 Gb/s 0 1 1 0 = Reserved 0 1 0 1 = 25 Gb/s 0 1 0 0 = 100 Gb/s 0 0 1 1 = 40 Gb/s 0 0 1 0 = 10/1 Gb/s 0 0 0 1 = 10PASS-TS/2BASE-TL 0 0 0 0 = 10 Gb/s | R/W |
| ... | | | |

^a RO = Read only, R/W = Read/Write, SC = Self-clearing

45.2.3.2 PCS status 1 register (Register 3.1)

45.2.3.2.7 PCS receive link status (3.1.2)

Change the third sentence of 45.2.3.2.7 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bq-2016) as follows:

When a 10/25/40/100/200/400GBASE-R, 10GBASE-W, or any MultiGBASE-T mode of operation is selected for the PCS using the PCS type selection field (3.7.3:0), this bit is a latching low version of bit 3.32.12.

45.2.3.4 PCS speed ability (Register 3.4)

Change the reserved row of Table 45-122 (as modified by IEEE Std 802.3bz-2016) as follows (unchanged rows not shown):

Table 45–122—PCS speed ability register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|----------|----------------------------|--|------------------|
| 3.4.15:8 | Reserved for future speeds | Value always 0 | RO |
| 3.4.9 | 400G capable | 1 = PCS is capable of operating at 400 Gb/s 0 = PCS is not capable of operating at 400 Gb/s | RO |
| 3.4.8 | 200G capable | 1 = PCS is capable of operating at 200 Gb/s 0 = PCS is not capable of operating at 200 Gb/s | RO |
| ... | | | |

^a RO = Read only

Insert 45.2.3.4.8 and 45.2.3.4.9 after 45.2.3.4.7 (as inserted by IEEE Std 802.3bz-2016) as follows:

45.2.3.4.8 200G capable (3.4.8)

When read as a one, bit 3.4.8 indicates that the PCS is able to operate at a data rate of 200 Gb/s. When read as a zero, bit 3.4.8 indicates that the PCS is not able to operate at a data rate of 200 Gb/s.

45.2.3.4.9 400G capable (3.4.9)

When read as a one, bit 3.4.9 indicates that the PCS is able to operate at a data rate of 400 Gb/s. When read as a zero, bit 3.4.9 indicates that the PCS is not able to operate at a data rate of 400 Gb/s.

45.2.3.6 PCS control 2 register (Register 3.7)

Change Table 45–123 (as modified by IEEE Std 802.3by-2016, IEEE Std 802.3bq-2016, and IEEE Std 802.3bz-2016) as follows:

Table 45–123—PCS control 2 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|----------|--------------------|---|------------------|
| 3.7.15:4 | Reserved | Value always 0 | RO |
| 3.7.3:0 | PCS type selection | 3 2 1 0 1 1 1 x = reserved 1 1 0 1 = Select 400GBASE-R PCS type 1 1 0 0 = Select 200GBASE-R PCS type 1 0 1 1 = Select 5GBASE-T PCS type 1 0 1 0 = Select 2.5GBASE-T PCS type 1 0 0 1 = Select 25GBASE-T PCS type 1 0 0 0 = reserved 0 1 1 1 = Select 25GBASE-R PCS type 0 1 1 0 = Select 40GBASE-T PCS type 0 1 0 1 = Select 100GBASE-R PCS type 0 1 0 0 = Select 40GBASE-R PCS type 0 0 1 1 = Select 10GBASE-T PCS type 0 0 1 0 = Select 10GBASE-W PCS type 0 0 0 1 = Select 10GBASE-X PCS type 0 0 0 0 = Select 10GBASE-R PCS type | R/W |

^a RO = Read only, R/W = Read/Write

45.2.3.6.1 PCS type selection (3.7.3.0)

Change the second sentence of 45.2.3.6.1 (as modified by IEEE Std 802.3bq-2016) as follows:

The PCS type abilities of the PCS are advertised in bits 3.8.9, ~~and 3.8.7:0,~~ and 3.9.1:0.

Insert 45.2.3.7a after 45.2.3.7 as follows:

45.2.3.7a PCS status 3 register (Register 3.9)

The assignment of bits in the PCS status 3 register is shown in Table 45–124a. All the bits in the PCS status 3 register are read only; a write to the PCS status 3 register shall have no effect.

Table 45–124a—PCS status 3 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|----------|--------------------|--|------------------|
| 3.9.15:2 | Reserved | Value always 0 | RO |
| 3.9.1 | 400GBASE-R capable | 1 = PCS is able to support 400GBASE-R PCS type 0 = PCS is not able to support 400GBASE-R PCS type | RO |
| 3.9.0 | 200GBASE-R capable | 1 = PCS is able to support 200GBASE-R PCS type 0 = PCS is not able to support 200GBASE-R PCS type | RO |

^a RO = Read only

45.2.3.7a.1 400GBASE-R capable (3.9.1)

When read as a one, bit 3.9.1 indicates that the PCS is able to support the 400GBASE-R PCS type. When read as a zero, bit 3.9.1 indicates that the PCS is not able to support the 400GBASE-R PCS type.

45.2.3.7a.2 200GBASE-R capable (3.9.0)

When read as a one, bit 3.9.0 indicates that the PCS is able to support the 200GBASE-R PCS type. When read as a zero, bit 3.9.0 indicates that the PCS is not able to support the 200GBASE-R PCS type.

45.2.3.9a EEE control and capability 2 (Register 3.21)

Change the reserved row for 3.21.15:3 in Table 45–125a (as inserted by IEEE Std 802.3bq-2016) as follows (unchanged rows not shown):

Table 45–125a—EEE control and capability 2 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|--|-----------------------------|--|------------------|
| 3.21.15:3 <u>3.21.15:6</u> | Reserved | Value always 0 | RO |
| <u>3.21.5</u> | <u>400GBASE-R fast wake</u> | <u>1 = EEE fast wake is supported for 400GBASE-R</u> <u>0 = EEE fast wake is not supported for 400GBASE-R</u> | <u>RO</u> |
| <u>3.21.4</u> | <u>Reserved</u> | <u>Value always 0</u> | <u>RO</u> |

Table 45–125a—EEE control and capability 2 register bit definitions (continued)

| Bit(s) | Name | Description | R/W ^a |
|--------|----------------------|--|------------------|
| 3.21.3 | 200GBASE-R fast wake | 1 = EEE fast wake is supported for 200GBASE-R 0 = EEE fast wake is not supported for 200GBASE-R | RO |
| ... | | | |

^a RO = Read only

Insert 45.2.3.9a.a and 45.2.3.9a.b before 45.2.3.9a.1 (as inserted by IEEE Std 802.3bq-2016) as follows:

45.2.3.9a.a 400GBASE-R EEE fast wake supported (3.21.5)

If the PCS supports EEE fast wake operation for 400GBASE-R, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.3.9a.b 200GBASE-R EEE fast wake supported (3.21.3)

If the PCS supports EEE fast wake operation for 200GBASE-R, this bit shall be set to a one; otherwise this bit shall be set to a zero.

45.2.3.13 BASE-R and MultiGBASE-T PCS status 1 register (Register 3.32)**45.2.3.13.1 BASE-R and MultiGBASE-T receive link status (3.32.12)**

Change the last sentence of 45.2.3.13.1 (as modified by IEEE Std 802.3by-2016, IEEE Std 802.3bq-2016, and IEEE Std 802.3bz-2016) as follows:

This bit is a reflection of the PCS_status variable defined in 49.2.14.1 for 10/25GBASE-R, in 126.3.7.1 for 2.5GBASE-T and 5GBASE-T, in 55.3.7.1 for 10GBASE-T, in 113.3.7.1 for 25GBASE-T and 40GBASE-T, and in 82.3.1 for 40/100GBASE-R, and in 119.3 for 200/400GBASE-R.

45.2.3.17 BASE-R PCS test-pattern control register (Register 3.42)

Change the fifth sentence of 45.2.3.17 (as modified by IEEE Std 802.3by-2016) as follows:

Scrambled idle test patterns are defined for 25/40/100/200/400GBASE-R PCS only.

45.2.3.46 Lane 0 mapping register (Register 3.400)

Change the text of 45.2.3.46 as follows:

The assignment of bits in the Lane 0 mapping register is shown in Table 45–160. When the multi-lane PCS described in Clause 82 or Clause 119 detects and locks the alignment marker for service interface lane 0, the detected PCS lane number is recorded in this register. The contents of the Lane 0 mapping register is valid when Lane 0 aligned bit (3.52.0) is set to one and is invalid otherwise.

Insert 45.2.3.47h through 45.2.3.47p after 45.2.3.47g (as inserted by IEEE Std 802.3bv-2017) as follows:

45.2.3.47h PCS FEC symbol error counter lane 0 (Register 3.600, 3.601)

The assignment of bits in the PCS FEC symbol error counter lane 0 register is shown in Table 45–160h. Symbol errors detected in PCS lane 0 are counted and shown in register 3.600.15:0 and 3.601.15:0. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 3.600, 3.601 are used to read the value of a 32-bit counter. When registers 3.600 and 3.601 are used to read the 32-bit counter value, the register 3.600 is read first, the value of the register 3.601 is latched when (and only when) register 3.600 is read, and reads of register 3.601 return the latched value rather than the current value of the counter.

Table 45–160h—PCS FEC symbol error counter lane 0 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-------------------------------------|-----------------------------------|------------------|
| 3.600.15:0 | PCS FEC symbol errors, lane 0 lower | FEC_symbol_error_counter_0[15:0] | RO, NR |
| 3.601.15:0 | PCS FEC symbol errors, lane 0 upper | FEC_symbol_error_counter_0[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.3.47i PCS FEC symbol error counter lane 1 through 15 (Registers 3.602 through 3.631)

The behavior of the PCS FEC symbol error counters, lane 1 through 15 is identical to that described for PCS lane 0 in 45.2.3.47h. Errors detected in each PCS lane are counted and shown in the corresponding register. PCS lane 1, lower 16 bits are shown in register 3.602; PCS lane 1, upper 16 bits are shown in register 3.603; PCS lane 2, lower 16 bits are shown in register 3.604; through register 3.631 for PCS lane 15, upper 16 bits.

45.2.3.47j PCS FEC control register (Register 3.800)

The assignment of bits in the PCS FEC control register is shown in Table 45–160i.

Table 45–160i—PCS FEC control register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|----------------------------------|--|------------------|
| 3.800.15:3 | Reserved | Value always 0 | RO |
| 3.800.2 | PCS FEC degraded SER enable | 1 = FEC decoder signals degraded SER 0 = FEC decoder does not signal degraded SER | R/W |
| 3.800.1 | PCS FEC bypass indication enable | 1 = FEC decoder does not indicate errors 0 = FEC decoder indicates errors | R/W |
| 3.800.0 | Reserved | Value always 0 | RO |

^a R/W = Read/Write, RO = Read only

45.2.3.47j.1 PCS FEC degraded SER enable (3.800.2)

This bit enables the PCS FEC decoder to signal the presence of a degraded SER. When set to a one, this bit enables degraded SER signaling. When set to a zero, degraded SER signaling is disabled. Writes to this bit are ignored and reads return a zero if the PCS FEC does not have the ability to signal the presence of a degraded SER (see 119.2.5.3).

45.2.3.47j.2 PCS FEC bypass indication enable (3.800.1)

This bit enables the PCS FEC decoder to bypass error indication to the upper layers through the sync bits. When set to a one, this bit enables bypass of the error indication. When set to a zero, errors are indicated to the upper layers through the sync bits. Writes to this bit are ignored and reads return a zero if the PCS FEC does not have the ability to bypass indicating decoding errors (see 119.2.5.3).

45.2.3.47k PCS FEC status register (Register 3.801)

The assignment of bits in the PCS FEC status register is shown in Table 45–160j.

Table 45–160j—PCS FEC status register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-----------------------------------|--|------------------|
| 3.801.15:7 | Reserved | Value always 0 | RO |
| 3.801.6 | Local degraded SER received | 1 = local degraded SER received 0 = no local degraded SER received | RO |
| 3.801.5 | Remote degraded SER received | 1 = remote degraded SER received 0 = no remote degraded SER received | RO |
| 3.801.4 | PCS FEC degraded SER | 1 = SER is degraded 0 = SER is not degraded | RO |
| 3.801.3 | PCS FEC degraded SER ability | 1 = FEC decoder has the ability to signal the presence of a degraded SER 0 = FEC decoder does not have the ability to signal the presence of a degraded SER | RO |
| 3.801.2 | PCS FEC high SER | 1 = FEC errors have exceeded threshold 0 = FEC errors have not exceeded threshold | RO/LH |
| 3.801.1 | PCS FEC bypass indication ability | 1 = FEC decoder has the ability to bypass error indication 0 = FEC decoder does not have the ability to bypass error indication | RO |
| 3.801.0 | Reserved | Value always 0 | RO |

^a RO = Read only, LH = Latching high

45.2.3.47k.1 Local degraded SER received (3.801.6)

When read as a one, bit 3.801.6 indicates that the local degraded SER signal has been received. This bit reflects the state of rx_local_degraded (see 119.2.6.2.2).

45.2.3.47k.2 Remote degraded SER received (3.801.5)

When read as a one, bit 3.801.5 indicates that the remote degraded SER signal has been received. This bit reflects the state of rx_rm_degraded (see 119.2.6.2.2).

45.2.3.47k.3 PCS FEC degraded SER (3.801.4)

When read as a one, bit 3.801.4 indicates that the local PCS has detected a degradation of the received signal. This bit reflects the state of the variable FEC_degraded_SER (see 119.2.5.3). The value of bit 3.801.4 is unspecified if the value of the PCS FEC degraded SER activate threshold (registers 3.806 and 3.807) is less than the value of the PCS FEC degraded SER deactivate threshold (registers 3.808 and 3.809).

45.2.3.47k.4 PCS FEC degraded SER ability (3.801.3)

The PCS FEC decoder may have the option to signal the presence of a degraded SER (see 119.2.5.3). This bit is set to one to indicate that the decoder has the ability to signal the presence of a degraded SER. The bit is set to zero if this ability is not supported.

45.2.3.47k.5 PCS FEC high SER (3.801.2)

When PCS FEC bypass indication enable (bit 3.800.1) is set to one, this bit is set to one if the number of FEC symbol errors in a window of 8192 codewords exceeds the threshold (see 119.2.5.3) and is set to zero otherwise. The bit is set to zero if PCS FEC bypass indication enable (bit 3.800.1) is set to zero. This bit shall be implemented with latching high behavior.

45.2.3.47k.6 PCS FEC bypass indication ability (3.801.1)

The PCS FEC decoder may have the option to perform error detection without error indication (see 119.2.5.3) to reduce the delay contributed by the FEC decoder. This bit is set to one to indicate that the decoder has this ability to bypass error indication. The bit is set to zero if this ability is not supported.

45.2.3.47l PCS FEC corrected codewords counter (Register 3.802, 3.803)

The assignment of bits in the PCS FEC corrected codewords counter register is shown in Table 45–160k. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 3.802, 3.803 are used to read the value of a 32-bit counter. When registers 3.802 and 3.803 are used to read the 32-bit counter value, the register 3.802 is read first, the value of the register 3.803 is latched when (and only when) register 3.802 is read, and reads of register 3.803 return the latched value rather than the current value of the counter.

Table 45–160k—PCS FEC corrected codewords counter register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-------------------------------|---------------------------------|------------------|
| 3.802.15:0 | FEC corrected codewords lower | FEC_corrected_cw_counter[15:0] | RO, NR |
| 3.803.15:0 | FEC corrected codewords upper | FEC_corrected_cw_counter[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.3.47m PCS FEC uncorrected codewords counter (Register 3.804, 3.805)

The assignment of bits in the PCS FEC uncorrected codewords counter register is shown in Table 45–160l. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 3.804, 3.805 are used to read the value of a 32-bit counter. When registers 3.804 and 3.805 are used to read the 32-bit counter value, the register 3.804 is read first, the value of the register 3.805 is latched when (and only when) register 3.804 is read, and reads of register 3.805 return the latched value rather than the current value of the counter.

Table 45–160l—PCS FEC uncorrected codewords counter register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|---------------------------------|-----------------------------------|------------------|
| 3.804.15:0 | FEC uncorrected codewords lower | FEC_uncorrected_cw_counter[15:0] | RO, NR |
| 3.805.15:0 | FEC uncorrected codewords upper | FEC_uncorrected_cw_counter[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.3.47n PCS FEC degraded SER activate threshold register (Register 3.806, 3.807)

The assignment of bits in the PCS FEC degraded SER activate threshold register is shown in Table 45–160m. The value controls the threshold used to set the PCS FEC degraded SER bit (3.801.4) as defined in 119.2.5.3.

Table 45–160m—PCS FEC degraded SER activate threshold register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|---|--|------------------|
| 3.806.15:0 | PCS FEC degraded SER activate threshold lower | FEC_degraded_SER_activate_threshold[15:0] | R/W |
| 3.807.15:0 | PCS FEC degraded SER activate threshold upper | FEC_degraded_SER_activate_threshold[31:16] | R/W |

^a R/W = Read/Write

45.2.3.47o PCS FEC degraded SER deactivate threshold register (Register 3.808, 3.809)

The assignment of bits in the PCS FEC degraded SER deactivate threshold register is shown in Table 45–160n. The value controls the threshold used to clear the PCS FEC degraded SER bit (3.801.4) as defined in 119.2.5.3.

Table 45–160n—PCS FEC degraded SER deactivate threshold register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|---|--|------------------|
| 3.808.15:0 | PCS FEC degraded SER deactivate threshold lower | FEC_degraded_SER_deactivate_threshold[15:0] | R/W |
| 3.809.15:0 | PCS FEC degraded SER deactivate threshold upper | FEC_degraded_SER_deactivate_threshold[31:16] | R/W |

^a R/W = Read/Write

45.2.3.47p PCS FEC degraded SER interval register (Register 3.810, 3.811)

The assignment of bits in the PCS FEC degraded SER interval register is shown in Table 45–160o. The value controls the interval used to set and clear the PCS FEC degraded SER bit (3.801.4) as defined in 119.2.5.3.

Table 45–160o—PCS FEC degraded SER interval register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-------------------------------------|----------------------------------|------------------|
| 3.810.15:0 | PCS FEC degraded SER interval lower | FEC_degraded_SER_interval[15:0] | R/W |
| 3.811.15:0 | PCS FEC degraded SER interval upper | FEC_degraded_SER_interval[31:16] | R/W |

^a R/W = Read/Write

45.2.4 PHY XS registers

Change the reserved row for 4.26 through 4.1799 in Table 45-164 as follows (unchanged rows not shown):

Table 45–164—PHY XS registers

| Register address | Register name | Subclause |
|-------------------------------------|--|--|
| ... | | |
| 4.26 through 4.31 4.1799 | Reserved | |
| <u>4.32</u> | <u>BASE-R PHY XS status 1</u> | <u>45.2.4.11a</u> |
| <u>4.33 through 4.41</u> | <u>Reserved</u> | |
| <u>4.42</u> | <u>BASE-R PHY XS test pattern control</u> | <u>45.2.4.11b</u> |
| <u>4.43 through 4.49</u> | <u>Reserved</u> | |
| <u>4.50</u> | <u>Multi-lane BASE-R PHY XS alignment status 1</u> | <u>45.2.4.11c</u> |
| <u>4.51</u> | <u>Reserved</u> | |
| <u>4.52</u> | <u>Multi-lane BASE-R PHY XS alignment status 3</u> | <u>45.2.4.11d</u> |
| <u>4.53</u> | <u>Multi-lane BASE-R PHY XS alignment status 4</u> | <u>45.2.4.11e</u> |
| <u>4.54</u> | <u>Reserved</u> | |
| <u>4.400 through 4.415</u> | <u>PHY XS lane mapping, lane 0 through 15</u> | <u>45.2.4.11f</u> <u>45.2.4.11g</u> |
| <u>4.415 through 4.599</u> | <u>Reserved</u> | |
| <u>4.600 through 4.631</u> | <u>PHY XS FEC symbol error counter, lane 0 to 15</u> | <u>45.2.4.11h</u> <u>45.2.4.11i</u> |
| <u>4.632 through 4.799</u> | <u>Reserved</u> | |
| <u>4.800</u> | <u>PHY XS FEC control</u> | <u>45.2.4.11j</u> |
| <u>4.801</u> | <u>PHY XS FEC status</u> | <u>45.2.4.11k</u> |
| <u>4.802, 4.803</u> | <u>PHY XS FEC corrected codewords counter</u> | <u>45.2.4.11l</u> |

Table 45–164—PHY XS registers (continued)

| Register address | Register name | Subclause |
|-----------------------------|---|-------------------|
| <u>4.804, 4.805</u> | <u>PHY XS FEC uncorrected codewords counter</u> | <u>45.2.4.11m</u> |
| <u>4.806, 4.807</u> | <u>PHY XS FEC degraded SER activate threshold</u> | <u>45.2.4.11n</u> |
| <u>4.808, 4.809</u> | <u>PHY XS FEC degraded SER deactivate threshold</u> | <u>45.2.4.11o</u> |
| <u>4.810, 4.811</u> | <u>PHY XS FEC degraded SER interval</u> | <u>45.2.4.11p</u> |
| <u>4.812 through 4.1799</u> | <u>Reserved</u> | |
| ... | | |

45.2.4.1 PHY XS control 1 register (Register 4.0)

Change the row for 4.0.5:2 in Table 45-165 as follows (unchanged rows not shown):

Table 45–165—PHY XS control 1 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|---------|-----------------|---|------------------|
| ... | | | |
| 4.0.5:2 | Speed selection | <div> <div>5 4 3 2</div> <div> <div>1 * * *</div> <div>* 1 * *</div> <div>* * 1 *</div> <div>1 1 x x</div> <div>1 0 1 1</div> <div>1 0 1 0</div> <div>1 0 0 1</div> <div>1 0 0 0</div> <div>0 1 x x</div> <div>0 0 1 x</div> <div>0 0 0 1</div> <div>0 0 0 0</div> </div> <div> <div>= Reserved</div> <div>= Reserved</div> <div>= Reserved</div> <div>= Reserved</div> <div>= Reserved</div> <div>= 400 Gb/s</div> <div>= 200 Gb/s</div> <div>= Reserved</div> <div>= Reserved</div> <div>= Reserved</div> <div>= Reserved</div> </div> </div> | R/W |
| ... | | | |

^a R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.4.4 PHY XS speed ability (Register 4.4)

Change Table 45-167 as follows:

Table 45–167—PHY XS speed ability register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------------|-----------------------------------|---|------------------|
| <u>4.4.15:10</u> | <u>Reserved for future speeds</u> | Value always 0 | RO |
| <u>4.4.9</u> | <u>400G capable</u> | <div>1 = PHY XS is capable of operating at 400 Gb/s</div> <div>0 = PHY XS is not capable of operating at 400 Gb/s</div> | RO |
| <u>4.4.8</u> | <u>200G capable</u> | <div>1 = PHY XS is capable of operating at 200 Gb/s</div> <div>0 = PHY XS is not capable of operating at 200 Gb/s</div> | RO |

Table 45–167—PHY XS speed ability register bit definitions (continued)

| Bit(s) | Name | Description | R/W ^a |
|---------|-------------|--|------------------|
| 4.4.7:1 | Reserved | Value always 0 | RO |
| 4.4.0 | 10G capable | 1 = PHY XS is capable of operating at 10 Gb/s 0 = PHY XS is not capable of operating at 10 Gb/s | RO |

^a RO = Read only

Insert 45.2.4.4.a and 45.2.4.4.b before 45.2.4.4.1 as follows:

45.2.4.4.a 400G capable (4.4.9)

When read as a one, bit 4.4.9 indicates that the PHY XS is able to operate at a data rate of 400 Gb/s. When read as a zero, bit 4.4.9 indicates that the PHY-XS is not able to operate at a data rate of 400 Gb/s.

45.2.4.4.b 200G capable (4.4.8)

When read as a one, bit 4.4.8 indicates that the PHY XS is able to operate at a data rate of 200 Gb/s. When read as a zero, bit 4.4.8 indicates that the PHY-XS is not able to operate at a data rate of 200 Gb/s.

Insert 45.2.4.11a through 45.2.4.11p after 45.2.4.11 as follows:

45.2.4.11a BASE-R PHY XS status 1 register (Register 4.32)

The assignment of bits in the BASE-R PHY XS status 1 register is shown in Table 45–171a. A PHY XS device that does not implement BASE-R shall return a zero for all bits in the BASE-R PHY XS status 1 register. 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.

Table 45–171a—BASE-R PHY XS status 1 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-----------------------------------|--|------------------|
| 4.32.15:13 | Reserved | Value always 0 | RO |
| 4.32.12 | BASE-R PHY XS receive link status | 1 = BASE-R PHY XS receive link up 0 = BASE-R PHY XS receive link down | RO |
| 4.32.11:0 | Reserved | Value always 0 | RO |

^a RO = Read only

45.2.4.11a.1 BASE-R PHY XS receive link status (4.32.12)

When read as a one, bit 4.32.12 indicates that the PHY XS is in a fully operational state. When read as a zero, bit 4.32.12 indicates that the PHY XS is not fully operational. This bit is a reflection of the PHY XS equivalent of the PCS_status variable defined in 119.3 for 200/400GBASE-R.

45.2.4.11b BASE-R PHY XS test-pattern control register (Register 4.42)

The assignment of bits in the BASE-R PHY XS test-pattern control register is shown in Table 45–171b. This register is only required when the 200/400GBASE-R capability is supported. The test-pattern methodology is described in 119.2.4.9.

Table 45–171b—BASE-R PHY XS test-pattern control register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|------------------------------|---|------------------|
| 4.42.15:4 | Reserved | Value always 0 | RO |
| 4.42.3 | Transmit test-pattern enable | 1 = Enable transmit test pattern 0 = Disable transmit test pattern | R/W |
| 4.42.2:0 | Reserved | Value always 0 | RO |

^a RO = Read only, R/W = Read/Write

45.2.4.11b.1 Transmit test-pattern enable (4.42.3)

When bit 4.42.3 is set to a one, pattern testing is enabled on the transmit path. When bit 4.42.3 is set to a zero, pattern testing is disabled on the transmit path.

The default value for bit 4.42.3 is zero.

45.2.4.11c Multi-lane BASE-R PHY XS alignment status 1 register (Register 4.50)

The assignment of bits in the multi-lane BASE-R PHY XS alignment status 1 register is shown in Table 45–171c. A PHY XS device that does not implement multi-lane BASE-R PHY XS shall return a zero for all bits in the multi-lane BASE-R PHY XS alignment status 1 register. 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.

Table 45–171c—Multi-lane BASE-R PHY XS alignment status 1 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|------------------------------|--|------------------|
| 4.50.15:13 | Reserved | Value always 0 | RO |
| 4.50.12 | PHY XS lane alignment status | 1 = PHY XS receive lanes locked and aligned 0 = PHY XS receive lanes not locked and aligned | RO |
| 4.50.11:0 | Reserved | Value always 0 | RO |

^a RO = Read only

45.2.4.11c.1 PHY XS lane alignment status (4.50.12)

When read as a one, bit 4.50.12 indicates that the PHY XS has locked and aligned all receive lanes. When read as a zero, bit 4.50.12 indicates that the PHY XS has not locked and aligned all receive lanes.

45.2.4.11d Multi-lane BASE-R PHY XS alignment status 3 register (Register 4.52)

The assignment of bits in the multi-lane BASE-R PHY XS alignment status 3 register is shown in Table 45–171d. A PHY XS device that does not implement multi-lane BASE-R PHY XS shall return a zero for all bits in the multi-lane BASE-R PHY XS alignment status 3 register. A device that implements multi-lane BASE-R PHY XS shall return a zero for all bits in the multi-lane BASE-R PHY XS alignment status 3 register that are not required for the PHY XS 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.

Table 45–171d—Multi-lane BASE-R PHY XS alignment status 3 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|----------------|--|------------------|
| 4.52.15:8 | Reserved | Value always 0 | RO |
| 4.52.7 | Lane 7 aligned | 1 = Lane 7 alignment marker is locked 0 = Lane 7 alignment marker is not locked | RO |
| 4.52.6 | Lane 6 aligned | 1 = Lane 6 alignment marker is locked 0 = Lane 6 alignment marker is not locked | RO |
| 4.52.5 | Lane 5 aligned | 1 = Lane 5 alignment marker is locked 0 = Lane 5 alignment marker is not locked | RO |
| 4.52.4 | Lane 4 aligned | 1 = Lane 4 alignment marker is locked 0 = Lane 4 alignment marker is not locked | RO |
| 4.52.3 | Lane 3 aligned | 1 = Lane 3 alignment marker is locked 0 = Lane 3 alignment marker is not locked | RO |
| 4.52.2 | Lane 2 aligned | 1 = Lane 2 alignment marker is locked 0 = Lane 2 alignment marker is not locked | RO |
| 4.52.1 | Lane 1 aligned | 1 = Lane 1 alignment marker is locked 0 = Lane 1 alignment marker is not locked | RO |
| 4.52.0 | Lane 0 aligned | 1 = Lane 0 alignment marker is locked 0 = Lane 0 alignment marker is not locked | RO |

^a RO = Read only

45.2.4.11d.1 Lane 7 aligned (4.52.7)

When read as a one, bit 4.52.7 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 7. When read as a zero, bit 4.52.7 indicates that the PHY XS receiver lane 7 has not achieved alignment marker lock. This bit reflects the state of amps_lock[7] (see 119.2.6.2.2).

45.2.4.11d.2 Lane 6 aligned (4.52.6)

When read as a one, bit 4.52.6 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 6. When read as a zero, bit 4.52.6 indicates that the PHY XS receiver lane 6 has not achieved alignment marker lock. This bit reflects the state of amps_lock[6] (see 119.2.6.2.2).

45.2.4.11d.3 Lane 5 aligned (4.52.5)

When read as a one, bit 4.52.5 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 5. When read as a zero, bit 4.52.5 indicates that the PHY XS receiver lane 5 has not achieved alignment marker lock. This bit reflects the state of amps_lock[5] (see 119.2.6.2.2).

45.2.4.11d.4 Lane 4 aligned (4.52.4)

When read as a one, bit 4.52.4 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 4. When read as a zero, bit 4.52.4 indicates that the PHY XS receiver lane 4 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[4]` (see 119.2.6.2.2).

45.2.4.11d.5 Lane 3 aligned (4.52.3)

When read as a one, bit 4.52.3 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 3. When read as a zero, bit 4.52.3 indicates that the PHY XS receiver lane 3 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[3]` (see 119.2.6.2.2).

45.2.4.11d.6 Lane 2 aligned (4.52.2)

When read as a one, bit 4.52.2 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 2. When read as a zero, bit 4.52.2 indicates that the PHY XS receiver lane 2 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[2]` (see 119.2.6.2.2).

45.2.4.11d.7 Lane 1 aligned (4.52.1)

When read as a one, bit 4.52.1 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 1. When read as a zero, bit 4.52.1 indicates that the PHY XS receiver lane 1 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[1]` (see 119.2.6.2.2).

45.2.4.11d.8 Lane 0 aligned (4.52.0)

When read as a one, bit 4.52.0 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 0. When read as a zero, bit 4.52.0 indicates that the PHY XS receiver lane 0 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[0]` (see 119.2.6.2.2).

45.2.4.11e Multi-lane BASE-R PHY XS alignment status 4 register (Register 4.53)

The assignment of bits in the multi-lane BASE-R PHY XS alignment status 4 register is shown in Table 45–171e. A PHY XS device that does not implement multi-lane BASE-R PHY XS shall return a zero for all bits in the multi-lane BASE-R PHY XS alignment status 4 register. A device that implements multi-lane BASE-R PHY XS shall return a zero for all bits in the multi-lane BASE-R PHY XS alignment status 4 register that are not required for the PHY XS 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.

**Table 45–171e—Multi-lane BASE-R PHY XS alignment status 4 register
bit definitions**

| Bit(s) | Name | Description | R/W ^a |
|-----------|-----------------|--|------------------|
| 4.53.15:8 | Reserved | Value always 0 | RO |
| 4.53.7 | Lane 15 aligned | 1 = Lane 15 alignment marker is locked 0 = Lane 15 alignment marker is not locked | RO |
| 4.53.6 | Lane 14 aligned | 1 = Lane 14 alignment marker is locked 0 = Lane 14 alignment marker is not locked | RO |
| 4.53.5 | Lane 13 aligned | 1 = Lane 13 alignment marker is locked 0 = Lane 13 alignment marker is not locked | RO |

**Table 45–171e—Multi-lane BASE-R PHY XS alignment status 4 register
 bit definitions (*continued*)**

| Bit(s) | Name | Description | R/W ^a |
|--------|-----------------|--|------------------|
| 4.53.4 | Lane 12 aligned | 1 = Lane 12 alignment marker is locked 0 = Lane 12 alignment marker is not locked | RO |
| 4.53.3 | Lane 11 aligned | 1 = Lane 11 alignment marker is locked 0 = Lane 11 alignment marker is not locked | RO |
| 4.53.2 | Lane 10 aligned | 1 = Lane 10 alignment marker is locked 0 = Lane 10 alignment marker is not locked | RO |
| 4.53.1 | Lane 9 aligned | 1 = Lane 9 alignment marker is locked 0 = Lane 9 alignment marker is not locked | RO |
| 4.53.0 | Lane 8 aligned | 1 = Lane 8 alignment marker is locked 0 = Lane 8 alignment marker is not locked | RO |

^a RO = Read only

45.2.4.11e.1 Lane 15 aligned (4.53.7)

When read as a one, bit 4.53.7 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 15. When read as a zero, bit 4.53.7 indicates that the PHY XS receiver lane 15 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[15]` (see 119.2.6.2.2).

45.2.4.11e.2 Lane 14 aligned (4.53.6)

When read as a one, bit 4.53.6 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 14. When read as a zero, bit 4.53.6 indicates that the PHY XS receiver lane 14 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[14]` (see 119.2.6.2.2).

45.2.4.11e.3 Lane 13 aligned (4.53.5)

When read as a one, bit 4.53.5 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 13. When read as a zero, bit 4.53.5 indicates that the PHY XS receiver lane 13 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[13]` (see 119.2.6.2.2).

45.2.4.11e.4 Lane 12 aligned (4.53.4)

When read as a one, bit 4.53.4 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 12. When read as a zero, bit 4.53.4 indicates that the PHY XS receiver lane 12 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[12]` (see 119.2.6.2.2).

45.2.4.11e.5 Lane 11 aligned (4.53.3)

When read as a one, bit 4.53.3 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 11. When read as a zero, bit 4.53.3 indicates that the PHY XS receiver lane 11 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[11]` (see 119.2.6.2.2).

45.2.4.11e.6 Lane 10 aligned (4.53.2)

When read as a one, bit 4.53.2 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 10. When read as a zero, bit 4.53.2 indicates that the PHY XS receiver lane 10 has not achieved alignment marker lock. This bit reflects the state of amps_lock[10] (see 119.2.6.2.2).

45.2.4.11e.7 Lane 9 aligned (4.53.1)

When read as a one, bit 4.53.1 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 9. When read as a zero, bit 4.53.1 indicates that the PHY XS receiver lane 9 has not achieved alignment marker lock. This bit reflects the state of amps_lock[9] (see 119.2.6.2.2).

45.2.4.11e.8 Lane 8 aligned (4.53.0)

When read as a one, bit 4.53.0 indicates that the PHY XS receiver has achieved alignment marker lock for service interface lane 8. When read as a zero, bit 4.53.0 indicates that the PHY XS receiver lane 8 has not achieved alignment marker lock. This bit reflects the state of amps_lock[8] (see 119.2.6.2.2).

45.2.4.11f PHY XS lane mapping, lane 0 register (Register 4.400)

The assignment of bits in the PHY XS lane mapping, lane 0 register is shown in Table 45–171f. When the multi-lane PHY XS described in Clause 118 detects and locks the alignment marker for service interface lane 0, the detected PCS lane number is recorded in this register. The contents of the PHY XS lane mapping, lane 0 register is valid when Lane 0 aligned bit (4.52.0) is set to one and is invalid otherwise.

Table 45–171f—PHY XS lane mapping, lane 0 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|----------------|--|------------------|
| 4.400.15:5 | Reserved | Value always 0 | RO |
| 4.400.4:0 | Lane 0 mapping | PHY XS lane received in service interface lane 0 | RO |

^a RO = Read only

45.2.4.11g PHY XS lane mapping, lane 1 through lane 15 registers (Registers 4.401 through 4.415)

The definition of the PHY XS lane mapping, lane 1 through 15 registers is identical to that described for lane 0 in 45.2.4.11f. The lane mapping for lane 1 is in register 4.401; lane 2 is in register 4.402; etc.

45.2.4.11h PHY XS FEC symbol error counter lane 0 (Register 4.600, 4.601)

The assignment of bits in the PHY XS FEC symbol error counter lane 0 register is shown in Table 45–171g. Symbol errors detected in PHY XS FEC lane 0 are counted and shown in register 4.600.15:0 and 4.601.15:0. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 4.600, 4.601 are used to read the value of a 32-bit counter. When registers 4.600 and 4.601 are used to read the 32-bit counter value, the register 4.600 is read first, the value of the register 4.601 is latched when (and only when) register 4.600 is read, and reads of register 4.601 return the latched value rather than the current value of the counter.

Table 45–171g—PHY XS FEC symbol error counter lane 0 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--|-----------------------------------|------------------|
| 4.600.15:0 | PHY XS FEC symbol errors, lane 0 lower | FEC_symbol_error_counter_0[15:0] | RO, NR |
| 4.601.15:0 | PHY XS FEC symbol errors, lane 0 upper | FEC_symbol_error_counter_0[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.4.11i PHY XS FEC symbol error counter lane 1 through 15 (Registers 4.602 through 4.631)

The behavior of the PHY XS FEC symbol error counters, lane 1 through 15 is identical to that described for PHY XS FEC lane 0 in 45.2.4.11h. Errors detected in each PHY XS FEC lane are counted and shown in the corresponding register. PHY XS FEC lane 1, lower 16 bits are shown in register 4.602; PHY XS FEC lane 1, upper 16 bits are shown in register 4.603; PHY XS FEC lane 2, lower 16 bits are shown in register 4.604; through register 4.631 for PHY XS FEC lane 15, upper 16 bits.

45.2.4.11j PHY XS FEC control register (Register 4.800)

The assignment of bits in the PHY XS FEC control register is shown in Table 45–171h.

Table 45–171h—PHY XS FEC control register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-------------------------------------|--|------------------|
| 4.800.15:3 | Reserved | Value always 0 | RO |
| 4.800.2 | PHY XS FEC degraded SER enable | 1 = FEC decoder signals degraded SER 0 = FEC decoder does not signal degraded SER | R/W |
| 4.800.1 | PHY XS FEC bypass indication enable | 1 = FEC decoder does not indicate errors 0 = FEC decoder indicates errors | R/W |
| 4.800.0 | Reserved | Value always 0 | RO |

^a R/W = Read/Write, RO = Read only

45.2.4.11j.1 PHY XS FEC degraded SER enable (4.800.2)

This bit enables the PHY XS FEC decoder to signal the presence of a degraded SER. When set to a one, this bit enables degraded SER signaling. When set to a zero, degraded SER signaling is disabled. Writes to this bit are ignored and reads return a zero if the PHY XS FEC does not have the ability to signal the presence of a degraded SER (see 119.2.5.3 for equivalent PCS behavior).

45.2.4.11j.2 PHY XS FEC bypass indication enable (4.800.1)

This bit enables the PHY XS FEC decoder to bypass error indication to the upper layers through the sync bits. When set to a one, this bit enables bypass of the error indication. When set to a zero, errors are indicated to the upper layers through the sync bits. Writes to this bit are ignored and reads return a zero if the PHY XS FEC does not have the ability to bypass indicating decoding errors (see 119.2.5.3 for equivalent PCS behavior).

45.2.4.11k PHY XS FEC status register (Register 4.801)

The assignment of bits in the PHY XS FEC status register is shown in Table 45–171i.

Table 45–171i—PHY XS FEC status register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--------------------------------------|--|------------------|
| 4.801.15:6 | Reserved | Value always 0 | RO |
| 4.801.5 | Remote degraded SER received | 1 = remote degraded SER received 0 = no remote degraded SER received | RO |
| 4.801.4 | PHY XS FEC degraded SER | 1 = SER is degraded 0 = SER is not degraded | RO |
| 4.801.3 | PHY XS FEC degraded SER ability | 1 = FEC decoder has the ability to signal the presence of a degraded SER 0 = FEC decoder does not have the ability to signal the presence of a degraded SER | RO |
| 4.801.2 | PHY XS FEC high SER | 1 = FEC errors have exceeded threshold 0 = FEC errors have not exceeded threshold | RO/LH |
| 4.801.1 | PHY XS FEC bypass indication ability | 1 = FEC decoder has the ability to bypass error indication 0 = FEC decoder does not have the ability to bypass error indication | RO |
| 4.801.0 | Reserved | Value always 0 | RO |

^a RO = Read only, LH = Latching high

45.2.4.11k.1 Remote degraded SER received (4.801.5)

When read as a one, bit 4.801.5 indicates that the remote degraded SER signal has been received. This bit reflects the state of rx_rm_degraded for the PHY XS (see 119.2.6.2.2 for equivalent PCS behavior).

45.2.4.11k.2 PHY XS FEC degraded SER (4.801.4)

When PHY XS FEC degraded SER enable (bit 4.800.2) is set to one, bit 4.801.4 is set to one if the number of FEC symbol errors in a window of PHY XS FEC degraded SER interval (registers 4.810 and 4.811) codewords exceeds the PHY XS FEC degraded SER activate threshold (registers 4.806 and 4.807) and is cleared if the number of FEC symbol errors in the same window is below the PHY XS FEC degraded SER deactivate threshold (registers 4.808 and 4.809). If the number of FEC symbol errors in the window is between the two thresholds, then bit 4.801.4 remains in its previous state. The bit is set to zero if PHY XS FEC degraded SER enable (bit 4.800.2) is set to zero (see 119.2.5.3 for equivalent PCS behavior). The value of bit 4.801.4 is undefined if the value of the PHY XS FEC degraded SER activate threshold is less than the value of the PHY XS FEC degraded SER deactivate threshold.

45.2.4.11k.3 PHY XS FEC degraded SER ability (4.801.3)

The PHY XS FEC decoder may have the option to signal the presence of a degraded SER (see 119.2.5.3 for equivalent PCS behavior). This bit is set to one to indicate that the decoder has the ability to signal the presence of a degraded SER. The bit is set to zero if this ability is not supported.

45.2.4.11k.4 PHY XS FEC high SER (4.801.2)

When PHY XS FEC bypass indication enable (bit 4.800.1) is set to one, this bit is set to one if the number of FEC symbol errors in a window of 8192 codewords exceeds the threshold (see 119.2.5.3 for equivalent PCS behavior) and is set to zero otherwise. The bit is set to zero if PHY XS FEC bypass indication enable (bit 4.800.1) is set to zero. This bit shall be implemented with latching high behavior.

45.2.4.11k.5 PHY XS FEC bypass indication ability (4.801.1)

The PHY XS FEC decoder may have the option to perform error detection without error indication (see 119.2.5.3 for equivalent PCS behavior) to reduce the delay contributed by the FEC decoder. This bit is set to one to indicate that the decoder has this ability to bypass error indication. The bit is set to zero if this ability is not supported.

45.2.4.11l PHY XS FEC corrected codewords counter (Register 4.802, 4.803)

The assignment of bits in the PHY XS FEC corrected codewords counter register is shown in Table 45–171j. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 4.802, 4.803 are used to read the value of a 32-bit counter. When registers 4.802 and 4.803 are used to read the 32-bit counter value, the register 4.802 is read first, the value of the register 4.803 is latched when (and only when) register 4.802 is read, and reads of register 4.803 return the latched value rather than the current value of the counter.

Table 45–171j—PHY XS FEC corrected codewords counter register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-------------------------------|---------------------------------|------------------|
| 4.802.15:0 | FEC corrected codewords lower | FEC_corrected_cw_counter[15:0] | RO, NR |
| 4.803.15:0 | FEC corrected codewords upper | FEC_corrected_cw_counter[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.4.11m PHY XS FEC uncorrected codewords counter (Register 4.804, 4.805)

The assignment of bits in the PHY XS FEC uncorrected codewords counter register is shown in Table 45–171k. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 4.804, 4.805 are used to read the value of a 32-bit counter. When registers 4.804 and 4.805 are used to read the 32-bit counter value, the register 4.804 is read first, the value of the register 4.805 is latched when (and only when) register 4.804 is read, and reads of register 4.805 return the latched value rather than the current value of the counter.

Table 45–171k—PHY XS FEC uncorrected codewords counter register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|---------------------------------|-----------------------------------|------------------|
| 4.804.15:0 | FEC uncorrected codewords lower | FEC_uncorrected_cw_counter[15:0] | RO, NR |
| 4.805.15:0 | FEC uncorrected codewords upper | FEC_uncorrected_cw_counter[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.4.11n PHY XS FEC degraded SER activate threshold register (Register 4.806, 4.807)

The assignment of bits in the PHY XS FEC degraded SER activate threshold register is shown in Table 45–171l. The value controls the threshold used to set the PHY XS FEC degraded SER bit (4.801.4) in an equivalent manner to that defined for the FEC_degraded_SER bit in 119.2.5.3.

Table 45–171l—PHY XS FEC degraded SER activate threshold register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--|--|------------------|
| 4.806.15:0 | PHY XS FEC degraded SER activate threshold lower | FEC_degraded_SER_activate_threshold[15:0] | R/W |
| 4.807.15:0 | PHY XS FEC degraded SER activate threshold upper | FEC_degraded_SER_activate_threshold[31:16] | R/W |

^a R/W = Read/Write

45.2.4.11o PHY XS FEC degraded SER deactivate threshold register (Register 4.808, 4.809)

The assignment of bits in the PHY XS FEC degraded SER deactivate threshold register is shown in Table 45–171m. The value controls the threshold used to clear the PHY XS FEC degraded SER bit (4.801.4) in an equivalent manner to that defined for the FEC_degraded_SER bit in 119.2.5.3.

Table 45–171m—PHY XS FEC degraded SER deactivate threshold register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--|--|------------------|
| 4.808.15:0 | PHY XS FEC degraded SER deactivate threshold lower | FEC_degraded_SER_deactivate_threshold[15:0] | R/W |
| 4.809.15:0 | PHY XS FEC degraded SER deactivate threshold upper | FEC_degraded_SER_deactivate_threshold[31:16] | R/W |

^a R/W = Read/Write

45.2.4.11p PHY XS FEC degraded SER interval register (Register 4.810, 4.811)

The assignment of bits in the PHY XS FEC degraded SER interval register is shown in Table 45–171n. The value controls the interval used to set and clear the PHY XS FEC degraded SER bit (4.801.4) in an equivalent manner to that defined for the FEC_degraded_SER bit in 119.2.5.3.

Table 45–171n—PHY XS FEC degraded SER interval register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--|----------------------------------|------------------|
| 4.810.15:0 | PHY XS FEC degraded SER interval lower | FEC_degraded_SER_interval[15:0] | R/W |
| 4.811.15:0 | PHY XS FEC degraded SER interval upper | FEC_degraded_SER_interval[31:16] | R/W |

^a R/W = Read/Write

45.2.5 DTE XS registers

Change the reserved row for 5.26 through 5.1799 in Table 45-175 as follows (unchanged rows not shown):

Table 45-175—DTE XS registers

| Register address | Register name | Subclause |
|-------------------------------------|--|--|
| ... | | |
| 5.26 through 5.31 5.1799 | Reserved | |
| <u>5.32</u> | <u>BASE-R DTE XS status 1</u> | <u>45.2.5.11a</u> |
| <u>5.33 through 5.41</u> | <u>Reserved</u> | |
| <u>5.42</u> | <u>BASE-R DTE XS test pattern control</u> | <u>45.2.5.11b</u> |
| <u>5.43 through 5.49</u> | <u>Reserved</u> | |
| <u>5.50</u> | <u>Multi-lane BASE-R DTE XS alignment status 1</u> | <u>45.2.5.11c</u> |
| <u>5.51</u> | <u>Reserved</u> | |
| <u>5.52</u> | <u>Multi-lane BASE-R DTE XS alignment status 3</u> | <u>45.2.5.11d</u> |
| <u>5.53</u> | <u>Multi-lane BASE-R DTE XS alignment status 4</u> | <u>45.2.5.11e</u> |
| <u>5.54</u> | <u>Reserved</u> | |
| <u>5.400 through 5.415</u> | <u>DTE XS lane mapping, lane 0 through 15</u> | <u>45.2.5.11f</u> <u>45.2.5.11g</u> |
| <u>5.415 through 5.599</u> | <u>Reserved</u> | |
| <u>5.600 through 5.631</u> | <u>DTE XS FEC symbol error counter, lane 0 to 15</u> | <u>45.2.5.11h</u> <u>45.2.5.11i</u> |
| <u>5.632 through 5.799</u> | <u>Reserved</u> | |
| <u>5.800</u> | <u>DTE XS FEC control</u> | <u>45.2.5.11j</u> |
| <u>5.801</u> | <u>DTE XS FEC status</u> | <u>45.2.5.11k</u> |
| <u>5.802, 5.803</u> | <u>DTE XS FEC corrected codewords counter</u> | <u>45.2.5.11l</u> |
| <u>5.804, 5.805</u> | <u>DTE XS FEC uncorrected codewords counter</u> | <u>45.2.5.11m</u> |
| <u>5.806, 5.807</u> | <u>DTE XS FEC degraded SER activate threshold</u> | <u>45.2.5.11n</u> |
| <u>5.808, 5.809</u> | <u>DTE XS FEC degraded SER deactivate threshold</u> | <u>45.2.5.11o</u> |
| <u>5.810, 5.811</u> | <u>DTE XS FEC degraded SER interval</u> | <u>45.2.5.11p</u> |
| <u>5.812 through 5.1799</u> | <u>Reserved</u> | |
| ... | | |

45.2.5.1 DTE XS control 1 register (Register 5.0)

Change the row for 5.0.5:2 in Table 45-176 as follows (unchanged rows not shown):

Table 45-176—DTE XS control 1 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|---------|-----------------|--|------------------|
| ... | | | |
| 5.0.5:2 | Speed selection | <div> <div>5 4 3 2</div> <div>1 * * * = Reserved</div> <div>* 1 * * = Reserved</div> <div>* * 1 * = Reserved</div> <div>1 1 x x = Reserved</div> <div>1 0 1 1 = Reserved</div> <div>1 0 1 0 = 400 Gb/s</div> <div>1 0 0 1 = 200 Gb/s</div> <div>1 0 0 0 = Reserved</div> <div>0 1 x x = Reserved</div> <div>0 0 1 x = Reserved</div> <div>0 0 0 1 = Reserved</div> <div>0 0 0 0 = 10 Gb/s</div> </div> | R/W |
| ... | | | |

^a R/W = Read/Write, SC = Self-clearing, RO = Read only

45.2.5.4 DTE XS speed ability (Register 5.4)

Change Table 45-178 as follows:

Table 45-178— DTE XS speed ability register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|----------------------------|---|------------------|
| 5.4.15:10 | Reserved for future speeds | Value always 0 | RO |
| 5.4.9 | 400G capable | <div>1 = DTE XS is capable of operating at 400 Gb/s</div> <div>0 = DTE XS is not capable of operating at 400 Gb/s</div> | RO |
| 5.4.8 | 200G capable | <div>1 = DTE XS is capable of operating at 200 Gb/s</div> <div>0 = DTE XS is not capable of operating at 200 Gb/s</div> | RO |
| 5.4.7:1 | Reserved | Value always 0 | RO |
| 5.4.0 | 10G capable | <div>1 = DTE XS is capable of operating at 10 Gb/s</div> <div>0 = DTE XS is not capable of operating at 10 Gb/s</div> | RO |

^a RO = Read only

Insert 45.2.5.4.a and 45.2.5.4.b before 45.2.5.4.1 as follows:

45.2.5.4.a 400G capable (5.4.9)

When read as a one, bit 5.4.9 indicates that the DTE XS is able to operate at a data rate of 400 Gb/s. When read as a zero, bit 5.4.9 indicates that the DTE XS is not able to operate at a data rate of 400 Gb/s.

45.2.5.4.b 200G capable (5.4.8)

When read as a one, bit 5.4.8 indicates that the DTE XS is able to operate at a data rate of 200 Gb/s. When read as a zero, bit 5.4.8 indicates that the DTE XS is not able to operate at a data rate of 200 Gb/s.

Insert 45.2.5.11a through 45.2.5.11p after 45.2.5.11 as follows:

45.2.5.11a BASE-R DTE XS status 1 register (Register 5.32)

The assignment of bits in the BASE-R DTE XS status 1 register is shown in Table 45–182a. A DTE XS device that does not implement BASE-R shall return a zero for all bits in the BASE-R DTE XS status 1 register. 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.

Table 45–182a—BASE-R DTE XS status 1 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-----------------------------------|--|------------------|
| 5.32.15:13 | Reserved | Value always 0 | RO |
| 5.32.12 | BASE-R DTE XS receive link status | 1 = BASE-R DTE XS receive link up 0 = BASE-R DTE XS receive link down | RO |
| 5.32.11:0 | Reserved | Value always 0 | RO |

^a RO = Read only

45.2.5.11a.1 BASE-R DTE XS receive link status (5.32.12)

When read as a one, bit 5.32.12 indicates that the DTE XS is in a fully operational state. When read as a zero, bit 5.32.12 indicates that the DTE XS is not fully operational. This bit is a reflection of the DTE XS equivalent of the PCS_status variable defined in 119.3 for 200/400GBASE-R.

45.2.5.11b BASE-R DTE XS test-pattern control register (Register 5.42)

The assignment of bits in the BASE-R DTE XS test-pattern control register is shown in Table 45–182b. This register is only required when the 200/400GBASE-R capability is supported. The test-pattern methodology is described in 119.2.4.9.

Table 45–182b—BASE-R DTE XS test-pattern control register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|------------------------------|---|------------------|
| 5.42.15:4 | Reserved | Value always 0 | RO |
| 5.42.3 | Transmit test-pattern enable | 1 = Enable transmit test pattern 0 = Disable transmit test pattern | R/W |
| 5.42.2:0 | Reserved | Value always 0 | RO |

^a RO = Read only, R/W = Read/Write

45.2.5.11b.1 Transmit test-pattern enable (5.42.3)

When bit 5.42.3 is set to a one, pattern testing is enabled on the transmit path. When bit 5.42.3 is set to a zero, pattern testing is disabled on the transmit path.

The default value for bit 5.42.3 is zero.

45.2.5.11c Multi-lane BASE-R DTE XS alignment status 1 register (Register 5.50)

The assignment of bits in the multi-lane BASE-R DTE XS alignment status 1 register is shown in Table 45–182c. A DTE XS device that does not implement multi-lane BASE-R DTE XS shall return a zero for all bits in the multi-lane BASE-R DTE XS alignment status 1 register. 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.

Table 45–182c—Multi-lane BASE-R DTE XS alignment status 1 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|------------------------------|--|------------------|
| 5.50.15:13 | Reserved | Value always 0 | RO |
| 5.50.12 | DTE XS lane alignment status | 1 = DTE XS receive lanes locked and aligned 0 = DTE XS receive lanes not locked and aligned | RO |
| 5.50.11:0 | Reserved | Value always 0 | RO |

^a RO = Read only

45.2.5.11c.1 DTE XS lane alignment status (5.50.12)

When read as a one, bit 5.50.12 indicates that the DTE XS has locked and aligned all receive lanes. When read as a zero, bit 5.50.12 indicates that the DTE XS has not locked and aligned all receive lanes.

45.2.5.11d Multi-lane BASE-R DTE XS alignment status 3 register (Register 5.52)

The assignment of bits in the multi-lane BASE-R DTE XS alignment status 3 register is shown in Table 45–182d. A DTE XS device that does not implement multi-lane BASE-R DTE XS shall return a zero for all bits in the multi-lane BASE-R DTE XS alignment status 3 register. A device that implements multi-lane BASE-R DTE XS shall return a zero for all bits in the multi-lane BASE-R DTE XS alignment status 3 register that are not required for the DTE XS 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.

Table 45–182d—Multi-lane BASE-R DTE XS alignment status 3 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|-----------|----------------|--|------------------|
| 5.52.15:8 | Reserved | Value always 0 | RO |
| 5.52.7 | Lane 7 aligned | 1 = Lane 7 alignment marker is locked 0 = Lane 7 alignment marker is not locked | RO |
| 5.52.6 | Lane 6 aligned | 1 = Lane 6 alignment marker is locked 0 = Lane 6 alignment marker is not locked | RO |

**Table 45–182d—Multi-lane BASE-R DTE XS alignment status 3 register
 bit definitions (*continued*)**

| Bit(s) | Name | Description | R/W ^a |
|--------|----------------|--|------------------|
| 5.52.5 | Lane 5 aligned | 1 = Lane 5 alignment marker is locked 0 = Lane 5 alignment marker is not locked | RO |
| 5.52.4 | Lane 4 aligned | 1 = Lane 4 alignment marker is locked 0 = Lane 4 alignment marker is not locked | RO |
| 5.52.3 | Lane 3 aligned | 1 = Lane 3 alignment marker is locked 0 = Lane 3 alignment marker is not locked | RO |
| 5.52.2 | Lane 2 aligned | 1 = Lane 2 alignment marker is locked 0 = Lane 2 alignment marker is not locked | RO |
| 5.52.1 | Lane 1 aligned | 1 = Lane 1 alignment marker is locked 0 = Lane 1 alignment marker is not locked | RO |
| 5.52.0 | Lane 0 aligned | 1 = Lane 0 alignment marker is locked 0 = Lane 0 alignment marker is not locked | RO |

^a RO = Read only

45.2.5.11d.1 Lane 7 aligned (5.52.7)

When read as a one, bit 5.52.7 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 7. When read as a zero, bit 5.52.7 indicates that the DTE XS receiver lane 7 has not achieved alignment marker lock. This bit reflects the state of amps_lock[7] (see 119.2.6.2.2).

45.2.5.11d.2 Lane 6 aligned (5.52.6)

When read as a one, bit 5.52.6 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 6. When read as a zero, bit 5.52.6 indicates that the DTE XS receiver lane 6 has not achieved alignment marker lock. This bit reflects the state of amps_lock[6] (see 119.2.6.2.2).

45.2.5.11d.3 Lane 5 aligned (5.52.5)

When read as a one, bit 5.52.5 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 5. When read as a zero, bit 5.52.5 indicates that the DTE XS receiver lane 5 has not achieved alignment marker lock. This bit reflects the state of amps_lock[5] (see 119.2.6.2.2).

45.2.5.11d.4 Lane 4 aligned (5.52.4)

When read as a one, bit 5.52.4 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 4. When read as a zero, bit 5.52.4 indicates that the DTE XS receiver lane 4 has not achieved alignment marker lock. This bit reflects the state of amps_lock[4] (see 119.2.6.2.2).

45.2.5.11d.5 Lane 3 aligned (5.52.3)

When read as a one, bit 5.52.3 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 3. When read as a zero, bit 5.52.3 indicates that the DTE XS receiver lane 3 has not achieved alignment marker lock. This bit reflects the state of amps_lock[3] (see 119.2.6.2.2).

45.2.5.11d.6 Lane 2 aligned (5.52.2)

When read as a one, bit 5.52.2 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 2. When read as a zero, bit 5.52.2 indicates that the DTE XS receiver lane 2 has not achieved alignment marker lock. This bit reflects the state of amps_lock[2] (see 119.2.6.2.2).

45.2.5.11d.7 Lane 1 aligned (5.52.1)

When read as a one, bit 5.52.1 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 1. When read as a zero, bit 5.52.1 indicates that the DTE XS receiver lane 1 has not achieved alignment marker lock. This bit reflects the state of amps_lock[1] (see 119.2.6.2.2).

45.2.5.11d.8 Lane 0 aligned (5.52.0)

When read as a one, bit 5.52.0 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 0. When read as a zero, bit 5.52.0 indicates that the DTE XS receiver lane 0 has not achieved alignment marker lock. This bit reflects the state of amps_lock[0] (see 119.2.6.2.2).

45.2.5.11e Multi-lane BASE-R DTE XS alignment status 4 register (Register 5.53)

The assignment of bits in the multi-lane BASE-R DTE XS alignment status 4 register is shown in Table 45–182e. A DTE XS device that does not implement multi-lane BASE-R DTE XS shall return a zero for all bits in the multi-lane BASE-R DTE XS alignment status 4 register. A device that implements multi-lane BASE-R DTE XS shall return a zero for all bits in the multi-lane BASE-R DTE XS alignment status 4 register that are not required for the DTE XS 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.

**Table 45–182e—Multi-lane BASE-R DTE XS alignment status 4 register
bit definitions**

| Bit(s) | Name | Description | R/W ^a |
|-----------|-----------------|--|------------------|
| 5.53.15:8 | Reserved | Value always 0 | RO |
| 5.53.7 | Lane 15 aligned | 1 = Lane 15 alignment marker is locked 0 = Lane 15 alignment marker is not locked | RO |
| 5.53.6 | Lane 14 aligned | 1 = Lane 14 alignment marker is locked 0 = Lane 14 alignment marker is not locked | RO |
| 5.53.5 | Lane 13 aligned | 1 = Lane 13 alignment marker is locked 0 = Lane 13 alignment marker is not locked | RO |
| 5.53.4 | Lane 12 aligned | 1 = Lane 12 alignment marker is locked 0 = Lane 12 alignment marker is not locked | RO |
| 5.53.3 | Lane 11 aligned | 1 = Lane 11 alignment marker is locked 0 = Lane 11 alignment marker is not locked | RO |
| 5.53.2 | Lane 10 aligned | 1 = Lane 10 alignment marker is locked 0 = Lane 10 alignment marker is not locked | RO |
| 5.53.1 | Lane 9 aligned | 1 = Lane 9 alignment marker is locked 0 = Lane 9 alignment marker is not locked | RO |
| 5.53.0 | Lane 8 aligned | 1 = Lane 8 alignment marker is locked 0 = Lane 8 alignment marker is not locked | RO |

^a RO = Read only

45.2.5.11e.1 Lane 15 aligned (5.53.7)

When read as a one, bit 5.53.7 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 15. When read as a zero, bit 5.53.7 indicates that the DTE XS receiver lane 15 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[15]` (see 119.2.6.2.2).

45.2.5.11e.2 Lane 14 aligned (5.53.6)

When read as a one, bit 5.53.6 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 14. When read as a zero, bit 5.53.6 indicates that the DTE XS receiver lane 14 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[14]` (see 119.2.6.2.2).

45.2.5.11e.3 Lane 13 aligned (5.53.5)

When read as a one, bit 5.53.5 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 13. When read as a zero, bit 5.53.5 indicates that the DTE XS receiver lane 13 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[13]` (see 119.2.6.2.2).

45.2.5.11e.4 Lane 12 aligned (5.53.4)

When read as a one, bit 5.53.4 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 12. When read as a zero, bit 5.53.4 indicates that the DTE XS receiver lane 12 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[12]` (see 119.2.6.2.2).

45.2.5.11e.5 Lane 11 aligned (5.53.3)

When read as a one, bit 5.53.3 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 11. When read as a zero, bit 5.53.3 indicates that the DTE XS receiver lane 11 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[11]` (see 119.2.6.2.2).

45.2.5.11e.6 Lane 10 aligned (5.53.2)

When read as a one, bit 5.53.2 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 10. When read as a zero, bit 5.53.2 indicates that the DTE XS receiver lane 10 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[10]` (see 119.2.6.2.2).

45.2.5.11e.7 Lane 9 aligned (5.53.1)

When read as a one, bit 5.53.1 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 9. When read as a zero, bit 5.53.1 indicates that the DTE XS receiver lane 9 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[9]` (see 119.2.6.2.2).

45.2.5.11e.8 Lane 8 aligned (5.53.0)

When read as a one, bit 5.53.0 indicates that the DTE XS receiver has achieved alignment marker lock for service interface lane 8. When read as a zero, bit 5.53.0 indicates that the DTE XS receiver lane 8 has not achieved alignment marker lock. This bit reflects the state of `amps_lock[8]` (see 119.2.6.2.2).

45.2.5.11f DTE XS lane mapping, lane 0 register (Register 5.400)

The assignment of bits in the DTE XS lane mapping, lane 0 register is shown in Table 45–182f. When the multi-lane DTE XS described in Clause 118 detects and locks the alignment marker for service interface lane 0, the detected PCS lane number is recorded in this register. The contents of the DTE XS lane mapping, lane 0 register is valid when Lane 0 aligned bit (5.52.0) is set to one and is invalid otherwise.

Table 45–182f—DTE XS lane mapping, lane 0 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|----------------|--|------------------|
| 5.400.15:5 | Reserved | Value always 0 | RO |
| 5.400.4:0 | Lane 0 mapping | DTE XS lane received in service interface lane 0 | RO |

^a RO = Read only

45.2.5.11g DTE XS lane mapping, lane 1 through lane 15 registers (Registers 5.401 through 5.415)

The definition of the DTE XS lane mapping, lane 1 through 15 registers is identical to that described for lane 0 in 45.2.5.11f. The lane mapping for lane 1 is in register 5.401; lane 2 is in register 5.402; etc.

45.2.5.11h DTE XS FEC symbol error counter lane 0 (Register 5.600, 5.601)

The assignment of bits in the DTE XS FEC symbol error counter lane 0 register is shown in Table 45–182g. Symbol errors detected in DTE XS FEC lane 0 are counted and shown in register 5.600.15:0 and 5.601.15:0. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 5.600, 5.601 are used to read the value of a 32-bit counter. When registers 5.600 and 5.601 are used to read the 32-bit counter value, the register 5.600 is read first, the value of the register 5.601 is latched when (and only when) register 5.600 is read, and reads of register 5.601 return the latched value rather than the current value of the counter.

Table 45–182g—DTE XS FEC symbol error counter lane 0 register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--|-----------------------------------|------------------|
| 5.600.15:0 | DTE XS FEC symbol errors, lane 0 lower | FEC_symbol_error_counter_0[15:0] | RO, NR |
| 5.601.15:0 | DTE XS FEC symbol errors, lane 0 upper | FEC_symbol_error_counter_0[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.5.11i DTE XS FEC symbol error counter lane 1 through 15 (Registers 5.602 through 5.631)

The behavior of the DTE XS FEC symbol error counters, lane 1 through 15 is identical to that described for DTE XS FEC lane 0 in 45.2.5.11h. Errors detected in each DTE XS FEC lane are counted and shown in the corresponding register. DTE XS FEC lane 1, lower 16 bits are shown in register 5.602; DTE XS FEC lane 1, upper 16 bits are shown in register 5.603; DTE XS FEC lane 2, lower 16 bits are shown in register 5.604; through register 5.631 for DTE XS FEC lane 15, upper 16 bits.

45.2.5.11j DTE XS FEC control register (Register 5.800)

The assignment of bits in the DTE XS FEC control register is shown in Table 45–182h.

Table 45–182h—DTE XS FEC control register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-------------------------------------|--|------------------|
| 5.800.15:2 | Reserved | Value always 0 | RO |
| 5.800.2 | DTE XS FEC degraded SER enable | 1 = FEC decoder signals degraded SER 0 = FEC decoder does not signal degraded SER | R/W |
| 5.800.1 | DTE XS FEC bypass indication enable | 1 = FEC decoder does not indicate errors 0 = FEC decoder indicates errors | R/W |
| 5.800.0 | Reserved | Value always 0 | RO |

^a R/W = Read/Write, RO = Read only

45.2.5.11j.1 DTE XS FEC degraded SER enable (5.800.2)

This bit enables the DTE XS FEC decoder to signal the presence of a degraded SER. When set to a one, this bit enables degraded SER signaling. When set to a zero, degraded SER signaling is disabled. Writes to this bit are ignored and reads return a zero if the DTE XS FEC does not have the ability to signal the presence of a degraded SER (see 119.2.5.3 for equivalent PCS behavior).

45.2.5.11j.2 DTE XS FEC bypass indication enable (5.800.1)

This bit enables the DTE XS FEC decoder to bypass error indication to the upper layers through the sync bits. When set to a one, this bit enables bypass of the error indication. When set to a zero, errors are indicated to the upper layers through the sync bits. Writes to this bit are ignored and reads return a zero if the DTE XS FEC does not have the ability to bypass indicating decoding errors (see 119.2.5.3 for equivalent PCS behavior).

45.2.5.11k DTE XS FEC status register (Register 5.801)

The assignment of bits in the DTE XS FEC status register is shown in Table 45–182i.

Table 45–182i—DTE XS FEC status register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|------------------------------|---|------------------|
| 5.801.15:7 | Reserved | Value always 0 | RO |
| 5.801.6 | Local degraded SER received | 1 = local degraded SER received 0 = no local degraded SER received | RO |
| 5.801.5 | Remote degraded SER received | 1 = remote degraded SER received 0 = no remote degraded SER received | RO |
| 5.801.4 | DTE XS FEC degraded SER | 1 = SER is degraded 0 = SER is not degraded | RO |

Table 45–182i—DTE XS FEC status register bit definitions (continued)

| Bit(s) | Name | Description | R/W ^a |
|---------|--------------------------------------|--|------------------|
| 5.801.3 | DTE XS FEC degraded SER ability | 1 = FEC decoder has the ability to signal the presence of a degraded SER 0 = FEC decoder does not have the ability to signal the presence of a degraded SER | RO |
| 5.801.2 | DTE XS FEC high SER | 1 = FEC errors have exceeded threshold 0 = FEC errors have not exceeded threshold | RO/LH |
| 5.801.1 | DTE XS FEC bypass indication ability | 1 = FEC decoder has the ability to bypass error indication 0 = FEC decoder does not have the ability to bypass error indication | RO |
| 5.801.0 | Reserved | Value always 0 | RO |

^a RO = Read only, LH = Latching high

45.2.5.11k.1 Local degraded SER received (5.801.6)

When read as a one, bit 5.801.6 indicates that the local degraded SER signal has been received. This bit reflects the state of rx_local_degraded (see 118.2.1) for the DTE XS.

45.2.5.11k.2 Remote degraded SER received (5.801.5)

When read as a one, bit 5.801.5 indicates that the remote degraded SER signal has been received. This bit reflects the state of rx_rm_degraded for the DTE XS (see 119.2.6.2.2 for equivalent PCS behavior).

45.2.5.11k.3 DTE XS FEC degraded SER (5.801.4)

When DTE XS FEC degraded SER enable (bit 5.800.2) is set to one, bit 5.801.4 is set to one if the number of FEC symbol errors in a window of DTE XS FEC degraded SER interval (registers 5.810 and 5.811) codewords exceeds the DTE XS FEC degraded SER activate threshold (registers 5.806 and 5.807) and is cleared if the number of FEC symbol errors in the same window is below the DTE XS FEC degraded SER deactivate threshold (registers 5.808 and 5.809). If the number of FEC symbol errors in the window is between the two thresholds, then bit 5.801.4 remains in its previous state. The bit is set to zero if DTE XS FEC degraded SER enable (bit 5.800.2) is set to zero (see 119.2.5.3 for equivalent PCS behavior). The value of bit 5.801.4 is undefined if the value of the DTE XS FEC degraded SER activate threshold is less than the value of the DTE XS FEC degraded SER deactivate threshold.

45.2.5.11k.4 DTE XS FEC degraded SER ability (5.801.3)

The DTE XS FEC decoder may have the option to signal the presence of a degraded SER (see 119.2.5.3 for equivalent PCS behavior). This bit is set to one to indicate that the decoder has the ability to signal the presence of a degraded SER. The bit is set to zero if this ability is not supported.

45.2.5.11k.5 DTE XS FEC high SER (5.801.2)

When DTE XS FEC bypass indication enable (bit 5.800.1) is set to one, this bit is set to one if the number of FEC symbol errors in a window of 8192 codewords exceeds the threshold (see 119.2.5.3) and is set to zero otherwise. The bit is set to zero if DTE XS FEC bypass indication enable (bit 5.800.1) is set to zero. This bit shall be implemented with latching high behavior.

45.2.5.11k.6 DTE XS FEC bypass indication ability (5.801.1)

The DTE XS FEC decoder may have the option to perform error detection without error indication (see 119.2.5.3) to reduce the delay contributed by the FEC decoder. This bit is set to one to indicate that the decoder has this ability to bypass error indication. The bit is set to zero if this ability is not supported.

45.2.5.11l DTE XS FEC corrected codewords counter (Register 5.802, 5.803)

The assignment of bits in the DTE XS FEC corrected codewords counter register is shown in Table 45–182j. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 5.802, 5.803 are used to read the value of a 32-bit counter. When registers 5.802 and 5.803 are used to read the 32-bit counter value, the register 5.802 is read first, the value of the register 5.803 is latched when (and only when) register 5.802 is read, and reads of register 5.803 return the latched value rather than the current value of the counter.

Table 45–182j—DTE XS FEC corrected codewords counter register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|-------------------------------|---------------------------------|------------------|
| 5.802.15:0 | FEC corrected codewords lower | FEC_corrected_cw_counter[15:0] | RO, NR |
| 5.803.15:0 | FEC corrected codewords upper | FEC_corrected_cw_counter[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.5.11m DTE XS FEC uncorrected codewords counter (Register 5.804, 5.805)

The assignment of bits in the DTE XS FEC uncorrected codewords counter register is shown in Table 45–182k. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 5.804, 5.805 are used to read the value of a 32-bit counter. When registers 5.804 and 5.805 are used to read the 32-bit counter value, the register 5.804 is read first, the value of the register 5.805 is latched when (and only when) register 5.804 is read, and reads of register 5.805 return the latched value rather than the current value of the counter.

Table 45–182k—DTE XS FEC uncorrected codewords counter register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|---------------------------------|-----------------------------------|------------------|
| 5.804.15:0 | FEC uncorrected codewords lower | FEC_uncorrected_cw_counter[15:0] | RO, NR |
| 5.805.15:0 | FEC uncorrected codewords upper | FEC_uncorrected_cw_counter[31:16] | RO, NR |

^a RO = Read only, NR = Non Roll-over

45.2.5.11n DTE XS FEC degraded SER activate threshold register (Register 5.806, 5.807)

The assignment of bits in the DTE XS FEC degraded SER activate threshold register is shown in Table 45–182l. The value controls the threshold used to set the DTE XS FEC degraded SER bit (5.801.4) in an equivalent manner to that defined for the FEC_degraded_SER bit in 119.2.5.3.

Table 45–182l—DTE XS FEC degraded SER activate threshold register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--|--|------------------|
| 5.806.15:0 | DTE XS FEC degraded SER activate threshold lower | FEC_degraded_SER_activate_threshold[15:0] | R/W |
| 5.807.15:0 | DTE XS FEC degraded SER activate threshold upper | FEC_degraded_SER_activate_threshold[31:16] | R/W |

^a R/W = Read/Write**45.2.5.11o DTE XS FEC degraded SER deactivate threshold register (Register 5.808, 5.809)**

The assignment of bits in the DTE XS FEC degraded SER deactivate threshold register is shown in Table 45–182m. The value controls the threshold used to clear the DTE XS FEC degraded SER bit (5.801.4) in an equivalent manner to that defined for the FEC_degraded_SER bit in 119.2.5.3.

Table 45–182m—DTE XS FEC degraded SER deactivate threshold register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--|--|------------------|
| 5.808.15:0 | DTE XS FEC degraded SER deactivate threshold lower | FEC_degraded_SER_deactivate_threshold[15:0] | R/W |
| 5.809.15:0 | DTE XS FEC degraded SER deactivate threshold upper | FEC_degraded_SER_deactivate_threshold[31:16] | R/W |

^a R/W = Read/Write**45.2.5.11p DTE XS FEC degraded SER interval register (Register 5.810, 5.811)**

The assignment of bits in the DTE XS FEC degraded SER interval register is shown in Table 45–182n. The value controls the interval used to set and clear the DTE XS FEC degraded SER bit (5.801.4) in an equivalent manner to that defined for the FEC_degraded_SER bit in 119.2.5.3.

Table 45–182n—DTE XS FEC degraded SER interval register bit definitions

| Bit(s) | Name | Description | R/W ^a |
|------------|--|----------------------------------|------------------|
| 5.810.15:0 | DTE XS FEC degraded SER interval lower | FEC_degraded_SER_interval[15:0] | R/W |
| 5.811.15:0 | DTE XS FEC degraded SER interval upper | FEC_degraded_SER_interval[31:16] | R/W |

^a R/W = Read/Write

78. Energy-Efficient Ethernet (EEE)

78.1 Overview

Change the third paragraph of 78.1 (as modified by IEEE Std 802.3by-2016) as follows:

EEE supports operation over twisted-pair cabling systems, twinax cable, electrical backplanes, optical fiber, the XGXS for 10 Gb/s PHYs, the 25GAUI for 25 Gb/s PHYs, the XLAUI for 40 Gb/s PHYs, ~~and the CAUI-10 or CAUI-4 for 100 Gb/s PHYs, the 200GAUI-n and 200GXS for 200 Gb/s PHYs, and the 400GAUI-n and 400GXS for 400 Gb/s PHYs.~~ Table 78–1 lists the supported PHYs and interfaces and their associated clauses.

78.1.4 PHY types optionally supporting EEE

Insert new rows at the end of Table 78-1 (as modified by IEEE Std 802.3by-2016) as follows (unchanged rows not shown):

Table 78–1—Clauses associated with each PHY or interface type

| PHY or interface type | Clause |
|----------------------------|---------------|
| ... | |
| 200GBASE-DR4 ^b | 119, 120, 121 |
| 200GBASE-FR4 ^b | 119, 120, 122 |
| 200GBASE-LR4 ^b | 119, 120, 122 |
| 400GBASE-SR16 ^b | 119, 120, 123 |
| 400GBASE-DR4 ^b | 119, 120, 124 |
| 400GBASE-FR8 ^b | 119, 120, 122 |
| 400GBASE-LR8 ^b | 119, 120, 122 |

^b The deep sleep mode of EEE is not supported for this PHY.

78.5 Communication link access latency

Insert three new rows (including footnote c) at the end of Table 78-4 as follows (unchanged rows not shown):

Table 78-4—Summary of the LPI timing parameters for supported PHYs or interfaces

| PHY or interface type | Case | $T_{w_sys_tx}$ (min) (μs) | T_{w_phy} (min) (μs) | $T_{phy_shrink_tx}$ (max) (μs) | $T_{phy_shrink_rx}$ (max) (μs) | $T_{w_sys_rx}$ (min) (μs) |
|----------------------------|------|-----------------------------------|-------------------------------|--|--|-----------------------------------|
| ... | | | | | | |
| 200GXS/400GXS ^c | | 0.34 | | | | |
| 200GBASE-R fast wake | | 0.34 | 0.3 | 0 | 0 | 0.25 |
| 400GBASE-R fast wake | | 0.34 | 0.3 | 0 | 0 | 0.25 |

^c The minimum $T_{w_sys_tx}$ of a PHY is increased by the indicated period if there is a 200GXS/400GXS in the transmit path. A PHY that includes a 200GXS/400GXS in the receive path may require an increase of $T_{w_sys_tx}$ on the link partner; this may be negotiated using LLDP (see 79.3.5).

Change the title of 78.5.1 as follows:

78.5.1 ~~40 Gb/s~~ PHY extension using extender sublayers ~~XGXS~~

Insert a new paragraph at the end of 78.5.1 as follows:

The 200GXS or 400GXS (see Clause 118) can be inserted between the RS and a 200 Gb/s or 400 Gb/s PHY, respectively, to transparently extend the physical reach of the 200GMII or 400GMII. The LPI signaling can operate through the 200GXS or 400GXS with the PHY timing parameters modified as described in Table 78-4.

90. Ethernet support for time synchronization protocols

90.1 Introduction

Change the second paragraph of 90.1 (as modified by IEEE Std 802.3by-2016) as follows:

The TSSI is defined for the full-duplex mode of operation only. It supports MAC operation at various data rates. The MII (Clause 22), GMII (Clause 35), XGMII (Clause 46), 25GMII (Clause 106), XLGMII (Clause 81), ~~and CGMII (Clause 81), 200GMII (Clause 117), and 400GMII (Clause 117)~~ specifications are all compatible with the gRS sublayer defined in 90.5.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd.10:2019

Insert Clause 116 to Clause 124 as follows:

116. Introduction to 200 Gb/s and 400 Gb/s networks

116.1 Overview

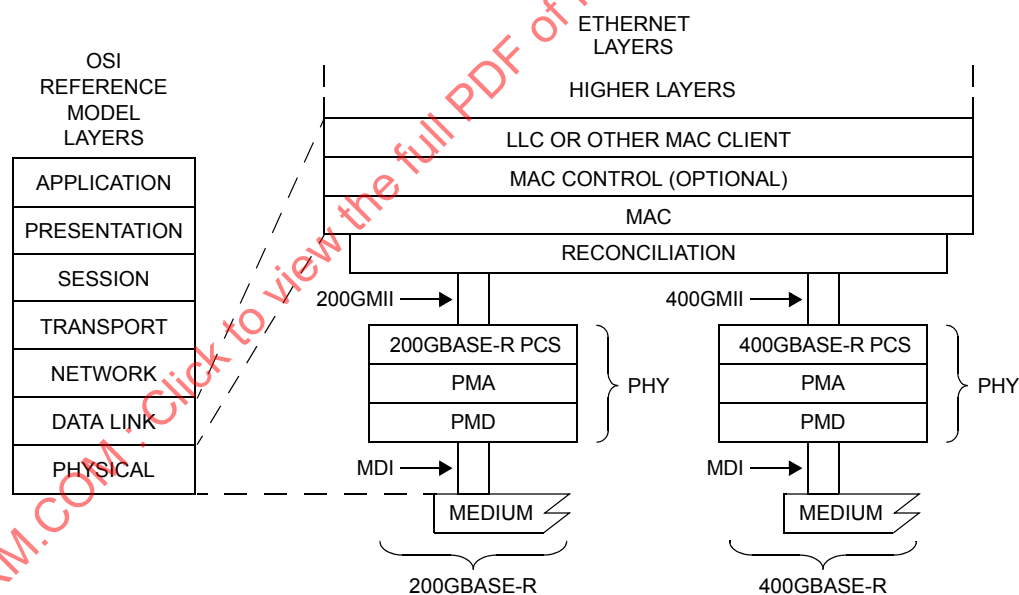
116.1.1 Scope

This clause describes the general requirements for 200 Gigabit and 400 Gigabit Ethernet.

200 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer operating at a data rate of 200 Gb/s, coupled with any IEEE 802.3 200GBASE Physical Layer implementation and is defined for full duplex operation only. 400 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer operating at a data rate of 400 Gb/s, coupled with any IEEE 802.3 400GBASE Physical Layer implementation and is defined for full duplex operation only.

116.1.2 Relationship of 200 Gigabit and 400 Gigabit Ethernet to the ISO OSI reference model

200 Gigabit and 400 Gigabit Ethernet couples the IEEE 802.3 MAC to a family of 200 Gb/s and 400 Gb/s Physical Layers. The relationships among 200 Gigabit and 400 Gigabit Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in Figure 116–1.



200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
 400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT

Figure 116–1—Architectural positioning of 200 Gigabit and 400 Gigabit Ethernet

While this specification defines interfaces in terms of bits, octets, and frames, implementations may choose other data-path widths for implementation convenience. The only exceptions are as follows:

- a) The 200GMII and 400GMII, which, when implemented as logical interconnection points between the MAC sublayer and the Physical Layer (PHY), use a 64-bit wide data path as specified in Clause 117. Physical instantiations of this interface may use other data-path widths.
- b) The management interface, which, when physically implemented as the MDIO/MDC (Management Data Input/Output and Management Data Clock) at an observable interconnection port, uses a bit-wide data path as specified in Clause 45.
- c) The PMA service interface, which, when physically implemented as 400GAUI-16 (400 Gb/s sixteen-lane Attachment Unit Interface) at an observable interconnection port, uses a 16-lane data path as specified in Annex 120B or Annex 120C.
- d) The PMA service interface, which, when physically implemented as 200GAUI-8 (200 Gb/s eight-lane Attachment Unit Interface) or 400GAUI-8 (400 Gb/s eight-lane Attachment Unit Interface) at an observable interconnection port, uses an 8-lane data path as specified in Annex 120B, Annex 120C, Annex 120D, or Annex 120E.
- e) The PMA service interface, which, when physically implemented as 200GAUI-4 (200 Gb/s four-lane Attachment Unit Interface) at an observable interconnection port, uses a 4-lane data path as specified in Annex 120D or Annex 120E.
- f) The MDI as specified in Clause 123 for 400GBASE-SR16 uses a 16-lane data path.
- g) The MDI as specified in Clause 122 for 400GBASE-FR8 and 400GBASE-LR8 uses an 8-lane data path.
- h) The MDIs as specified in Clause 121 for 200GBASE-DR4, in Clause 122 for 200GBASE-FR4 and 200GBASE-LR4, and in Clause 124 for 400GBASE-DR4, all use a 4-lane data path.

116.1.3 Nomenclature

The nomenclature employed by the 200 Gb/s and 400 Gb/s Physical Layer is explained as follows.

The alpha-numeric prefix 200GBASE in the port type (e.g., 200GBASE-R) represents a family of Physical Layer devices operating at a speed of 200 Gb/s. The alpha-numeric prefix 400GBASE in the port type (e.g., 400GBASE-R) represents a family of Physical Layer devices operating at a speed of 400 Gb/s.

200GBASE-R represents a family of Physical Layer devices using the Physical Coding Sublayer for 200 Gb/s operation over multiple PCS lanes (see Clause 119). Physical Layer devices listed in Table 116–1 are defined for operation at 200 Gb/s.

Table 116–1—200 Gb/s PHYs

| Name | Description |
|--------------|---|
| 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) |

400GBASE-R represents a family of Physical Layer devices using the Physical Coding Sublayer for 400 Gb/s operation over multiple PCS lanes (see Clause 119). Physical Layer devices listed in Table 116–2 are defined for operation at 400 Gb/s.

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-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-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) |

116.1.4 Physical Layer signaling systems

This standard specifies a family of Physical Layer implementations. Table 116–3 and Table 116–4 specify the correlation between PHY types and clauses. Implementations conforming to one or more PHY types must meet the requirements of the corresponding clauses.

Table 116–3—PHY type and clause correlation (200GBASE optical)

| PHY type | Clause ^a | | | | | | | | | | | | |
|--------------|---------------------|-----|---------|------------------|----------------|----------------|---------------|---------------|---------------|---------------|------------------|------------------|------------------|
| | 78 | 117 | | 118 | 119 | 120 | 120B | 120C | 120D | 120E | 121 | 122 | 122 |
| | EEE | RS | 200GMII | 200GMII Extender | 200GBASE-R PCS | 200GBASE-R PMA | 200GAUI-8 C2C | 200GAUI-8 C2M | 200GAUI-4 C2C | 200GAUI-4 C2M | 200GBASE-DR4 PMD | 200GBASE-FR4 PMD | 200GBASE-LR4 PMD |
| 200GBASE-DR4 | O | M | O | O | M | M | O | O | O | O | M | | |
| 200GBASE-FR4 | O | M | O | O | M | M | O | O | O | O | | M | |
| 200GBASE-LR4 | O | M | O | O | M | M | O | O | O | O | | | M |

^a O = Optional, M = Mandatory.

Table 116–4—PHY type and clause correlation (400GBASE optical)

| PHY type | Clause ^a | | | | | | | | | | | | |
|---------------|---------------------|-----|---------|------------------|----------------|----------------|----------------|----------------|---------------|---------------|-------------------|------------------|------------------|
| | 78 | 117 | | 118 | 119 | 120 | 120B | 120C | 120D | 120E | 123 | 124 | 122 |
| | EEE | RS | 400GMII | 400GMII Extender | 400GBASE-R PCS | 400GBASE-R PMA | 400GAUI-16 C2C | 400GAUI-16 C2M | 400GAUI-8 C2C | 400GAUI-8 C2M | 400GBASE-SR16 PMD | 400GBASE-DR4 PMD | 400GBASE-FR8 PMD |
| 400GBASE-SR16 | O | M | O | O | M | M | O | O | O | O | M | | |
| 400GBASE-DR4 | O | M | O | O | M | M | O | O | O | O | | M | |
| 400GBASE-FR8 | O | M | O | O | M | M | O | O | O | O | | | M |
| 400GBASE-LR8 | O | M | O | O | M | M | O | O | O | O | | | M |

^a O = Optional, M = Mandatory.

116.2 Summary of 200 Gigabit and 400 Gigabit Ethernet sublayers

116.2.1 Reconciliation Sublayer (RS) and Media Independent Interface

The Media Independent Interface (Clause 117) provides a logical interconnection between the MAC sublayer and Physical Layer entities (PHY). The Media Independent Interface is not intended to be physically instantiated, rather it can logically connect layers within a device.

The 200GMII supports 200 Gb/s operation and the 400GMII supports 400 Gb/s operation through its 64-bit-wide transmit and receive data paths. The Reconciliation Sublayer (RS) provides a mapping between the signals provided at the Media Independent Interface (200GMII and 400GMII) and the MAC/PLS service definition.

While the 200GMII and 400GMII are optional interfaces, they are used extensively in this standard as a basis for functional specification and provides a common service interface for the physical coding sublayer defined in Clause 119.

116.2.2 200GMII and 400GMII Extender Sublayers (200GXS and 400GXS)

The 200 Gigabit 200GMII Extender Sublayer (200GXS) is part of the 200GMII Extender (Clause 118). It is identical in function to the 200GBASE-R PCS in Clause 119 with the exceptions defined in Clause 118. The 400 Gigabit 400GMII Extender Sublayer (400GXS) is part of the 400GMII Extender (Clause 118). It is identical in function to the 400GBASE-R PCS in Clause 119 with the exceptions defined in Clause 118.

The optional 200GMII Extender (Clause 118) can be inserted between the Reconciliation Sublayer and the PHY to transparently extend the reach of the 200GMII. The optional 400GMII Extender (Clause 118) can be inserted between the Reconciliation Sublayer and the PHY to transparently extend the reach of the 400GMII.

116.2.3 Physical Coding Sublayer (PCS)

The terms 200GBASE-R and 400GBASE-R refer to a specific family of Physical Layer implementations based upon the 64B/66B data coding method specified in Clause 119 and the PMA specifications defined in Clause 120. Clause 119 PCSs perform encoding (decoding) of data from (to) the 200GMII or 400GMII to 256B/257B code blocks, apply FEC, distribute the data to multiple lanes, and transfer the encoded data to the PMA.

The 200GBASE-R PCS has almost the same functionality as the 200GXS, and therefore may be configured as a 200GXS in order to implement part of the optional 200GMII Extender (see Clause 118). The 400GBASE-R PCS has almost the same functionality as the 400GXS, and therefore may be configured as a 400GXS in order to implement part of the optional 400GMII Extender (see Clause 118).

116.2.4 Physical Medium Attachment (PMA) sublayer

The PMA provides a medium-independent means for the PCS to support the use of a range of physical media. The 200GBASE-R and 400GBASE-R PMAs perform the mapping of transmit and receive data streams between the PCS and PMA via the PMA service interface, and the mapping and multiplexing of transmit and receive data streams between the PMA and PMD via the PMD service interface. In addition, the PMA performs retiming of the received data stream when appropriate, optionally provides data loopback at the PMA or PMD service interface, and optionally provides test pattern generation and checking.

The 200GBASE-R and 400GBASE-R PMAs are specified in Clause 120.

116.2.5 Physical Medium Dependent (PMD) sublayer

The Physical Medium Dependent sublayer is responsible for interfacing to the transmission medium. The PMD is located just above the Medium Dependent Interface (MDI). The MDI, logically subsumed within each PMD subclause, is the actual medium attachment for the various supported media.

The 200GBASE-R PMDs and their corresponding media are specified in Clause 121 and Clause 122. The 400GBASE-R PMDs and their corresponding media are specified in Clause 122 through Clause 124.

116.2.6 Management interface (MDIO/MDC)

The optional MDIO/MDC management interface (Clause 45) provides an interconnection between MDIO Manageable Devices (MMDs) and Station Management (STA) entities.

116.2.7 Management

Managed objects, attributes, and actions are defined for all 200 Gigabit and 400 Gigabit Ethernet components. These items are defined in Clause 30.

116.3 Service interface specification method and notation

The service interface specification for the 200GBASE-R and 400GBASE-R Physical Layers is as per the definition in 1.2.2. Note that the 200GBASE-R and 400GBASE-R inter-sublayer service interfaces use multiple scalar REQUEST and INDICATION primitives, to indicate the transfer of multiple independent streams of data units, as defined in 116.3.1 through 116.3.3.

116.3.1 Inter-sublayer service interface

The inter-sublayer service interface is described in an abstract manner and does not imply any particular implementation. The inter-sublayer service interface primitives are defined as follows:

IS_UNITDATA_ *i*.request
 IS_UNITDATA_ *i*.indication
 IS_SIGNAL.indication

The IS_UNITDATA_ *i*.request (where $i = 0$ to $n - 1$, and n is the number of streams of data units) primitive is used to define the transfer of multiple streams of data units from a sublayer to the next lower (closer to the medium) sublayer. The IS_UNITDATA_ *i*.indication (where $i = 0$ to $n - 1$, and n is the number of streams of data units) primitive is used to define the transfer of multiple streams of data units from a sublayer to the next higher (closer to the MAC) sublayer. The IS_SIGNAL.indication primitive is used to define the transfer of signal status from a sublayer to the next higher sublayer.

116.3.2 Instances of the Inter-sublayer service interface

The inter-sublayer interface can be instantiated between different sublayers, hence a prefix notation is defined to identify a specific instance of an inter-sublayer service interface. The following prefixes are defined:

- a) PMD:—for primitives issued on the interface between the PMD sublayer and the PMA sublayer called the PMD service interface.
- b) PMA:—for primitives issued on the interface between the PMA sublayer and the PCS, DTE 200GXS, DTE 400GXS, or PMA sublayer above called the PMA service interface.
- c) PHY XS:—for primitives issued on the interface between the PMA sublayer and the PHY 200GXS sublayer or PHY 400GXS sublayer called the PHY XS service interface.

Examples of inter-sublayer service interfaces for 200GBASE-R and 400GBASE-R with their corresponding instance names are illustrated in Figure 116-2 and Figure 116-3, respectively. For example, the primitives for one instance of the inter-sublayer service interface, named the PMD service interface, are identified as follows:

PMD:IS_UNITDATA_ *i*.request
 PMD:IS_UNITDATA_ *i*.indication
 PMD:IS_SIGNAL.indication.

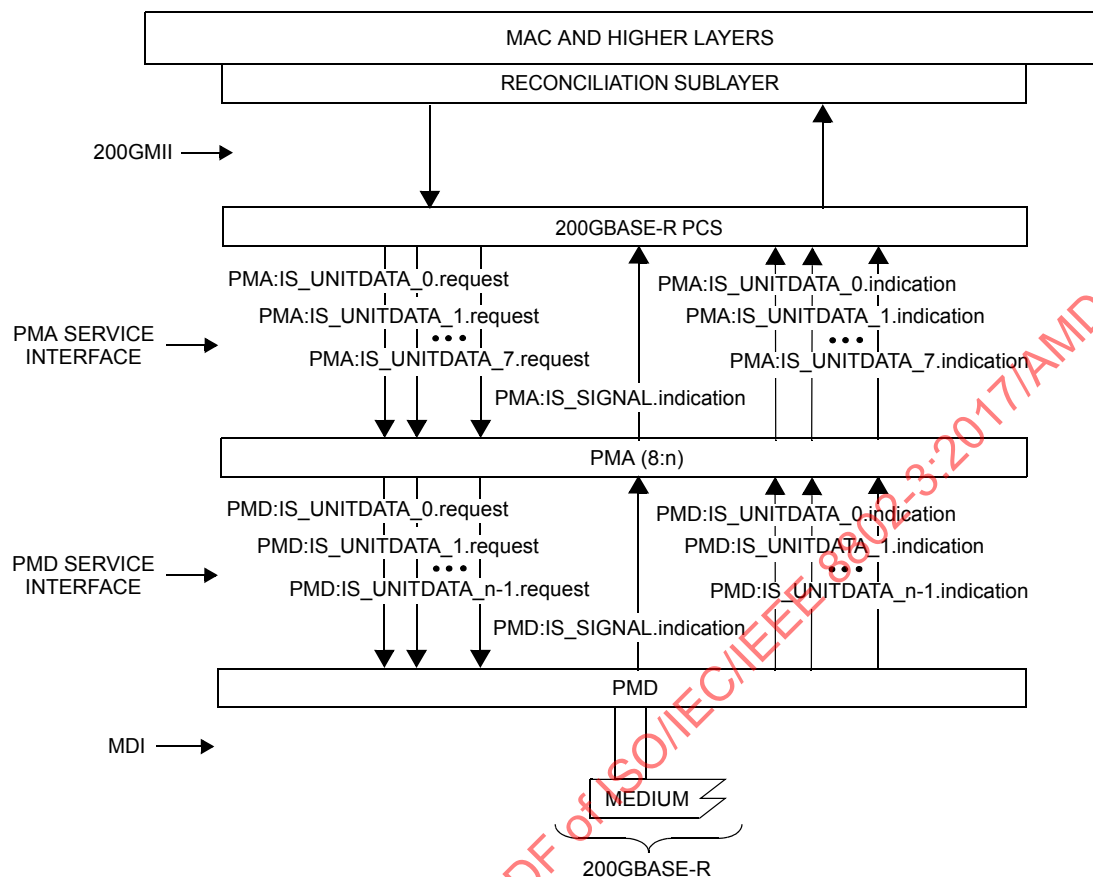
Primitives for other instances, of inter-sublayer interfaces, are represented in a similar manner as described above.

116.3.3 Semantics of inter-sublayer service interface primitives

The semantics of the inter-sublayer service interface primitives for the 200GBASE-R and 400GBASE-R sublayers are described in 116.3.3.1 through 116.3.3.3.

116.3.3.1 IS_UNITDATA_ *i*.request

The IS_UNITDATA_ *i*.request (where $i = 0$ to $n - 1$) primitive is used to define the transfer of multiple streams of data units from a sublayer to the next lower sublayer, where n is the number of parallel streams of data units.



200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER

PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

n = NUMBER OF PARALLEL STREAMS OF DATA UNITS

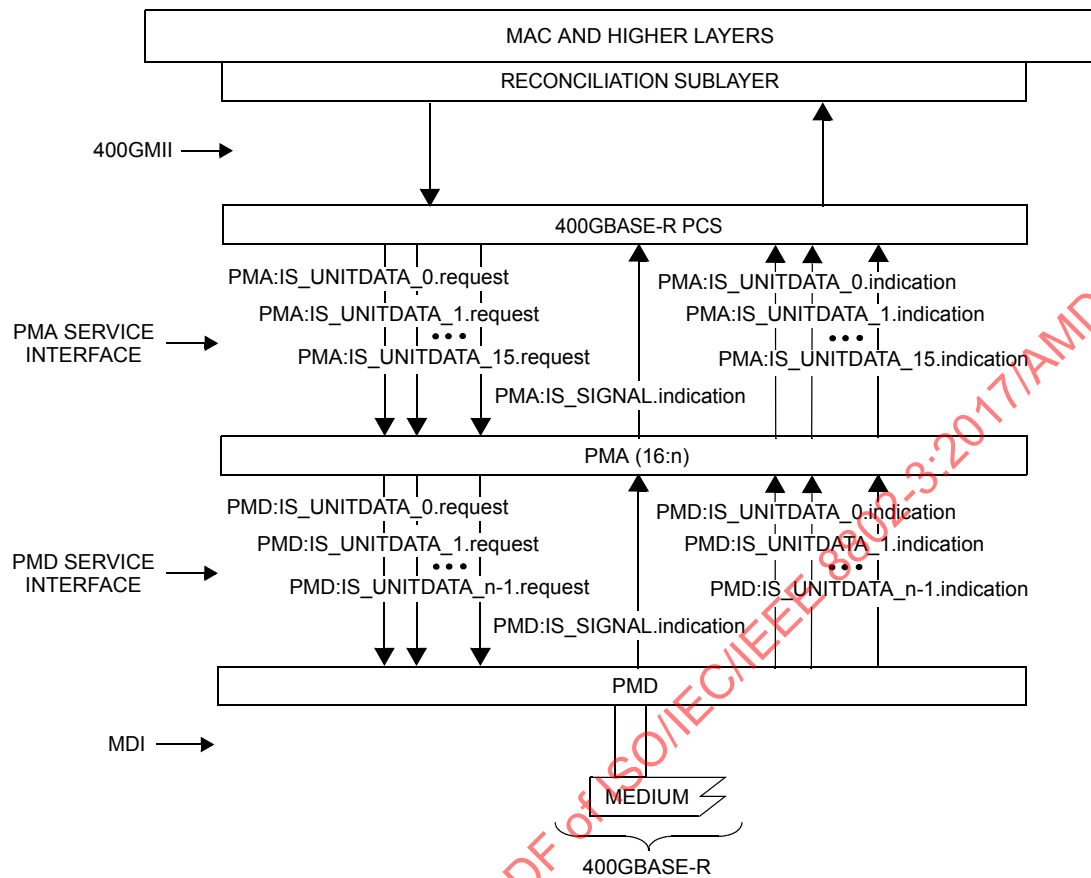
Figure 116-2—200GBASE-R inter-sublayer service interfaces

116.3.3.1.1 Semantics of the service primitive

IS_UNITDATA_0.request(tx_symbol)
IS_UNITDATA_1.request(tx_symbol)
...
IS_UNITDATA_n-1.request(tx_symbol)

The data conveyed by IS_UNITDATA_0.request to IS_UNITDATA_n-1.request consists of n parallel continuous streams of encoded symbols, one stream for each lane. Depending on the specific instance of the inter-sublayer service interface each of the tx_symbol parameters can either take one of two values: zero or one; or take one of four values: zero, one, two, or three.

IEEE Std 802.3bs-2017
Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
for 200 Gb/s and 400 Gb/s Operation



400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER

PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

n = NUMBER OF PARALLEL STREAMS OF DATA UNITS

Figure 116-3—400GBASE-R inter-sublayer service interfaces

116.3.3.1.2 When generated

The sublayer continuously sends n parallel symbol streams $IS_UNITDATA_i.request(tx_symbol)$ to the next lower sublayer, each at a nominal signaling rate defined by a specific instance of the inter-sublayer service interface.

116.3.3.1.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

116.3.3.2 IS_UNITDATA_i.indication

The IS_UNITDATA_i.indication (where $i = 0$ to $n - 1$) primitive is used to define the transfer of multiple streams of data units from the sublayer to the next higher sublayer, where n is the number of parallel streams of data units.

116.3.3.2.1 Semantics of the service primitive

IS_UNITDATA_0.indication(rx_symbol)
IS_UNITDATA_1.indication(rx_symbol)
...
IS_UNITDATA_n-1.indication(rx_symbol)

The data conveyed by IS_UNITDATA_0.indication to IS_UNITDATA_n-1.indication consists of n parallel continuous streams of encoded symbols, one stream for each lane. Depending on the specific instance of the inter-sublayer service interface each of the rx_symbol parameters can either take one of two values: zero or one; or take one of four values: zero, one, two, or three.

116.3.3.2.2 When generated

The sublayer continuously sends n parallel symbol streams IS_UNITDATA_i.indication(rx_symbol) to the next higher sublayer, each at a nominal signaling rate defined by a specific instance of the inter-sublayer service interface.

116.3.3.2.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

116.3.3.3 IS_SIGNAL.indication

The IS_SIGNAL.indication primitive is generated by the sublayer to the next higher sublayer to indicate the status of the receive process. This primitive is generated by the receive process to propagate the detection of severe error conditions (e.g., no valid signal being received by the sublayer that generates this primitive) to the next higher sublayer.

116.3.3.3.1 Semantics of the service primitive

IS_SIGNAL.indication(SIGNAL_OK)

The SIGNAL_OK parameter can take on one of two values: OK or FAIL. A value of FAIL denotes that invalid data is being presented (rx_symbol parameters undefined) by the sublayer to the next higher sublayer. A value of OK does not guarantee valid data is being presented by the sublayer to the next higher sublayer.

116.3.3.3.2 When generated

The sublayer generates the IS_SIGNAL.indication primitive to the next higher sublayer whenever there is change in the value of the SIGNAL_OK parameter.

116.3.3.3.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

116.4 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 implementers must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. Table 116–5 and Table 116–6 contain the values of maximum sublayer delay (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 200 Gigabit and 400 Gigabit Ethernet, respectively. 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.

Table 116–5—Sublayer delay constraints (200GBASE)

| Sublayer | Maximum (bit time) ^a | Maximum (pause_quanta) ^b | Maximum (ns) | Notes ^c |
|---------------------------------------|---------------------------------|-------------------------------------|--------------|-------------------------------------|
| 200G MAC, RS, and MAC Control | 49 152 | 96 | 245.76 | See 117.1.4. |
| 200GBASE-R PCS or 200GXS ^d | 160 256 | 313 | 801.28 | See 119.5. |
| 200GBASE-R PMA | 18 432 | 36 | 92.16 | See 120.5.4. |
| 200GBASE-DR4 PMD | 4 096 | 8 | 20.48 | Includes 2 m of fiber. See 121.3.1. |
| 200GBASE-FR4 PMD | 4 096 | 8 | 20.48 | Includes 2 m of fiber. See 122.3.1. |
| 200GBASE-LR4 PMD | 4 096 | 8 | 20.48 | Includes 2 m of fiber. See 122.3.1. |

^a For 200GBASE-R, 1 bit time (BT) is equal to 5 ps. (See 1.4.117 for the definition of bit time.)

^b For 200GBASE-R, 1 pause_quantum is equal to 2.56 ns. (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 200GMII extender, the delay associated with the 200GMII extender includes two 200GXS sublayers.

Table 116–6—Sublayer delay constraints (400GBASE)

| Sublayer | Maximum (bit time) ^a | Maximum (pause_quanta) ^b | Maximum (ns) | Notes ^c |
|---------------------------------------|---------------------------------|-------------------------------------|--------------|-------------------------------------|
| 400G MAC, RS, and MAC Control | 98 304 | 192 | 245.76 | See 117.1.4. |
| 400GBASE-R PCS or 400GXS ^d | 320 000 | 625 | 800 | See 119.5. |
| 400GBASE-R PMA | 36 864 | 72 | 92.16 | See 120.5.4. |
| 400GBASE-SR16 PMD | 8 192 | 16 | 20.48 | Includes 2 m of fiber. See 123.3.1. |
| 400GBASE-DR4 PMD | 8 192 | 16 | 20.48 | Includes 2 m of fiber. See 124.3.1. |
| 400GBASE-FR8 PMD | 8 192 | 16 | 20.48 | Includes 2 m of fiber. See 122.3.1. |
| 400GBASE-LR8 PMD | 8 192 | 16 | 20.48 | Includes 2 m of fiber. See 122.3.1. |

^a For 400GBASE-R, 1 bit time (BT) is equal to 2.5 ps. (See 1.4.117 for the definition of bit time.)

^b For 400GBASE-R, 1 pause_quantum is equal to 1.28 ns. (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 400GMII extender, the delay associated with the 400GMII extender includes two 400GXS sublayers.

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

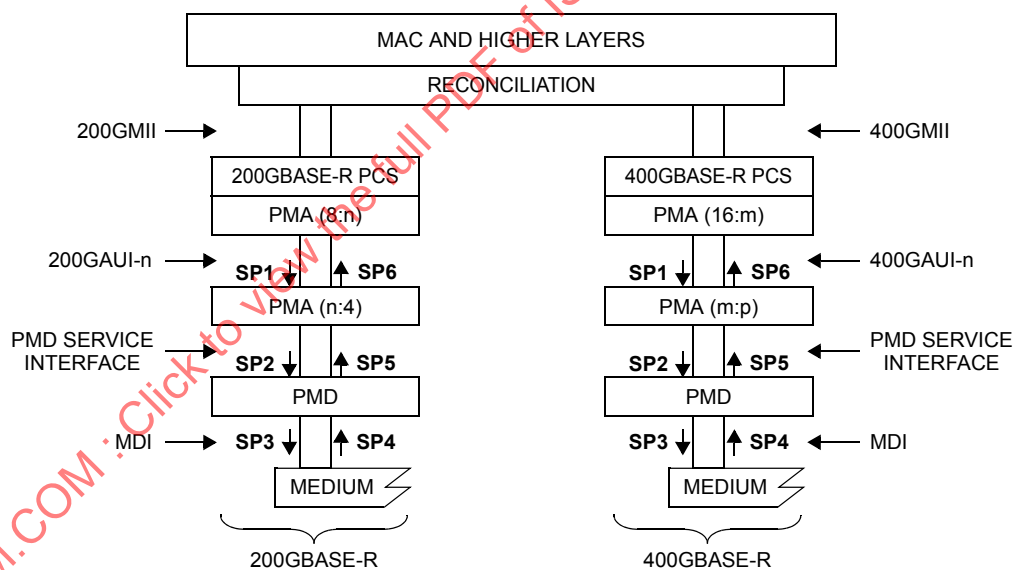
See 31B.3.7 for PAUSE reaction timing constraints for stations at operating speeds of 200 Gb/s and 400 Gb/s.

116.5 Skew constraints

Skew (or relative delay) can be introduced between lanes by both active and passive elements of a 200GBASE-R or 400GBASE-R link. Skew is defined as the difference between the times of the earliest PCS lane and latest PCS lane for the one to zero transition of the alignment marker sync bits. The PCS deskew function (see 119.2.5.1) compensates for all lane-to-lane Skew observed at the receiver. The Skew between the lanes must be kept within limits as shown in Table 116–7 so that the transmitted information on the lanes can be reassembled by the receive PCS.

Skew Variation may be introduced due to variations in electrical, thermal or environmental characteristics. Skew Variation is defined as the change in Skew between any PCS lane and any other PCS lane over the entire time that the link is in operation. From the time the link is brought up, Skew Variation is limited so that each PCS lane always traverses the same lane between any pair of adjacent sublayers while the link remains in operation.

The maximum Skew and Skew Variation at physically instantiated interfaces is specified at Skew points SP1, SP2, and SP3 for the transmit direction and SP4, SP5, and SP6 for the receive direction as illustrated in Figure 116–4 (single 200GAUI-n or 400GAUI-n interface) and Figure 116–5 (multiple 200GAUI-n or 400GAUI-n interfaces).

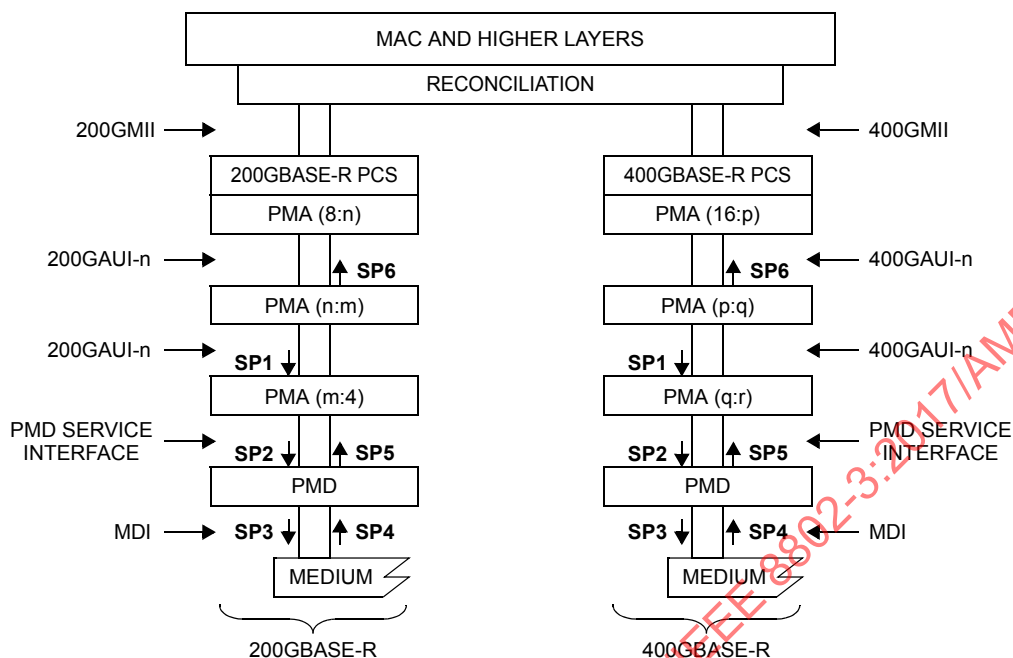


200GAUI-n = 200 Gb/s ATTACHMENT UNIT INTERFACE
 200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
 400GAUI-n = 400 Gb/s ATTACHMENT UNIT INTERFACE
 400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 n = 8 or 4
 m = 16 or 8
 p = 16, 8, or 4

**Figure 116–4—200GBASE-R and 400GBASE-R Skew points
 for single 200GAUI-n or 400GAUI-n**

IEEE Std 802.3bs-2017
Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
for 200 Gb/s and 400 Gb/s Operation



200GAUI-n = 200 Gb/s ATTACHMENT UNIT INTERFACE
200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
400GAUI-n = 400 Gb/s ATTACHMENT UNIT INTERFACE
400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER

PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
n = 8 or 4
m = 8 or 4
p = 16 or 8
q = 16 or 8
r = 16, 8, or 4

**Figure 116-5—200GBASE-R and 400GBASE-R Skew points
for multiple 200GAUI-n or 400GAUI-n**

In the transmit direction, the Skew points are defined in the following locations:

- SP1 on the 200GAUI-n/400GAUI-n interface, at the input of the PMA closest to the PMD;
- SP2 on the PMD service interface, at the input of the PMD;
- SP3 at the output of the PMD, at the MDI.

In the receive direction, the Skew points are defined in the following locations:

- SP4 at the MDI, at the input of the PMD;
- SP5 on the PMD service interface, at the output of the PMD;
- SP6 on the 200GAUI-n/400GAUI-n interface, at the output of the PMA closest to the 200GBASE-R/400GBASE-R PCS or DTE 200GXS/400GXS.

The allowable limits for Skew are shown in Table 116-7 and the allowable limits for Skew Variation are shown in Table 116-8.

The Skew requirements for the PCS, PMA and PMD sublayers are specified in the respective clauses as noted in Table 116-7 and Table 116-8.

Table 116–7—Summary of Skew constraints

| Skew points | Maximum Skew (ns) ^a | Maximum Skew for 200GBASE-R or 400GBASE-R PCS lane (UI) ^b | Notes ^c |
|----------------|--------------------------------|--|--|
| SP1 | 29 | » 770 | See 120.5.3.1 |
| SP2 | 43 | » 1142 | See 120.5.3.3, 121.3.2, 122.3.2, 123.3.2, or 124.3.2 |
| SP3 | 54 | » 1434 | See 121.3.2, 122.3.2, 123.3.2, or 124.3.2 |
| SP4 | 134 | » 3559 | See 121.3.2, 122.3.2, 123.3.2, or 124.3.2 |
| SP5 | 145 | » 3852 | See 121.3.2, 122.3.2, 123.3.2, or 124.3.2 |
| SP6 | 160 | » 4250 | See 120.5.3.5 |
| At PCS receive | 180 | » 4781 | See 119.2.5.1 |

^a The Skew limit includes 1 ns allowance for PCB traces that are associated with the Skew points.

^b The symbol » indicates approximate equivalent of maximum Skew in UI based on 1 UI equals 37.64706 ps at PCS lane signaling rate of 26.5625 GBd.

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

Table 116–8—Summary of Skew Variation constraints

| Skew points | Maximum Skew Variation (ns) | Maximum Skew Variation for 26.5625 GBd PMD lane (UI) ^a | Maximum Skew Variation for 53.125 GBd PMD lane (UI) ^b | Notes ^c |
|----------------|-----------------------------|---|--|--|
| SP1 | 0.2 | » 5 | N/A | See 120.5.3.1 |
| SP2 | 0.4 | » 11 | » 21 | See 120.5.3.3, 121.3.2, 122.3.2, 123.3.2, or 124.3.2 |
| SP3 | 0.6 | » 16 | » 32 | See 121.3.2, 122.3.2, 123.3.2, or 124.3.2 |
| SP4 | 3.4 | » 90 | » 181 | See 121.3.2, 122.3.2, 123.3.2, or 124.3.2 |
| SP5 | 3.6 | » 96 | » 191 | See 121.3.2, 122.3.2, 123.3.2, or 124.3.2 |
| SP6 | 3.8 | » 101 | N/A | See 120.5.3.5 |
| At PCS receive | 4 | » 106 | N/A | See 119.2.5.1 |

^a The symbol » indicates approximate equivalent of maximum Skew Variation in UI based on 1 UI equals 37.64706 ps at PMD lane signaling rate of 26.5625 GBd.

^b The symbol » indicates approximate equivalent of maximum Skew Variation in UI based on 1 UI equals 18.82353 ps at PMD lane signaling rate of 53.125 GBd.

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

116.6 FEC Degrade

FEC degrade is an optional feature allowing for the detection of a non-service affecting link degradation condition based on exceeding a threshold for FEC corrected errors. If there are multiple FEC decoders in a given direction of transmission between the MAC sublayers at each end of the link, a Local Degrade condition is cascaded in that direction of transmission to convey the fact that one or more FEC decoders in the path have exceeded their threshold of FEC corrected errors. If any FEC decoder in a given direction of transmission exceeds its provisioned threshold for FEC corrected errors, a Remote Degrade condition is indicated in the opposite direction of transmission from the PCS or XS closest to the MAC.

Figure 116–6 illustrates the signaling of the Remote Degraded condition in the case that there is no Clause 118 extender sublayer present between the MAC and the PCS. Note that the PCS will not initiate the signaling for local degrade in this configuration as there are no additional FEC decoders in the receive direction between the PCS and the MAC.

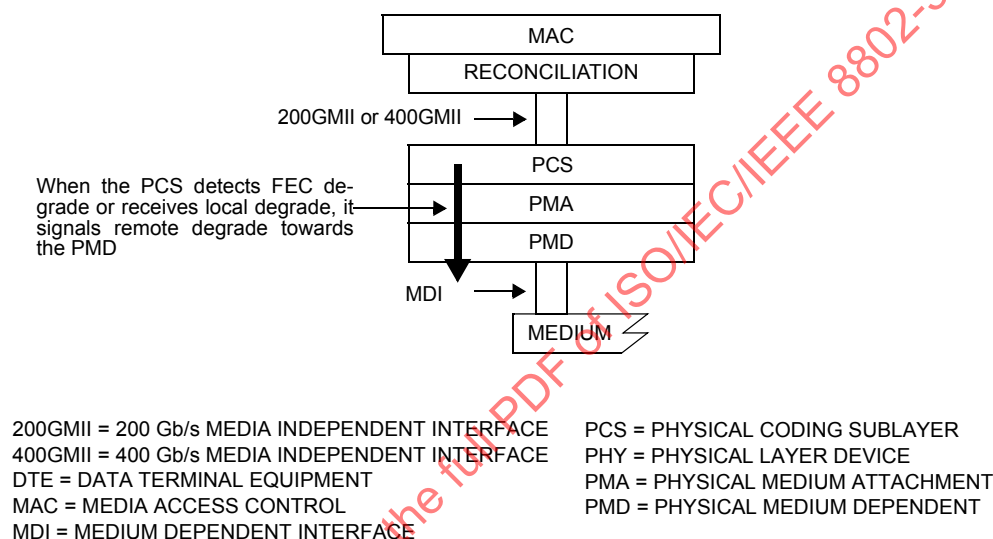


Figure 116–6—Remote Degrade signaling without Extender Sublayer

Figure 116–7 illustrates the signaling of the Local Degraded condition in the case that a Clause 118 extender sublayer is present between the MAC and the PCS.

Figure 116–8 illustrates the signaling of the Remote Degraded condition in the case that a Clause 118 extender sublayer is present between the MAC and the PCS.

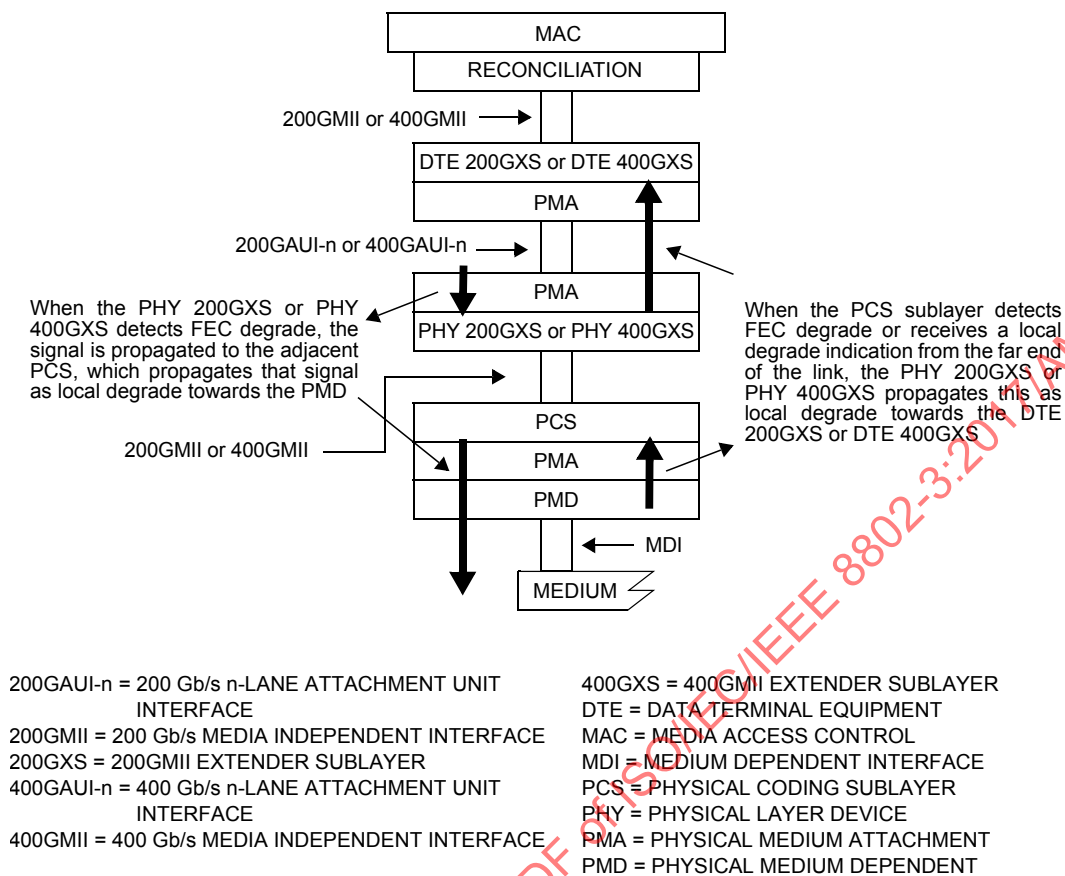


Figure 116-7—Local Degrad Signaling with Extender Sublayer

116.7 State diagrams

State diagrams take precedence over text.

The conventions of 1.2 are adopted, along with the extensions listed in 21.5.

Multiple states of a function that have a transition to a common state utilizing different qualifiers (for example, multiple exit conditions to an IDLE or WAIT state) may be indicated by a shared arrow. An exit transition arrow must connect to the shared arrow, and the qualifier must be met prior to termination of the transition arrow on the shared arrow. The shared arrow has no qualifier.

IEEE Std 802.3bs-2017
Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
for 200 Gb/s and 400 Gb/s Operation

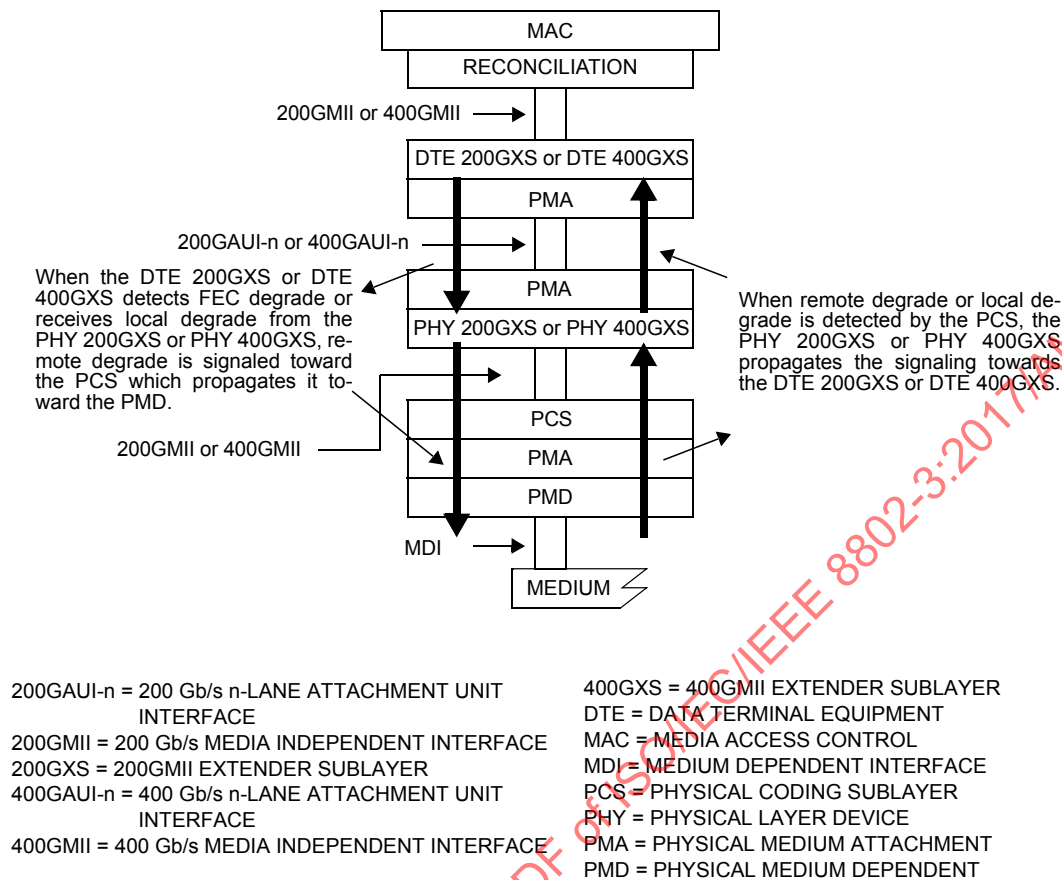


Figure 116-8—Remote Degrad signaling with Extender Sublayer

116.8 Protocol implementation conformance statement (PICS) proforma

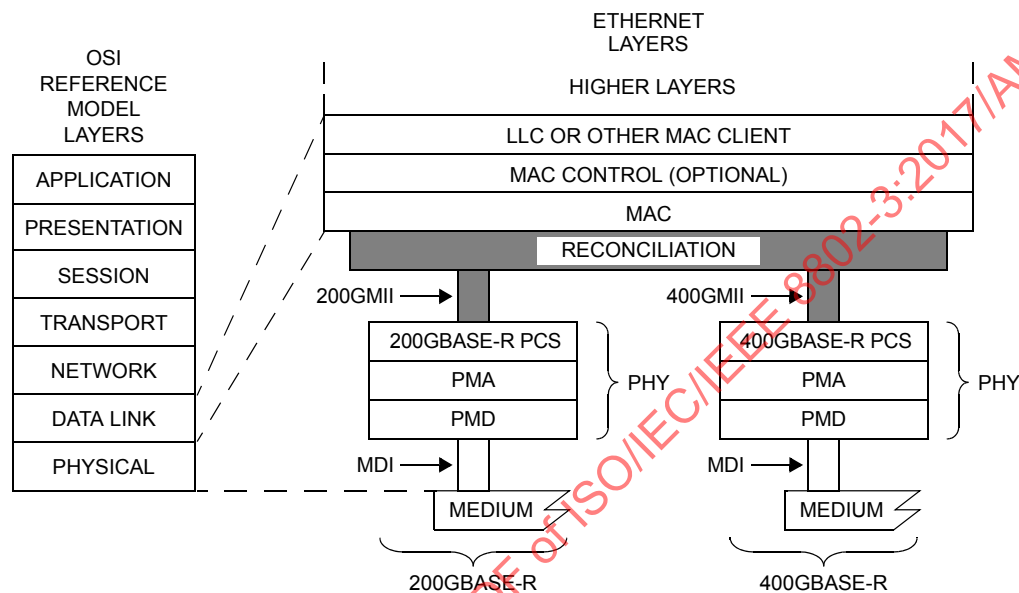
The supplier of a protocol implementation that is claimed to conform to any part of IEEE Std 802.3, Clause 117 through Clause 124, and related annexes demonstrates compliance by completing a protocol implementation conformance statement (PICS) proforma.

A completed PICS proforma is the PICS for the implementation in question. The PICS is a statement of which capabilities and options of the protocol have been implemented. A PICS is included at the end of each clause as appropriate. Each of the 200 Gigabit and 400 Gigabit Ethernet PICS conforms to the same notation and conventions used in 21.6.

117. Reconciliation Sublayer (RS) and Media Independent Interface for 200 Gb/s and 400 Gb/s operation (200GMII and 400GMII)

117.1 Overview

This clause defines the characteristics of the Reconciliation Sublayer (RS) and the Media Independent Interface between Ethernet media access controllers and various PHYs. Figure 117–1 shows the relationship of the RS and Media Independent Interface to the ISO/IEC OSI reference model.



200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
 400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT

Figure 117–1—RS and MII relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

The 200GMII and the 400GMII are optional logical interfaces between the MAC sublayer and the Physical Layer (PHY). The 200GXS/400GXS sublayer in conjunction with the 200GAUI-n/400GAUI-n interface may be used to optionally extend the 200GMII/400GMII.

The RS adapts the bit serial protocols of the MAC to the parallel format of the PCS service interface. Though the 200GMII and 400GMII are optional interfaces, they are used in this standard as a basis for specification. The Physical Coding Sublayer (PCS) is specified to the 200GMII/400GMII, so if not implemented, a conforming implementation shall behave functionally as if the RS and 200GMII/400GMII were implemented.

The 200GMII/400GMII have the following characteristics:

- The 200GMII supports a speed of 200 Gb/s.
- The 400GMII supports a speed of 400 Gb/s.
- Data and delimiters are synchronous to a clock reference.

- d) They provide independent 64-bit wide transmit and receive data paths.
- e) They supports full duplex operation only.

117.1.1 Summary of major concepts

The following are the major concepts of the 200GMII/400GMII:

- a) The 200GMII/400GMII are functionally similar to other media independent interfaces that have been defined for lower speeds, as they all define an interface allowing independent development of MAC and PHY logic.
- b) The RS converts between the MAC serial data stream and the parallel data paths of the 200GMII/400GMII.
- c) The RS maps the signal set provided at the 200GMII/400GMII to the PLS service primitives provided at the MAC.
- d) Each direction of data transfer is independent and serviced by data, control, and clock signals.
- e) The RS generates continuous data or control characters on the transmit path and expects continuous data or control characters on the receive path.
- f) The RS participates in link fault detection and reporting by monitoring the receive path for status reports that indicate an unreliable link, and generating status reports on the transmit path to report detected link faults to the DTE on the remote end of the connecting link.
- g) The 200GMII/400GMII may also support Low Power Idle (LPI) signaling for PHY types supporting Energy Efficient Ethernet (EEE) (see Clause 78).
- h) The 200GMII/400GMII can be extended through the use of a pair of 200GXS/400GXS sublayers (DTE XS and PHY XS) with a 200GAUI-n/400GAUI-n between them.

117.1.2 Application

This clause applies to the interface between the MAC and PHY. This logical interface is used to provide media independence so that an identical media access controller may be used with supported PHY types.

117.1.3 Rate of operation

The 200GMII is specified to support 200 Gb/s operation. The 400GMII is specified to support 400 Gb/s operation.

117.1.4 Delay constraints

The maximum cumulative MAC Control, MAC, and RS delay (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 117–1. A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 116.4 and its references.

Table 117–1—Delay constraints

| Sublayer | Maximum (bit time) | Maximum (pause_quanta) | Maximum (ns) |
|-----------------------------------|-----------------------|---------------------------|--------------|
| 200 Gb/s MAC, RS, and MAC Control | 49 152 | 96 | 245.76 |
| 400 Gb/s MAC, RS, and MAC Control | 98 304 | 192 | 245.76 |

117.1.5 Allocation of functions

The allocation of functions at the 200GMII/400GMII balances the need for media independence with interface simplicity. The 200GMII/400GMII provides media independence by separating the Data Link and Physical Layers of the OSI seven-layer reference model.

117.1.6 200GMII/400GMII structure

The 200GMII/400GMII structure is identical to the CGMII structure specified in 81.1.6.

117.1.7 Mapping of 200GMII/400GMII signals to PLS service primitives

The Reconciliation Sublayer (RS) shall map the signals provided at the 200GMII/400GMII to the PLS service primitives defined in Clause 6. The PLS service primitives provided by the RS and described here behave in exactly the same manner as defined in Clause 6. Full duplex operation only is implemented at either 200 Gb/s or 400 Gb/s, therefore, PLS service primitives supporting CSMA/CD operation are not mapped through the RS to the 200GMII/400GMII. This behavior and restrictions are the same as described in 22.7, with the details of the signaling described in 117.3. LPI_REQUEST shall not be set to ASSERT unless the attached link has been operational for at least one second (i.e., link_status = OK, according to the underlying PCS/PMA).

EEE capability requires the use of the MAC defined in Annex 4A for simplified full duplex operation (with carrier sense deferral). This provides full duplex operation but uses the carrier sense signal to defer transmission until the system returns from its low power state.

Mappings for the following primitives are defined for 200 Gb/s and 400 Gb/s operation:

PLS_DATA.request
PLS_DATA.indication
PLS_CARRIER.indication
PLS_SIGNAL.indication
PLS_DATA_VALID.indication

The RS maps all primitives in an identical manner as the CGMII does and as specified in 81.1.7.

117.2 200GMII/400GMII data stream

The 200GMII/400GMII data stream has the same characteristics as the CGMII data stream described in 81.2.

117.3 200GMII/400GMII functional specifications

The 200GMII/400GMII functions identically to the CGMII specified in 81.3.

117.4 LPI Assertion and Detection

LPI assertion and detection function identically to the CGMII specified in 81.4, with the single exception that the PMA stop signaling described in 81.4.4 is not applicable.

117.5 Protocol implementation conformance statement (PICS) proforma for Clause 117, Reconciliation Sublayer (RS) and Media Independent Interface for 200 Gb/s and 400 Gb/s operation (200GMII and 400GMII)²

117.5.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 117, Reconciliation Sublayer (RS) and Media Independent Interface for 200 Gb/s and 400 Gb/s operation (200GMII and 400GMII), shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

117.5.2 Identification

117.5.2.1 Implementation identification

| | |
|--|--|
| Supplier ¹ | |
| Contact point for inquiries about the PICS ¹ | |
| Implementation Name(s) and Version(s) ^{1,3} | |
| Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ² | |
| NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model). | |

117.5.2.2 Protocol summary

| | |
|---|---|
| Identification of protocol standard | IEEE Std 802.3bs-2017, Clause 117, Reconciliation Sublayer (RS) and Media Independent Interface for 200 Gb/s and 400 Gb/s operation (200GMII and 400GMII) |
| Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS | |
| Have any Exception items been required? No <input type="checkbox"/> Yes <input type="checkbox"/> (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3bs-2017.) | |

| | |
|-------------------|--|
| Date of Statement | |
|-------------------|--|

²Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

117.5.3 Major capabilities/options

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|--------------|---------------|--------|-------------------|
| *MII | PHY support of either 200GMII or 400GMII | 117.2, 117.3 | | O | Yes [] No [] |
| *LPI | Implementation of LPI | 117.1.7 | | O | Yes [] No [] |

117.5.4 PICS proforma tables for Reconciliation Sublayer (RS) and Media Independent Interface for 200 Gb/s and 400 Gb/s operation (200GMII and 400GMII)**117.5.4.1 General**

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|-----------------|--------|---------|
| G1 | Cumulative MAC Control, MAC and RS delay | 117.1.4 | Per Table 117-1 | M | Yes [] |
| G3 | Lane structure | 117.1.6 | Per Table 81-2 | M | Yes [] |

117.5.4.2 Mapping of PLS service primitives

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--------------------------------------|-----------|---|--------|--------------------|
| PL1 | Mapping to Clause 6 | 117.1.7 | RS implements mapping to Clause 6 PLS service primitives | MII:M | Yes [] N/A [] |
| PL2 | Mapping of PLS_DATA.requests | 117.1.7 | In sequence TXD<0> to TXD<63> | MII:M | Yes [] N/A [] |
| PL3 | Start control character creation | 117.1.7 | First octet of preamble converted to Start control character | MII:M | Yes [] N/A [] |
| PL4 | TXD and TXC generation | 117.1.7 | For each 64 PLS_DATA.requests | MII:M | Yes [] N/A [] |
| PL5 | Terminate control character creation | 117.1.7 | DATA_COMPLETE causes creation of Terminate control character in next lane in sequence | MII:M | Yes [] N/A [] |
| PL6 | Mapping RXD to PLS_DATA.indications | 117.1.7 | Create PLS_DATA.indications in sequence from RXD<0> to RXD<63> | MII:M | Yes [] N/A [] |
| PL7 | PLS_DATA.indication generation | 117.1.7 | Generate 64 PLS_DATA.indications for each RXD<63:0> until Terminate then generating 0, 8, 16, 24, 32, 40, 48, or 56 | MII:M | Yes [] N/A [] |
| PL8 | Start control character conversion | 117.1.7 | Convert valid Start control character to preamble before generating PLS_DATA.indications | MII:M | Yes [] N/A [] |

IEEE Std 802.3bs-2017
 Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 200 Gb/s and 400 Gb/s Operation

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---|--------|--------------------|
| PL9 | Terminate control character | 117.1.7 | No PLS_DATA.indications generated | MII:M | Yes [] N/A [] |
| PL10 | PLS_DATA_VALID.indication generation | 117.1.7 | On change of value of DATA_VALID_STATUS | MII:M | Yes [] N/A [] |
| PL11 | DATA_VALID_STATUS | 117.1.7 | Value of DATA_VALID on a lane 0 Start control character preceded by eight idles, a Sequence ordered set, or a Terminate character | MII:M | Yes [] N/A [] |
| PL12 | DATA_VALID_STATUS | 117.1.7 | Value of DATA_NOT_VALID on any control character but Error | MII:M | Yes [] N/A [] |
| PL13 | Frame not ending with Terminate control character | 117.1.7 | Ensure MAC detects CRC error | MII:M | Yes [] N/A [] |

117.5.4.3 Data stream structure

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|----------------------|-----------|------------------------------------|--------|--------------------|
| DS1 | Frame transfer | 117.2 | Within 200GMII/400GMII data stream | MII:M | Yes [] N/A [] |
| DS2 | Bit mapping | 117.2 | Per Figure 81–4 | MII:M | Yes [] N/A [] |
| DS3 | Content of <data> | 117.2 | Consist of data octets | MII:M | Yes [] N/A [] |
| DS4 | Recognition of <efd> | 117.2 | Terminate recognized in any lane | MII:M | Yes [] N/A [] |

117.5.4.4 200GMII/400GMII signal functional specifications

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--------------------------------|-----------|--|--------|--------------------|
| FS1 | TX_CLK active edge | 117.3 | TXD and TXC sampled on the rising edge of TX_CLK | MII:M | Yes [] N/A [] |
| FS2 | TX_CLK frequency | 117.3 | One-sixty-fourth of the MAC transmit data rate | MII:M | Yes [] N/A [] |
| FS3 | TXC assertion and de-assertion | 117.3 | De-asserted for data, asserted for control character | MII:M | Yes [] N/A [] |
| FS4 | TXC clock | 117.3 | Synchronous to TX_CLK | MII:M | Yes [] N/A [] |
| FS5 | TXD encoding | 117.3 | Per Table 81–3 | MII:M | Yes [] N/A [] |
| FS6 | TXD clock | 117.3 | Synchronous to TX_CLK | MII:M | Yes [] N/A [] |
| FS7 | Start alignment | 117.3 | Start control character aligned to lane 0 | MII:M | Yes [] N/A [] |

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|----------------------------------|-----------|---|--------|--------------------|
| FS8 | RX_CLK active edge | 117.3 | RXD and RXC sampled on the rising edge of RX_CLK | MII:M | Yes [] N/A [] |
| FS9 | RX_CLK frequency | 117.3 | One-sixty-fourth of the MAC receive data rate | MII:M | Yes [] N/A [] |
| FS10 | Loss of receive signal | 117.3 | Source RX_CLK from nominal clock | MII:M | Yes [] N/A [] |
| FS11 | RXC assertion and de-assertion | 117.3 | De-asserted for data, asserted for control character | MII:M | Yes [] N/A [] |
| FS12 | RXC clock | 117.3 | Synchronous to RX_CLK | MII:M | Yes [] N/A [] |
| FS13 | RXD decoding | 117.3 | Per Table 81–4 | MII:M | Yes [] N/A [] |
| FS14 | RXD clock | 117.3 | Synchronous to RX_CLK | MII:M | Yes [] N/A [] |
| FS15 | Received Error control character | 117.3 | RS cause MAC FrameCheckError | MII:M | Yes [] N/A [] |
| FS16 | DATA_VALID assertion | 117.3 | RS not assert DATA_VALID unless Start control character in lane 0 | MII:M | Yes [] N/A [] |

117.5.4.5 Link fault signaling state diagram

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|------------------------------------|-----------|--|--------|--------------------|
| LF1 | Link fault signaling state diagram | 117.3 | Implement per Figure 81–11 | MII:M | Yes [] N/A [] |
| LF2 | link_fault = OK and MAC frames | 117.3 | RS services MAC frame transmission requests | MII:M | Yes [] N/A [] |
| LF3 | link_fault = OK and no MAC frames | 117.3 | In absence of MAC frames, RS transmits Idle control characters | MII:M | Yes [] N/A [] |
| LF4 | link_fault = Local Fault | 117.3 | RS transmits continuous Remote Fault Sequence ordered sets | MII:M | Yes [] N/A [] |
| LF5 | link_fault = Remote Fault | 117.3 | RS transmits continuous Idle control characters | MII:M | Yes [] N/A [] |

117.5.4.6 LPI functions

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|----------------------------------|-----------|--------------------------|--------|--------------------|
| L1 | Assertion of LPI in Tx direction | 117.3 | As defined in Table 81–3 | LPI:M | Yes [] N/A [] |
| L2 | Assertion of LPI in Rx direction | 117.3 | As defined in Table 81–4 | LPI:M | Yes [] N/A [] |

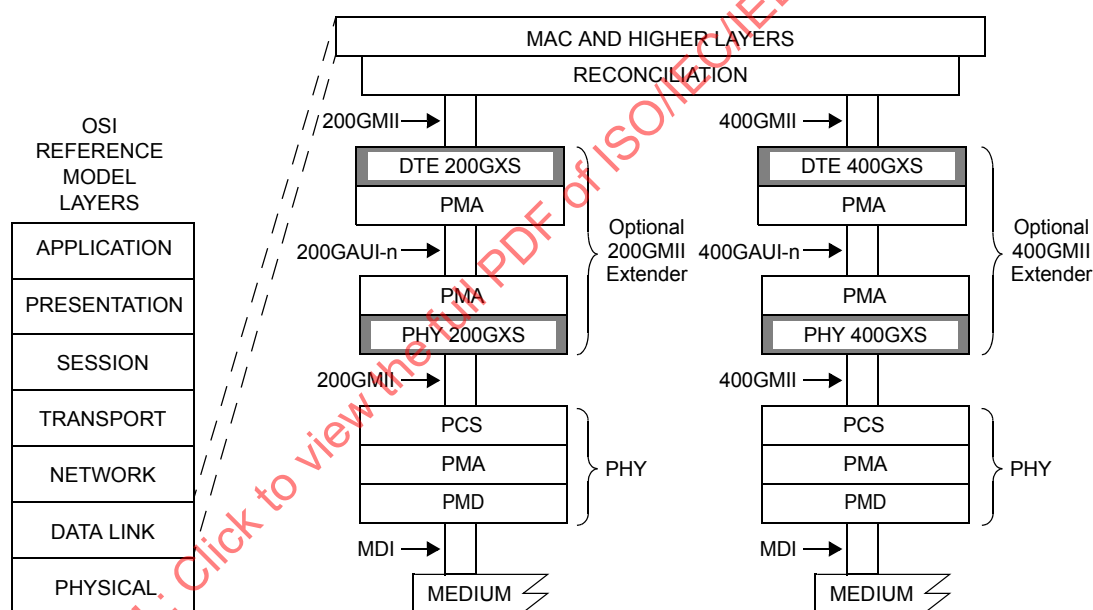
118. 200GMII Extender, 400GMII Extender, 200GMII Extender Sublayer (200GXS), and 400GMII Extender Sublayer (400GXS)

118.1 Overview

This clause defines the functional characteristics for the optional 200GMII Extender and 200GMII Extender Sublayer (200GXS), and also for the optional 400GMII Extender and 400GMII Extender Sublayer (400GXS). Figure 118–1 shows the relationship of the 200GMII/400GMII Extender and 200GXS/400GXS sublayer with other sublayers to the ISO Open System Interconnection (OSI) reference model.

The 200GMII/400GMII Extender allows the extension of the 200GMII/400GMII to the PCS via a physical instantiation. The 200GMII/400GMII Extender is composed of a 200GXS/400GXS at the RS end, a 200GXS/400GXS at the PHY end with a physical instantiation of 200GAUI-n/400GAUI-n between two adjacent PMA sublayers.

A 200GMII/400GMII Extender with the optional Energy-Efficient Ethernet (EEE) capability (see Clause 78) encodes and decodes Low Power Idle (LPI) signals. The assertion of LPI at the 200GMII/400GMII is encoded in the transmitted symbols. Detection of LPI encoding in the received symbols is indicated as LPI at the 200GMII/400GMII.



200GAUI-n = 200 Gb/s n-LANE ATTACHMENT UNIT
INTERFACE

200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE

200GXS = 200GMII EXTENDER SUBLAYER

400GAUI-n = 400 Gb/s n-LANE ATTACHMENT UNIT
INTERFACE

400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE

400GXS = 400GMII EXTENDER SUBLAYER

DTE = DATA TERMINAL EQUIPMENT

MAC = MEDIA ACCESS CONTROL

MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE

PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT

Figure 118–1—200GXS and 400GXS relationship to the ISO/IEC Open System Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

118.1.1 Summary of major concepts

The following is a list of the major concepts of the 200GMII/400GMII Extender:

- a) Simple signal mapping to the 200GMII/400GMII
- b) The optional 200GMII/400GMII Extender can be inserted between the Reconciliation Sublayer and the PHY to transparently extend the reach of the 200GMII/400GMII
- c) Independent transmit and receive data paths
- d) The 200GXS/400GXS leverages all functions in the Clause 119 PCS and supports physical instantiations of the 200GAUI-n/400GAUI-n
- e) Optionally extends LPI signaling to the PHY for EEE

118.1.2 200GXS/400GXS Sublayer

The 200GXS, if implemented, shall be identical in function to the 200GBASE-R PCS in Clause 119 with the addition of the functions defined in 118.2. A single device may be configured as either a 200GXS or the 200GBASE-R PCS and may be managed through different optional management registers.

The 400GXS, if implemented, shall be identical in function to the 400GBASE-R PCS in Clause 119 with the addition of the functions defined in 118.2. A single device may be configured as either a 400GXS or the 400GBASE-R PCS and may be managed through different optional management registers.

118.1.3 200GAUI-n/400GAUI-n

A 200GMII Extender may use any of the following physical instantiations of the 200GAUI-n:

- 200GAUI-8 chip-to-chip (Annex 120B)
- 200GAUI-8 chip-to-module (Annex 120C)
- 200GAUI-4 chip-to-chip (Annex 120D)
- 200GAUI-4 chip-to-module (Annex 120E)

A 400GMII Extender may use any of the following physical instantiations of the 400GAUI-n:

- 400GAUI-16 chip-to-chip (Annex 120B)
- 400GAUI-16 chip-to-module (Annex 120C)
- 400GAUI-8 chip-to-chip (Annex 120D)
- 400GAUI-8 chip-to-module (Annex 120E)

118.2 FEC Degradation

The propagation of FEC degradation signaling across PCS and XS sublayers is described in 116.6 and is based on the optional FEC degradation 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 Degradation signaling

The variable tx_am_sf is set as follows:

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

118.2.2 PHY XS FEC Degradе signaling

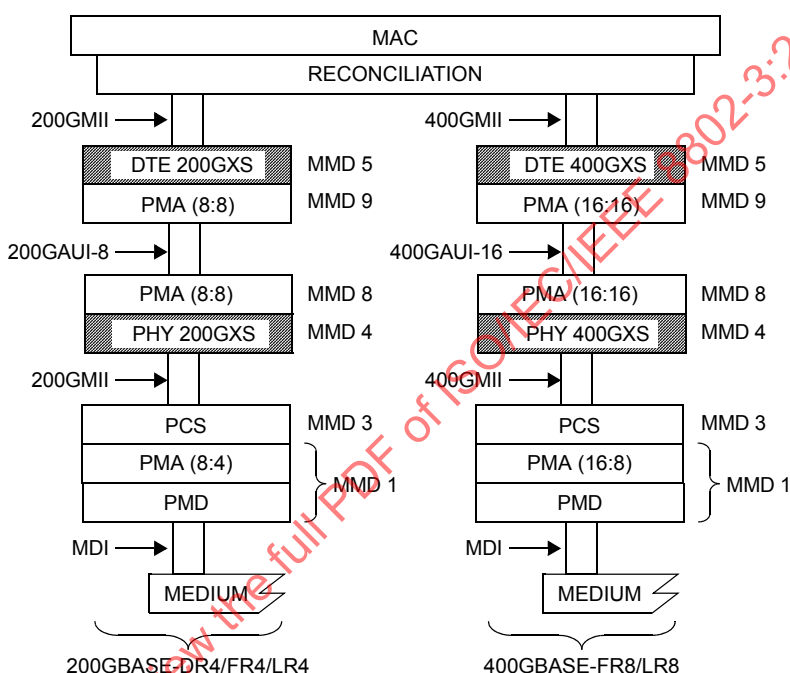
The variable tx_am_sf is set as follows:

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

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.

118.3 200GXS and 400GXS partitioning example

A partitioning example and MMD numbering using the 200GXS and 400GXS is shown in Figure 118–2.



200GAUI = 200 Gb/s ATTACHMENT UNIT INTERFACE
 200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
 200GXS = 200 Gb/s EXTENDER SUBLAYER
 400GAUI = 400 Gb/s ATTACHMENT UNIT INTERFACE
 400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
 400GXS = 400 Gb/s EXTENDER SUBLAYER
 DTE = DATA TERMINAL EQUIPMENT

MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 MMD = MDIO MANAGEABLE DEVICE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT

Figure 118–2—Example 200GBASE-DR4/FR4/LR4 and 400GBASE-FR8/LR8 PMA layering with 200GXS and 400GXS

118.4 200GXS and 400GXS MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the 200GXS or 400GXS. If MDIO is implemented, it shall map MDIO PHY XS and DTE XS control bits to Clause 119 control variables as shown in Table 118–1 and Table 118–3, respectively. Similarly, if MDIO is implemented, it shall map MDIO PHY XS and DTE XS status bits to Clause 119 status variables as shown in Table 118–2 and Table 118–4, respectively.

Table 118–1—MDIO PHY XS to Clause 119 control variable mapping

| MDIO control variable | PHY XS register name | Register/ bit number | Clause 119 control variable |
|--|---|----------------------|---------------------------------------|
| Reset | PHY XS control 1 register | 4.0.15 | reset |
| Loopback | PHY XS control 1 register | 4.0.14 | Loopback |
| Transmit test-pattern enable | BASE-R PHY XS test-pattern control register | 4.42.3 | tx_test_mode |
| PHY XS FEC bypass indication enable | PHY XS FEC control register | 4.800.1 | FEC_bypass_indication_enable |
| PHY XS FEC degraded SER enable | PHY XS FEC control register | 4.800.2 | FEC_degraded_SER_enable |
| PHY XS FEC degraded SER activate threshold | PHY XS FEC degraded SER activate threshold register | 4.806, 4.807 | FEC_degraded_SER_activate_threshold |
| PHY XS FEC degraded SER deactivate threshold | PHY XS FEC degraded SER deactivate threshold register | 4.808, 4.809 | FEC_degraded_SER_deactivate_threshold |
| PHY XS FEC degraded SER interval | PHY XS FEC degraded SER interval | 4.810, 4.811 | FEC_degraded_SER_interval |

Table 118–2—MDIO PHY XS to Clause 119 status variable mapping

| MDIO status variable | PHY XS register name | Register/ bit number | Clause 119 status variable |
|---|--|----------------------|------------------------------------|
| BASE-R PHY XS receive link status | BASE-R PHY XS status 1 | 4.32.12 | PCS_status |
| Lane <i>x</i> aligned | Multi-lane BASE-R PHY XS alignment status 3 and 4 | 4.52.7:0 4.53.7:0 | am_lock< <i>x</i> > |
| PHY XS lane alignment status | Multi-lane BASE-R PHY XS alignment status 1 | 4.50.12 | align_status |
| Lane <i>x</i> mapping | PHY XS lane mapping, lane 0 through lane 15 | 4.400 through 4.415 | pcs_lane_mapping< <i>x</i> > |
| PHY XS FEC bypass indication ability | PHY XS FEC status | 4.801.1 | FEC_bypass_indication_ability |
| FEC corrected codewords | PHY XS FEC corrected codewords counter | 4.802, 4.803 | FEC_corrected_cw_counter |
| FEC uncorrected codewords | PHY XS FEC uncorrected codewords counter | 4.804, 4.805 | FEC_uncorrected_cw_counter |
| PHY XS FEC symbol errors, lane 0 to lane 15 | PHY XS FEC symbol error counter, lane 0 to lane 15 | 4.600 to 4.631 | FEC_symbol_error_counter_ <i>i</i> |
| Tx LPI indication | PHY XS status 1 | 4.1.9 | Tx LPI indication |
| Tx LPI received | PHY XS status 1 | 4.1.11 | Tx LPI received |
| Rx LPI indication | PHY XS status 1 | 4.1.8 | Rx LPI indication |
| Rx LPI received | PHY XS status 1 | 4.1.10 | Rx LPI received |
| EEE wake error counter | EEE wake error counter | 4.22 | Wake_error_counter |

Table 118–2—MDIO PHY XS to Clause 119 status variable mapping (continued)

| MDIO status variable | PHY XS register name | Register/ bit number | Clause 119 status variable |
|---------------------------------|----------------------------|----------------------|----------------------------|
| PHY XS FEC degraded SER ability | PHY XS FEC status register | 4.801.3 | FEC_degraded_SER_ability |
| PHY XS FEC degraded SER | PHY XS FEC status register | 4.801.4 | FEC_degraded_SER |
| Remote degraded SER received | PHY XS FEC status register | 4.801.5 | rx_rm_degraded |

Table 118–3—MDIO DTE XS to Clause 119 control variable mapping

| MDIO control variable | DTE XS register name | Register/ bit number | Clause 119 control variable |
|--|---|----------------------|---------------------------------------|
| Reset | DTE XS control 1 register | 5.0.15 | reset |
| Loopback | DTE XS control 1 register | 5.0.14 | Loopback |
| Transmit test-pattern enable | BASE-R DTE XS test-pattern control register | 5.42.3 | tx_test_mode |
| DTE XS FEC bypass indication enable | DTE XS FEC control register | 5.800.1 | FEC_bypass_indication_enable |
| DTE XS FEC degraded SER enable | DTE XS FEC control register | 5.800.2 | FEC_degraded_SER_enable |
| DTE XS FEC degraded SER activate threshold | DTE XS FEC degraded SER activate threshold register | 5.806, 5.807 | FEC_degraded_SER_activate_threshold |
| DTE XS FEC degraded SER deactivate threshold | DTE XS FEC degraded SER deactivate threshold register | 5.808, 5.809 | FEC_degraded_SER_deactivate_threshold |
| DTE XS FEC degraded SER interval | DTE XS FEC degraded SER interval | 5.810, 5.811 | FEC_degraded_SER_interval |

Table 118–4—MDIO DTE XS to Clause 119 status variable mapping

| MDIO status variable | DTE XS register name | Register/ bit number | Clause 119 status variable |
|--------------------------------------|---|----------------------|-------------------------------|
| BASE-R DTE XS receive link status | BASE-R DTE XS status 1 | 5.32.12 | PCS_status |
| Lane <i>x</i> aligned | Multi-lane BASE-R DTE XS alignment status 3 and 4 | 5.52.7:0 5.53.7:0 | am_lock< <i>x</i> > |
| DTE XS lane alignment status | Multi-lane BASE-R DTE XS alignment status 1 | 5.50.12 | align_status |
| Lane <i>x</i> mapping | DTE XS lane mapping, lane 0 through lane 15 | 5.400 through 5.415 | pcs_lane_mapping< <i>x</i> > |
| DTE XS FEC bypass indication ability | DTE XS FEC status | 5.801.1 | FEC_bypass_indication_ability |
| FEC corrected codewords | DTE XS FEC corrected codewords counter | 5.802, 5.803 | FEC_corrected_cw_counter |

Table 118–4—MDIO DTE XS to Clause 119 status variable mapping (continued)

| MDIO status variable | DTE XS register name | Register/ bit number | Clause 119 status variable |
|---|--|----------------------|----------------------------|
| FEC uncorrected codewords | DTE XS FEC uncorrected codewords counter | 5.804, 5.805 | FEC_uncorrected_cw_counter |
| DTE XS FEC symbol errors, lane 0 to lane 15 | DTE XS FEC symbol error counter, lane 0 to lane 15 | 5.600 to 5.631 | FEC_symbol_error_counter_i |
| Tx LPI indication | DTE XS status 1 | 5.1.9 | Tx LPI indication |
| Tx LPI received | DTE XS status 1 | 5.1.11 | Tx LPI received |
| Rx LPI indication | DTE XS status 1 | 5.1.8 | Rx LPI indication |
| Rx LPI received | DTE XS status 1 | 5.1.10 | Rx LPI received |
| EEE wake error counter | EEE wake error counter | 5.22 | Wake_error_counter |
| DTE XS FEC degraded SER ability | DTE XS FEC status register | 5.801.3 | FEC_degraded_SER_ability |
| DTE XS FEC degraded SER | DTE XS FEC status register | 5.801.4 | FEC_degraded_SER |
| Remote degraded received | DTE XS FEC status register | 5.801.5 | rx_rm_degraded |
| Local degraded received | DTE XS FEC status register | 5.801.6 | rx_local_degraded |

118.5 Protocol implementation conformance statement (PICS) proforma for Clause 118, 200GMII Extender, 400GMII Extender, 200GMII Extender Sublayer (200GXS), and 400GMII Extender Sublayer (400GXS)³

118.5.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 118, 200GMII Extender, 400GMII Extender, 200GMII Extender Sublayer (200GXS), and 400GMII Extender Sublayer (400GXS), shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

118.5.2 Identification

118.5.2.1 Implementation identification

| | |
|--|--|
| Supplier ¹ | |
| Contact point for inquiries about the PICS ¹ | |
| Implementation Name(s) and Version(s) ^{1,3} | |
| Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ² | |
| NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model). | |

118.5.2.2 Protocol summary

| | |
|---|---|
| Identification of protocol standard | IEEE Std 802.3bs-2017, Clause 118, 200GMII Extender, 400GMII Extender, 200GMII Extender Sublayer (200GXS), and 400GMII Extender Sublayer (400GXS) |
| Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS | |
| Have any Exception items been required? No <input type="checkbox"/> Yes <input type="checkbox"/> (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3bs-2017.) | |

| | |
|-------------------|--|
| Date of Statement | |
|-------------------|--|

³Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

118.5.3 Major capabilities/options

| Item | Feature | Subclause | Value/Comment | Status | Support |
|---------|---|-----------------------|---|--------|-------------------|
| MII | 200GMII or 400GMII logical interface | 117, 118.1 | Logical interface is supported | O | Yes [] No [] |
| *PHYXS | PHY 200GXS or PHY 400GXS | 118.1 | | O/2 | Yes [] No [] |
| *DTEXS | DTE 200GXS or DTE 400GXS | 118.1 | | O/2 | Yes [] No [] |
| *200GXS | Using the XS for 200GBASE-R | 118.1.2 | | O.1 | Yes [] No [] |
| *400GXS | Using the XS for 400GBASE-R | 118.1.2 | | O.1 | Yes [] No [] |
| *MD | MDIO | 45, 118.4 | Registers and interface supported | O | Yes [] No [] |
| *BI | Bypass indication | 119.2.5.3 | Capability is supported | O | Yes [] No [] |
| DC | Delay constraints | 119.5 | Conforms to delay constraints specified in 119.5 | M | Yes [] |
| EEE | EEE capability | 119.2.3.3 | Capability is supported | O | Yes [] No [] |
| JTM | Supports test-pattern mode | 119.2.1, 119.2.4.9 | | M | Yes [] |
| FDD | Support for optional FEC degraded detection | 119.2.5.3 | In the FEC decoder can optionally detect FEC SER degraded at a programmable threshold | O | Yes [] No [] |

118.5.4 PICS proforma tables for 200GMII Extender, 400GMII Extender, 200GMII Extender Sublayer (200GXS), and 400GMII Extender Sublayer (400GXS)

118.5.4.1 Transmit function

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---------------------------------|-----------|--|--------|---------|
| TF1 | 64B/66B to 256B/257B transcoder | 119.2.4.2 | tx_xcoded<256:0> constructed per 119.2.4.2 | M | Yes [] |
| TF2 | Transmission bit ordering | 119.2.4.8 | First bit transmitted is bit 0 | M | Yes [] |
| TF3 | Pad value | 119.2.4.4 | PRBS9 | M | Yes [] |
| TF4 | Alignment marker insertion | 119.2.4.4 | | M | Yes [] |
| TF5 | Pre-FEC distribution | 119.2.4.5 | Distribute the data to two FEC codewords | M | Yes [] |
| TF6 | Reed-Solomon encoder | 119.2.4.6 | RS(544,514) | M | Yes [] |
| TF7 | Symbol distribution | 119.2.4.7 | Distribution is based on 10b symbols | M | Yes [] |

118.5.4.2 Receive function

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---|--------|--------------------|
| RF1 | Skew tolerance | 119.2.5.1 | Maximum Skew of 180 ns between PCS lanes and a maximum Skew Variation of 4 ns | M | Yes [] |
| RF2 | Lane reorder and de-interleave | 119.2.5.2 | Order the PCS lanes according to the PCS lane number and de-interleave the FEC codewords | M | Yes [] |
| RF3 | Reed-Solomon decoder | 119.2.5.3 | Corrects any combination of up to $t=15$ symbol errors in a codeword | M | Yes [] |
| RF4 | Reed-Solomon decoder | 119.2.5.3 | Capable of indicating when a codeword was not corrected. | M | Yes [] |
| RF5 | Error monitoring while error indication is bypassed | 119.2.5.3 | When the number of symbol errors in a block of 8192 codewords exceeds 5560 assert hi_ser for 60 ms to 75 ms | BI:M | Yes [] N/A [] |
| RF6 | 256B/257B to 64B/66B transcoder | 119.2.5.7 | rx_coded $j < 65:0$, $j=0$ to 3 constructed per 119.2.5.7 | M | Yes [] |

118.5.4.3 64B/66B coding rules

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-------------------------|--|--------|---------|
| C1 | Encoder (and ENCODE function) implements the code as specified | 119.2.3, 119.2.6.2.3 | | M | Yes [] |
| C2 | Decoder (and DECODE function) implements the code as specified | 119.2.3, 119.2.6.2.3 | | M | Yes [] |
| C3 | Only valid block types are transmitted | 119.2.3.2 | | M | Yes [] |
| C4 | Invalid block types are treated as an error | 119.2.3.2 | | M | Yes [] |
| C5 | Only valid control characters are transmitted | 119.2.3.3 | | M | Yes [] |
| C6 | Invalid control characters are treated as an error | 119.2.3.3 | | M | Yes [] |
| C7 | Idles do not interrupt data | 119.2.3.5 | | M | Yes [] |
| C8 | IDLE control code insertion and deletion | 119.2.3.5 | Insertion or Deletion in groups of 8 /I/s | M | Yes [] |
| C9 | Sequence ordered set deletion | 119.2.3.8 | Only one whole ordered set of two consecutive sequence ordered sets may be deleted | M | Yes [] |

118.5.4.4 Scrambler and descrambler

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|-------------|-----------|-----------------------------------|--------|---------|
| S1 | Scrambler | 119.2.4.3 | Performs as shown in Figure 49–8 | M | Yes [] |
| S2 | Descrambler | 119.2.5.6 | Performs as shown in Figure 49–10 | M | Yes [] |

118.5.4.5 Alignment markers

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|----------------------------|-----------|---|--------|--------------------|
| AM1 | Alignment marker insertion | 119.2.4.4 | Alignment markers are inserted periodically as in 119.2.4.4 | M | Yes [] |
| AM2 | Alignment marker form | 119.2.4.4 | Alignment markers are formed as described in 119.2.4.4 | M | Yes [] |
| AM3 | Lane mapping | 119.2.4.4 | PCS lane number is captured | MD:M | Yes [] N/A [] |
| AM4 | Alignment marker removal | 119.2.5.5 | Alignment markers are removed prior to descrambling as described in 119.2.5.5 | M | Yes [] |

118.5.5 Test-pattern modes

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---------------|--------|---------|
| JT1 | Scrambled idle transmit test-pattern generator is implemented | 119.2.4.9 | | M | Yes [] |

118.5.6 Bit order

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--------------------|-----------|---|--------|---------|
| B1 | Transmit bit order | 119.2.4.8 | Placement of bits into the PCS lanes as shown in Figure 119–10 or Figure 119–11 | M | Yes [] |

118.5.7 Management

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|---------------------------------|----------------|--------------------|
| M1 | Alternate access to XS management objects is provided | 119.3 | | O | Yes [] No [] |
| M2 | Mapping of MDIO control bits and MDIO status bits for PHY 200GXS or PHY 400GXS | 118.4 | See Table 118–1 and Table 118–2 | MD*PHY XS:M | Yes [] N/A [] |
| M3 | Mapping of MDIO control bits and MDIO status bits for DTE 200GXS or DTE 400GXS | 118.4 | See Table 118–3 and Table 118–4 | MD*DTEX S:M | Yes [] N/A [] |

118.5.7.1 State diagrams

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-------------|--|--------------|--------------------|
| SM1 | Alignment marker lock | 119.2.6 | Implements 8 alignment marker lock processes as depicted in Figure 119–12 | 200GXS: M | Yes [] N/A [] |
| SM2 | Alignment marker lock | 119.2.6 | Implements 16 alignment marker lock processes as depicted in Figure 119–12 | 400GXS: M | Yes [] N/A [] |
| SM3 | The SLIP function evaluates all possible block positions | 119.2.6.2.3 | | M | Yes [] |
| SM4 | PCS synchronization state diagram | 119.2.6 | Meets the requirements of Figure 119–13 | M | Yes [] |
| SM5 | Transmit process | 119.2.6 | Meets the requirements of Figure 119–14 | M | Yes [] |
| SM6 | Receive process | 119.2.6 | Meets the requirements of Figure 119–15 | M | Yes [] |

118.5.7.2 Loopback

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---------------|--------|---------|
| L1 | Supports loopback | 119.4 | | M | Yes [] |
| L2 | When in loopback, transmits what it receives from the 200GMII/400GMII | 119.4 | | M | Yes [] |
| L3 | When in loopback, ignore all data presented by the PMA sublayer | 119.4 | | M | Yes [] |

118.5.7.3 Delay constraints

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|----------------------|-----------|---|--------------|--------------------|
| TIM1 | PCS delay constraint | 119.5 | No more than 160 256 BT for sum of transmit and receive path delays for 200GBASE-R. | 200GXS :M | Yes [] N/A [] |
| TIM2 | PCS delay constraint | 119.5 | No more than 320 000 BT for sum of transmit and receive path delays for 400GBASE-R. | 400GXS :M | Yes [] N/A [] |

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd.10:2019

119. Physical Coding Sublayer (PCS) for 64B/66B, type 200GBASE-R and 400GBASE-R

119.1 Overview

119.1.1 Scope

This clause specifies the Physical Coding Sublayer (PCS) that is common to two families of Physical Layer implementation known as 200GBASE-R and 400GBASE-R. The 200GBASE-R PCS is a sublayer of the 200 Gb/s PHYs listed in Table 116–1. The 400GBASE-R PCS is a sublayer of the 400 Gb/s PHYs listed in Table 116–2. The terms 200GBASE-R and 400GBASE-R are used when referring generally to Physical Layers using the PCS defined in this clause. Both 200GBASE-R and 400GBASE-R are based on a 64B/66B code. The 64B/66B code supports transmission of data and control characters. The 64B/66B code is then transcoded to 256B/257B encoding to reduce the overhead and make room for Forward Error Correction (FEC). The 256B/257B encoded data is then FEC encoded before being transmitted. Data distribution is introduced to support multiple lanes in the Physical Layer. Part of the distribution includes the periodic insertion of an alignment marker, which allows the receive PCS to align data from multiple lanes.

119.1.2 Relationship of 200GBASE-R and 400GBASE-R to other standards

Figure 119–1 depicts the relationship of the 200GBASE-R and 400GBASE-R sublayers (shown shaded), the Ethernet MAC and Reconciliation Sublayers, and the higher layers.

This clause borrows heavily from [Clause 82](#) and [Clause 91](#). 64B/66B encoding is reused with the addition of transcoding and PCS FEC.

119.1.3 Physical Coding Sublayer (PCS)

The PCS service interface is the Media Independent Interface (200GMII/400GMII), which is defined in Clause 117. The 200GMII provides a uniform interface to the Reconciliation Sublayer for all 200 Gb/s PHY implementations. The 400GMII provides a uniform interface to the Reconciliation Sublayer for all 400 Gb/s PHY implementations.

The 200GBASE-R and 400GBASE-R PCSs provide all services required by the 200GMII/400GMII, including the following:

- a) Encoding (decoding) of eight 200GMII/400GMII data octets to (from) 66-bit blocks (64B/66B).
- b) Transcoding from 66-bit blocks to (from) 257-bit blocks.
- c) Reed-Solomon encoding (decoding) the 257-bit blocks.
- d) Transferring encoded data to (from) the PMA.
- e) Compensation for any rate differences caused by the insertion or deletion of alignment markers or due to any rate difference between the 200GMII/400GMII and PMA through the insertion or deletion of idle control characters.
- f) Determining when a functional link has been established and informing the management entity via the MDIO when the PHY is ready for use.

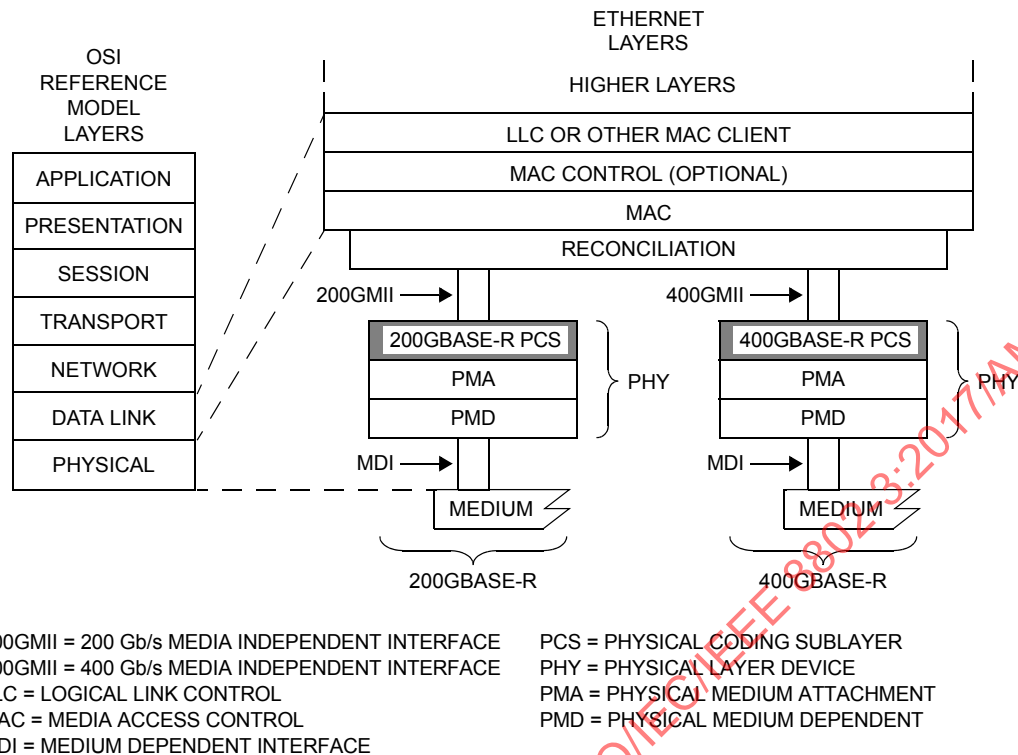


Figure 119-1—200GBASE-R and 400GBASE-R PCS relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

119.1.4 Inter-sublayer interfaces

The upper interface of the PCS may connect to the Reconciliation Sublayer through the 200GMII/400GMII. The lower interface of the PCS connects to the PMA sublayer to support a PMD. The 200GBASE-R PCS has a nominal rate at the PMA service interface of 26.5625 Gtransfers/s on each of 8 PCS lanes, which provides capacity for the MAC data rate of 200 Gb/s. The 400GBASE-R PCS has a nominal rate at the PMA service interface of 26.5625 Gtransfers/s on each of 16 PCS lanes, which provides capacity for the MAC data rate of 400 Gb/s.

It is important to note that, while this specification defines interfaces in terms of bits, octets, and frames, implementations may choose other data-path widths for implementation convenience.

119.1.4.1 PCS service interface (200GMII/400GMII)

The PCS service interface allows the 200GBASE-R or 400GBASE-R PCS to transfer information to and from a PCS client. The PCS client is the Reconciliation Sublayer. The PCS Service Interface is precisely defined as the Media Independent Interface (200GMII/400GMII) in Clause 117.

119.1.4.2 Physical Medium Attachment (PMA) service interface

The PMA service interface for the PCS is described in an abstract manner and does not imply any particular implementation. The PMA Service Interface supports the exchange of encoded data between the PCS and PMA sublayer. The PMA service interface is defined in 120.3 and is an instance of the inter-sublayer service interface definition in 116.3.

119.1.5 Functional block diagram

Figure 119–2 provides a functional block diagram of the 200GBASE-R and 400GBASE-R PCSs.

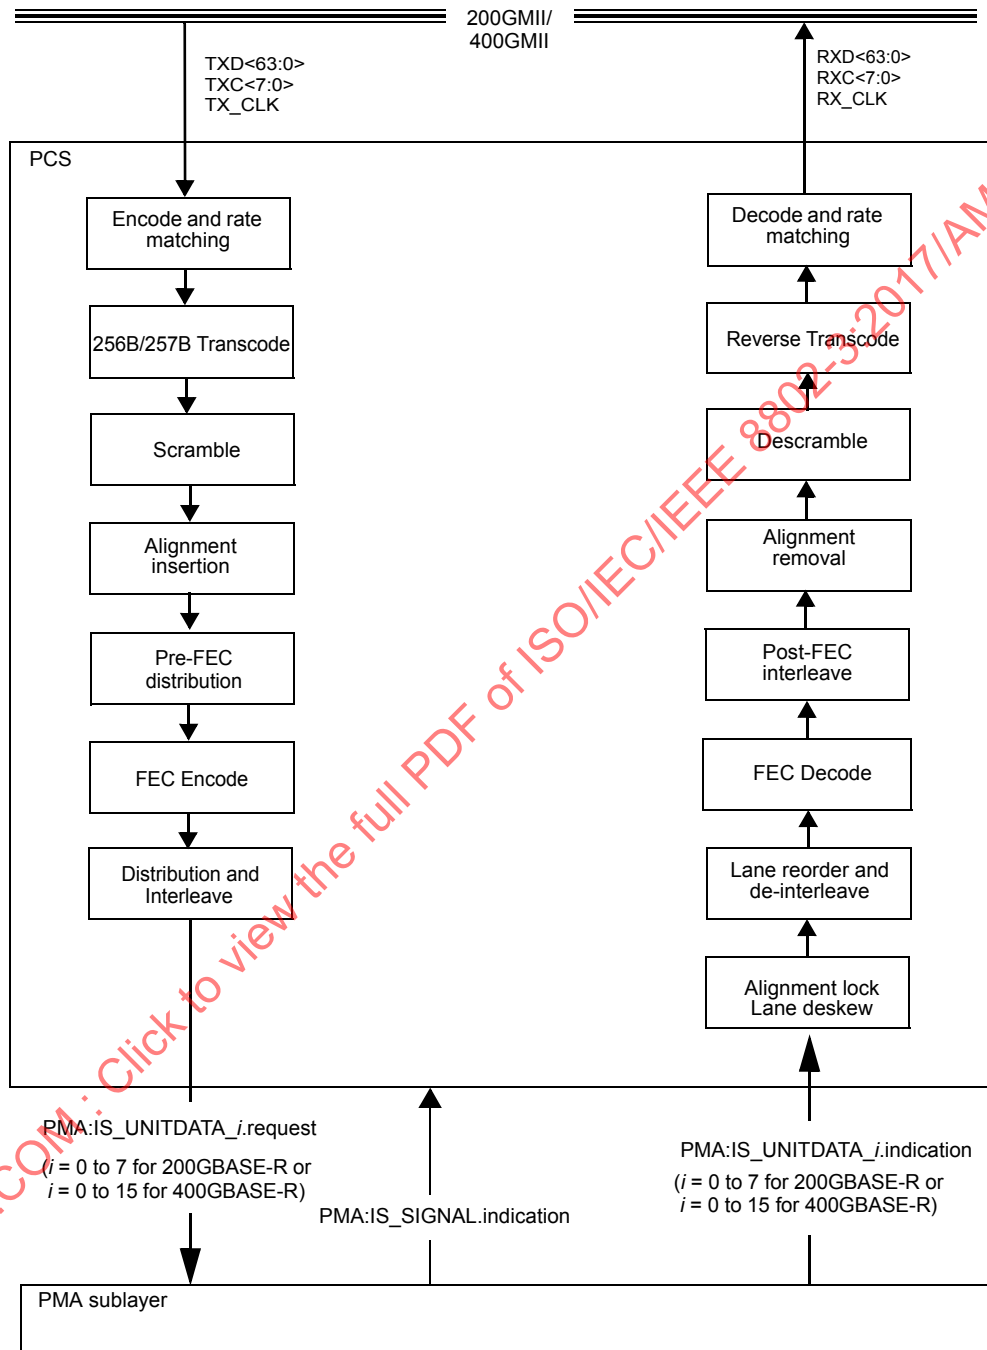


Figure 119–2—Functional block diagram

119.2 Physical Coding Sublayer (PCS)

119.2.1 Functions within the PCS

The 200GBASE-R and 400GBASE-R PCS are composed of the PCS Transmit and PCS Receive processes. The PCS shields the Reconciliation Sublayer (and MAC) from the specific nature of the underlying channel. The PCS transmit channel can operate in normal mode or test-pattern mode.

When communicating with the 200GMII/400GMII, the PCS uses an eight octet-wide, synchronous data path, with packet delineation being provided by transmit control signals ($\text{TXC}_{<n>} = 1$) and receive control signals ($\text{RXC}_{<n>} = 1$). When communicating with the PMA, the 200GBASE-R PCS uses 8 encoded bit streams (also known as PCS lanes) and the 400GBASE-R PCS uses 16 encoded bit streams. Per direction (RX or TX), the PCS lanes originate from a common clock but may vary in phase and skew dynamically. The PMA sublayer operates independently of block and packet boundaries. The PCS provides the functions necessary to map packets between the 200GMII/400GMII format and the PMA service interface format.

When the transmit channel is in normal mode, the PCS Transmit process continuously generates 66-bit blocks based upon the $\text{TXD}_{<63:0>}$ and $\text{TXC}_{<7:0>}$ signals on the 200GMII/400GMII. The 66-bit blocks are transcoded to 257-bit blocks to reduce the line coding overhead. The transcoded blocks are then scrambled to control clock content and baseline wander. Alignment markers are periodically added to the data stream. The data stream is distributed to two 5140-bit blocks and then FEC encoded to control errors. The two FEC codewords are then interleaved before data is distributed to individual PCS lanes.

Transmit data-units are sent to the service interface via the PMA:IS_UNITDATA_i .request primitive.

When the transmit channel is in test-pattern mode, a test pattern is packed into the transmit data-units that are sent to the PMA service interface via the PMA:IS_UNITDATA_i .request primitive.

The PCS Synchronization process continuously monitors $\text{PMA:IS_SIGNAL.indication}(\text{SIGNAL_OK})$. When SIGNAL_OK indicates OK, then the PCS synchronization process accepts data-units via the PMA:IS_UNITDATA_i .indication primitive. It attains alignment marker lock based on the common marker (CM) portion that is periodically transmitted on every PCS lane. After alignment markers are found on all PCS lanes, the individual PCS lanes are identified using the unique marker portion (UM) and then re-ordered and deskewed. Note that a particular transmit PCS lane can be received on any receive lane of the service interface due to the skew and multiplexing that occurs in the path.

The PCS deskew process deskews, aligns and reorders the individual PCS lanes, forms a single stream, and sets the align_status flag to indicate whether the PCS has obtained alignment. The PCS then de-interleaves and processes the FEC codewords. Next the PCS re-interleaves the corrected FEC codewords on a 10-bit basis to form a single stream. The PCS then removes alignment markers, descrambles the data, transcodes the data back to 64B/66B and then decodes the 64B/66B encoded data.

The PCS shall provide transmit test-pattern mode for the scrambled idle pattern (see 119.2.4.9).

119.2.2 Use of blocks

The PCS maps 200GMII/400GMII signals into 66-bit blocks, and vice versa, using a 64B/66B coding scheme. The PCS functions ENCODE and DECODE generate, manipulate, and interpret blocks as defined in 119.2.3.

119.2.3 64B/66B code

The PCS uses 64B/66B coding to support transmission of control and data characters. The 64B/66B codestream is then transcoded into a 256B/257B stream and FEC bits are added in this PCS before transmission.

119.2.3.1 Notation conventions

For values shown as binary, the leftmost bit is the first transmitted bit.

64B/66B encodes 8 data octets or control characters into a block. Blocks containing control characters also contain a block type field. Data octets are labeled D_0 to D_7 . The control characters /I/ and /E/ are labeled C_0 to C_7 . The control characters, /Q/ and /Fsig/, for ordered sets are labeled as O_0 since they are only valid on the first octet of the 200GMII/400GMII. The control character for start is labeled as S_0 for the same reason. The control character for terminate is labeled as T_0 to T_7 . The four trailing zero data octets in ordered sets are labeled as Z_4 to Z_7 .

One 200GMII/400GMII transfer provides eight characters that are encoded into one 66-bit transmission block. The subscript in the above labels indicates the position of the character in the eight characters from the 200GMII/400GMII transfer.

Contents of block type fields, data octets, and control characters are shown as hexadecimal values. The LSB of the hexadecimal value represents the first transmitted bit. For instance, the block type field 0x1E is sent from left to right as 01111000. The bits of a transmitted or received block are labeled $TxB<65:0>$ and $RxB<65:0>$, respectively, where $TxB<0>$ and $RxB<0>$ represent the first transmitted bit. The value of the sync header is shown as a binary value. Binary values are shown with the first transmitted bit (the LSB) on the left.

119.2.3.2 64B/66B block structure

Blocks consist of 66 bits. The first two bits of a block are the synchronization header (sync header). Blocks are either data blocks or control blocks. The sync header is 01 for data blocks and 10 for control blocks. The remainder of the block contains the payload.

Data blocks contain eight data characters. Control blocks begin with an 8-bit block type field that indicates the format of the remainder of the block. For control blocks containing a Start, Terminate, or ordered set, that character is implied by the block type field. Other control characters are encoded in a 7-bit control code. Each control block encodes eight characters.

The format of the blocks is as shown in Figure 82-5. In the figure, the column labeled Input Data shows, in abbreviated form, the eight characters used to create the 66-bit block. These characters are either data characters or control characters and, when transferred across the 200GMII/400GMII, the corresponding TXC or RXC bit is set accordingly. Within the Input Data column, D_0 through D_7 are data octets and are transferred with the corresponding TXC or RXC bit set to zero. All other characters are control characters and are transferred with the corresponding TXC or RXC bit set to one. The single bit fields (thin rectangles with no label in the figure) are sent as zero and ignored upon receipt.

All values of block type field not listed in Figure 82-5 are invalid; they shall not be transmitted and shall be considered an error if received.

119.2.3.3 Control codes

The same set of control characters are supported by the 200GMII/400GMII and the PCS. The representations of the control characters are the control codes. The 200GMII/400GMII encodes a control

character into an octet (an eight bit value). The PCS encodes the start and terminate control characters implicitly by the block type field. The PCS encodes the ordered set control codes using the block type field. The PCS encodes each of the other control characters into a 7-bit control code.

The control characters and their mappings to 200GBASE-R/400GBASE-R control codes and 200GMII/400GMII control codes are specified in Table 82-1. All 200GMII/400GMII and 200GBASE-R/400GBASE-R control code values that do not appear in the table shall not be transmitted and shall be considered an error if received. The ability to transmit or receive Low Power Idle (LPI) is required for PHYs that support EEE (see Clause 78). If EEE has not been negotiated, LPI shall not be transmitted and shall be treated as an error if received.

119.2.3.4 Valid and invalid blocks

The valid and invalid blocks are identical to those in 82.2.3.5 with the exception that it is a 200GMII/400GMII data stream instead of XLGMII/CGMII.

119.2.3.5 Idle (/I/)

Idle and LPI control characters are identical to those in 82.2.3.6.

119.2.3.6 Start (/S/)

Start control characters are identical to those in 82.2.3.7.

119.2.3.7 Terminate (/T/)

The terminate control characteristics are identical to those in 82.2.3.8.

119.2.3.8 Ordered set (/O/)

Ordered sets are specified identically as in 82.2.3.9.

119.2.3.9 Error (/E/)

In both the 64B/66B encoder and decoder, the /E/ is generated whenever an /E/ is detected. The /E/ is also generated when invalid blocks are detected. The /E/ allows the PCS to propagate detected errors. See R_TYPE and T_TYPE function definitions in 119.2.6.2.3 for further information.

119.2.4 Transmit

119.2.4.1 Encode and rate matching

The PCS generates 66-bit blocks based upon the TXD<63:0> and TXC<7:0> signals received from the 200GMII/400GMII. One 200GMII/400GMII data transfer is encoded into one 66-bit block. If the transmit PCS spans multiple clock domains, it may also perform clock rate compensation via the deletion of idle control characters or sequence ordered sets or the insertion of idle (or LPI) control characters.

Idle (or LPI) control characters or sequence ordered sets are removed, if necessary, to accommodate the insertion of the alignment markers. See 119.2.3.5 and 119.2.3.8 for the deletion and insertion rules, and 119.2.4.4 for more details on alignment markers.

The transmit PCS generates blocks as specified in the transmit state diagram as shown in Figure 119-14. The contents of each block are contained in a vector tx_coded<65:0>, which is passed to the transcoder. tx_coded<1:0> contains the sync header and the remainder of the bits contain the block payload.

Note—The stream of 66-bit blocks generated by this process, together with the FEC_degraded_SER and rx_local_degraded bits should be used as the reference signal for mapping to OTN.

119.2.4.2 64B/66B to 256B/257B transcoder

The transcoder constructs a 257-bit block, tx_xcoded<256:0>, from a group of four 66-bit blocks, tx_coded_j<65:0> where $j=0$ to 3. For each group of four 66-bit blocks, $j=3$ corresponds to the most recently received block. Bit 0 in each 66-bit block is the first bit received and corresponds to the first bit of the synchronization header.

Note—This transcoder differs from that described in 91.5.2.5, there is no scrambling of the first 5 bits since this clause scrambles the complete 257-bit blocks after transcoding.

If for all $j=0$ to 3, tx_coded_j<0>=0 and tx_coded_j<1>=1, tx_xcoded<256:0> shall be constructed as follows:

- a) tx_xcoded<0> = 1
- b) tx_xcoded<(64j+64):(64j+1)> = tx_coded_j<65:2> for $j=0$ to 3

If for all $j=0$ to 3, tx_coded_j<0> ≠ tx_coded_j<1> (valid synchronization header) and for any $j=0$ to 3, tx_coded_j<0>=1 and tx_coded_j<1>=0, tx_xcoded<256:0> shall be constructed as follows:

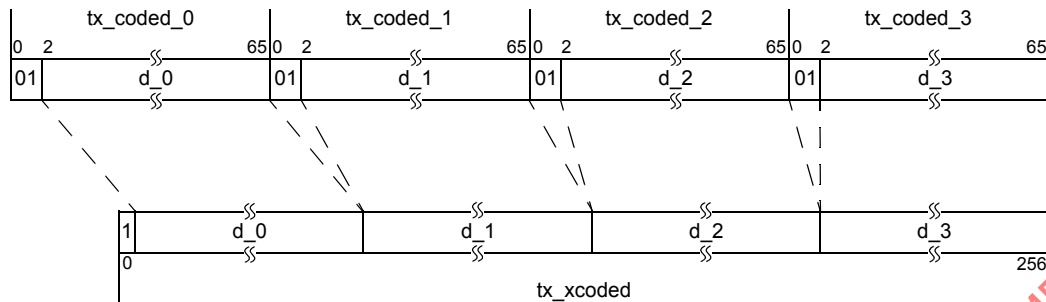
- a1) tx_xcoded<0> = 0
- b1) tx_xcoded<j+1> = tx_coded_j<1> for $j=0$ to 3
- c1) Let c be the smallest value of j such that tx_coded_c<0>=1. In other words, tx_coded_c is the first 66-bit control block that was received in the current group of four blocks.
- d1) Let tx_payloads<(64j+63):64j> = tx_coded_j<65:2> for $j=0$ to 3
- e1) Omit tx_coded_c<9:6>, which is the second nibble (based on transmission order) of the block type field for tx_coded_c, from tx_xcoded per the following expressions.
tx_xcoded<(64c+8):5> = tx_payloads<(64c+3):0>
tx_xcoded<256:(64c+9)> = tx_payloads<255:(64c+8)>

If for any $j=0$ to 3, tx_coded_j<0> = tx_coded_j<1> (invalid synchronization header), tx_xcoded<256:0> shall be constructed as follows:

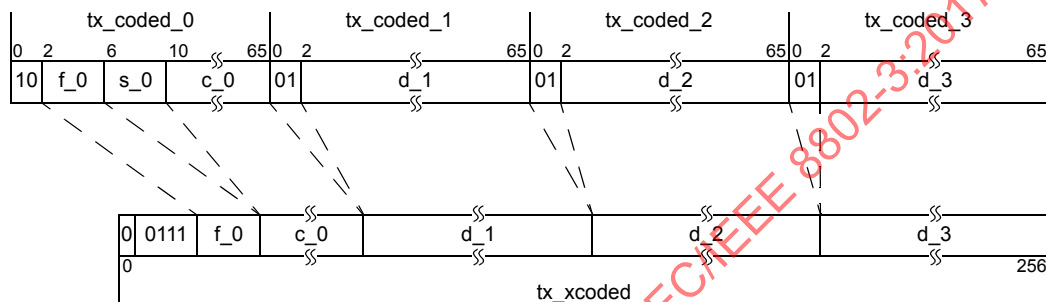
- a2) tx_xcoded<0> = 0
- b2) tx_xcoded<j+1> = 1 for $j=0$ to 3
- c2) Let tx_payloads<(64j+63):64j> = tx_coded_j<65:2> for $j=0$ to 3
- d2) Omit the second nibble (based on transmission order) of tx_coded_0 per the following expressions.
tx_xcoded<8:5> = tx_payloads<3:0>
tx_xcoded<256:9> = tx_payloads<255:8>

Several examples of the construction of tx_xcoded<256:0> are shown in Figure 119–3. In Figure 119–3, d_j indicates the jth 66-bit block contains only data octets, c_j indicates the jth 66-bit block contains one or more control characters, f_j denotes the first nibble of the block type field for 66-bit block j, and s_j denotes the second nibble of the block type field for 66-bit block j.

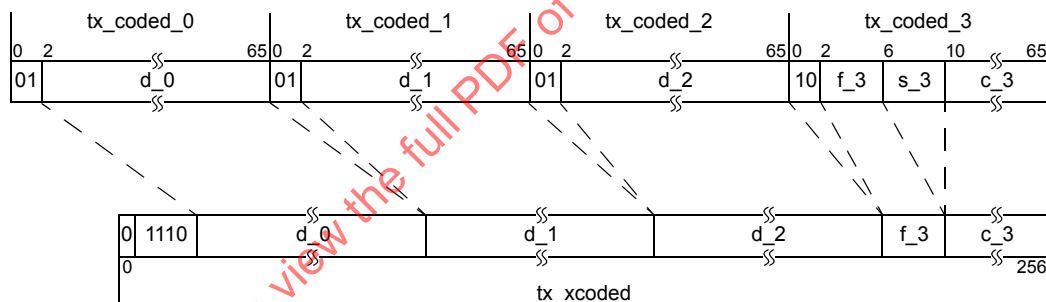
For each 257-bit block, bit 0 shall be the first bit transmitted.



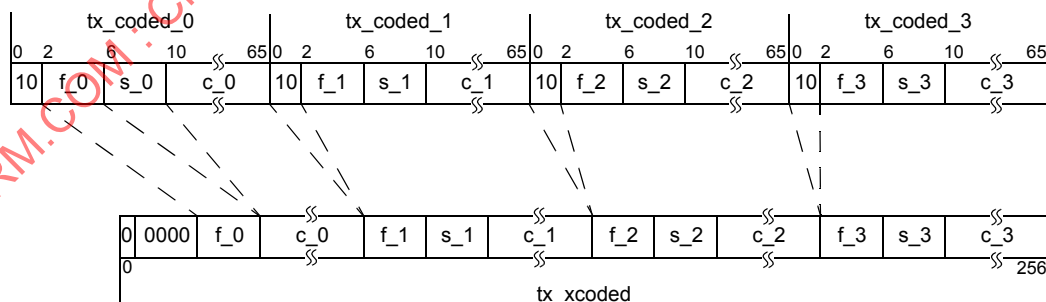
Example 1: All data blocks



Example 2: Control block followed by three data blocks



Example 3: Three data blocks followed by a control block



Example 4: All control blocks

Figure 119–3—Examples of the construction of tx_xcoded

119.2.4.3 Scrambler

The transcoded 257-bit block, $tx_xcoded<256:0>$, shall be scrambled with a self-synchronizing scrambler to generate $tx_scrambled<256:0>$. The scrambler polynomial is identical to that in [Clause 49](#), see [Equation \(49-1\)](#) for the definition of the polynomial.

119.2.4.4 Alignment marker mapping and insertion

In order to support deskew and reordering of the individual PCS lanes at the receive PCS, alignment markers corresponding to PCS lanes are periodically inserted after being processed by the alignment marker mapping function. The alignment marker mapping function compensates for the operation of the symbol distribution function and rearranges the alignment marker bits so that they appear on the PCS lanes intact and in the desired sequence. This preserves the properties of the alignment markers (e.g. DC balance, transition density) and provides a deterministic pattern for the purpose of synchronization.

For the 200GBASE-R PCS, an alignment marker group is composed of the alignment markers for all 8 PCS lanes plus an additional 65-bit pad and a 3-bit status field to yield the equivalent of four 257-bit blocks. For the 400GBASE-R PCS, an alignment marker group is composed of the alignment markers for all 16 PCS lanes plus an additional 133-bit pad and a 3-bit status field to yield the equivalent of eight 257-bit blocks. The alignment marker group is aligned to the beginning of two FEC messages, and interrupts any data transfer that is already in progress. The pad bits at the end of the alignment marker group shall be set to a free running PRBS9 pattern, defined by the polynomial $x^9 + x^5 + 1$. The initial value of the PRBS9 pattern generator may be any pattern other than all zeros.

Room for the alignment marker group is created by the transmit PCS (see 119.2.4.1). Special properties of the alignment marker group are that it is not scrambled, does not conform to the encoding rules as outlined in [Figure 82-5](#) and is not transcoded. This is possible because the alignment marker group is added after encoding, transcoding, and scrambling, and removed before descrambling, transcoding, and 64B/66B decoding. The alignment marker group is not scrambled, which allows the receiver to directly search for and find the individual alignment markers, deskew the PCS lanes, and reassemble the aggregate stream before FEC decoding is performed. The alignment markers are formed from a known pattern that is defined to be balanced and with many transitions and therefore scrambling is not necessary.

The format of each PCS lane's alignment marker is shown in Figure 119–4. There is a portion that is common across all alignment markers (designated as CM_0 to CM_5), a unique portion per PCS lane (designated as UM_0 to UM_5), and finally a unique pad per PCS lane (designated as UP_0 to UP_2). Common synchronization logic independent of the received PCS lane number can be used with the common portion of the alignment marker. The unique pad (UP_0 to UP_2) within the alignment markers and the PRBS9 pad at the end of the alignment marker group are ignored on receive.

The content of the alignment markers shall be as shown in Table 119–1 for the 200GBASE-R PCS and as shown in Table 119–2 for the 400GBASE-R PCS. The contents depend on the PCS lane number and the octet number, with CM_0 through CM_5 being identical across all alignment markers to allow for common synchronization across lanes. The format shown in Table 119–1 defines how the alignment markers appear on a given PCS lane. In the FEC codewords, they appear in a permuted format due to the codeword interleaving that occurs before FEC codewords are distributed to PCS lanes.

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.

119.2.4.4.1 AM creation for the 200GBASE-R PCS

For the 200GBASE-R PCS, the alignment marker mapping function creates a set of 8 alignment markers, and in combination with an additional 65-bit PRBS9 pad and a 3-bit status field, the PCS generates an alignment marker group. Let $am_x<119:0>$ be the alignment marker for PCS lane x , $x=0$ to 7, where bit 0 is the first bit transmitted. The alignment markers shall be mapped to $am_mapped<959:0>$ in a manner that yields the same result as the following process.

For $x=0$ to 7, $am_x<119:0>$ is constructed as follows:

$am_x<119:0>$ is set to $CM_0, CM_1, CM_2, UP_0, CM_3, CM_4, CM_5, UP_1, UM_0, UM_1, UM_2, UP_2, UM_3, UM_4$ and UM_5 , as shown in Figure 119–4 (bits 119:0) using the values in Table 119–1 for PCS lane number x .




Figure 119–4—Alignment marker format

Table 119–1—200GBASE-R alignment marker encodings

| PCS lane number | Encoding ^a { $CM_0, CM_1, CM_2, UP_0, CM_3, CM_4, CM_5, UP_1, UM_0, UM_1, UM_2, UP_2, UM_3, UM_4, UM_5$ } |
|-----------------|---|
| 0 | 0x9A, 0x4A, 0x26, 0x05, 0x65, 0xB5, 0xD9, 0xD6, 0xB3, 0xC0, 0x8C, 0x29, 0x4C, 0x3F, 0x73 |
| 1 | 0x9A, 0x4A, 0x26, 0x04, 0x65, 0xB5, 0xD9, 0x67, 0x5A, 0xDE, 0x7E, 0x98, 0xA5, 0x21, 0x81 |
| 2 | 0x9A, 0x4A, 0x26, 0x46, 0x65, 0xB5, 0xD9, 0xFE, 0x3E, 0xF3, 0x56, 0x01, 0xC1, 0x0C, 0xA9 |
| 3 | 0x9A, 0x4A, 0x26, 0x5A, 0x65, 0xB5, 0xD9, 0x84, 0x86, 0x80, 0xD0, 0x7B, 0x79, 0x7F, 0x2F |
| 4 | 0x9A, 0x4A, 0x26, 0xE1, 0x65, 0xB5, 0xD9, 0x19, 0x2A, 0x51, 0xF2, 0xE6, 0xD5, 0xAE, 0x0D |
| 5 | 0x9A, 0x4A, 0x26, 0xF2, 0x65, 0xB5, 0xD9, 0x4E, 0x12, 0x4F, 0xD1, 0xB1, 0xED, 0xB0, 0x2E |
| 6 | 0x9A, 0x4A, 0x26, 0x3D, 0x65, 0xB5, 0xD9, 0xEE, 0x42, 0x9C, 0xA1, 0x11, 0xBD, 0x63, 0x5E |
| 7 | 0x9A, 0x4A, 0x26, 0x22, 0x65, 0xB5, 0xD9, 0x32, 0xD6, 0x76, 0x5B, 0xCD, 0x29, 0x89, 0xA4 |

^a Each octet is transmitted LSB to MSB.

As an example, the variable am_0 is sent as (left most bit sent first, showing the first 32 bits transmitted of am_0):

01011001 01010010 01100100 10100000

The variable am_mapped is then derived from 10-bit interleaving the group of 8 alignment markers am_x per the following procedure:

For all $k=0$ to 11

For all $j=0$ to 3

if even(k)

$\text{am_mapped}\langle 80k+20j+9:80k+20j \rangle = \text{am_}\{2j\}\langle 10k+9:10k \rangle$

$\text{am_mapped}\langle 80k+20j+19:80k+20j+10 \rangle = \text{am_}\{2j+1\}\langle 10k+9:10k \rangle$

else

$\text{am_mapped}\langle 80k+20j+9:80k+20j \rangle = \text{am_}\{2j+1\}\langle 10k+9:10k \rangle$

$\text{am_mapped}\langle 80k+20j+19:80k+20j+10 \rangle = \text{am_}\{2j\}\langle 10k+9:10k \rangle$

The additional 65-bit pad is appended to variable am_mapped as follows:

$\text{am_mapped}\langle 1024:960 \rangle = \text{PRBS9}\langle 64:0 \rangle$

In this expression, $\text{PRBS9}\langle 0 \rangle$ is the first PRBS9 bit output of the 65-bit pad.

The 3-bit transmit alignment marker status field is then appended to the variable am_mapped as follows:

$\text{am_mapped}\langle 1027:1025 \rangle = \text{tx_am_sf}\langle 2:0 \rangle$

Alignment marker mapping is shown in Figure 119–5.

| PCS lane, i | am_mapped 10-bit symbol index, k | | | | | | | | | | | | |
|------------------|------------------------------------|---|---|---|---|---|---|---|---|---|----|----|----|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
| 0 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 1 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 2 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 3 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 4 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 5 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 6 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 7 | B | A | B | A | B | A | B | A | B | A | B | A | B |




 = 65-bit pad
  = 3-bit status field
  = Resumption of 257-bit blocks
 A = from FEC codeword A B = from FEC codeword B

Figure 119–5—200GBASE-R alignment marker mapping to PCS lanes

The alignment marker group $\text{am_mapped}\langle 1027:0 \rangle$ shall be inserted so it appears in the output stream every $81\,920 \times 257$ -bit blocks. The variable $\text{tx_scrambled_am}\langle 10279:0 \rangle$ is constructed in one of two ways. Let the set of vectors $\text{tx_scrambled_i}\langle 256:0 \rangle$ represent consecutive values of $\text{tx_scrambled}\langle 256:0 \rangle$.

For a 10280-bit block with an alignment marker group inserted:

$\text{tx_scrambled_am}\langle 1027:0 \rangle = \text{am_mapped}\langle 1027:0 \rangle$

For all $i=0$ to 35

$\text{tx_scrambled_am}\langle 257i+1284:257i+1028 \rangle = \text{tx_scrambled_i}\langle 256:0 \rangle$

For a 10280-bit block without an alignment marker group:

For all $i=0$ to 39

$\text{tx_scrambled_am}\langle 257i+256:257i \rangle = \text{tx_scrambled_i}\langle 256:0 \rangle$

Alignment marker repetition rate is shown in Figure 119–6.




Figure 119–6—200GBASE-R alignment marker insertion period

119.2.4.4.2 AM creation for the 400GBASE-R PCS

For the 400GBASE-R PCS, the alignment marker mapping function creates a set of 16 alignment markers, and in combination with an additional 133-bit PRBS9 pad and a 3-bit status field, the PCS generates an alignment marker group.

Let $am_x<119:0>$ be the alignment marker for PCS lane x , $x=0$ to 15, where bit 0 is the first bit transmitted. The alignment markers shall be mapped to $am_mapped<1919:0>$ in a manner that yields the same result as the following process.

For $x=0$ to 15, $am_x<119:0>$ is constructed as follows:

$am_x<119:0>$ is set to $CM_0, CM_1, CM_2, UP_0, CM_3, CM_4, CM_5, UP_1, UM_0, UM_1, UM_2, UP_2, UM_3, UM_4$ and UM_5 , as shown in Figure 119–4 (bits 119:0) using the values in Table 119–2 for PCS lane number x .

Table 119–2—400GBASE-R alignment marker encodings

| PCS lane number | Encoding ^a { $CM_0, CM_1, CM_2, UP_0, CM_3, CM_4, CM_5, UP_1, UM_0, UM_1, UM_2, UP_2, UM_3, UM_4, UM_5$ } |
|-----------------|---|
| 0 | 0x9A, 0x4A, 0x26, 0xB6, 0x65, 0xB5, 0xD9, 0xD9, 0x01, 0x71, 0xF3, 0x26, 0xFE, 0x8E, 0x0C |
| 1 | 0x9A, 0x4A, 0x26, 0x04, 0x65, 0xB5, 0xD9, 0x67, 0x5A, 0xDE, 0x7E, 0x98, 0xA5, 0x21, 0x81 |
| 2 | 0x9A, 0x4A, 0x26, 0x46, 0x65, 0xB5, 0xD9, 0xFE, 0x3E, 0xF3, 0x56, 0x01, 0xC1, 0x0C, 0xA9 |
| 3 | 0x9A, 0x4A, 0x26, 0x5A, 0x65, 0xB5, 0xD9, 0x84, 0x86, 0x80, 0xD0, 0x7B, 0x79, 0x7F, 0x2F |
| 4 | 0x9A, 0x4A, 0x26, 0xE1, 0x65, 0xB5, 0xD9, 0x19, 0x2A, 0x51, 0xF2, 0xE6, 0xD5, 0xAE, 0x0D |
| 5 | 0x9A, 0x4A, 0x26, 0xF2, 0x65, 0xB5, 0xD9, 0x4E, 0x12, 0x4F, 0xD1, 0xB1, 0xED, 0xB0, 0x2E |
| 6 | 0x9A, 0x4A, 0x26, 0x3D, 0x65, 0xB5, 0xD9, 0xEE, 0x42, 0x9C, 0xA1, 0x11, 0xBD, 0x63, 0x5E |
| 7 | 0x9A, 0x4A, 0x26, 0x22, 0x65, 0xB5, 0xD9, 0x32, 0xD6, 0x76, 0x5B, 0xCD, 0x29, 0x89, 0xA4 |
| 8 | 0x9A, 0x4A, 0x26, 0x60, 0x65, 0xB5, 0xD9, 0x9F, 0xE1, 0x73, 0x75, 0x60, 0x1E, 0x8C, 0x8A |
| 9 | 0x9A, 0x4A, 0x26, 0x6B, 0x65, 0xB5, 0xD9, 0xA2, 0x71, 0xC4, 0x3C, 0x5D, 0x8E, 0x3B, 0xC3 |
| 10 | 0x9A, 0x4A, 0x26, 0xFA, 0x65, 0xB5, 0xD9, 0x04, 0x95, 0xEB, 0xD8, 0xFB, 0x6A, 0x14, 0x27 |
| 11 | 0x9A, 0x4A, 0x26, 0x6C, 0x65, 0xB5, 0xD9, 0x71, 0x22, 0x66, 0x38, 0x8E, 0xDD, 0x99, 0xC7 |
| 12 | 0x9A, 0x4A, 0x26, 0x18, 0x65, 0xB5, 0xD9, 0x5B, 0xA2, 0xF6, 0x95, 0xA4, 0x5D, 0x09, 0x6A |
| 13 | 0x9A, 0x4A, 0x26, 0x14, 0x65, 0xB5, 0xD9, 0xCC, 0x31, 0x97, 0xC3, 0x33, 0xCE, 0x68, 0x3C |
| 14 | 0x9A, 0x4A, 0x26, 0xD0, 0x65, 0xB5, 0xD9, 0xB1, 0xCA, 0xFB, 0xA6, 0x4E, 0x35, 0x04, 0x59 |
| 15 | 0x9A, 0x4A, 0x26, 0xB4, 0x65, 0xB5, 0xD9, 0x56, 0xA6, 0xBA, 0x79, 0xA9, 0x59, 0x45, 0x86 |

^a Each octet is transmitted LSB to MSB.

As an example, the variable am_0 is sent as (left most bit sent first, showing the first 32 bits transmitted of am_0):

01011001 01010010 01100100 01101101

The variable am_mapped is then derived from 10-bit interleaving the group of 16 alignment markers am_x per the following procedure:

For all $k=0$ to 11

For all $j=0$ to 7

if even(k)

$am_mapped<160k+20j+9:160k+20j> = am_{{2j}}<10k+9:10k>$

$am_mapped<160k+20j+19:160k+20j+10> = am_{{2j+1}}<10k+9:10k>$

else

$am_mapped<160k+20j+9:160k+20j> = am_{{2j+1}}<10k+9:10k>$

$am_mapped<160k+20j+19:160k+20j+10> = am_{{2j}}<10k+9:10k>$

The additional 133-bit pad is appended to variable am_mapped as follows:

$am_mapped<2052:1920> = PRBS9<132:0>$

In this expression, $PRBS9<0>$ is the first PRBS9 bit output of the 133-bit pad.

The 3-bit transmit alignment marker status field is then appended to the variable am_mapped as follows:

$am_mapped<2055:2053> = tx_am_sf<2:0>$

Alignment marker mapping is shown in Figure 119–7.

| PCS lane, i | am_mapped 10-bit symbol index, k | | | | | | | | | | | | |
|---------------|------------------------------------|---|---|---|---|---|---|---|---|---|----|----|----|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
| 0 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 1 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 2 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 3 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 4 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 5 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 6 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 7 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 8 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 9 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 10 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 11 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 12 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 13 | B | A | B | A | B | A | B | A | B | A | B | A | B |
| 14 | A | B | A | B | A | B | A | B | A | B | A | B | A |
| 15 | B | A | B | A | B | A | B | A | B | A | B | A | B |



= 133-bit pad



= 3-bit status field



= Resumption of 257-bit blocks

A = from FEC codeword A

B = from FEC codeword B

Figure 119–7—400GBASE-R alignment marker mapping to PCS lanes

The alignment marker group $\text{am_mapped}\langle 2055:0 \rangle$ shall be inserted so it appears in the output stream every $163\,840 \times 257$ -bit blocks. The variable $\text{tx_scrambled_am}\langle 10279:0 \rangle$ is constructed in one of two ways. Let the set of vectors $\text{tx_scrambled_i}\langle 256:0 \rangle$ represent consecutive values of $\text{tx_scrambled}\langle 256:0 \rangle$.

For a 10280-bit block with an alignment marker group inserted:

$\text{tx_scrambled_am}\langle 2055:0 \rangle = \text{am_mapped}\langle 2055:0 \rangle$

For all $i=0$ to 31

$\text{tx_scrambled_am}\langle 257i+2312:257i+2056 \rangle = \text{tx_scrambled_i}\langle 256:0 \rangle$

For a 10280-bit block without an alignment marker group:

For all $i=0$ to 39

$\text{tx_scrambled_am}\langle 257i+256:257i \rangle = \text{tx_scrambled_i}\langle 256:0 \rangle$

Alignment marker repetition rate is shown in Figure 119–8.




Figure 119–8—400GBASE-R alignment marker insertion period

119.2.4.5 Pre-FEC distribution

In order to improve error correction capability, each set of two consecutive Reed-Solomon FEC codewords is interleaved before being distributed to form the PCS lanes. To enable this interleaving, the Pre-FEC distribution function receives a 10280-bit block tx_scrambled_am , and shall perform a 10-bit symbol round robin distribution to form two 514-symbol FEC messages, m_A and m_B , which are subsequently each encoded by the RS FEC. The following describes the 10-bit round robin distribution process.

For all $i=0$ to 513

$m_A\langle 513-i \rangle = \text{tx_scrambled_am}\langle (20i+9):(20i) \rangle$

$m_B\langle 513-i \rangle = \text{tx_scrambled_am}\langle (20i+19):(20i+10) \rangle$

119.2.4.6 Reed-Solomon encoder

The PCS sublayer employs a Reed-Solomon code operating over the Galois Field $\text{GF}(2^{10})$ where the symbol size is 10 bits. The encoder processes k message symbols to generate $2t$ parity symbols, which are then appended to the message to produce a codeword of $n=k+2t$ symbols. For the purposes of this clause, a particular Reed-Solomon code is denoted $\text{RS}(n,k)$.

The PCS shall implement an $\text{RS}(544,514)$ based FEC encoder. The PCS distributes a group of 40×257 -bit blocks from tx_scrambled_am on a 10-bit round robin basis into two 5140-bit message blocks, m_A and m_B , as described in 119.2.4.5. These are then encoded using $\text{RS}(544,514)$ encoder into codeword A and codeword B, respectively. The $\text{RS}(544,514)$ code is based on the generating polynomial given by Equation (119–1).

$$g(x) = \prod_{j=0}^{2t-1} (x - \alpha^j) = g_{2t}x^{2t} + g_{2t-1}x^{2t-1} + \dots + g_1x + g_0 \quad (119-1)$$

In Equation (119-1), α is a primitive element of the finite field defined by the polynomial $x^{10} + x^3 + 1$.

Equation (119-2) defines the message polynomial $m(x)$ whose coefficients are the message symbols m_{k-1} to m_0 .

$$m(x) = m_{k-1}x^{n-1} + m_{k-2}x^{n-2} + \dots + m_1x^{2t+1} + m_0x^{2t} \quad (119-2)$$

Each message symbol m_i is the bit vector $(m_{i,9}, m_{i,8}, \dots, m_{i,1}, m_{i,0})$, which is identified with the element $m_{i,9}\alpha^9 + m_{i,8}\alpha^8 + \dots + m_{i,1}\alpha + m_{i,0}$ of the finite field. The message symbols are composed of the variable m_A for codeword A and m_B for codeword B (including a group of alignment markers when appropriate). The first symbol input to the encoder is m_{k-1} .

Equation (119-3) defines the parity polynomial $p(x)$ whose coefficients are the parity symbols p_{2t-1} to p_0 .

$$p(x) = p_{2t-1}x^{2t-1} + p_{2t-2}x^{2t-2} + \dots + p_1x + p_0 \quad (119-3)$$

The parity polynomial is the remainder from the division of $m(x)$ by $g(x)$. This may be computed using the shift register implementation illustrated in Figure 119-9. The outputs of the delay elements are initialized to zero prior to the computation of the parity for a given message. After the last message symbol, m_0 , is processed by the encoder, the outputs of the delay elements are the parity symbols for that message.

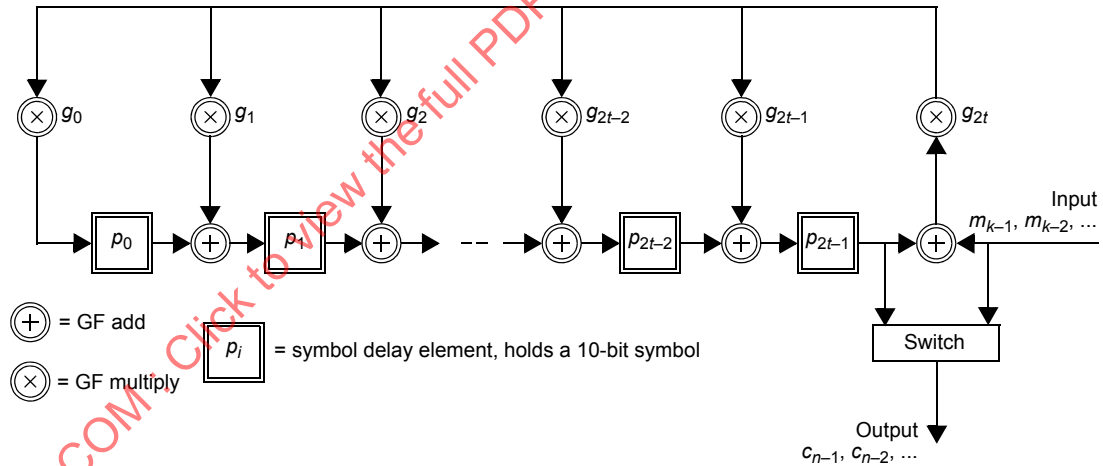


Figure 119-9—Reed-Solomon encoder functional model

The codeword polynomial $c(x)$ is then the sum of $m(x)$ and $p(x)$ where the coefficient of the highest power of x , $c_{n-1} = m_{k-1}$ is first and the coefficient of the lowest power of x , $c_0 = p_0$ is last.

The coefficients of the generator polynomial for the RS(544,514) code are presented in Table 119-3. Example codewords for the RS(544,514) code are provided in Annex 119A.

Table 119–3—Coefficients of the generator polynomial g_i (decimal)

| i | g_i | i | g_i | i | g_i |
|-----|-------|-----|-------|-----|-------|
| 0 | 523 | 11 | 883 | 22 | 565 |
| 1 | 834 | 12 | 503 | 23 | 108 |
| 2 | 128 | 13 | 942 | 24 | 1 |
| 3 | 158 | 14 | 385 | 25 | 552 |
| 4 | 185 | 15 | 495 | 26 | 230 |
| 5 | 127 | 16 | 720 | 27 | 187 |
| 6 | 392 | 17 | 94 | 28 | 552 |
| 7 | 193 | 18 | 132 | 29 | 575 |
| 8 | 610 | 19 | 593 | 30 | 1 |
| 9 | 788 | 20 | 249 | | |
| 10 | 361 | 21 | 282 | | |

119.2.4.7 Symbol distribution

Once the data has been FEC encoded, two FEC codewords ($c_A<543:0>$ and $c_B<543:0>$) are interleaved on a 10-bit basis before the data is distributed to each PCS lane.

The interleaving of two codewords for the 200GBASE-R PCS shall follow this procedure:

For all $k=0$ to 135

For all $j=0$ to 3

if even(k)

$tx_out<8k+2j> = c_A<543-4k-j>$

$tx_out<8k+2j+1> = c_B<543-4k-j>$

else

$tx_out<8k+2j> = c_B<543-4k-j>$

$tx_out<8k+2j+1> = c_A<543-4k-j>$

The interleaving of two codewords for the 400GBASE-R PCS shall follow this procedure:

For all $k=0$ to 67

For all $j=0$ to 7

if even(k)

$tx_out<16k+2j> = c_A<543-8k-j>$

$tx_out<16k+2j+1> = c_B<543-8k-j>$

else

$tx_out<16k+2j> = c_B<543-8k-j>$

$tx_out<16k+2j+1> = c_A<543-8k-j>$

Once the data has been Reed-Solomon encoded and interleaved, it shall be distributed to 8 PCS lanes for a 200GBASE-R PCS and to 16 PCS lanes for a 400GBASE-R PCS, one 10-bit symbol at a time, from the lowest to the highest PCS lane. The first bit transmitted from each 10-bit symbol is bit 0.

119.2.4.8 Transmit bit ordering and distribution

The transmit bit ordering and distribution are illustrated in Figure 119–10 for a 200GBASE-R PCS.

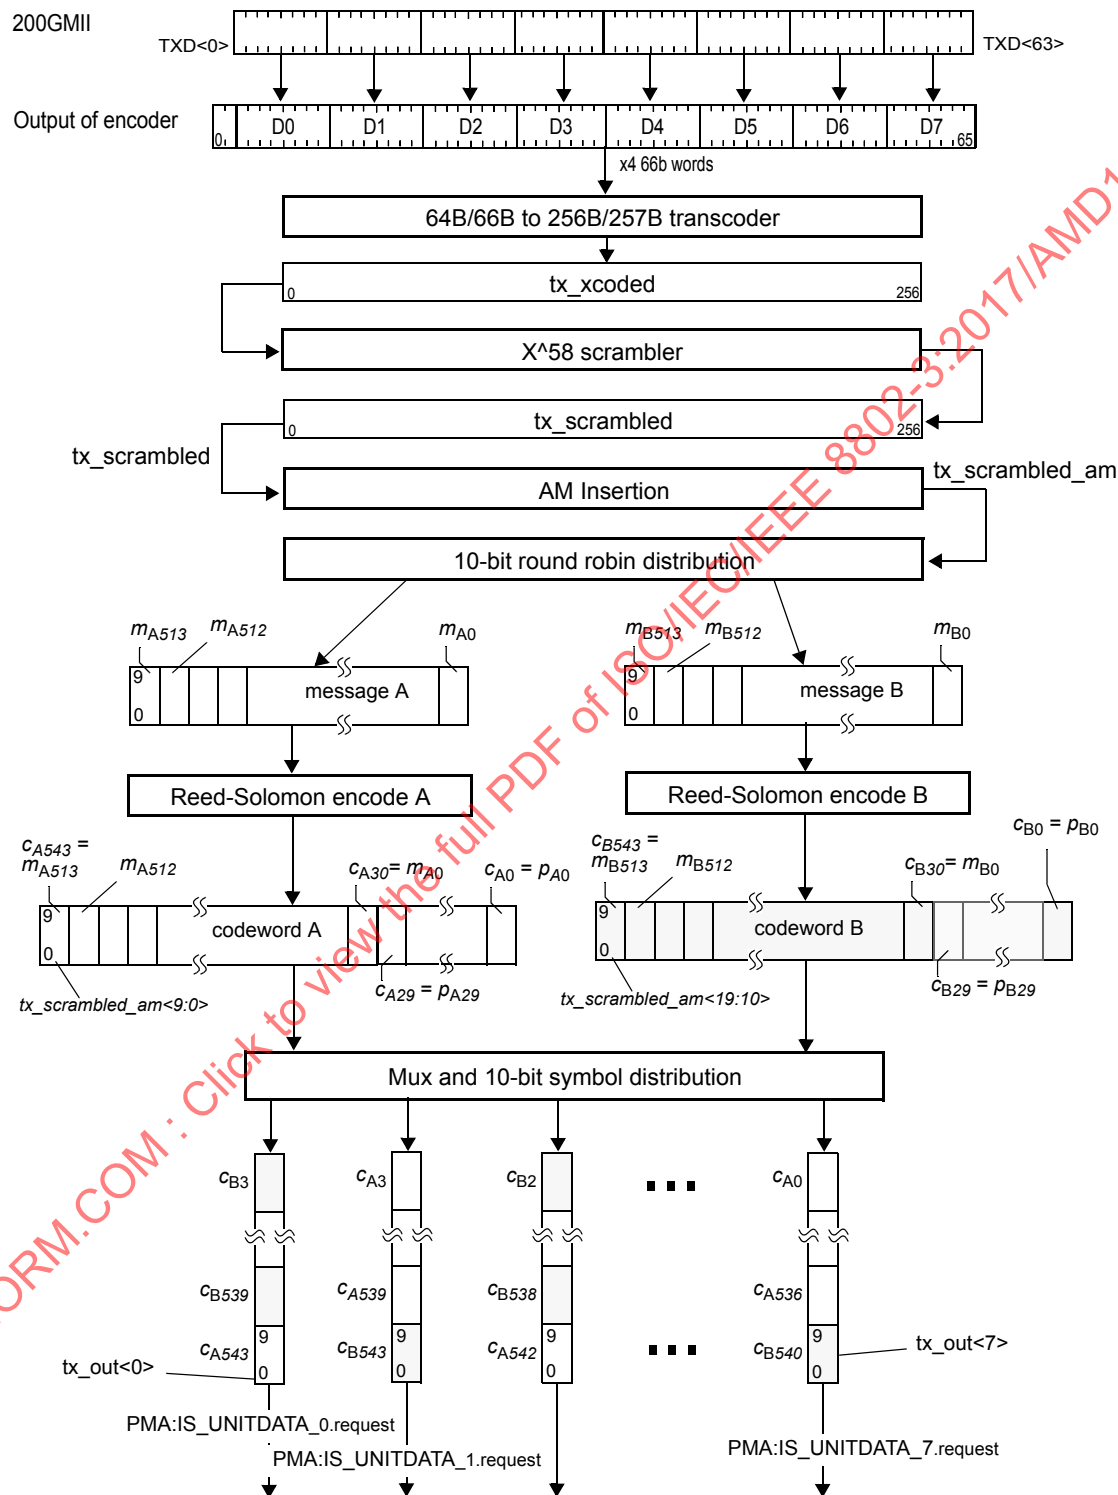


Figure 119–10—200GBASE-R Transmit bit ordering and distribution

The transmit bit ordering and distribution is illustrated in Figure 119–11 for a 400GBASE-R PCS.

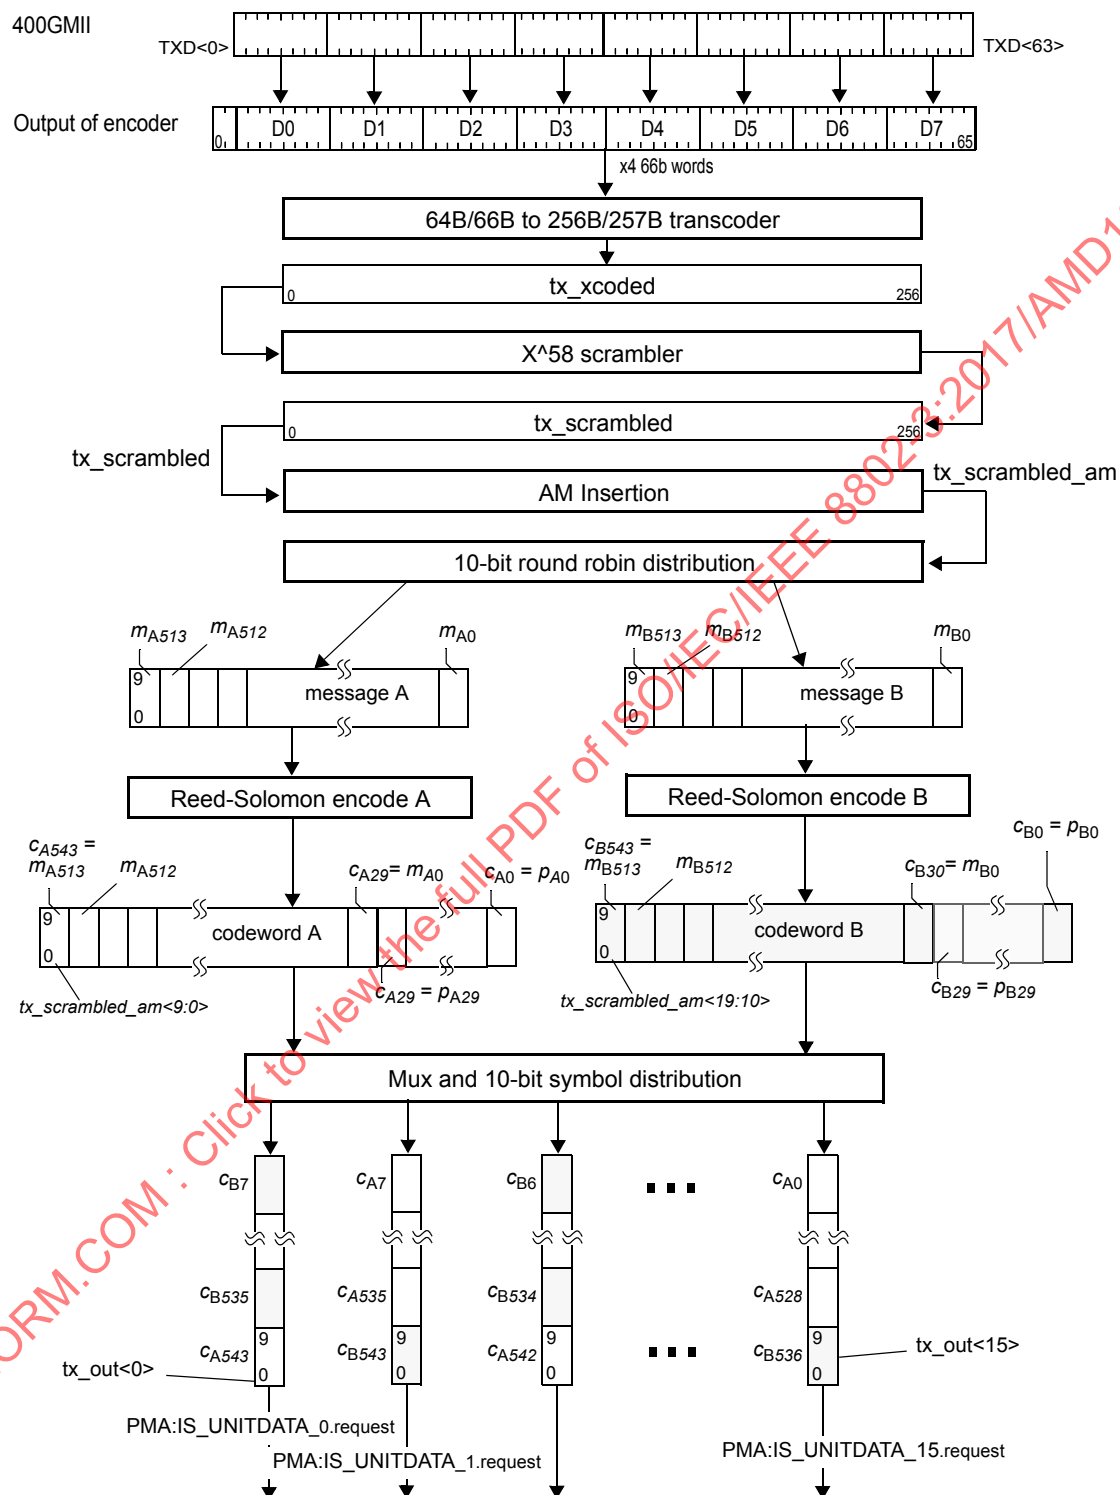


Figure 119–11—400GBASE-R Transmit bit ordering and distribution

119.2.4.9 Test-pattern generators

The PCS shall have the ability to generate a scrambled idle test pattern which is suitable for receiver tests and for certain transmitter tests. When a scrambled idle pattern is enabled, the test pattern is generated by the PCS. The test pattern is an idle control block (block type=0x1E) with all idles as defined in Figure 82-5. The test pattern is sent continuously and is transcoded, scrambled, alignment markers are inserted and finally encapsulated by the FEC.

When the transmit channel is operating in test-pattern mode, the encoded bit stream is distributed to the PCS Lanes as in normal operation (see 119.2.4.7).

If a Clause 45 MDIO is implemented, then control of the test-pattern generation is from the BASE-R PCS test-pattern control register (bit 3.42.3).

119.2.5 Receive function

119.2.5.1 Alignment lock and deskew

The receive PCS forms n separate bit streams by concatenating the bits from each of the n PMA:IS_UNITDATA_ i .indication primitives in the order they are received (where $n = 8$ for a 200GBASE-R PCS and $n = 16$ for a 400GBASE-R PCS). It obtains lock to the alignment markers as specified by the alignment marker lock state diagram shown in Figure 119-12. Note that alignment marker lock is achieved before FEC codewords are processed and therefore the alignment markers are processed in a high error probability environment.

After alignment marker lock is achieved on each of the n lanes (bit streams), all inter-lane Skew is removed as specified by the PCS synchronization state diagram shown in Figure 119-13. The PCS receive function shall support a maximum Skew of 180 ns, and maximum Skew Variation of 4 ns, between PCS lanes.

119.2.5.2 Lane reorder and de-interleave

PCS lanes can be received on different lanes of the service interface from which they were originally transmitted. The PCS receive function shall order the PCS lanes according to the PCS lane number. The PCS lane number is defined by the unique portion (UM₀ to UM₅) of the alignment marker that is mapped to each PCS lane (see 119.2.4.4).

After all PCS lanes are aligned, deskewed, and reordered, the two FEC codewords are de-interleaved to reconstruct the original stream of two FEC codewords.

119.2.5.3 Reed-Solomon decoder

The Reed-Solomon decoder extracts the message symbols from the codeword, corrects them as necessary, and discards the parity symbols.

The Reed-Solomon decoder shall be capable of correcting any combination of up to $t=15$ symbol errors in a codeword. The Reed-Solomon decoder shall also be capable of indicating when an errored codeword was not corrected. The probability that the decoder fails to indicate a codeword with $t+1$ errors as uncorrected is not expected to exceed 10^{-16} . This limit is also expected to apply for $t+2$ errors, $t+3$ errors, and so on.

If bypass error indication is not supported or not enabled, when the Reed-Solomon decoder determines that a codeword contains errors that were not corrected, it shall cause the PCS receive function to set every 66-bit block within the two associated codewords to an error block (set to EBLOCK_R). This may be achieved by setting the synchronization header to 11 for all 66-bit blocks created from these codewords by the 256B/257B to 64B/66B transcoder.

The Reed-Solomon decoder may optionally provide the ability to bypass the error indication feature to reduce the delay contributed by the FEC function. The presence of this option is indicated by the assertion of the `FEC_bypass_indication_ability` variable (see 119.3). When the option is provided it is enabled by the assertion of the `FEC_bypass_indication_enable` variable (see 119.3).

When `FEC_bypass_indication_enable` is asserted, additional error monitoring is performed by the Reed-Solomon decoder to reduce the likelihood that errors in a packet are not detected. The Reed-Solomon decoder counts the number of symbol errors detected on all PCS lanes in consecutive non-overlapping blocks of 8192 codewords. When the number of symbol errors in a block of 8192 codewords exceeds 5560, the Reed-Solomon decoder shall assert `hi_ser` for a period of 60 ms to 75 ms.

The Reed-Solomon decoder may optionally provide the ability to signal a degradation of the received signal. The presence of this option is indicated by the assertion of the `FEC_degraded_SER_ability` variable (see 119.3). When the option is provided it is enabled by the assertion of the `FEC_degraded_SER_enable` variable (see 119.3).

When `FEC_degraded_SER_enable` is asserted, additional error monitoring is performed by the PCS. The Reed-Solomon decoder counts the number of symbol errors detected on all PCS lanes in consecutive non-overlapping blocks of `FEC_degraded_SER_interval` codewords (see 119.3.1), where the least significant bit of `FEC_degraded_SER_interval` is ignored (evaluated as 0) to make the number of codewords even. If the decoder determines that a codeword is uncorrectable, the number of symbol errors detected is increased by 16. When the number of symbol errors exceeds the threshold set in `FEC_degraded_SER_activate_threshold` (see 119.3.1), the `FEC_degraded_SER` bit (see 119.3.1) is set. At the end of each interval, if the number of symbol errors is less than `FEC_degraded_SER_deactivate_threshold`, the `FEC_degraded_SER` bit is cleared. If either `FEC_degraded_SER_ability` or `FEC_degraded_SER_enable` is de-asserted then the `FEC_degraded_SER` bit is cleared.

119.2.5.4 Post FEC interleave

After the Reed-Solomon decoder processes the data, data is interleaved on a 10-bit basis into `rx_scrambled_am` from two codewords corresponding to 40 transcoded blocks in order to recreate the transmitted data stream.

119.2.5.5 Alignment marker removal

For the 200GBASE-R PCS, every 81 920 x 257-bit blocks (corresponds to 4096 codewords) the first 1028 bits of `rx_scrambled_am` blocks is the vector `am_rx<1027:0>` where bit 0 is the first bit received. The specific codewords that include this vector are indicated by the alignment lock and deskew function. The 3-bit receive alignment marker status field is assigned from the variable `am_rx` as follows:

$$\text{rx_am_sf}<2:0> = \text{am_rx}<1027:1025>$$

For the 400GBASE-R PCS, every 163 840 x 257-bit blocks (corresponds to 8192 codewords) the first 2056 bits of `rx_scrambled_am` blocks is the vector `am_rx<2055:0>` where bit 0 is the first bit received. The specific codewords that include this vector are indicated by the alignment lock and deskew function. The 3-bit receive alignment marker status field is assigned from the variable `am_rx` as follows:

$$\text{rx_am_sf}<2:0> = \text{am_rx}<2055:2053>$$

The vector `am_rx` shall be removed from `rx_scrambled_am` to create `rx_scrambled` prior to descrambling.

119.2.5.6 Descrambler

The descrambler shall process `rx_scrambled<256:0>` to reverse the effect of the scrambler using the polynomial given in Equation (49-1).

119.2.5.7 256B/257B to 64B/66B transcoder

The transcoder extracts a group of four 66-bit blocks, $rx_coded_j<65:0>$ where $j=0$ to 3, from each 257-bit block $rx_xcoded<256:0>$. Bit 0 of the 257-bit block is the first bit received.

If $rx_xcoded<0>$ is 1, $rx_coded_j<65:0>$ for $j=0$ to 3 shall be derived as follows.

- a) $rx_coded_j<65:2> = rx_xcoded<(64j+64):(64j+1)>$ for $j=0$ to 3
- b) $rx_coded_j<0>=0$ and $rx_coded_j<1>=1$ for all $j=0$ to 3

If $rx_xcoded<0>$ is 0 and any $rx_xcoded<j+1>=0$ for $j=0$ to 3, $rx_coded_j<65:0>$ for $j=0$ to 3 shall be derived as follows.

- a1) Let c be the smallest value of j such that $rx_xcoded<j+1>=0$. In other words, rx_coded_c is the first 66-bit control block in the resulting group of four blocks.
- b1) Let $rx_payloads$ be a vector representing the payloads of the four 66-bit blocks. It is derived using the following expressions:
 $rx_payloads<(64c+3):0> = rx_xcoded<(64c+8):5>$
 $rx_payloads<(64c+7):(64c+4)>$ is set to a value derived by cross-referencing $rx_payloads<(64c+3):64c>$ using Figure 82-5. For example, if $rx_payloads<(64c+3):64c>$ is 0xE then $rx_payloads<(64c+7):(64c+4)>$ is 0x1. If no match to $rx_payloads<(64c+3):64c>$ is found, $rx_payloads<(64c+7):(64c+4)>$ is set to 0000
 $rx_payloads<255:(64c+8)> = rx_xcoded<256:(64c+9)>$
- c1) $rx_coded_j<65:2> = rx_payloads<(64j+63):64j>$ for $j=0$ to 3
- d1) If $rx_xcoded<j+1>=0$, $rx_coded_j<0>=1$ and $rx_coded_j<1>=0$ for $j=0$ to 3
- e1) If $rx_xcoded<j+1>=1$, $rx_coded_j<0>=0$ and $rx_coded_j<1>=1$ for $j=0$ to 3
- f1) If $rx_payloads<(64c+7):(64c+4)> = 0000$, $rx_coded_c<1>=1$ (invalidate synchronization header)

If $rx_xcoded<0>$ is 0 and all $rx_xcoded<j+1>=1$ for $j=0$ to 3, $rx_coded_j<65:0>$ for $j=0$ to 3 shall be derived as follows.

- a2) Set $c = 0$
- b2) Let $rx_payloads$ be a vector representing the payloads of the four 66-bit blocks. It is derived using the following expressions.
 $rx_payloads<(64c+3):0> = rx_xcoded<(64c+8):5>$
 $rx_payloads<(64c+7):(64c+4)> = 0000$
 $rx_payloads<255:(64c+8)> = rx_xcoded<256:(64c+9)>$
- c2) $rx_coded_j<65:2> = rx_payloads<(64j+63):(64j)>$ for $j=0$ to 3
- d2) $rx_coded_j<0>=0$ and $rx_coded_j<1>=0$ for $j=0$ and 2
- e2) $rx_coded_j<0>=1$ and $rx_coded_j<1>=1$ for $j=1$ and 3

The 66-bit blocks are transmitted in order from $j=0$ to 3. Bit 0 of each block is the first bit transmitted.

Note—The stream of 66-bit blocks generated by this process is used as the reference signal for de-mapping from OTN. See ITU-T G.709 [B50].

119.2.5.8 Decode and rate matching

The receive PCS decodes blocks to produce $RXD<63:0>$ and $RXC<7:0>$ for transmission to the 200GMII/400GMII. One 200GMII/400GMII data transfer is decoded from each block. The receive PCS may insert idle control characters to compensate for the removal of alignment markers. If the receive PCS spans multiple clock domains, it may also perform clock rate compensation via the deletion of idle control characters or sequence ordered sets or the insertion of idle control characters (see 82.2.3.6 and 82.2.3.9 for insertion and deletion rules).

The PCS receive decodes blocks as specified in the receive state diagram shown in Figure 119–15.

119.2.6 Detailed functions and state diagrams

119.2.6.1 State diagram conventions

The body of this subclause is composed of state diagrams, including the associated definitions of variables, functions, and counters. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams follows the conventions of 21.5. The notation ++ after a counter or integer variable indicates that its value is to be incremented.

119.2.6.2 State variables

119.2.6.2.1 Constants

- EBLOCK_R<71:0>**
72-bit vector to be sent to the 200GMII/400GMII containing /E/ in all the eight character locations.
- EBLOCK_T<65:0>**
66-bit vector to be sent to the transcoder containing /E/ in all the eight character locations.
- LBLOCK_R<71:0>**
72-bit vector to be sent to the 200GMII/400GMII containing one Local Fault ordered set. The Local Fault ordered set is defined in 119.3.
- LBLOCK_T<65:0>**
66-bit vector to be sent to the transcoder containing one Local Fault ordered set.

119.2.6.2.2 Variables

- align_status**
A variable set by the PCS synchronization process to reflect the status of PCS lane-to-lane alignment. Set to true when all lanes are synchronized and aligned and set to false when the PCS synchronization process is not complete.
- all_locked**
A Boolean variable that is set to true when `amps_lock<x>` is true for all x and is set to false when `amps_lock<x>` is false for any x .
- amp_counter_done**
Boolean variable that indicates that `amp_counter` has reached its terminal count.
- amp_match**
Boolean variable that holds the output of the function `AMP_COMPARE`.
- amp_valid**
Boolean variable that is set to true if the received 120-bit block is a valid alignment marker payload. The alignment marker payload, mapped to a PCS lane according to the process described in 119.2.4.4, consists of 96 known bits. The 48 bits of the common marker portion are compared on a nibble-wise basis (12 comparisons). If 9 or more nibbles in the candidate block match the corresponding known nibbles in the common portion of the alignment marker payload, the candidate block is considered a valid alignment marker payload.
- amps_lock<x>**
Boolean variable that is set to true when the receiver has detected the location of the alignment marker payload sequence for a given lane on the PMA service interface, where $x = 0:7$ for 200GBASE-R and $x = 0:15$ for 400GBASE-R.

current_pcs1

A variable that holds the PCS lane number corresponding to the current alignment marker payload that is recognized on a given lane of the PMA service interface. It is compared to the variable **first_pcs1** to confirm that the location of the alignment marker payload sequence has been detected. The PCS lane number is determined by the alignment marker payloads based on the mapping defined in 119.2.4.4. The 48 bits that are in the positions of the unique marker bits in the received alignment marker payload are compared to the expected values for a given payload position and PCS lane on a nibblewise basis (12 comparisons). If 9 or more nibbles in the candidate block match the corresponding known nibbles for any payload position on a given PCS lane, then the PCS lane number is assigned accordingly. If a match is not found, then the **AMP_COMPARE** function sets **amp_match** to false.

cw_A_bad

A Boolean variable that is set to true if the Reed-Solomon decoder (see 119.2.5.3) fails to correct the current FEC codeword A and is set to false otherwise.

cw_B_bad

A Boolean variable that is set to true if the Reed-Solomon decoder (see 119.2.5.3) fails to correct the current FEC codeword B and is set to false otherwise.

deskew_done

A Boolean variable that is set to true when **pcs_enable_deskew** is set to true and the deskew process is completed. Otherwise, this variable is set to false.

first_pcs1

A variable that holds the PCS lane number that corresponds to the first alignment marker payload that is recognized on a given lane of the PMA service interface. It is compared to the PCS lane number corresponding to the second alignment marker payload that is tested.

hi_ser

When **FEC_bypass_indication_enable** is set to one, this bit is set to one for a period of 60 ms to 75 ms if the number of RS-FEC symbol errors in a window of 8192 codewords exceeds the threshold (see 119.2.5.3), and is set to zero otherwise. When **FEC_bypass_indication_enable** is set to zero, this bit is set to zero. This variable is mapped to the bit defined in 45.2.3.47k (3.801.2).

pcs_alignment_valid

Boolean variable that is set to true if all PCS lanes are aligned. PCS lanes are considered to be aligned when **amps_lock_{<x>}** is true for all *x*, each PCS lane is locked to a unique alignment marker payload sequence (see 119.2.4.4), and the PCS lanes are deskewed. Otherwise, this variable is set to false.

pcs_enable_deskew

A Boolean variable that enables and disables the PCS synchronization process. Received bits may be discarded whenever deskew is enabled. It is set to true when deskew is enabled and set to false when deskew is disabled.

pcs_lane_mapping_{<x>}

A variable that holds the value of the **pcs_lane** received on physical lane *x*.

PCS_status

A boolean variable that is true when **align_status** is true and is false otherwise.

reset

Boolean variable that controls the resetting of the PCS sublayer. It is true whenever a reset is necessary including when reset is initiated from the MDIO, during power on, and when the MDIO has put the PCS sublayer into low-power mode.

restart_lock

Boolean variable that is set by the PCS synchronization process to restart the alignment marker lock process on all PCS lanes. It is set to true after 3 consecutive uncorrected codewords are received (3_BAD state) or when 5 Alignment Markers in a row fail to match (5_BAD state) and set to false upon entry into the **LOSS_OF_ALIGNMENT** state.

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_local_degraded

Boolean variable that is asserted true when the receiver detects rx_am_sf<1> asserted true for two consecutive alignment marker periods. It is deasserted when rx_am_sf<1> is deasserted for two consecutive alignment marker periods. If a Clause 45 MDIO is implemented, the status of this variable is reflected in bit 3.801.6.

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.

rx_rm_degraded

Boolean variable that is asserted true when the receiver detects rx_am_sf<2> asserted true for two consecutive alignment marker periods. It is deasserted when rx_am_sf<2> is deasserted for two consecutive alignment marker periods. If a Clause 45 MDIO is implemented, the status of this variable is reflected in bit 3.801.5.

signal_ok

Boolean variable that is set based on the most recently received value of PMA:IS_SIGNAL.indication(SIGNAL_OK). It is true if the value was OK and false if the value was FAIL.

slip_done

Boolean variable that is set to true when the SLIP requested by the alignment marker lock state diagram has been completed indicating that the next candidate 120-bit block position can be tested.

test_amp

Boolean variable that is set to true when a candidate block position is available for testing and false when the FIND_1ST state is entered.

test_cw

Boolean variable that is set to true when a new FEC codeword is available for decoding and is set to false when the TEST_CW state is entered.

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.

119.2.6.2.3 Functions

AMP_COMPARE

This function compares the values of first_pcs1 and current_pcs1 to determine if a valid alignment marker payload sequence has been detected and returns the result of the comparison using the variable amp_match. If current_pcs1 and first_pcs1 both found a match and indicate the same PCS lane number, amp_match is set to true.

DECODE(rx_coded<65:0>)

Decodes the 66-bit vector returning rx_raw<71:0>, which is sent to the 200GMII/400GMII. The DECODE function shall decode the block as specified in 119.2.3.

ENCODE(tx_raw<71:0>)

Encodes the 72-bit vector returning tx_coded<65:0>. The ENCODE function shall encode the block as specified in 119.2.3.

R_TYPE(rx_coded<65:0>)

This function classifies the current rx_coded<65:0> vector as belonging to one of the following types, depending on its contents. The classification results are returned via the r_block_type variable.

Values: C; The vector contains a sync header of 10 and one of the following:

- a) A block type field of 0x1E and eight valid control characters other than /E/ or /LI/;
- b) A block type field of 0x4B.

LI; For EEE capability, the LI type is supported where the vector contains a sync header of 10, a block type field of 0x1E and eight control characters of 0x06 (/LI/).

S; The vector contains a sync header of 10 and the following:

- a) A block type field of 0x78.

T; The vector contains a sync header of 10, a block type field of 0x87, 0x99, 0xAA, 0xB4, 0xCC, 0xD2, 0xE1 or 0xFF and all control characters are valid.

D; The vector contains a sync header of 01.

E; The vector does not meet the criteria for any other value.

Valid control characters are specified in Table 82-1.

NOTE—A PCS that does not support EEE classifies vectors containing one or more /LI/ control characters as type E.

R_TYPE_NEXT

This function classifies the 66-bit rx_coded vector that immediately follows the current rx_coded<65:0> vector as belonging to one of the five types defined in R_TYPE, depending on its contents. It is intended to perform a prescient end of packet check. The classification results are returned via the r_block_type_next variable.

SLIP

Causes the next candidate block position to be tested. The precise method for determining the next candidate block position is not specified and is implementation dependent. However, an implementation shall ensure that all possible block positions are evaluated.

T_TYPE = (tx_raw<71:0>)

This function classifies each 72-bit tx_raw vector as belonging to one of the following types depending on its contents. The classification results are returned via the t_block_type variable.

Values: C; The vector contains one of the following:

- a) Eight valid control characters other than /O/, /S/, /T/, /LI/, and /E/;
- b) One valid ordered set.

LI; For EEE capability, this vector contains eight /LI/ characters.

S; The vector contains an /S/ in its first character, and all characters following the /S/ are data characters.

T; The vector contains a /T/ in one of its characters, all characters before the /T/ are data characters, and all characters following the /T/ are valid control characters other than /O/, /S/ and /T/.

D; The vector contains eight data characters.

E; The vector does not meet the criteria for any other value.

A tx_raw character is a control character if its associated TXC bit is asserted. A valid control character is one containing an 200GMII/400GMII control code specified in Table 82-1. A valid ordered set consists of a valid /O/ character in the first character and data characters in the seven characters following the /O/. A valid /O/ is any character with a value for O code in Table 82-1.

NOTE—A PCS that does not support EEE classifies vectors containing one or more /LI/ control characters as type E.

119.2.6.2.4 Counters

amp_bad_count

Counts the number of consecutive alignment markers that don't match the expected values for a given PCS lane.

amp_counter

This counter counts the interval of i FEC codewords containing normal alignment marker payload sequences (where i is 4096 for the 200GBASE-R PCS, and 8192 for the 400GBASE-R PCS).

cw_A_bad_count

Counts the number of consecutive uncorrected FEC codewords for codeword A. This counter is set to zero when an FEC codeword A is received and cw_A_bad is false.

cw_B_bad_count

Counts the number of consecutive uncorrected FEC codewords for codeword B. This counter is set to zero when an FEC codeword B is received and cw_B_bad is false.

119.2.6.3 State diagrams

The 200GBASE-R PCS shall implement eight alignment marker lock processes and the 400GBASE-R PCS shall implement sixteen alignment marker lock processes as depicted in Figure 119–12. An alignment marker lock process operates independently on each lane. The alignment marker lock state diagram shown in Figure 119–12 determines when the PCS has obtained alignment marker lock to the received bit stream for a given lane of the service interface. Each alignment marker lock process looks for two valid alignment markers 278 528 × 10-bit Reed-Solomon symbols apart, on a per PCS lane basis, to gain alignment marker lock. When the alignment marker lock process achieves lock for a lane, and if a Clause 45 MDIO is implemented, the PCS shall record the PCS lane number received on that lane of the service interface in the appropriate lane mapping register (3.400 to 3.415). Once in lock, a lane goes out of alignment marker lock only when restart_lock is signaled. This occurs when the PCS synchronization process determines that three uncorrectable codewords in a row are seen, or when the alignment marker lock process sees five alignment markers in a row that fail to match the expected pattern on a given lane.

The PCS shall run the synchronization process as depicted in Figure 119–13. The PCS synchronization process is responsible for determining if the PCS is capable of presenting coherent data to the 200GMII/400GMII. The synchronization process ensures that all PCS lanes have alignment marker lock, are locked to different alignment markers, and that the Skew is within the boundaries of what the PCS can deskew. The synchronization process, along with the alignment marker lock process, are restarted if three consecutive FEC codewords from the same codeword (A or B) are uncorrectable.

The Transmit state diagram shown in Figure 119–14 controls the encoding of 66-bit transmitted blocks. It makes exactly one transition for each 66-bit transmit block processed. Though the Transmit state diagram sends Local Fault ordered sets when reset is asserted, the scrambler is not guaranteed to be operational during reset. Thus, the Local Fault ordered sets may not appear on the PMA service interface.

The Receive state diagram shown in Figure 119–15 controls the decoding of received blocks. It makes exactly one transition for each receive 66-bit block processed.

The PCS shall perform the functions of alignment marker lock, PCS synchronization, Transmit, and Receive as specified in the respective state diagrams.

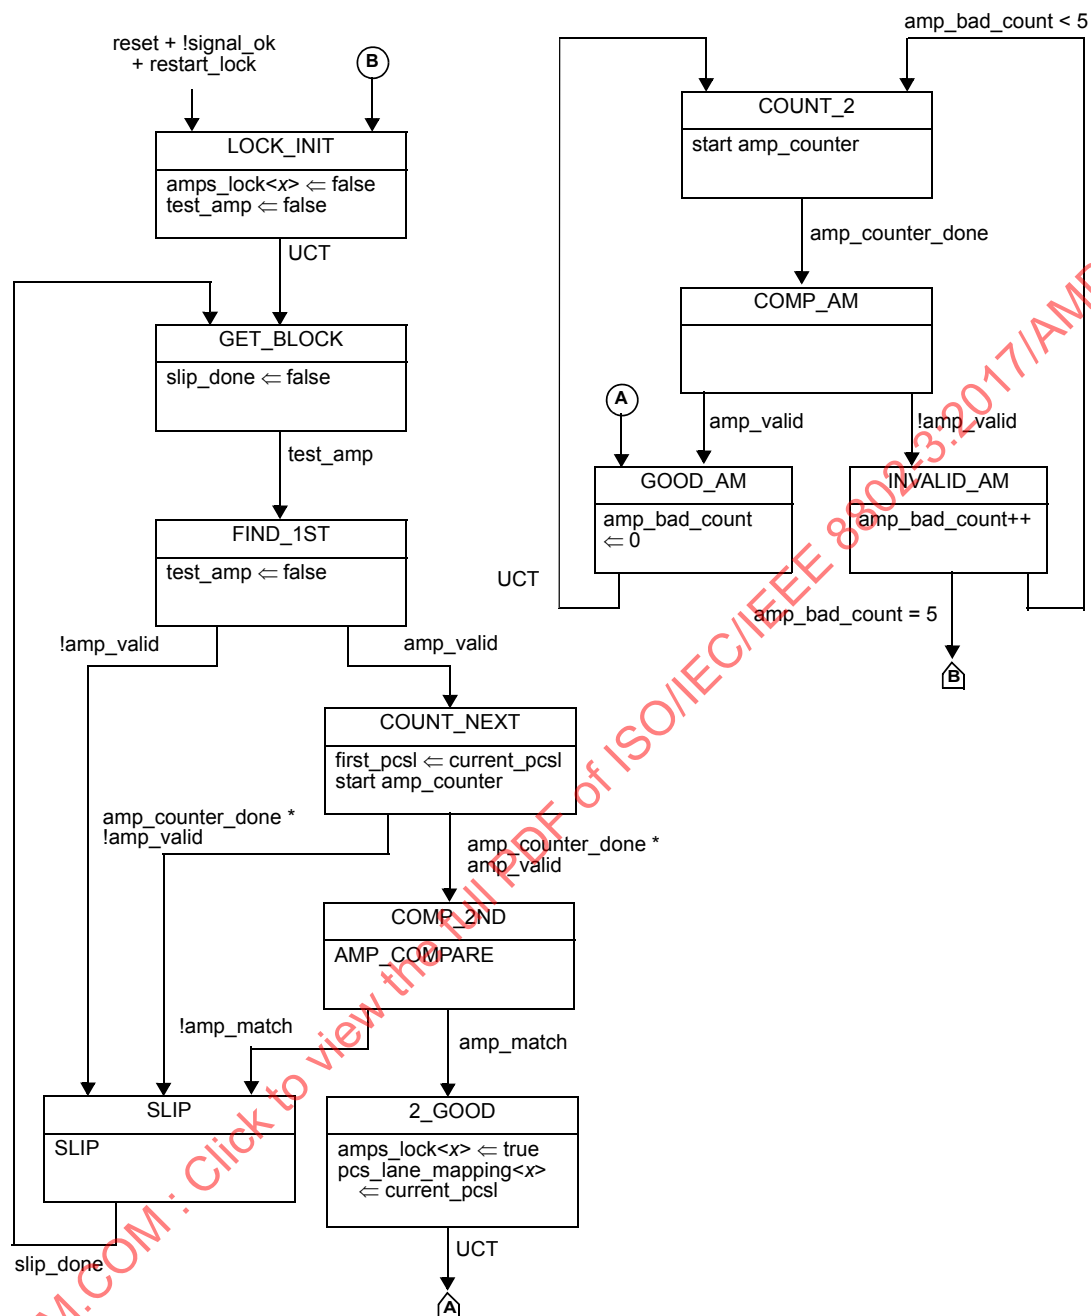


Figure 119-12—Alignment marker lock state diagram

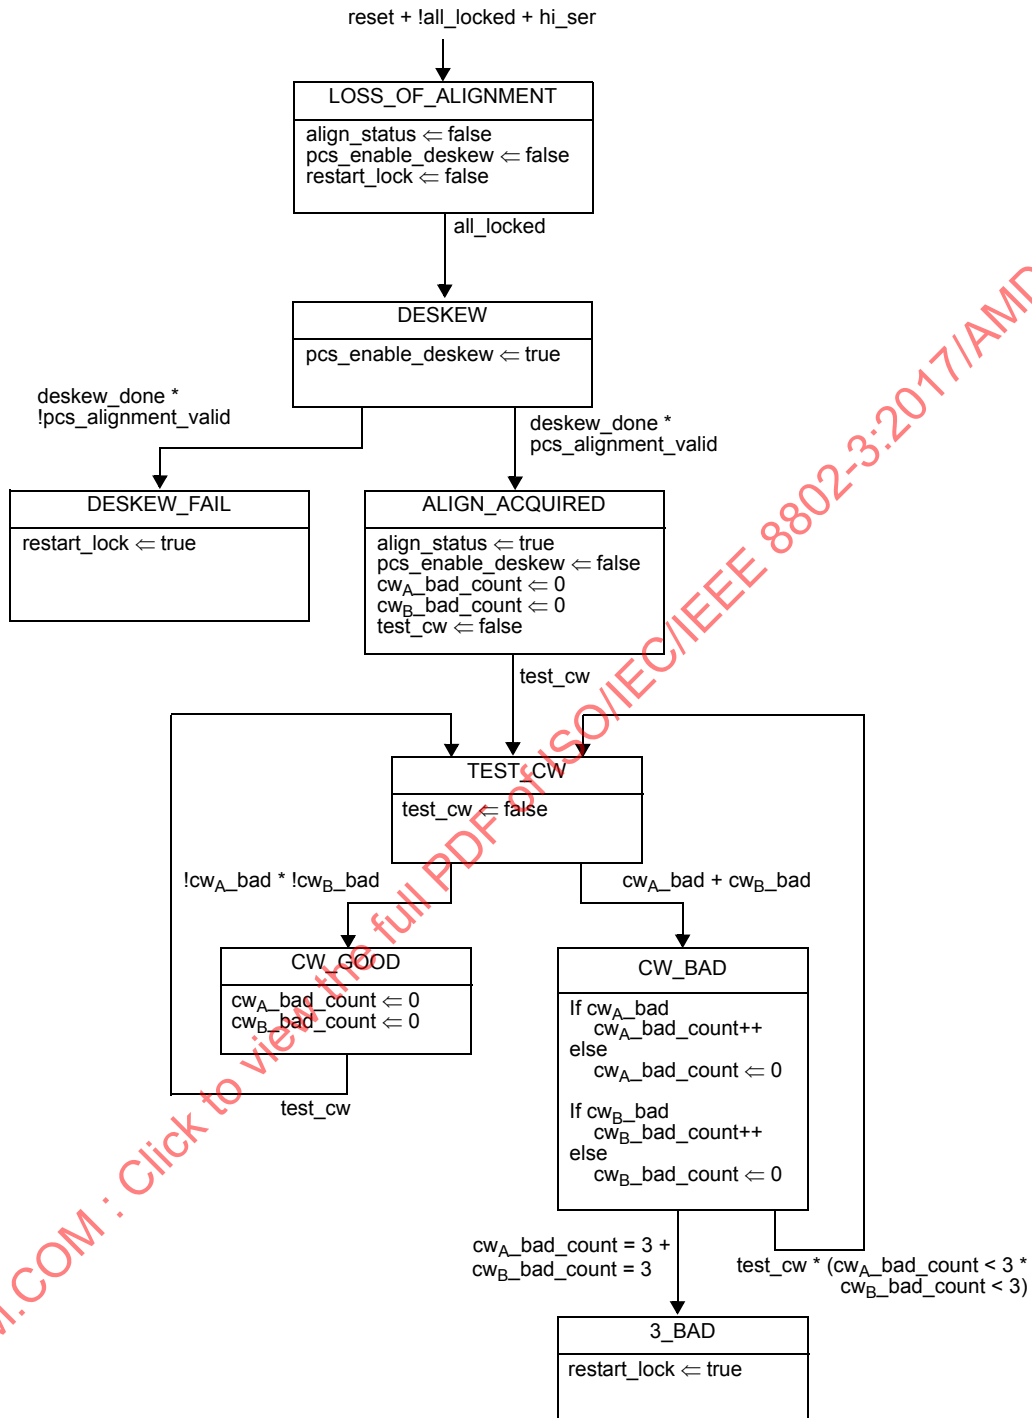


Figure 119–13—PCS synchronization state diagram

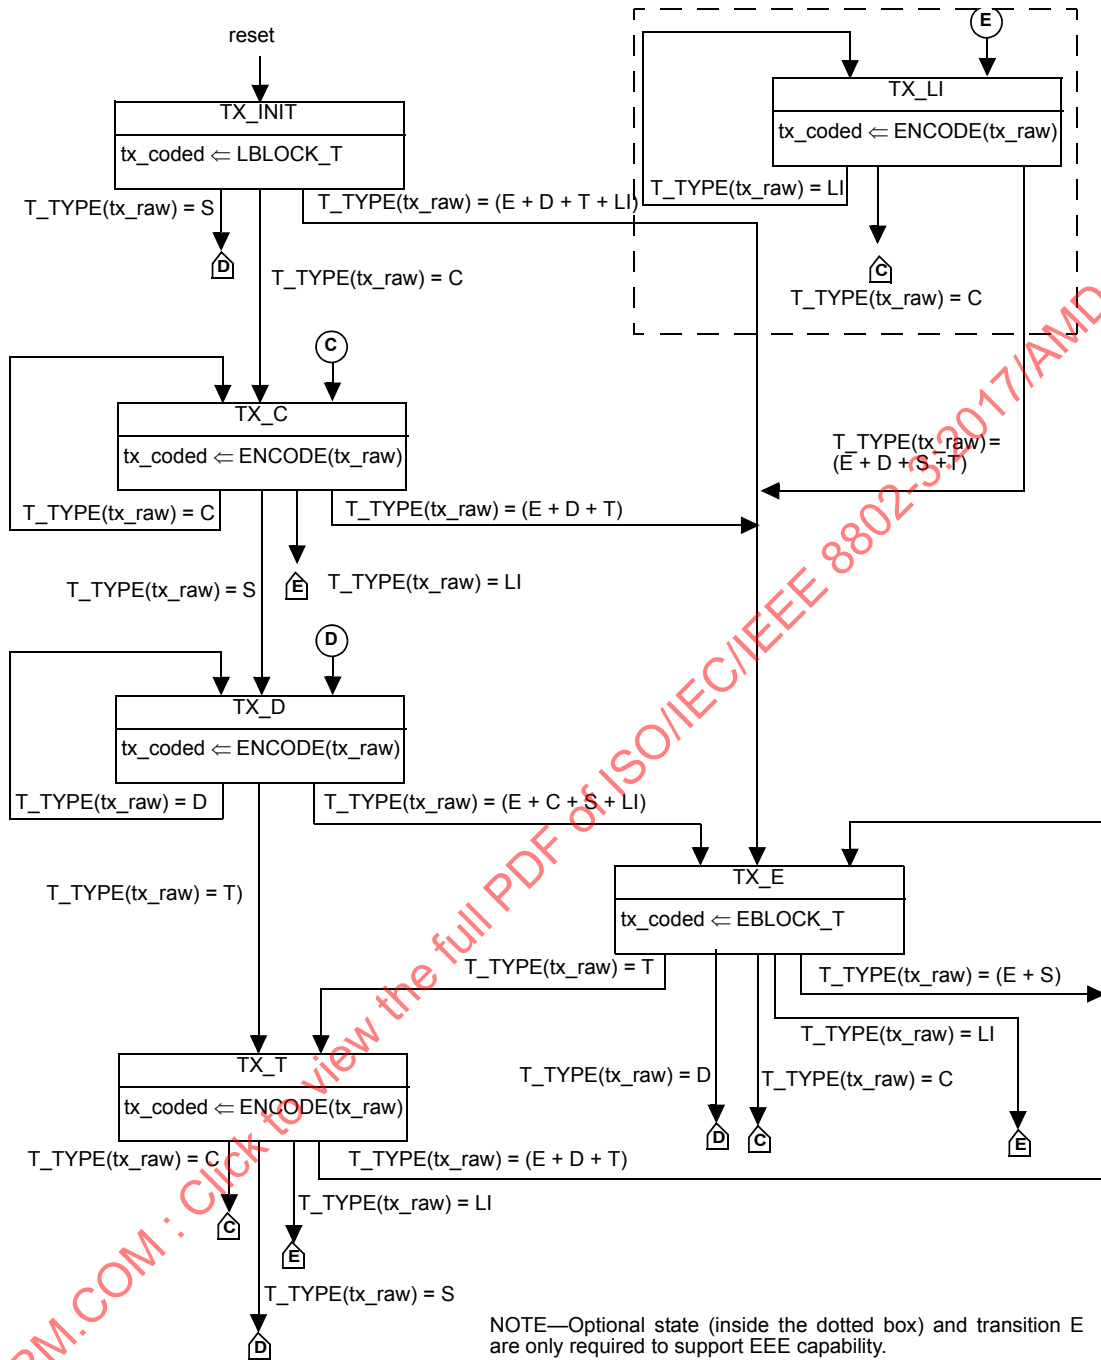


Figure 119–14—Transmit state diagram

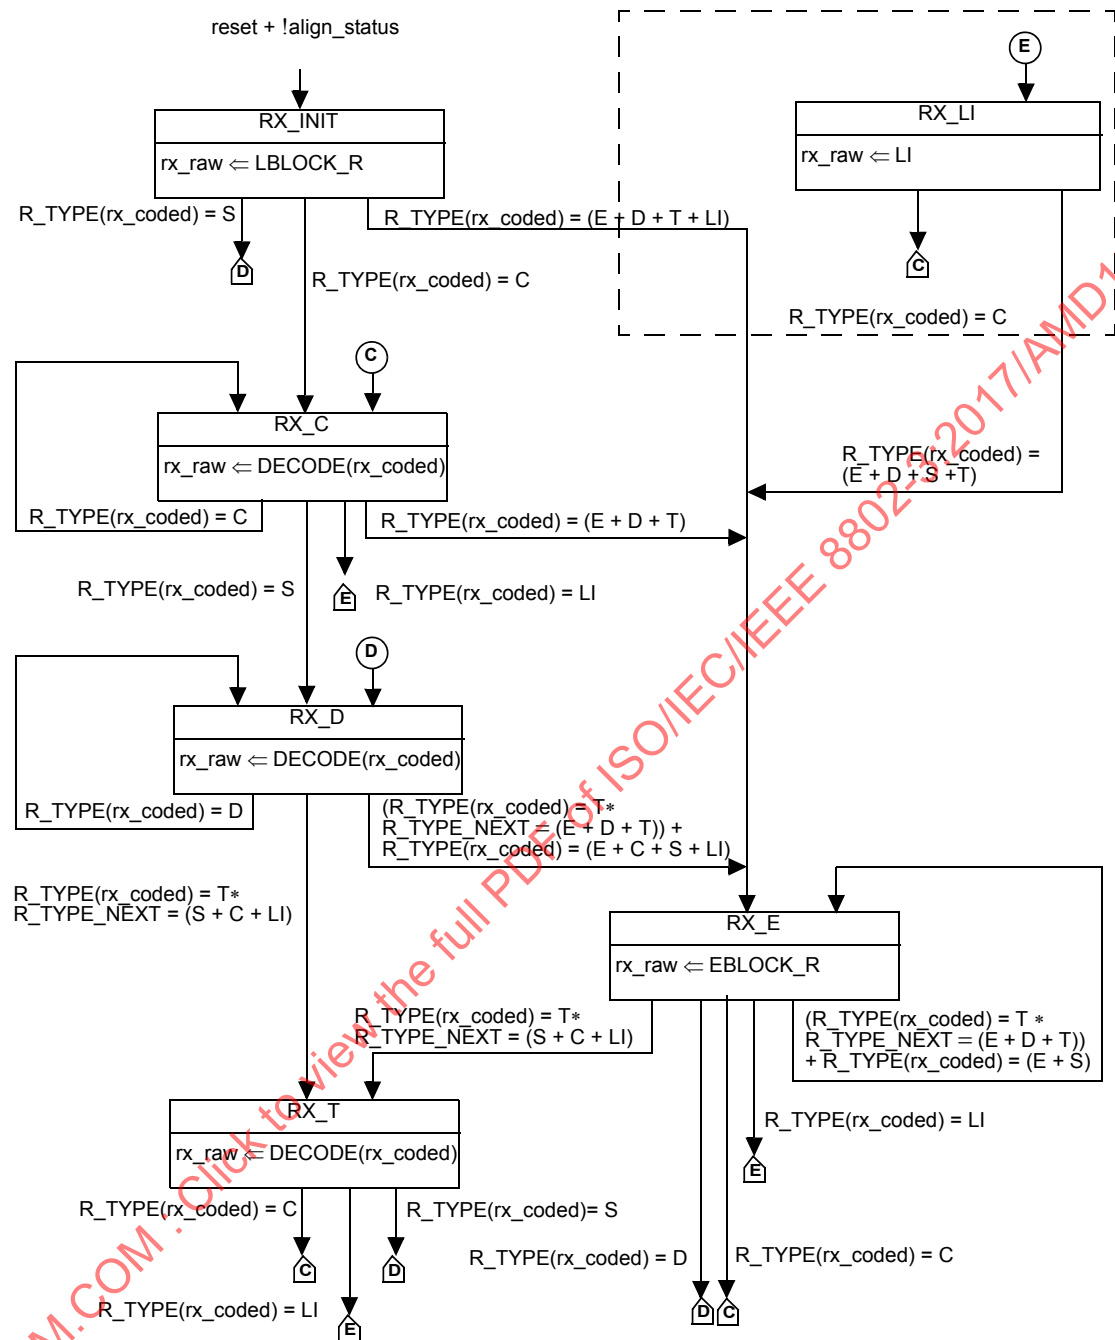


Figure 119–15—Receive state diagram

119.3 PCS management

The following objects apply to PCS management. If an MDIO Interface is provided (see Clause 45), they are accessed via that interface. If not, it is recommended that an equivalent access is provided.

119.3.1 PCS MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the PCS. If MDIO is implemented, it shall map MDIO control bits to PCS control variables, as shown in Table 119–4, and MDIO status bits to PCS status variables, as shown in Table 119–5.

Table 119–4—MDIO/PCS control variable mapping

| MDIO control variable | PCS register name | Register/ bit number | PCS control variable |
|---|--|----------------------|---------------------------------------|
| Reset | PCS control 1 register | 3.0.15 | reset |
| Loopback | PCS control 1 register | 3.0.14 | Loopback |
| Transmit test-pattern enable | BASE-R PCS test-pattern control register | 3.42.3 | tx_test_mode |
| LPI_FW | EEE control and capability | 3.20.0 | LPI_FW |
| PCS FEC bypass indication enable | PCS FEC control register | 3.800.1 | FEC_bypass_indication_enable |
| PCS FEC degraded SER enable | PCS FEC control register | 3.800.2 | FEC_degraded_SER_enable |
| PCS FEC degraded SER activate threshold | PCS FEC degraded SER activate threshold register | 3.806, 3.807 | FEC_degraded_SER_activate_threshold |
| PCS FEC degraded SER deactivate threshold | PCS FEC degraded SER deactivate threshold register | 3.808, 3.809 | FEC_degraded_SER_deactivate_threshold |
| PCS FEC degraded SER interval | PCS FEC degraded SER interval | 3.810, 3.811 | FEC_degraded_SER_interval |

Table 119–5—MDIO/PCS status variable mapping

| MDIO status variable | PCS register name | Register/ bit number | PCS status variable |
|---|--|----------------------|-------------------------------|
| BASE-R and MultiGBASE-T receive link status | BASE-R and MultiGBASE-T PCS status 1 register | 3.32.12 | PCS_status |
| Lane <i>x</i> aligned | Multi-lane BASE-R PCS alignment status 3 and 4 registers | 3.52.7:0 3.53.7:0 | am_lock< <i>x</i> > |
| PCS lane alignment status | Multi-lane BASE-R PCS alignment status 1 register | 3.50.12 | align_status |
| Lane <i>x</i> mapping | Lane <i>x</i> mapping register | 3.400 through 3.415 | pcs_lane_mapping< <i>x</i> > |
| PCS FEC bypass indication ability | PCS FEC status register | 3.801.1 | FEC_bypass_indication_ability |

Table 119–5—MDIO/PCS status variable mapping (continued)

| MDIO status variable | PCS register name | Register/ bit number | PCS status variable |
|---|---|----------------------|--|
| PCS FEC corrected codewords | PCS FEC corrected codewords counter register | 3.802, 3.803 | FEC_corrected_cw_counter |
| PCS FEC uncorrected codewords | PCS FEC uncorrected codewords counter register | 3.804, 3.805 | FEC_uncorrected_cw_counter |
| PCS FEC symbol errors, PCS lanes 0 to x | PCS FEC symbol error counter register, lanes 0 to x | 3.600 to 3.631 | FEC_symbol_error_counter _{i} |
| Tx LPI indication | PCS status 1 | 3.1.9 | Tx LPI indication |
| Tx LPI received | PCS status 1 | 3.1.11 | Tx LPI received |
| Rx LPI indication | PCS status 1 | 3.1.8 | Rx LPI indication |
| Rx LPI received | PCS status 1 | 3.1.10 | Rx LPI received |
| EEE wake error counter | EEE wake error counter | 3.22 | Wake_error_counter |
| PCS FEC degraded SER ability | PCS FEC status register | 3.801.3 | FEC_degraded_SER_ability |
| PCS FEC degraded SER | PCS FEC status register | 3.801.4 | FEC_degraded_SER |
| Local degraded SER received | PCS FEC status register | 3.801.6 | rx_local_degraded |
| Remote degraded SER received | PCS FEC status register | 3.801.5 | rx_rm_degraded |
| PCS FEC high SER | PCS FEC status register | 3.801.2 | hi_ser |

119.4 Loopback

When the PCS is in loopback, the PCS shall accept data on the transmit path from the 200GMII/400GMII and return it on the receive path to the 200GMII/400GMII. In addition, the PCS shall transmit what it receives from the 200GMII/400GMII to the PMA sublayer, and shall ignore all data presented to it by the PMA sublayer. If a Clause 45 MDIO is implemented, then the PCS is placed in the loopback when the loopback bit from the PCS control 1 register (bit 3.0.14) is set to a one.

119.5 Delay constraints

The maximum delay contributed by the 200GBASE-R PCS (sum of transmit and receive delays at one end of the link) shall be no more than 160 256 BT (313 pause_quanta or 801.28 ns). The maximum delay contributed by the 400GBASE-R PCS (sum of transmit and receive delays at one end of the link) shall be no more than 320 000 BT (625 pause_quanta or 800 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 116.4 and its references.

119.6 Protocol implementation conformance statement (PICS) proforma for Clause 119, Physical Coding Sublayer (PCS) for 64B/66B, type 200GBASE-R and 400GBASE-R⁴

119.6.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 119, Physical Coding Sublayer (PCS) for 64B/66B, type 200GBASE-R and 400GBASE-R, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

119.6.2 Identification

119.6.2.1 Implementation identification

| | |
|--|--|
| Supplier ¹ | |
| Contact point for inquiries about the PICS ¹ | |
| Implementation Name(s) and Version(s) ^{1,3} | |
| Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ² | |
| NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model). | |

119.6.2.2 Protocol summary

| | |
|---|---|
| Identification of protocol standard | IEEE Std 802.3bs-2017, Clause 119, Physical Coding Sublayer (PCS) 64B/66B, type 200GBASE-R and 400GBASE-R |
| Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS | |
| Have any Exception items been required? No <input type="checkbox"/> Yes <input type="checkbox"/> (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3bs-2017.) | |

| | |
|-------------------|--|
| Date of Statement | |
|-------------------|--|

⁴Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

119.6.3 Major capabilities/options

| Item | Feature | Subclause | Value/Comment | Status | Support |
|---------|---|-------------------|---|--------|-------------------|
| MII | 200GMII or 400GMII logical interface | 117, 119.1.4.1 | Logical interface is supported | O | Yes [] No [] |
| *PCS200 | PCS for 200GBASE-R | 119.1.1 | | O.1 | Yes [] No [] |
| *PCS400 | PCS for 400GBASE-R | 119.1.1 | | O.1 | Yes [] No [] |
| *MD | MDIO | 45, 119.3 | Registers and interface supported | O | Yes [] No [] |
| *BI | Bypass indication | 119.2.5.3 | Capability is supported | O | Yes [] No [] |
| DC | Delay constraints | 119.5 | Conforms to delay constraints specified in 119.5 | M | Yes [] |
| *EEE | EEE capability | 119.2.3.3 | Capability is supported | O | Yes [] No [] |
| JTM | Supports test-pattern mode | 119.2.1 | | M | Yes [] |
| FDD | Support for optional FEC degraded detection | 119.2.5.3 | In the FEC decoder can optionally detect FEC SER degraded at a programmable threshold | O | Yes [] No [] |

119.6.4 PICS proforma tables for Physical Coding Sublayer (PCS) 64B/66B, type 200GBASE-R and 400GBASE-R

119.6.4.1 Transmit function

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---------------------------------|-----------|--|--------|---------|
| TF1 | 64B/66B to 256B/257B transcoder | 119.2.4.2 | tx_xcoded<256:0> constructed per 119.2.4.2 | M | Yes [] |
| TF2 | Transmission bit ordering | 119.2.4.8 | First bit transmitted is bit 0 | M | Yes [] |
| TF3 | Pad value | 119.2.4.4 | PRBS9 | M | Yes [] |
| TF4 | Alignment marker insertion | 119.2.4.4 | | M | Yes [] |
| TF5 | Pre-FEC distribution | 119.2.4.5 | Distribute the data to two FEC codewords | M | Yes [] |
| TF6 | Reed-Solomon encoder | 119.2.4.6 | RS(544,514) | M | Yes [] |
| TF7 | Symbol distribution | 119.2.4.7 | Distribution is based on 10b symbols | M | Yes [] |

119.6.4.2 Receive function

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---|--------|--------------------|
| RF1 | Skew tolerance | 119.2.5.1 | Maximum Skew of 180 ns between PCS lanes and a maximum Skew Variation of 4 ns | M | Yes [] |
| RF2 | Lane reorder and de-interleave | 119.2.5.2 | Order the PCS lanes according to the PCS lane number and de-interleave the FEC codewords | M | Yes [] |
| RF3 | Reed-Solomon decoder | 119.2.5.3 | Corrects any combination of up to $t=15$ symbol errors in a codeword | M | Yes [] |
| RF4 | Reed-Solomon decoder | 119.2.5.3 | Capable of indicating when a codeword was not corrected. | M | Yes [] |
| RF5 | Error monitoring while error indication is bypassed | 119.2.5.3 | When the number of symbol errors in a block of 8192 codewords exceeds 5560 assert hi_ser for 60 ms to 75 ms | BI:M | Yes [] N/A [] |
| RF6 | 256B/257B to 64B/66B transcoder | 119.2.5.7 | rx_coded, $j < 65:0$, $j=0$ to 3 constructed per 119.2.5.7 | M | Yes [] |

119.6.4.3 64B/66B coding rules

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-------------------------|---------------|--------|--------------------|
| C1 | Encoder (and ENCODE function) implements the code as specified | 119.2.3, 119.2.6.2.3 | | M | Yes [] |
| C2 | Decoder (and DECODE function) implements the code as specified | 119.2.3, 119.2.6.2.3 | | M | Yes [] |
| C3 | Only valid block types are transmitted | 119.2.3.2 | | M | Yes [] |
| C4 | Invalid block types are treated as an error | 119.2.3.2 | | M | Yes [] |
| C5 | Only valid control characters are transmitted | 119.2.3.3 | | M | Yes [] |
| C6 | Invalid control characters are treated as an error | 119.2.3.3 | | M | Yes [] |
| C7 | If EEE has not been negotiated, LPI is not transmitted | 119.2.3.3 | | EEE:M | Yes [] N/A [] |
| C8 | If EEE has not been negotiated, LPI is treated as an error if received | 119.2.3.3 | | EEE:M | Yes [] N/A [] |

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|--|--------|---------|
| C9 | Idles do not interrupt data | 119.2.3.5 | | M | Yes [] |
| C10 | IDLE control code insertion and deletion | 119.2.3.5 | Insertion or Deletion in groups of 8 /I/s | M | Yes [] |
| C11 | Sequence ordered set deletion | 119.2.3.8 | Only one whole ordered set of two consecutive sequence ordered sets may be deleted | M | Yes [] |

119.6.4.4 Scrambler and descrambler

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|-------------|-----------|-----------------------------------|--------|---------|
| S1 | Scrambler | 119.2.4.3 | Performs as shown in Figure 49–8 | M | Yes [] |
| S2 | Descrambler | 119.2.5.6 | Performs as shown in Figure 49–10 | M | Yes [] |

119.6.4.5 Alignment markers

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|----------------------------|-----------|---|--------|--------------------|
| AM1 | Alignment marker insertion | 119.2.4.4 | Alignment markers are inserted periodically as in 119.2.4.4 | M | Yes [] |
| AM2 | Alignment marker form | 119.2.4.4 | Alignment markers are formed as described in 119.2.4.4 | M | Yes [] |
| AM3 | Lane mapping | 119.2.6.3 | PCS lane number is captured | MD:M | Yes [] N/A [] |
| AM4 | Alignment marker removal | 119.2.5.5 | Alignment markers are removed prior to descrambling as described in 119.2.5.5 | M | Yes [] |

119.6.4.6 Test-pattern modes

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---------------|--------|---------|
| JT1 | Scrambled idle transmit test-pattern generator is implemented | 119.2.4.9 | | M | Yes [] |

119.6.4.7 Bit order

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--------------------|-----------|---|--------|---------|
| B1 | Transmit bit order | 119.2.4.8 | Placement of bits into the PCS lanes as shown in Figure 119–10 or Figure 119–11 | M | Yes [] |

119.6.4.8 Management

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|-----------------------------|--------|--------------------|
| M1 | Alternate access to PCS management object is provided | 119.3 | | O | Yes [] No [] |
| M2 | Mapping of MDIO control bits and MDIO status bits | 119.3.1 | Table 119–4 and Table 119–5 | MD:M | Yes [] N/A [] |

119.6.4.9 State diagrams

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-------------|--|----------|--------------------|
| SM1 | Alignment marker lock | 119.2.6 | Implements 8 alignment marker lock processes as depicted in Figure 119–12 | PCS200:M | Yes [] N/A [] |
| SM2 | Alignment marker lock | 119.2.6 | Implements 16 alignment marker lock processes as depicted in Figure 119–12 | PCS400:M | Yes [] N/A [] |
| SM3 | The SLIP function evaluates all possible block positions | 119.2.6.2.3 | | M | Yes [] |
| SM4 | PCS synchronization state diagram | 119.2.6 | Meets the requirements of Figure 119–13 | M | Yes [] |
| SM5 | Transmit process | 119.2.6 | Meets the requirements of Figure 119–14 | M | Yes [] |
| SM6 | Receive process | 119.2.6 | Meets the requirements of Figure 119–15 | M | Yes [] |

119.6.4.10 Loopback

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---------------|--------|---------|
| L1 | Supports loopback | 119.4 | | M | Yes [] |
| L2 | When in loopback, transmits what it receives from the 200GMII/400GMII | 119.4 | | M | Yes [] |
| L3 | When in loopback, ignore all data presented by the PMA sublayer. | 119.4 | | M | Yes [] |

119.6.4.11 Delay constraints

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|----------------------|-----------|---|--------------|--------------------|
| TIM1 | PCS delay constraint | 119.5 | No more than 160 256 BT for sum of transmit and receive path delays for 200GBASE-R. | PCS200: M | Yes [] N/A [] |
| TIM2 | PCS delay constraint | 119.5 | No more than 320 000 BT for sum of transmit and receive path delays for 400GBASE-R. | PCS400: M | Yes [] N/A [] |

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd.10:2019

120. Physical Medium Attachment (PMA) sublayer, type 200GBASE-R and 400GBASE-R

120.1 Overview

120.1.1 Scope

This clause specifies the Physical Medium Attachment sublayer (PMA) for the families of 200 Gb/s and 400 Gb/s Physical Layer implementations, known as 200GBASE-R and 400GBASE-R. The PMA allows the PCS (specified in Clause 119) to connect in a media-independent way with a range of physical media, or for the DTE XS to connect to the PHY XS (specified in Clause 118). The 200GBASE-R PMA(s) can support any of the 200 Gb/s PMDs in Table 116–1, and the 400GBASE-R PMA(s) can support any of the 400 Gb/s PMDs in Table 116–2. The terms 200GBASE-R and 400GBASE-R are used when referring generally to Physical Layers using the PMA defined in this clause.

200GBASE-R and 400GBASE-R can be extended to support any full duplex medium requiring only that the PMD be compliant with the appropriate PMA interface.

The interfaces for the inputs of the 200GBASE-R and 400GBASE-R PMAs are defined in an abstract manner and do not imply any particular implementation. For 200GBASE-R PMAs, electrical interfaces connecting PMA sublayers, known as 200GAUI-n, are defined in Annex 120B, Annex 120C, Annex 120D, and Annex 120E. For 400GBASE-R PMAs, electrical interfaces connecting PMA sublayers, known as 400GAUI-n, are defined in Annex 120B, Annex 120C, Annex 120D, and Annex 120E.

120.1.2 Position of the PMA in the 200GBASE-R and 400GBASE-R sublayers

Figure 120–1 shows the relationship of the PMA sublayer (shown shaded) with other sublayers to the ISO Open System Interconnection (OSI) reference model.

120.1.3 Summary of functions

The following is a summary of the principal functions implemented (when required) by the PMA in both the transmit and receive directions:

- a) Adapt the PCSL formatted signal to the appropriate number of abstract or physical lanes.
- b) Provide per input-lane clock and data recovery.
- c) Provide bit-level multiplexing.
- d) Provide clock generation.
- e) Provide signal drivers.
- f) Optionally provide local loopback to/from the PMA service interface.
- g) Optionally provide remote loopback to/from the PMD service interface.
- h) Optionally provide test-pattern generation and detection.
- i) Tolerate Skew Variation.
- j) Perform PAM4 encoding and decoding for 200GBASE-R PMAs where the number of physical lanes is 4, and for 400GBASE-R PMAs where the number of physical lanes is 4 or 8.

In addition, the PMA provides receive link status information in the receive direction.

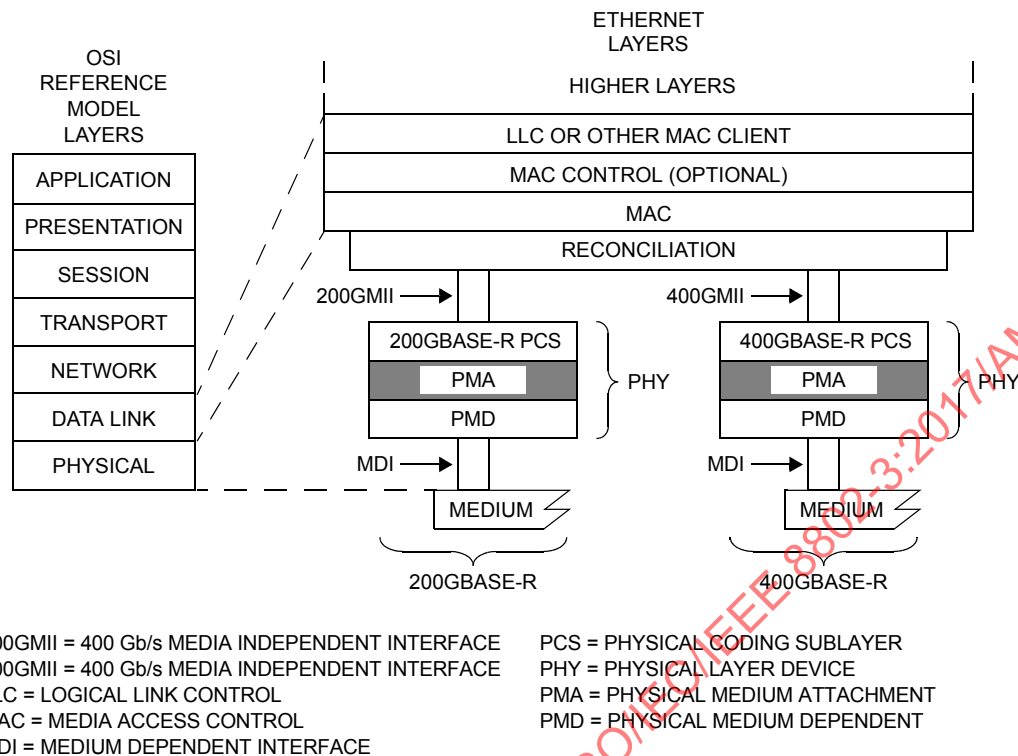


Figure 120-1—200GBASE-R and 400GBASE-R PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

120.1.4 PMA sublayer positioning

An implementation may use one or more PMA sublayers to adapt the number and rate of the PCS lanes to the number and rate of the PMD lanes. The number of PMA sublayers required depends on the partitioning of functionality for a particular implementation. An example is illustrated in Figure 120-2. Additional examples are illustrated in Annex 120A. Each PMA maps the PCSLs from p PMA input lanes to q PMA output lanes in the Tx direction, and from q PMA input lanes to p PMA output lanes in the Rx direction.

Management Data Input/Output (MDIO) Manageable Device (MMD) addresses 1, 8, 9, and 10 are available for addressing multiple instances of PMA sublayers (see Table 45-1 for MMD device addresses). If the PMA sublayer that is closest to the PMD is packaged with the PMD, it shares MMD 1 with the PMD. More addressable instances of PMA sublayers, each one separated from lower addressable instances by chip-to-chip interfaces, may be implemented and addressed allocating MMD addresses to PMAs in increasing numerical order going from the PMD toward the MAC. The example shown in Figure 120-2 could be implemented with three addressable instances: MMD 1 addressing the lowest PMA sublayer (co-packaged with the PMD), MMD 8 addressing the PMA sublayer above the 200GAUI-4 below the 200GAUI-8 or above the 400GAUI-8 below the 400GAUI-16, and MMD 9 addressing the PMA sublayer below the PCS.

The number of input lanes and the number of output lanes for a PMA are always divisors of the number of PCSLs. For PMA sublayers supporting 200GBASE-R PMDs, the number of PCSLs is 8. For PMA sublayers supporting 400GBASE-R PMDs, the number of PCSLs is 16.

IEEE Std 802.3bs-2017
Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
for 200 Gb/s and 400 Gb/s Operation

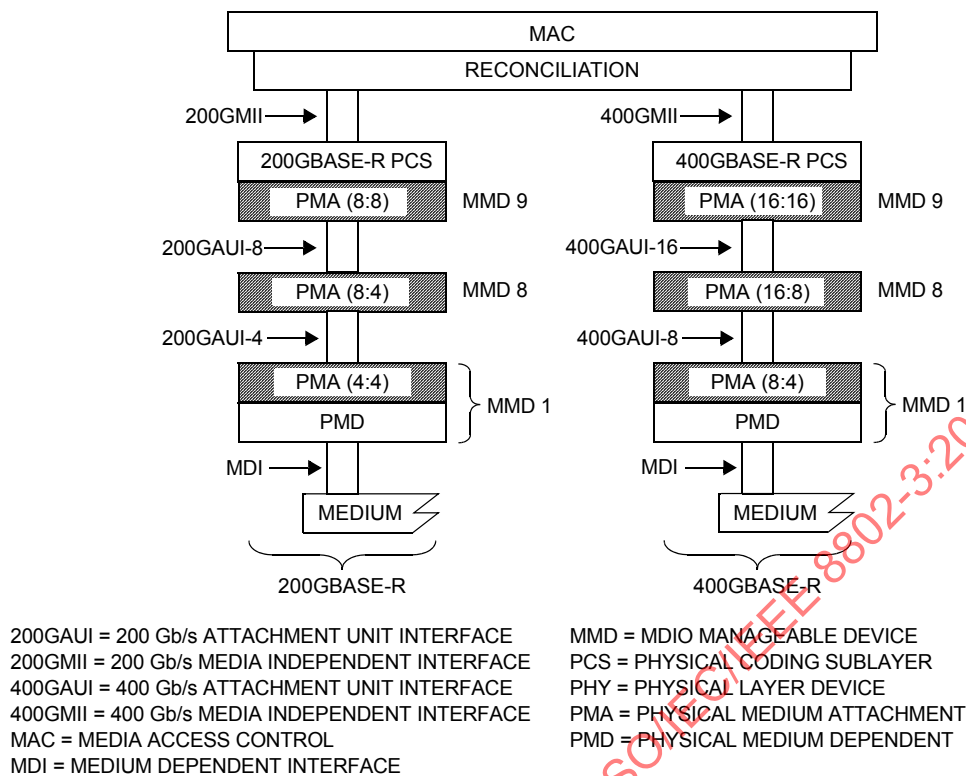


Figure 120-2—Example 200GBASE-R and 400GBASE-R PMA layering

The following guidelines apply to the partitioning of PMAs:

- The inter-sublayer service interface, defined in 116.3.1, is used for the PMA service interfaces supporting a flexible architecture with multiple PMA sublayers.
 - An instance of this interface can only connect service interfaces with the same number of lanes, where the lanes operate at the same rate.
- 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.
 - As physical instantiations, these define electrical and timing specification as well as requiring a receive re-timing function.
 - 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.
 - 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.
- The abstract inter-sublayer service interface can be physically instantiated as a 200GAUI-n or 400GAUI-n, using associated PMAs to map to the appropriate number of lanes.
- Opportunities for optional test-pattern generation, optional test-pattern detection, optional local loopback and optional remote loopback are dependent upon the location of the PMA sublayer in the implementation. See Figure 120-5.

- e) A minimum of one PMA sublayer is required in a PHY.
- f) A maximum of four PMA sublayers are addressable as MDIO MMDs.

120.2 PMA interfaces

All PMA variants for 200GBASE-R and 400GBASE-R are based on a generic specification of a bit mux function that applies to all input/output lane counts and each direction of transmission. Each direction of transmission may employ one or more such bit muxes to adapt from the appropriate number of input lanes to the appropriate number of output lanes as illustrated in Figure 120–3.




Figure 120–3—PMA bit mux used in both Tx and Rx directions

Conceptually, the PMA bit mux operates in one direction of transmission by demultiplexing PCSs from m PMA input lanes and remultiplexing them into n PMA output lanes. The mapping of PCSs from input to output lanes is not specified. See 120.5.2 and Figure 120–4 for details.

Figure 120–5 provides the functional block diagram of a PMA. The parameters of a PMA include the following:

- The numbers of input and output lanes in each direction.
- Whether the PMA is adjacent to a physically instantiated interface (200GAUI- n or 400GAUI- n above or below).
- Whether the PMA is adjacent to the PCS or DTE XS.
- Whether the PMA is adjacent to the PMD or PHY XS.

120.3 PMA service interface

The PMA service interface for 200GBASE-R and 400GBASE-R is an instance of the inter-sublayer service interface defined in 116.3.1. The PMA service interface primitives are summarized as follows:

PMA:IS_UNITDATA_ i .request(tx_symbol)
 PMA:IS_UNITDATA_ i .indication(rx_symbol)
 PMA:IS_SIGNAL.indication(SIGNAL_OK)

For a PMA service interface with p lanes, the primitives are defined for $i = 0$ to $p - 1$.

If the PMA client is the PCS or DTE XS, the PMA continuously sends eight (for 200GBASE-R) or sixteen (for 400GBASE-R) parallel bit streams to the PCS, each at the nominal signaling rate of the PCSL.

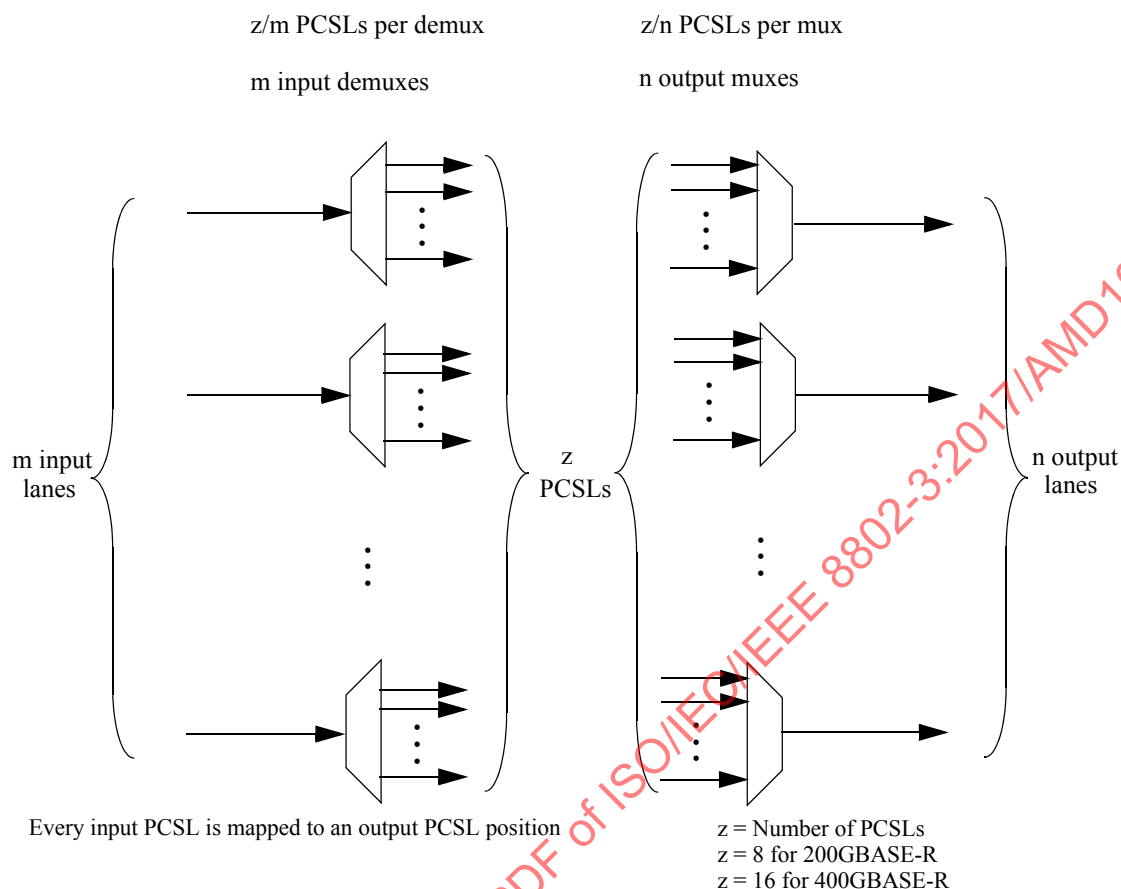
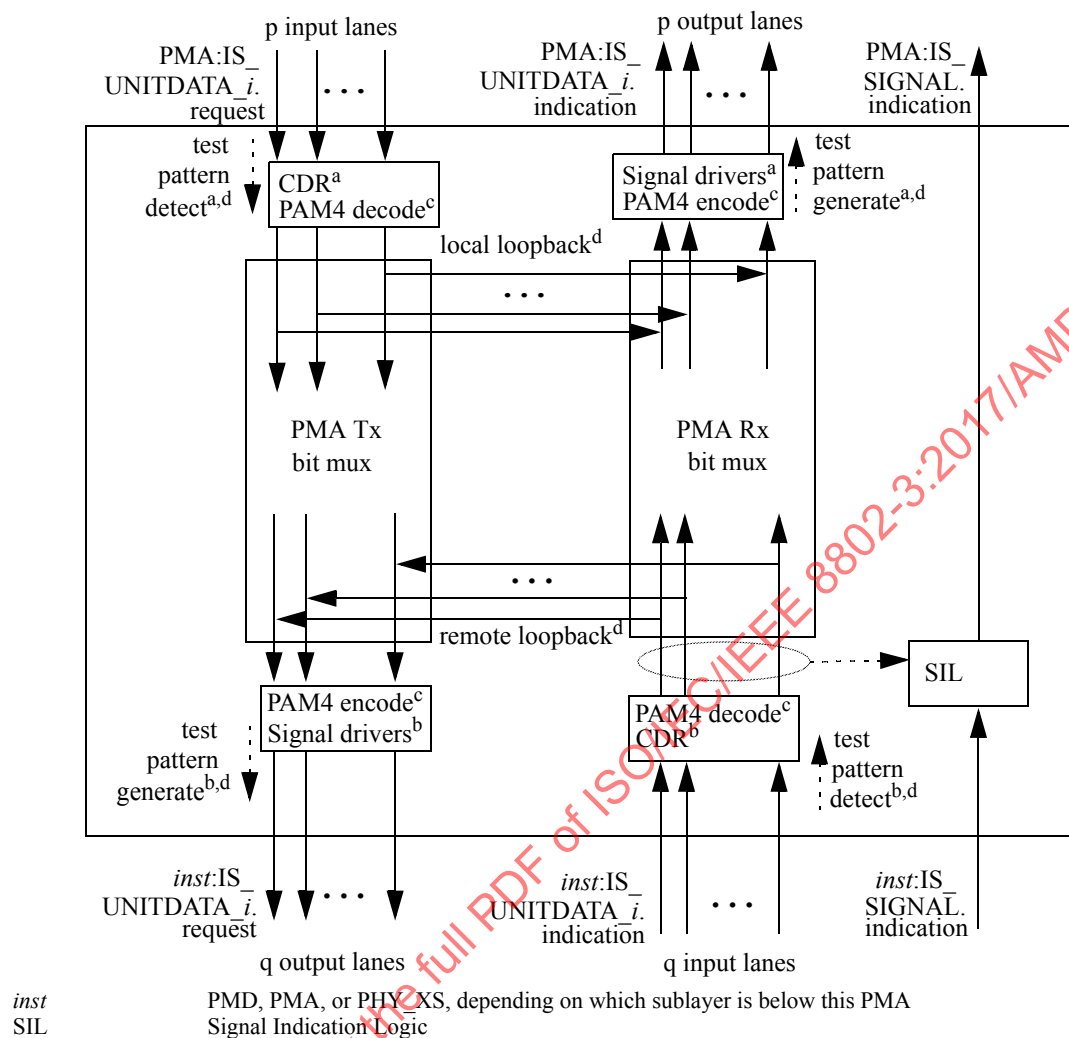


Figure 120-4—PMA bit mux operation used in both Tx and Rx directions

If a PMA client is another PMA, a 200GBASE-R PMA with $p=8$ or a 400GBASE-R PMA with $p=16$ physical input lanes receives bits from one PCSL on each of its input lanes. A 200GBASE-R PMA with $p=4$ physical input lanes receives PAM4 symbols on each of its input lanes at the PCSL rate, each symbol formed from two bits. A 400GBASE-R PMA with $p=4$ or $p=8$ physical input lanes receives PAM4 symbols on each of its input lanes at $8/p$ times the PCSL rate, each symbol formed from two bits. The bit stream represented by the input symbols carries z/p bit-multiplexed PCSLs on each physical input lane, where $z=8$ for 200GBASE-R and $z=16$ for 400GBASE-R. Skew may exist between different PCSLs received on the same physical input lane even though all PCSLs originate from the same synchronous source, so there is independence of arrival of bits from each PCSL on each physical input lane.

In the Tx direction, if the symbol from a PMA:IS_UNITDATA_ i .request primitive is received over a physically instantiated interface (200GAUI-n or 400GAUI-n), clock and data are recovered on the lane receiving the symbol. If necessary, PAM4 symbols are converted to pairs of bits. Bits are routed through the PMA to an output lane through a process that may demultiplex PCSLs from the input, perform any necessary buffering to tolerate Skew Variation across input lanes, and multiplex PCSLs to output lanes. If necessary, pairs of bits on output lanes are converted to PAM4 symbols. The symbols are sent on an output lane to the sublayer below using the *inst*:IS_UNITDATA_ k .request (k not necessarily equal to i) primitive (see 120.4).



^a If physically instantiated interface (200GAUI-n or 400GAUI-n) immediately above this PMA.

^b If physically instantiated interface (200GAUI-n or 400GAUI-n) immediately below this PMA, or if this is the closest PMA to the PMD.

^c If number of input or output lanes is 4 for a 200GBASE-R PMA, or 4 or 8 for a 400GBASE-R PMA.

^d Optional.

Figure 120-5—PMA Functional Block Diagram

The PMA passes symbols from the input lanes to the output lanes in the Rx direction when data is being received from every input lane from the sublayer below the PMA that has a PCSL that is routed to a particular output lane at the PMA service interface, and (if necessary), buffers are filled to allow tolerating the Skew Variation that may appear between the input lanes. PCSLs are demultiplexed from the input lanes, remultiplexed to the output lanes, and symbols are transferred over each output lane to the PMA client via the PMA:IS_UNITDATA_i.indication primitive. If necessary, PAM4 symbols are converted to pairs of bits on the input lanes and/or pairs of bits are converted to PAM4 symbols on the output lanes.

The PMA:IS_SIGNAL.indication primitive is generated through a set of Signal Indication Logic (SIL) that reports signal health based on receipt of the *inst:IS_SIGNAL.indication* from the sublayer below, data being received on all of the input lanes from the sublayer below, buffers filled (if necessary) to accommodate Skew Variation, and symbols being sent to the PMA client on all of the output lanes. When these conditions

are met, the SIGNAL_OK parameter sent to the PMA client via the PMA:IS_SIGNAL.indication primitive has the value OK. Otherwise, the SIGNAL_OK primitive has the value FAIL. In the case where the sublayer below the PMA is a PHY XS the PMA does not receive a PHY_XS:IS_SIGNAL.indication as an input to the SIL.

120.4 Service interface below PMA

Since the architecture supports multiple PMA sublayers for various PMD lane counts and device partitioning, there are several different sublayers that may appear below a PMA, including the PMD, an extender sublayer, or another PMA. The variable *inst* represents whichever sublayer appears below the PMA (e.g., another PMA or PMD).

The sublayer below the PMA utilizes the inter-sublayer service interface defined in 116.3. The service interface primitives provided to the PMA are summarized as follows:

```
inst:IS_UNITDATA_i.request(tx_symbol)
inst:IS_UNITDATA_i.indication(rx_symbol)
inst:IS_SIGNAL.indication(SIGNAL_OK)
```

The number of lanes *q* for the service interface matches the number of lanes expected by the PMA. The *inst:IS_UNITDATA_i* primitives are defined for each lane *i* = 0 to *q* - 1 of the service interface below the PMA. Note that electrical and timing specifications of the service interface are defined if the interface is physically instantiated (e.g., 200GAUI-n or 400GAUI-n), otherwise the service interface is specified only abstractly. The interface between the PMA and the sublayer below consists of *q* lanes for data transfer and a status indicating a good signal from the sublayer below the PMA (see Figure 120-5).

The PMA transfers symbols from the input lanes to the output lanes in the Tx direction when data is being received via the PMA:IS_UNITDATA_i.request primitive from every input lane from the PMA client at the PMA service interface (see 120.3) that has a PCSL that is routed to this output lane, and (if necessary), buffers are filled to provide the ability to tolerate the Skew Variation that may appear between the input lanes from the PMA client. PCSLs are demultiplexed from the input lanes, remultiplexed to the output lanes, and symbols are transferred over each output lane to the sublayer below the PMA.

In the Rx direction, if the symbol is received over a physically instantiated interface (200GAUI-n, 400GAUI-n, or physically instantiated PMD service interface), clock and data are recovered on the lane receiving the symbol. If necessary, PAM4 symbols received on the input lanes are converted to pairs of bits. The bits are routed through the PMA to an output lane toward the PMA client through a process that may demultiplex PCSLs from the input, perform any necessary buffering to tolerate Skew Variation across input lanes, and multiplex PCSLs to output lanes. If necessary, pairs of bits are converted to PAM4 symbols on the output lanes. Each symbol is sent on an output lane to the PMA client using the PMA:IS_UNITDATA_k.indication (*k* not necessarily equal to *i*) primitive at the PMA service interface.

In the case where the sublayer below the PMA is a PHY XS, there is an additional primitive:

```
PHY_XS:IS_SIGNAL.request(SIGNAL_OK)
```

The PHY_XS:IS_SIGNAL.request primitive is generated through a set of SIL that reports signal health based on data being received on all of the input lanes from the sublayer above, buffers filled (if necessary) to accommodate Skew Variation, and symbols being sent to the PHY XS on all of the output lanes. When these conditions are met, the SIGNAL_OK parameter sent to the PHY XS via the PHY_XS:IS_SIGNAL.request primitive has the value OK. Otherwise, the SIGNAL_OK primitive has the value FAIL.

120.5 Functions within the PMA

The purpose of the PMA is to adapt the PCSL formatted signal to an appropriate number of abstract or physical lanes, to recover clock from the received signal (if appropriate), and optionally to provide test signals and loopback. Each input (Tx direction) or output (Rx direction) lane between the PMA and the PMA client carries one or more PCSLs that are bit-multiplexed. All input and output lanes between the PMA and the PMA client carry the same number of PCSLs and operate at the same nominal signaling rate. Likewise, each input (Rx direction) or output (Tx direction) lane between the PMA and the sublayer below the PMA carries one or more PCSLs that are bit-multiplexed. All input and output lanes between the PMA and the sublayer below the PMA carry the same number of PCSLs and operate at the same nominal signaling rate. As described in 120.1.4, the number of input lanes and the number of output lanes for a given PMA are divisors of 8 for 200GBASE-R, or 16 for 400GBASE-R, which are the number of PCSLs for the respective PHYs. The symbols received from input lanes or sent on output lanes are bits for 200GBASE-R interfaces where the number of lanes is 8, or 400GBASE-R interfaces where the number of lanes is 16. The symbols received from input lanes or sent on output lanes are PAM4 symbols for 200GBASE-R interfaces where the number of lanes is 4, or 400GBASE-R interfaces where the number of lanes is 8 or 4.

120.5.1 Per input-lane clock and data recovery

If the interface between the PMA client and the PMA is physically instantiated as 200GAUI-n or 400GAUI-n, the PMA shall meet the electrical and timing specifications in Annex 120B, Annex 120C, Annex 120D, or Annex 120E as appropriate. If the interface between the sublayer below the PMA and the PMA is physically instantiated as 200GAUI-n or 400GAUI-n, the PMA shall meet the electrical and timing specifications at the service interface below the PMA as specified in Annex 120B, Annex 120C, Annex 120D, or Annex 120E as appropriate.

Test patterns that are intended for transmitter testing, such as square wave, may not be correctly recovered by an adjacent PMA.

120.5.2 Bit-level multiplexing

The PMA provides bit-level multiplexing in both the Tx and Rx directions. In the Tx direction, the function is performed among the bits received from the PMA client via the PMA:IS_UNITDATA_i.request primitives (for PMA client lanes $i = 0$ to $p - 1$) with the result sent to the service interface below the PMA using the inst:IS_UNITDATA_i.request primitives (for service interface lanes $i = 0$ to $q - 1$), referencing the functional block diagram shown in Figure 120–5. The bit multiplexing behavior is illustrated in Figure 120–4.

The aggregate signal carried by the group of input lanes or the group of output lanes is arranged as a set of PCSLs. The number of PCSLs z is 8 for 200GBASE-R interfaces and 16 for 400GBASE-R interfaces. The nominal bit rate R of each PCSL is 26.5625 Gb/s.

For a PMA with m input lanes (Tx or Rx direction), each input lane carries, bit multiplexed, z/m PCSLs. Each input lane has a nominal bit rate of $26.5625 \times z/m$ Gb/s. Note that the signaling (Baud) rate is equal to the bit rate when the number of physical lanes is 8 for 200GBASE-R or 16 for 400GBASE-R (bits are sent or received on the lanes). The Baud rate is equal to half of the bit rate when the number of physical lanes is 4 for 200GBASE-R or the number of physical lanes is 8 or 4 for 400GBASE-R (PAM4 symbols are sent or received on the lanes). If necessary, PAM4 symbols are converted to pairs of bits on the input lanes and/or pairs of bits are converted to PAM4 symbols on the output lanes. If bit x received on an input lane belongs to a particular PCSL, the next bit of that same PCSL is received on the same input lane at bit position $x + (z/m)$. The z/m PCSLs may arrive in any sequence on a given input lane.

For a PMA with n output lanes (Tx or Rx direction), each output lane carries, bit multiplexed, z/n PCSLs. Each output lane has a nominal signaling rate of $26.5625 \times z/n$ Gb/s. Each PCSL is mapped from a position

in the sequence on one of the m input lanes to a position in the sequence on one of the n output lanes. If bit x sent on an output lane belongs to a particular PCSL, the next bit of that same PCSL is sent on the same output lane at bit position $x + (z/n)$. The PMA shall maintain the chosen sequence of PCSLs on all output lanes while it is receiving a valid stream of bits on all input lanes.

Each PCSL received in any temporal position on an input lane is transferred into a temporal position on an output lane. As the PCS (see Clause 119) has fully flexible receive logic, an implementation is free to perform the mapping of PCSLs from input lanes to output lanes without constraint. Figure 120–6 illustrates one possible bit ordering for a 400GBASE-R 8:4 PMA bit mux. Other bit orderings are also valid.

Note that since the number of input lanes and output lanes for a 200GBASE-R or 400GBASE-R PMA is always a power of two, many PMAs converting between different numbers of lanes normally simply multiplex two or four input lanes to one output lane, or demultiplex two or four output lanes from one input lane. However, any PMA implementation which produces an allowable order of bits from all PCSLs on the output lanes is valid.

NOTE—PMA output lanes composed of some specific combinations of four PCSLs with specific skew offsets (e.g., 400GBASE-R PCSLs 0, 2, 4, and 10 with delays 0, 1, 0, and 2 bits, respectively) may have reduced transition density.

120.5.3 Skew and Skew Variation

The Skew (relative delay) between the PCSLs must be kept within limits so that the information on the lanes can be reassembled by the PCS.

Any PMA that combines PCSLs from different input lanes onto the same output lane must tolerate Skew Variation between the input lanes without changing the PCSL positions on the output. Skew and Skew Variation are defined in 116.5. The limits for Skew and Skew Variation at physically instantiated interfaces are specified at Skew points SP1 and SP2 in the transmit direction and SP5 and SP6 in the receive direction as defined in 116.5 and illustrated in Figure 116–4 and Figure 116–5.

120.5.3.1 Skew generation toward SP1

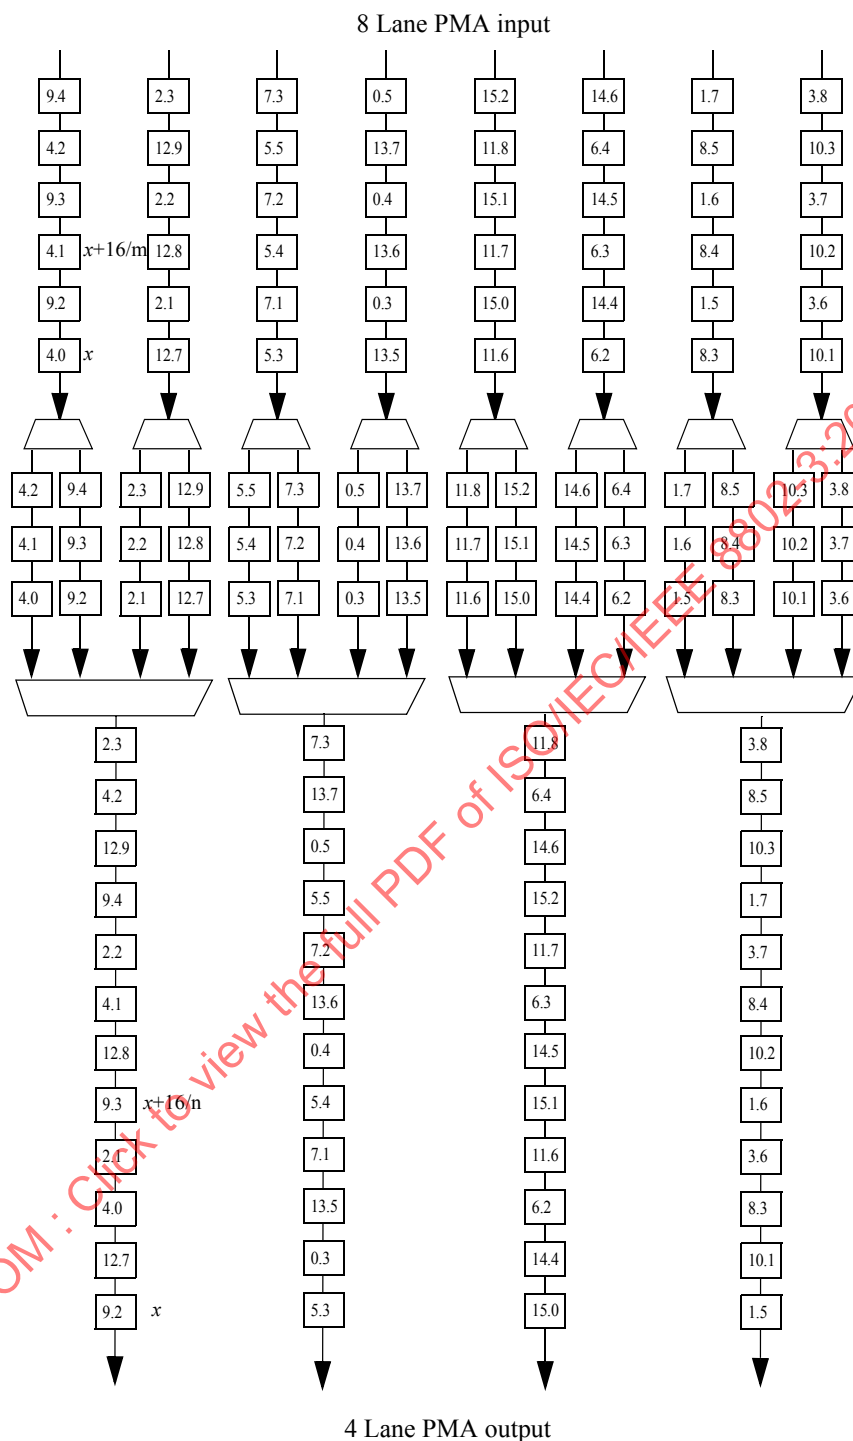
In an implementation with one or more physically instantiated 200GAUI-n or 400GAUI-n interfaces, the PMA that sends data in the transmit direction toward the 200GAUI-n or 400GAUI-n that is closest to the PMD (SP1 in Figure 116–4 and Figure 116–5) shall produce no more than 29 ns of Skew between PCSLs toward the 200GAUI-n or 400GAUI-n, and no more than 200 ps of Skew Variation.

120.5.3.2 Skew tolerance at SP1

In an implementation with one or more physically instantiated 200GAUI-n or 400GAUI-n interfaces, the PMA service interface that receives data in the transmit direction from the 200GAUI-n or 400GAUI-n (SP1 in Figure 116–4 and Figure 116–5) shall tolerate the maximum amount of Skew Variation allowed at SP1 (200 ps) between input lanes while maintaining the bit ordering and position of each PCSL on each PMA lane in the transmit direction (toward the PMD).

120.5.3.3 Skew generation toward SP2

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, the PMA adjacent to the PMD service interface shall generate no more than 43 ns of Skew, and no more than 400 ps of Skew Variation between output lanes toward the PMD service interface (SP2 in Figure 116–4 and Figure 116–5). If there is a physically instantiated 200GAUI-n or 400GAUI-n as well, then the Skew measured at SP1 is limited to no more than 29 ns of Skew and no more than 200 ps of Skew Variation.



NOTE—i.k indicates bit k on PCSL i. Skew may exist between PCSLs.

Figure 120–6—Example 8:4 400GBASE-R PMA bit mux

120.5.3.4 Skew tolerance at SP5

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, the PMA adjacent to the PMD service interface (SP5) shall tolerate the maximum amount of Skew Variation allowed at SP5 (3.6 ns) between output lanes from the PMD service interface while maintaining the bit ordering and position of each PCSL on each PMA lane in the receive direction (toward the MAC).

120.5.3.5 Skew generation at SP6

In an implementation with one or more physically instantiated 200GAUI-n or 400GAUI-n interfaces, at SP6 (the receive direction of the 200GAUI-n or 400GAUI-n closest to the PCS), the PMA or group of PMAs between the PMD and the 200GAUI-n or 400GAUI-n closest to the PCS shall deliver no more than 160 ns of Skew, and no more than 3.8 ns of Skew Variation between output lanes toward the 200GAUI-n or 400GAUI-n in the Rx direction. If there is a physically instantiated PMD service interface that allows the Skew to be measured, the Skew measured at SP5 is limited to no more than 145 ns of Skew and no more than 3.6 ns of Skew Variation. If there is no physically instantiated PMD service interface, the Skew measured at SP4 is limited to no more than 134 ns of Skew, and no more than 3.4 ns of Skew Variation.

120.5.3.6 Skew tolerance at SP6

In an implementation with one or more physically instantiated 200GAUI-n or 400GAUI-n interfaces, the PMA between the 200GAUI-n or 400GAUI-n closest to the PCS and the PCS shall tolerate the maximum amount of Skew Variation allowed at SP6 (3.8 ns) between input lanes while maintaining the bit order and position of PCSLs on lanes sent in the receive direction towards the MAC.

120.5.4 Delay constraints

The maximum cumulative delay contributed by up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 120–1. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 116.4 and its references.

Table 120–1—Delay constraints

| Sublayer | Maximum (bit time) | Maximum (pause_quanta) | Maximum (ns) |
|----------------|-----------------------|---------------------------|--------------|
| 200GBASE-R PMA | 18 432 | 36 | 92.16 |
| 400GBASE-R PMA | 36 864 | 72 | 92.16 |

120.5.5 Clocking architecture

A PMA with m input lanes and n output lanes shall clock the output lanes such that the bit rate of the output lanes is m/n times the bit rate of the input lanes. For input or output lanes encoded as PAM4 (for 200GBASE-R, where the number of input or output lanes is 4, or for 400GBASE-R, where the number of input or output lanes is 4 or 8), the symbol rate is half of the bit rate. This applies in both the Tx and Rx directions of transmission. In the case where the interfaces between the PMA client and the PMA and/or the PMA and the sublayer below the PMA are physically instantiated, the PMA may derive its input clock(s) from the electrical interface on one or more of the input lanes, and generate the output clock(s) with an appropriate PLL multiplier/divider circuit.

There is no requirement that the PMA clock all output lanes in unison. Examples of independent clocking of output lanes include the following:

- The case where the number of input and output lanes are equal (the PMA is provided for retiming and regeneration of the signal). This may be implemented without any rearrangement of PCSs between input lanes and output lanes (although rearrangements are allowed), and such a PMA may be implemented by driving each output lane using the clock recovered from the corresponding input lane.
- Since the number of PCSs is a power of 2, whenever the number of input and output lanes are unequal, each input lane could be mapped to multiple output lanes (in the case where the number of output lanes is greater), or multiple input lanes could be mapped to a single output lane.

120.5.6 Signal drivers

For cases where the interface between the PMA client and the PMA, or between the PMA and the sublayer below the PMA represent a physically instantiated interface, the PMA provides electrical signal drivers for that interface. The electrical and jitter/timing specifications for these interfaces appear in

- Annex 120B, which specifies the 200GAUI-8 and 400GAUI-16 interfaces for chip-to-chip applications.
- Annex 120C, which specifies the 200GAUI-8 and 400GAUI-16 interfaces for chip-to-module applications.
- Annex 120D, which specifies the 200GAUI-4 and 400GAUI-8 interfaces for chip-to-chip applications.
- Annex 120E, which specifies the 200GAUI-4 and 400GAUI-8 interfaces for chip-to-module applications.

For 200GAUI-8 or 400GAUI-16, the modulation format is NRZ. For 200GAUI-4 or 400GAUI-8, the modulation format is PAM4.

120.5.7 Gray mapping for PAM4 encoded lanes

For output lanes encoded as PAM4 (for 200GBASE-R, where the number of output lanes is 4, or for 400GBASE-R, where the number of output lanes is 4 or 8), the PMA transmit process shall map consecutive pairs of bits {A, B}, where A is the bit arriving first, to a Gray-coded symbol as follows:

{0, 0} maps to 0,
 {0, 1} maps to 1,
 {1, 1} maps to 2, and
 {1, 0} maps to 3.

For input lanes encoded as PAM4 (for 200GBASE-R, where the number of input lanes is 4, or for 400GBASE-R, where the number of input lanes is 4 or 8), the PMA receive process shall map Gray-coded PAM4 symbols to pairs of bits {A, B} where A is considered to be the first bit as follows:

0 maps to {0, 0},
 1 maps to {0, 1},
 2 maps to {1, 1}, and
 3 maps to {1, 0}.

120.5.8 Link status

The PMA shall provide link status information to the PMA client using the PMA:IS_SIGNAL.indication primitive. The PMA continuously monitors the link status reported by the service interface below from the

inst:IS_SIGNAL.indication primitive, and uses this as input to Signal Indication Logic (SIL) to determine the link status to report to the layer above. Other inputs to the SIL may include the status of clock and data recovery on the lanes from the service interface below the PMA and whether buffers/FIFOs have reached the required fill level to accommodate Skew Variation so that data is being sent on the output lanes.

120.5.9 PMA local loopback mode (optional)

PMA local loopback mode is optional. If it is implemented, it shall be as described in this subclause (120.5.9).

The PMA sublayer may provide a local loopback function. The function involves looping back each input lane to the corresponding output lane. Each bit received from the PMA:IS_UNITDATA_0.request(tx_symbol) primitive is looped back in the direction of the MAC using the PMA:IS_UNITDATA_0.indication(rx_symbol) primitive.

During local loopback, the PMA performs normal bit muxing of PCSs per 120.5.2 onto the lanes in the Tx direction toward the service interface below the PMA.

Ability to perform this function is indicated by the Local_loopback_ability status variable. If a Clause 45 MDIO is implemented, this variable is accessible through bit 1.8.0 (45.2.1.7.15). A device is placed in local loopback mode when the Local_loopback_enable control variable is set to one, and removed from local loopback mode when this variable is set to zero. If a Clause 45 MDIO is implemented, this variable is accessible through PMA/PMD control 1 register (bit 1.0.0, see 45.2.1.1.5).

120.5.10 PMA remote loopback mode (optional)

PMA remote loopback mode is optional. If implemented, it shall be as described in this subclause (120.5.10).

Remote loopback, if provided, should be implemented in a PMA sublayer close enough to the PMD to maintain the bit sequence on each individual PMD lane. When remote loopback is enabled, each symbol received over a lane of the service interface below the PMA via *inst:IS_UNITDATA_0.indication* is looped back to the corresponding output lane toward the PMD via *inst:IS_UNITDATA_0.request*. Note that the service interface below the PMA can be provided by the PHY XS, PMD, or another PMA sublayer.

During remote loopback, the PMA performs normal bit muxing of PCSs per 120.5.2 onto the lanes in the Rx direction towards the PMA client.

The ability to perform this function is indicated by the 200G_Remote_loopback_ability and 400G_Remote_loopback_ability status variables for the 200GBASE-R PMA and 400GBASE-R PMA, respectively. If a Clause 45 MDIO is implemented, the 200G_Remote_loopback_ability and 400G_Remote_loopback_ability variables are accessible through bit 1.23.15 (45.2.1.14e.1) and bit 1.24.15 (45.2.1.14f.1), respectively. A device is placed in remote loopback mode when the Remote_loopback_enable control variable is set to one, and removed from remote loopback mode when this variable is set to zero. If a Clause 45 MDIO is implemented, this variable is accessible through PMA/PMD Control register 1 (bit 1.0.1, see 45.2.1.1.4).

120.5.11 PMA test patterns (optional)

Where the output lanes of the PMA appear on a physically instantiated interface 200GAUI-n, 400GAUI-n or the PMD service interface (whether or not it is physically instantiated), the PMA may optionally generate and detect test patterns. These test patterns are used to test adjacent layer interfaces for an individual PMA sublayer or to perform testing between a physically instantiated interface of a PMA sublayer and external testing equipment.

120.5.11.1 Test patterns for NRZ encoded signals

For a 200GBASE-R PMA with 8 output lanes or a 400GBASE-R PMA with 16 output lanes using 2-level NRZ encoding, the test patterns in this clause may be supported.

The ability to generate each of the respective test patterns in each direction of transmission are indicated by the `PRBS9_Tx_generator_ability`, `PRBS9_Rx_generator_ability`, `PRBS31_Tx_generator_ability`, `PRBS31_Rx_generator_ability`, and `Square_wave_ability` status variables, which if a Clause 45 MDIO is implemented are accessible through bits 1.1500.5, 1.1500.4, 1.1500.3, 1.1500.1, and 1.1500.12, respectively (see 45.2.1.123).

120.5.11.1.1 PRBS31 test pattern

The ability to check PRBS31 test patterns in each direction of transmission are indicated by the `PRBS31_Tx_checker_ability` and `PRBS31_Rx_checker_ability` status variables, which if a Clause 45 MDIO is implemented, are accessible through bits 1.1500.2 and 1.1500.0, respectively (see 45.2.1.123).

If supported, when send Tx PRBS31 test pattern is enabled by the `PRBS31_enable` and `PRBS_Tx_gen_enable` control variables, the PMA shall generate a PRBS31 pattern (as defined in 49.2.8) on each of the lanes toward the service interface below the PMA via the `inst:IS_UNITDATA_i.request` primitive. To avoid correlated crosstalk, it is highly recommended that the PRBS31 patterns generated on each lane be generated from independent, random seeds or at a minimum offset of 20 000 UI between the PRBS31 sequence on any lane and any other lane. When send Tx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 120.5.2. If a Clause 45 MDIO is implemented, the `PRBS31_enable` and `PRBS_Tx_gen_enable` control variables are accessible through bits 1.1501.7 and 1.1501.3 (see 45.2.1.124).

If supported, when send Rx PRBS31 test pattern is enabled by the `PRBS31_enable` and `PRBS_Rx_gen_enable` control variables, the PMA shall generate a PRBS31 pattern on each of the lanes toward the PMA client via the `PMA:IS_UNITDATA_i.indication` primitive. While this test pattern is enabled, the PMA also generates `PMA:IS_SIGNAL.indication(SIGNAL_OK)` toward the PMA client independent of the link status at the service interface below the PMA. When send Rx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 120.5.2. If a Clause 45 MDIO is implemented, the `PRBS31_enable` and `PRBS_Rx_gen_enable` control variables are accessible through bits 1.1501.7 and 1.1501.1 (see 45.2.1.124).

If supported, when check Tx PRBS31 test pattern mode is enabled by the `PRBS31_enable` and `PRBS_Tx_check_enable` control variables, the PMA shall check for the PRBS31 pattern on each of the lanes received from the PMA client via the `PMA:IS_UNITDATA_i.request` primitive. The checker shall increment the test-pattern error counter by one for each incoming bit error in the PRBS31 pattern for isolated single bit errors. Implementations should be capable of counting at least one error whenever one or more errors occur in a sliding 1000-bit window. If a Clause 45 MDIO is implemented, the `PRBS31_enable` and `PRBS_Tx_check_enable` control variables are accessible through bits 1.1501.7 and 1.1501.2 (see 45.2.1.124). The Tx test-pattern error counters `Ln0_PRBS_Tx_test_err_counter` through `Ln15_PRBS_Tx_test_err_counter` count, per lane, errors in detecting the PRBS31 pattern on the lanes from the PMA client. If a Clause 45 MDIO is implemented, these counters are accessible through registers 1.1600 through 1.16015 (see 45.2.1.126). When check Tx PRBS31 test pattern is disabled, the PMA expects normal traffic and test-pattern error counting does not continue. While in check Tx PRBS31 test-pattern mode, bit multiplexing continues as described in 120.5.2. Note that bit multiplexing of per-lane PRBS31 may produce a signal which is not meaningful for downstream sublayers.

If supported, when check Rx PRBS31 test-pattern mode is enabled by the `PRBS31_enable` and `PRBS_Rx_check_enable` control variables, the PMA shall check for the PRBS31 pattern on each of the lanes received from the service interface below the PMA via the `inst:IS_UNITDATA_i.indication` primitive.

If a Clause 45 MDIO is implemented, the PRBS31_enable and PRBS_Rx_check_enable control variables are accessible through bits 1.1501.7 and 1.1501.0 (see 45.2.1.124). The Rx test-pattern error counters Ln0_PRBS_Rx_test_err_counter through Ln15_PRBS_Rx_test_err_counter count, per lane, errors in detecting the PRBS31 pattern on the lanes from the service interface below the PMA. If a Clause 45 MDIO is implemented, these counters are accessible through registers 1.1700 through 1.1715 (see 45.2.1.127). While in check Rx PRBS31 test-pattern mode, the PMA:IS_SIGNAL.indication primitive does not indicate a valid signal. When check Rx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 120.5.2.

120.5.11.1.2 PRBS9 test pattern

If supported, when send Tx PRBS9 test-pattern mode is enabled by the PRBS9_enable and Tx_gen_enable control variables, the PMA shall generate a PRBS9 pattern (as defined in Table 68–6) on each lane toward the service interface below the PMA via the *inst:IS_UNITDATA_i.request* primitive. If a Clause 45 MDIO is implemented, the PRBS9_enable and Tx_gen_enable control variables are accessible through bits 1.1501.6 and 1.1501.3 (see 45.2.1.124). When send Tx PRBS9 test-pattern mode is disabled, the PMA returns to normal operation performing bit multiplexing as described in 120.5.2.

If supported, when send Rx PRBS9 test-pattern mode is enabled by the PRBS9_enable and Rx_gen_enable control variables, the PMA shall generate a PRBS9 pattern on each lane toward the PMA client via the PMA:IS_UNITDATA_i.indication primitive. The PMA also generates PMA:IS_SIGNAL.indication(SIGNAL_OK) toward the PMA client independent of the link status at the service interface below the PMA. If a Clause 45 MDIO is implemented, the PRBS9_enable and Rx_gen_enable control variables are accessible through bits 1.1501.6 and 1.1501.1 (see 45.2.1.124). When send Rx PRBS9 test-pattern mode is disabled, the PMA returns to normal operation performing bit multiplexing as described in 120.5.2.

Note that PRBS9 is intended to be checked by external test gear, and no PRBS9 checking function is provided within the PMA.

120.5.11.1.3 Square wave test pattern

Transmit square wave test-pattern mode optionally applies to each lane of the Tx direction PMA towards a physically instantiated 200GAUI-8, 400GAUI-16 or towards an NRZ PMD service interface whether or not it is physically instantiated. The ability to perform this function is indicated by the Square_wave_ability status variable. If a Clause 45 MDIO is implemented, the Square_wave_ability status variable is accessible through the Square wave test ability bit 1.1500.12 (see 45.2.1.123). If implemented, the transmit square wave test-pattern mode is enabled by control variables Square_wave_enable_0 through Square_wave_enable_15. If a Clause 45 MDIO is implemented, these control variables are accessible through the square wave testing control and status register bits 1.1510.0 through 1.1510.15 (limited to the number of lanes of the service interface below the PMA, see 45.2.1.125). When enabled, the PMA shall generate a square wave test pattern (8 ones followed by 8 zeros) on the square wave enabled lanes toward the service interface below the PMA via the *inst:IS_UNITDATA_i.request* primitive. Lanes for which square wave is not enabled transmit normal data resulting from the bit multiplexing operations described in 120.5.2 or test patterns as determined by other registers. When transmit square wave test pattern is disabled for all lanes, the PMA performs normal operation performing bit multiplexing as described in 120.5.2 or transmit test patterns as determined by other registers.

120.5.11.2 Test patterns for PAM4 encoded signals

For a 200GBASE-R PMA with 4 output lanes or a 400GBASE-R PMA with 4 or 8 output lanes using PAM4 encoding, the test patterns described in this clause may optionally be supported.

The patterns PRBS13Q and square wave (quaternary) can be enabled on a lane-by-lane basis. The patterns PRBS31Q and SSPRQ can be enabled on all lanes of an interface at once. If per-lane pattern(s) are enabled for a subset of the lanes and a per-interface pattern is also enabled, the per-lane patterns are generated only on the indicated lanes and the per-interface pattern is generated on the remaining lanes. The behavior if more than one per-lane pattern is enabled for the same lane or more than one per-interface pattern is enabled is not defined.

120.5.11.2.1 PRBS13Q test pattern

A PMA may optionally include a PRBS13Q pattern generator as specified in this subclause. The ability to perform this function is indicated by the PRBS13Q_gen_Tx_ability and PRBS13Q_gen_Rx_ability status variables, reflecting the ability to send this test pattern in the direction towards the PMD and towards the MAC, respectively. If a Clause 45 MDIO is implemented, the PRBS13Q_gen_Tx_ability and PRBS13Q_gen_Rx_ability status variables are accessible through the PRBS13Q Tx generator ability and PRBS13Q Rx generator ability bits 1.1500.11 and 1.1500.10 (see 45.2.1.123).

When the PRBS13Q test pattern enabled, it replaces the signal on the output lane(s) for which it is enabled. The PRBS13Q test pattern is a repeating 8191-symbol sequence formed by Gray coding pairs of bits from two repetitions of the PRBS13 pattern into PAM4 symbols as described in 120.5.7. The PRBS pattern generator produces the same result as the implementation shown in Figure 94-6, which implements the generator polynomial shown in Equation (94-3). Since the PRBS13 pattern is an odd number of bits in length, bits which are mapped as the first bit of a PAM4 symbol during one repetition of the PRBS13 sequence are mapped as the second bit of a PAM4 symbol during the next repetition of the PRBS13 sequence, and bits which are mapped as the second bit of a PAM4 symbol are mapped as the first bit of the following symbol in the next repetition of the PRBS13 sequence. For example, if the PRBS13 generator used to create the PRBS13Q sequence is initialized to a seed value of 0000010101011 (with the leftmost bit in S0 and the rightmost in S12), the PRBS13Q sequence begins with the following Gray coded PAM4 symbols, transmitted left to right: 1031320220111130103121231210012102121023131112.

The PRBS13Q test pattern is enabled in the transmit direction by the Tx_gen_enable and PRBS13Q_enable_0 through PRBS13Q_enable_7 control variables (in the downstream direction towards the PMD), and in the receive direction by the Rx_gen_enable and PRBS13Q_enable_0 through PRBS13Q_enable_7 control variables (in the upstream direction towards the MAC). If the optional Clause 45 MDIO is implemented, the Tx_gen_enable and Rx_gen_enable control variables are accessible through the Tx generator enable and Rx generator enable bits 1.1501.3 and 1.1501.1 (see 45.2.1.124), and the PRBS13_enable_0 through PRBS13_enable_7 control variables are accessible through bits 1.1512.0 through 1.1512.7 (see 45.2.1.125a).

120.5.11.2.2 PRBS31Q test pattern

A PMA may optionally include a PRBS31Q pattern generator as specified in this subclause. The ability to generate PRBS31Q patterns in each direction of transmission are indicated by the PRBS31Q_gen_Tx_ability and PRBS31Q_gen_Rx_ability status variables, reflecting the ability to send this test pattern in the direction towards the PMD and towards the MAC, respectively. The ability to check PRBS31Q patterns in each direction of transmission are indicated by the PRBS31Q_Tx_checker_ability and PRBS31Q_Rx_checker_ability status variables. If a Clause 45 MDIO is implemented, the PRBS31Q_gen_Tx_ability, PRBS31Q_gen_Rx_ability, PRBS31Q_Tx_checker_ability and PRBS31Q_Rx_checker_ability status variables are accessible through the PRBS31Q Tx generator ability, PRBS31Q Rx generator ability, PRBS31Q Tx checker ability, and PRBS31Q Rx checker ability bits 1.1500.9, 1.1500.7, 1.1500.8, and 1.1500.6 (see 45.2.1.123).

The PRBS31Q test pattern is a repeating $2^{31}-1$ -symbol sequence formed by Gray coding pairs of bits from two repetitions of the PRBS31 pattern defined in 49.2.8 into PAM4 symbols as described in 120.5.7. Since the PRBS31 pattern is an odd number of bits in length, bits which are mapped as the first bit of a PAM4

symbol during one repetition of the PRBS31 sequence are mapped as the second bit of a PAM4 symbol during the next repetition of the PRBS31 sequence, and bits which are mapped as the second bit of a PAM4 symbol during one repetition of the PRBS31 sequence are mapped as the first bit of the following symbol in the next repetition of the PRBS31 sequence. For example, if the PRBS31 generator used to create the PRBS31Q sequence is initialized to a seed value of all ones, the PRBS31Q sequence begins with the following Gray coded PAM4 symbols, transmitted left to right: 222222222222012222222222220002222222222201201222. To avoid correlated crosstalk, it is highly recommended that the PRBS31 patterns used to generate the PRBS31Q pattern on each lane are generated from independent, random seeds, or at a minimum offset of 20 000 UI between the PRBS31 sequence used to generate the PRBS31Q pattern on any lane and any other lane. A PRBS31Q pattern checker operates by converting PAM4 symbols received on each input lane to pairs of bits as described in 120.5.7 and then using a PRBS31 pattern checker on the resulting bit stream. The checker shall increment the test-pattern error counter by one for each incoming bit error in the PRBS31 pattern for isolated single bit errors. Implementations should be capable of counting at least one error whenever one or more errors occur in a sliding 1000-bit window.

If supported, when send Tx PRBS31Q test pattern is enabled by the PRBS31Q_enable and PRBS_tx_gen enable control variables, the PMA shall generate a PRBS31Q pattern on each of the lanes toward the service interface below the PMA via the *inst:IS_UNITDATA_i.request* primitive. When send Tx PRBS31Q test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 120.5.2. If a Clause 45 MDIO is implemented, the PRBS31Q_enable and PRBS_Tx_gen_enable control variables are accessible through bits 1.1501.13 and 1.1501.3 (see 45.2.1.124).

If supported, when send Rx PRBS31Q test pattern is enabled by the PRBS31Q_enable and PRBS_rx_gen enable control variables, the PMA shall generate a PRBS31Q pattern on each of the lanes toward the PMA client via the *PMA:IS_UNITDATA_i.indication* primitive. The PMA also generates *PMA:IS_SIGNAL.indication(SIGNAL_OK)* toward the PMA client independent of the link status at the service interface below the PMA. If a Clause 45 MDIO is implemented, the PRBS31Q_enable and PRBS_Rx_gen_enable control variables are accessible through bits 1.1501.13 and 1.1501.1 (see 45.2.1.124). When send Rx PRBS31Q test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 120.5.2.

If supported, when check Tx PRBS31Q test pattern mode is enabled by the PRBS31Q_enable and PRBS_Tx_check_enable control variables, the PMA checks for the PRBS31Q pattern on each of the lanes received from the PMA client via the *PMA:IS_UNITDATA_i.request* primitive. If a Clause 45 MDIO is implemented, the PRBS31Q_enable and PRBS_Tx_check_enable control variables are accessible through bits 1.1501.13 and 1.1501.2 (see 45.2.1.124). The Tx test-pattern error counters Ln0_PRBS_Tx_test_err_counter through either Ln3_PRBS_Tx_test_error_counter or Ln7_PRBS_Tx_test_error_counter (depending on whether the number of lanes is 4 or 8) count, per lane, errors in detecting the PRBS31 pattern resulting from converting the PAM4 symbols received on each lane to pairs of bits. If a Clause 45 MDIO is implemented, these counters are accessible through registers 1.1600 through 1.1603 or 1.1607 (depending on whether the number of lanes is 4 or 8) (see 45.2.1.126). When check Tx PRBS31Q test pattern is disabled, the PMA expects normal traffic and test pattern error counting does not continue. While in check Tx PRBS31Q mode, bit multiplexing continues as described in 120.5.2. Note that bit multiplexing of per-lane PRBS31Q may produce a signal which is not meaningful for downstream sublayers.

If supported, when check Rx PRBS31Q test pattern mode is enabled by the PRBS31Q_enable and PRBS_Rx_check_enable control variables, the PMA checks for the PRBS31Q pattern on each of the lanes received from the service interface below the PMA via the *inst:IS_UNITDATA_i.indication* primitive. If a Clause 45 MDIO is implemented, the PRBS31Q_enable and PRBS_Rx_check_enable control variables are accessible through bits 1.1501.13 and 1.1501.0 (see 45.2.1.124). The Rx test-pattern error counters Ln0_PRBS_Rx_test_err_counter through either Ln3_PRBS_Rx_test_error_counter or Ln7_PRBS_Rx_test_error_counter (depending on whether the number of lanes is 4 or 8) count, per lane,

errors in detecting the PRBS31 pattern resulting from converting the PAM4 symbols received on each lane to pairs of bits. If a Clause 45 MDIO is implemented, these counters are accessible through registers 1.1700 through 1.1703 or 1.1707 (depending on whether the number of lanes is 4 or 8) (see 45.2.1.127). While in check Rx PRBS31Q mode, the PMA:IS_SIGNAL.indication primitive does not indicate a valid signal. When check Rx PRBS31Q test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 120.5.2.

120.5.11.2.3 SSPRQ test pattern

A PMA may optionally include a short stress pattern random quaternary (SSPRQ) test pattern generator as specified in this subclause. The ability to perform this function is indicated by the SSPRQ_gen_Tx_ability status variable, reflecting the ability to send this test pattern in the direction towards the PMD. If a Clause 45 MDIO is implemented, the SSPRQ_gen_Tx_ability status variable is accessible through the SSPRQ Tx generator ability bit 1.1500.13 (see 45.2.1.123).

The SSPRQ pattern is a repeating $2^{16}-1$ PAM4 symbol sequence constructed as follows:

- Bit sequence A is a 32768-bit sequence composed of three sections of the PRBS31 binary sequence (see 120.5.11.1.1) according to Table 120–2.

Table 120–2—SSPRQ bit sequence A

| Pattern | Seed | Length |
|---------|------------|------------|
| PRBS31 | 0x00000002 | 10924 bits |
| | 0x34013FF7 | 10922 bits |
| | 0x0CCCCCCC | 10922 bits |

- Each section of PRBS31 is generated as if produced by the shift register implementation shown in Figure 49-9 and the seed is a 31-bit hexadecimal value used to preset S30 through S0 (S30 is set to the MSB and S0 is set to the LSB) prior to the generation of the PRBS31 sequence for the indicated length of bits.
- Bit sequence B is a 65534-bit sequence formed by removing the first and last binary bit from a sequence consisting of two repetitions of bit sequence A.
- PAM4 sequence 1 is a 16384-symbol sequence formed by Gray coding bit sequence A according to 120.5.7.
- PAM4 sequence 2 is a 16384-symbol sequence formed by Gray coding bit sequence A according to 120.5.7, and then inverting each PAM4 symbol by replacing each PAM4 symbol n with the value $3-n$.
- PAM4 sequence 3 is a 16383-symbol sequence formed by Gray coding the first 32766 bits of bit-sequence B according to 120.5.7.
- PAM4 sequence 4 is a 16384-symbol sequence formed by Gray coding the last 32768 bits of bit-sequence B according to 120.5.7, and then inverting each PAM4 symbol by replacing each PAM4 symbol n with the value $3-n$.
- The repeating SSPRQ pattern formed by concatenating PAM4 sequences 1, 2, 3 and 4.

NOTE—A CSV file containing the entire PAM4 symbol sequence for the SSPRQ test pattern is available at <http://standards.ieee.org/downloads/802.3/>.

If supported, when send SSPRQ test pattern is enabled by the SSPRQ_enable control variable, the PMA shall generate an SSPRQ pattern on each of its lanes in the Tx direction towards the PMD with at least a 31 UI delay between the SSPRQ pattern on one lane and any other lane. If a Clause 45 MDIO is

implemented, the SSPRQ_pattern_enable control variable is accessible through the SSPRQ pattern enable bit 1.1501.14 (see 45.2.1.124).

120.5.11.2.4 Square wave (quaternary) test pattern

Transmit square wave (quaternary) test-pattern mode optionally applies to each lane of the Tx direction PMA towards a physically instantiated 200GAUI-4, 400GAUI-8 or towards a PAM4 PMD service interface whether or not it is physically instantiated.

The ability to perform this function is indicated by the Square_wave_ability status variable. If a Clause 45 MDIO is implemented, the Square_wave_ability status variable is accessible through the Square wave test ability bit 1.1500.12 (see 45.2.1.123). If implemented, the transmit square wave test-pattern mode is enabled by control variables Square_wave_enable_0 through Square_wave_enable_8. If a Clause 45 MDIO is implemented, these control variables are accessible through the square wave testing control and status register bits 1.1510.0 through 1.1510.15 (limited to the number of lanes of the service interface below the PMA, see 45.2.1.125). When enabled, the PMA shall generate a square wave test pattern (8 threes followed by 8 zeros) on the square wave enabled lanes toward the service interface below the PMA via the *inst:IS_UNITDATA_i.request* primitive. Lanes for which square wave is not enabled transmit normal data resulting from the bit multiplexing operations described in 120.5.2 or test patterns as determined by other registers. When transmit square wave test pattern is disabled for all lanes, the PMA performs normal operation performing bit multiplexing as described in 120.5.2 or transmit test patterns as determined by other registers.

NOTE—A square wave transmitted over a 200GAUI-4 or 400GAUI-8 may not be correctly forwarded to the PMD transmitter output.

120.6 PMA MDIO function mapping

The optional MDIO capability described in Clause 45 describes several variables that provide control and status information for and about the PMA. Since a given implementation may employ more than one PMA sublayer, the PMA control and status information is organized into multiple addressable instances, one for each possible PMA sublayer. See 45.2.1 and 120.1.4 for the allocation of MMD addresses to PMA sublayers. Control and status registers for MMD 8, 9, and 10 use the Extended PMA control and status registers at identical locations to those for MMD 1.

Mapping of MDIO control variables to PMA control variables is shown in Table 120–3. Mapping of MDIO status variables to PMA status variables is shown in Table 120–4. Mapping of MDIO counter to PMA counters is shown in Table 120–5. These tables provide the register and bit numbers for the PMA addressed as MMD 1. For implementations with multiple PMA sublayers, additional PMA sublayers use the corresponding register and bit numbers in MMDs 8, 9, and 10 as necessary.

Table 120–3—MDIO/PMA control variable mapping

| MDIO variable | PMA/PMD register name | Register/ bit number | PMA control variable |
|------------------------|------------------------------|-------------------------|------------------------|
| PMA remote loopback | PMA/PMD control 1 | 1.0.1 | Remote_loopback_enable |
| PMA local loopback | PMA/PMD control 1 | 1.0.0 | Local_loopback_enable |
| PRBS31 pattern enable | PRBS pattern testing control | 1.1501.7 | PRBS31_enable |
| PRBS9 pattern enable | PRBS pattern testing control | 1.1501.6 | PRBS9_enable |
| PRBS31Q pattern enable | PRBS pattern testing control | 1.1501.13 | PRBS31Q_pattern_enable |

Table 120–3—MDIO/PMA control variable mapping (continued)

| MDIO variable | PMA/PMD register name | Register/ bit number | PMA control variable |
|-----------------------|------------------------------|-------------------------|-----------------------|
| SSPRQ pattern enable | PRBS pattern testing control | 1.1501.14 | SSPRQ_pattern_enable |
| Tx generator enable | PRBS pattern testing control | 1.1501.3 | PRBS_Tx_gen_enable |
| Tx checker enable | PRBS pattern testing control | 1.1501.2 | PRBS_Tx_check_enable |
| Rx generator enable | PRBS pattern testing control | 1.1501.1 | PRBS_Rx_gen_enable |
| Rx checker enable | PRBS pattern testing control | 1.1501.0 | PRBS_Rx_check_enable |
| Lane 0 SW enable | Square wave testing control | 1.1510.0 | Square_wave_enable_0 |
| Lane 1 SW enable | Square wave testing control | 1.1510.1 | Square_wave_enable_1 |
| Lane 2 SW enable | Square wave testing control | 1.1510.2 | Square_wave_enable_2 |
| Lane 3 SW enable | Square wave testing control | 1.1510.3 | Square_wave_enable_3 |
| Lane 4 SW enable | Square wave testing control | 1.1510.4 | Square_wave_enable_4 |
| Lane 5 SW enable | Square wave testing control | 1.1510.5 | Square_wave_enable_5 |
| Lane 6 SW enable | Square wave testing control | 1.1510.6 | Square_wave_enable_6 |
| Lane 7 SW enable | Square wave testing control | 1.1510.7 | Square_wave_enable_7 |
| Lane 8 SW enable | Square wave testing control | 1.1510.8 | Square_wave_enable_8 |
| Lane 9 SW enable | Square wave testing control | 1.1510.9 | Square_wave_enable_9 |
| Lane 10 SW enable | Square wave testing control | 1.1510.10 | Square_wave_enable_10 |
| Lane 11 SW enable | Square wave testing control | 1.1510.11 | Square_wave_enable_11 |
| Lane 12 SW enable | Square wave testing control | 1.1510.12 | Square_wave_enable_12 |
| Lane 13 SW enable | Square wave testing control | 1.1510.13 | Square_wave_enable_13 |
| Lane 14 SW enable | Square wave testing control | 1.1510.14 | Square_wave_enable_14 |
| Lane 15 SW enable | Square wave testing control | 1.1510.15 | Square_wave_enable_15 |
| Lane 0 PRBS13Q enable | PRBS13Q testing control | 1.1512.0 | PRBS13Q_enable_0 |
| Lane 1 PRBS13Q enable | PRBS13Q testing control | 1.1512.1 | PRBS13Q_enable_1 |
| Lane 2 PRBS13Q enable | PRBS13Q testing control | 1.1512.2 | PRBS13Q_enable_2 |
| Lane 3 PRBS13Q enable | PRBS13Q testing control | 1.1512.3 | PRBS13Q_enable_3 |
| Lane 4 PRBS13Q enable | PRBS13Q testing control | 1.1512.4 | PRBS13Q_enable_4 |
| Lane 5 PRBS13Q enable | PRBS13Q testing control | 1.1512.5 | PRBS13Q_enable_5 |
| Lane 6 PRBS13Q enable | PRBS13Q testing control | 1.1512.6 | PRBS13Q_enable_6 |
| Lane 7 PRBS13Q enable | PRBS13Q testing control | 1.1512.7 | PRBS13Q_enable_7 |

Table 120–4—MDIO/PMA status variable mapping

| MDIO status variable | PMA/PMD register name | Register/bit number | PMA status variable |
|------------------------------|--|---------------------|------------------------------|
| PMA remote loopback ability | 200G PMA/PMD extended ability register | 1.23.15 | 200G_Remote_loopback_ability |
| PMA remote loopback ability | 400G PMA/PMD extended ability register | 1.24.15 | 400G_Remote_loopback_ability |
| PMA local loopback ability | PMA/PMD status 2 register | 1.8.0 | Local_loopback_ability |
| PRBS9 Tx generator ability | Test-pattern ability register | 1.1500.5 | PRBS9_Tx_generator_ability |
| PRBS9 Rx generator ability | Test-pattern ability register | 1.1500.4 | PRBS9_Rx_generator_ability |
| PRBS31Tx generator ability | Test-pattern ability register | 1.1500.3 | PRBS31_Tx_generator_ability |
| PRBS31Tx checker ability | Test-pattern ability register | 1.1500.2 | PRBS31_Tx_checker_ability |
| PRBS31 Rx generator ability | Test-pattern ability register | 1.1500.1 | PRBS31_Rx_generator_ability |
| PRBS31 Rx checker ability | Test-pattern ability register | 1.1500.0 | PRBS31_Rx_checker_ability |
| Square wave test ability | Test-pattern ability register | 1.1500.12 | Square_wave_ability |
| PRBS13Q Tx generator ability | Test-pattern ability register | 1.1500.11 | PRBS13Q_gen_Tx_ability |
| PRBS13Q Rx generator ability | Test-pattern ability register | 1.1500.10 | PRBS13Q_gen_Rx_ability |
| PRBS31Q Tx generator ability | Test-pattern ability register | 1.1500.9 | PRBS31Q_gen_Tx_ability |
| PRBS31Q Rx generator ability | Test-pattern ability register | 1.1500.7 | PRBS31Q_gen_rx_ability |
| PRBS31Q Tx checker ability | Test-pattern ability register | 1.1500.8 | PRBS31Q_Tx_checker_ability |
| PRBS31Q Rx checker ability | Test-pattern ability register | 1.1500.6 | PRBS31Q_Rx_checker_ability |
| SSPRQ Tx generator ability | Test-pattern ability register | 1.1500.13 | SSPRQ_gen_Tx_ability |

Table 120–5—MDIO/PMA counters mapping

| MDIO variable | PMA/PMD register name | Register/bit number | PMA status variable |
|--------------------------|---|---------------------|------------------------------|
| Error counter Tx, lane 0 | PRBS Tx pattern testing error counter, lane 0 | 1.1600 | Ln0_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 1 | PRBS Tx pattern testing error counter, lane 1 | 1.1601 | Ln1_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 2 | PRBS Tx pattern testing error counter, lane 2 | 1.1602 | Ln2_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 3 | PRBS Tx pattern testing error counter, lane 3 | 1.1603 | Ln3_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 4 | PRBS Tx pattern testing error counter, lane 4 | 1.1604 | Ln4_PRBS_Tx_test_err_counter |

Table 120–5—MDIO/PMA counters mapping (*continued*)

| MDIO variable | PMA/PMD register name | Register/bit number | PMA status variable |
|---------------------------|--|---------------------|-------------------------------|
| Error counter Tx, lane 5 | PRBS Tx pattern testing error counter, lane 5 | 1.1605 | Ln5_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 6 | PRBS Tx pattern testing error counter, lane 6 | 1.1606 | Ln6_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 7 | PRBS Tx pattern testing error counter, lane 7 | 1.1607 | Ln7_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 8 | PRBS Tx pattern testing error counter, lane 8 | 1.1608 | Ln8_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 9 | PRBS Tx pattern testing error counter, lane 9 | 1.1609 | Ln9_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 10 | PRBS Tx pattern testing error counter, lane 10 | 1.1610 | Ln10_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 11 | PRBS Tx pattern testing error counter, lane 11 | 1.1611 | Ln11_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 12 | PRBS Tx pattern testing error counter, lane 12 | 1.1612 | Ln12_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 13 | PRBS Tx pattern testing error counter, lane 13 | 1.1613 | Ln13_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 14 | PRBS Tx pattern testing error counter, lane 14 | 1.1614 | Ln14_PRBS_Tx_test_err_counter |
| Error counter Tx, lane 15 | PRBS Tx pattern testing error counter, lane 15 | 1.1615 | Ln15_PRBS_Tx_test_err_counter |
| Error counter Rx, lane 0 | PRBS Rx pattern testing error counter, lane 0 | 1.1700 | Ln0_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 1 | PRBS Rx pattern testing error counter, lane 1 | 1.1701 | Ln1_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 2 | PRBS Rx pattern testing error counter, lane 2 | 1.1702 | Ln2_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 3 | PRBS Rx pattern testing error counter, lane 3 | 1.1703 | Ln3_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 4 | PRBS Rx pattern testing error counter, lane 4 | 1.1704 | Ln4_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 5 | PRBS Rx pattern testing error counter, lane 5 | 1.1705 | Ln5_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 6 | PRBS Rx pattern testing error counter, lane 6 | 1.1706 | Ln6_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 7 | PRBS Rx pattern testing error counter, lane 7 | 1.1707 | Ln7_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 8 | PRBS Rx pattern testing error counter, lane 8 | 1.1708 | Ln8_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 9 | PRBS Rx pattern testing error counter, lane 9 | 1.1709 | Ln9_PRBS_Rx_test_err_counter |

Table 120–5—MDIO/PMA counters mapping (*continued*)

| MDIO variable | PMA/PMD register name | Register/bit number | PMA status variable |
|---------------------------|--|---------------------|-------------------------------|
| Error counter Rx, lane 10 | PRBS Rx pattern testing error counter, lane 10 | 1.1710 | Ln10_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 11 | PRBS Rx pattern testing error counter, lane 11 | 1.1711 | Ln11_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 12 | PRBS Rx pattern testing error counter, lane 12 | 1.1712 | Ln12_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 13 | PRBS Rx pattern testing error counter, lane 13 | 1.1713 | Ln13_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 14 | PRBS Rx pattern testing error counter, lane 14 | 1.1714 | Ln14_PRBS_Rx_test_err_counter |
| Error counter Rx, lane 15 | PRBS Rx pattern testing error counter, lane 15 | 1.1715 | Ln15_PRBS_Rx_test_err_counter |

120.7 Protocol implementation conformance statement (PICS) proforma for Clause 120, Physical Medium Attachment (PMA) sublayer, type 200GBASE-R and 400GBASE-R⁵

120.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 120, Physical Medium Attachment (PMA) sublayer, type 200GBASE-R and 400GBASE-R, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

120.7.2 Identification

120.7.2.1 Implementation identification

| | |
|--|--|
| Supplier ¹ | |
| Contact point for inquiries about the PICS ¹ | |
| Implementation Name(s) and Version(s) ^{1,3} | |
| Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ² | |
| NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model). | |

120.7.2.2 Protocol summary

| | |
|---|--|
| Identification of protocol standard | IEEE Std 802.3bs-2017, Clause 120, Physical Medium Attachment (PMA) sublayer, type 200GBASE-R and 400GBASE-R |
| Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS | |
| Have any Exception items been required? No [] Yes [] (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3bs-2017.) | |

| | |
|-------------------|--|
| Date of Statement | |
|-------------------|--|

⁵Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

120.7.3 Major capabilities/options

| Item | Feature | Subclause | Value/Comment | Status | Support |
|--------------|--|-----------|--|--|--------------------------|
| *PMA200 | PMA for 200GBASE-R | 120.1.1 | | O.1 | Yes [] No [] |
| *PMA400 | PMA for 400GBASE-R | 120.1.1 | | O.1 | Yes [] No [] |
| *LNS_UP-STRM | Number of lanes in direction of MAC | 120.1.4 | Divisor of number of PCS lanes | PMA200: M | 8 [] 4 [] |
| | | | | PMA400:M | 16 [] 8 [] |
| *UP_NRZ | NRZ modulation used for PMA service interface | 120.1.4 | | (PMA200* LNS_UPSTRM=8 or PMA400* LNS_UPSTRM=16) :M | Yes [] N/A [] |
| *UP_PAM4 | PAM4 modulation used for PMA service interface | 120.1.4 | | (PMA200* LNS_UPSTRM=4 or PMA400* LNS_UPSTRM=8) :M | Yes [] N/A [] |
| *LNS_DN-STRM | Number of lanes in the direction of PMD | 120.1.4 | Divisor of number of PCS lanes | PMA200: M | 8 [] 4 [] |
| | | | | PMA400:M | 16 [] 8 [] 4 [] |
| *DN_NRZ | NRZ modulation used for service interface below the PMA | 120.1.4 | | (PMA200* LNS_DNSTRM=8 or PMA400* LNS_DNSTRM=16) :M | Yes [] N/A [] |
| *DN_PAM4 | PAM4 modulation used for service interface below the PMA | 120.1.4 | | (PMA200* LNS_DNSTRM=4 or PMA400* LNS_DNSTRM<16) :M | Yes [] N/A [] |
| RX_CLOCK | Signaling rate of output lanes in Rx direction | 120.5.5 | 26.5625 GBd | M | Yes [] |
| TX_CLOCK | Signaling rate of output lanes in Tx direction | 120.5.5 | 26.5625 GBd | PMA200:M | Yes [] N/A [] |
| | | | LNS_DNSTRM=8 or 16: 26.5625 GBd LNS_DNSTRM=4: 53.125 GBd | PMA400:M | Yes [] N/A [] |
| LANE MAPPING | Maintain lane mapping while link is in operation | 120.5.2 | Maintain sequence of PCSs on all output lanes | M | Yes [] |

IEEE Std 802.3bs-2017
Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
for 200 Gb/s and 400 Gb/s Operation

| Item | Feature | Subclause | Value/Comment | Status | Support |
|----------|---|------------------|---|-----------|--------------------|
| LNKS | PMA link status | 120.5.8 | Meets the requirements of 120.5.8 | M | Yes [] |
| *LBL | PMA local loopback | 120.5.9 | Supports local loopback | O | Yes [] No [] |
| *LBR | PMA remote loopback | 120.5.10 | Supports remote loopback | O | Yes [] No [] |
| MD | MDIO | 120.6 | Registers and interface supported | O | Yes [] No [] |
| *USP1SP6 | Physically instantiated 200GAUI-n or 400GAUI-n above (toward MAC) | 120.5.3 | | O | Yes [] No [] |
| *DSP1SP6 | Physically instantiated 200GAUI-n or 400GAUI-n below (toward PMD) | 120.5.3 | | O | Yes [] No [] |
| *SP2SP5 | Physically instantiated PMD service interface | 120.5.3 | | O | Yes [] No [] |
| UNAUI | Electrical and timing requirements of Annex 120B, Annex 120C, Annex 120D, or Annex 120E as appropriate met by 200GAUI-n or 400GAUI-n of the PMA service interface | 120.5.1, 120.5.5 | | USP1SP6:M | Yes [] N/A [] |
| DNAUI | Electrical and timing requirements of Annex 120B, Annex 120C, Annex 120D, or Annex 120E as appropriate met by 200GAUI-n or 400GAUI-n of the service interface below the PMA | 120.5.1, 120.5.5 | | DSP1SP6:M | Yes [] N/A [] |
| DELAY200 | Roundtrip delay limit for 200GBASE-R | 120.5.4 | No more than 18 432 BT or 36 pause_quanta | PMA200:M | Yes [] N/A [] |
| DELAY400 | Roundtrip delay limit for 400GBASE-R | 120.5.4 | No more than 36 864 BT or 72 pause_quanta | PMA400:M | Yes [] N/A [] |

120.7.4 Skew generation and tolerance

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|----------------|---------------|--------------------|
| S1 | Skew generation toward SP1 in Tx direction | 120.5.3.1 | ≤ 29 ns | DSP1S P6:M | Yes [] N/A [] |
| S2 | Skew variation generation toward SP1 in Tx direction | 120.5.3.1 | ≤ 200 ps | DSP1S P6:M | Yes [] N/A [] |
| S3 | Skew variation tolerance at SP1 | 120.5.3.2 | Minimum 200 ps | USP1S P6:M | Yes [] N/A [] |
| S4 | Skew generation toward SP2 in Tx direction | 120.5.3.3 | ≤ 43 ns | SP2SP 5:M | Yes [] N/A [] |
| S5 | Skew variation generation toward SP2 in Tx direction | 120.5.3.3 | ≤ 400 ps | SP2SP 5:M | Yes [] N/A [] |
| S6 | Skew variation tolerance at SP5 | 120.5.3.4 | Minimum 3.6 ns | SP2SP 5:M | Yes [] N/A [] |
| S7 | Skew generation toward SP6 in Rx direction | 120.5.3.5 | ≤ 160 ns | USP1S P6:M | Yes [] N/A [] |
| S8 | Skew variation generation toward SP6 in Rx direction | 120.5.3.5 | ≤ 3.8 ns | USP1S P6:M | Yes [] N/A [] |
| S9 | Skew variation tolerance at SP6 | 120.5.3.6 | Minimum 3.8 ns | DSP1S P6:M | Yes [] N/A [] |

120.7.5 Test patterns

| Item | Feature | Subclause | Value/Comment | Status | Support |
|-------|---|--------------|---------------|-------------------|------------------------------|
| *JTP1 | Physically Instantiated 200GAUI-n or 400GAUI-n between PMA and PMA client | 120.5.11 | | O | Yes [] No [] |
| *JTP2 | Physically Instantiated 200GAUI-n or 400GAUI-n between PMA and sublayer below the PMA, or adjacent to PMD | 120.5.11 | | O | Yes [] No [] |
| J1 | Send PRBS31 Tx | 120.5.11.1.1 | | JTP2* DN_NRZ:O | Yes [] No [] N/A [] |
| J2 | Send PRBS31 Rx | 120.5.11.1.1 | | JTP1* UP_NRZ:O | Yes [] No [] N/A [] |
| J3 | Check PRBS31 Tx | 120.5.11.1.1 | | JTP1* UP_NRZ:O | Yes [] No [] N/A [] |
| J4 | Check PRBS31 Rx | 120.5.11.1.1 | | JTP2* DN_NRZ:O | Yes [] No [] N/A [] |

IEEE Std 802.3bs-2017
Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
for 200 Gb/s and 400 Gb/s Operation

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|-------------------------------------|--------------|---------------|--------------------|---|
| J5 | Send PRBS9 Tx | 120.5.11.1.2 | | JTP2* DN_NRZ:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J6 | Send PRBS9 Rx | 120.5.11.1.2 | | JTP1* UP_NRZ:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J7 | Send square wave Tx | 120.5.11.1.3 | | JTP2* DN_NRZ:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J8 | Send PRBS13Q Tx | 120.5.11.2.1 | | JTP2* DN_PAM4:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J9 | Send PRBS13Q Rx | 120.5.11.2.1 | | JTP1* UP_PAM4:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J10 | Send PRBS31Q Tx | 120.5.11.2.2 | | JTP2* DN_PAM4:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J11 | Send PRBS31Q Rx | 120.5.11.2.2 | | JTP1* UP_PAM4:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J12 | Check PRBS31Q Tx | 120.5.11.2.2 | | JTP1* UP_PAM4:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J13 | Check PRBS31Q Rx | 120.5.11.2.2 | | JTP2* DN_PAM4:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J14 | Send SSPRQ Tx | 120.5.11.2.3 | | DN_PAM4:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |
| J15 | Send square wave (quaternary) Tx | 120.5.11.2.4 | | DN_PAM4:O | Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/> |

120.7.6 Loopback modes

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---------------------------------|-----------|------------------------------------|--------|--|
| LB1 | PMA local loopback implemented | 120.5.9 | Meets the requirements of 120.5.9 | LBL:M | Yes <input type="checkbox"/> N/A <input type="checkbox"/> |
| LB2 | PMA remote loopback implemented | 120.5.10 | Meets the requirements of 120.5.10 | LBR:M | Yes <input type="checkbox"/> N/A <input type="checkbox"/> |

121. Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-DR4

121.1 Overview

This clause specifies the 200GBASE-DR4 PMD together with the single-mode fiber medium. The optical signal generated by this PMD type is modulated using a 4-level pulse amplitude modulation (PAM4) format. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 121–1, to the medium through the MDI and optionally with the management functions that may be accessible through the management interface defined in Clause 45, or equivalent.

Table 121–1—Physical Layer clauses associated with the 200GBASE-DR4 PMD

| Associated clause | 200GBASE-DR4 |
|-------------------------------|--------------|
| 117—RS | Required |
| 117—200GMII ^a | Optional |
| 118—200GMII Extender | Optional |
| 119—PCS | Required |
| 120—PMA | Required |
| 120B—Chip-to-chip 200GAUI-8 | Optional |
| 120C—Chip-to-module 200GAUI-8 | Optional |
| 120D—Chip-to-chip 200GAUI-4 | Optional |
| 120E—Chip-to-module 200GAUI-4 | Optional |
| 78—Energy Efficient Ethernet | Optional |

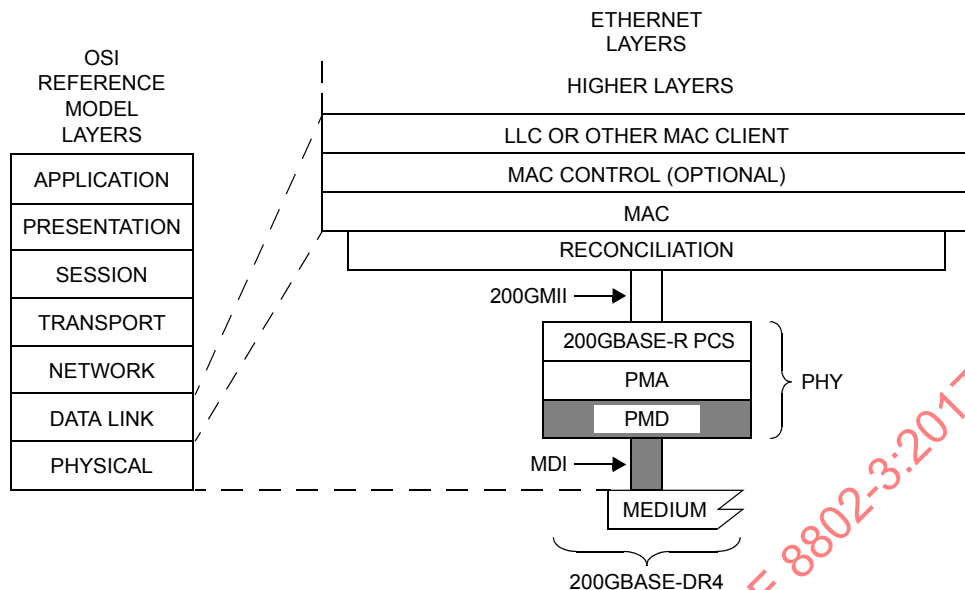
^a The 200GMII is an optional interface. However, if the 200GMII is not implemented, a conforming implementation must behave functionally as though the RS and 200GMII were present.

Figure 121–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 200 Gb/s Ethernet is introduced in Clause 116 and the purpose of each PHY sublayer is summarized in 116.2.

200GBASE-DR4 PHYs with the optional Energy Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

121.1.1 Bit error ratio

The bit error ratio (BER) when processed according to Clause 120 shall be less than 2.4×10^{-4} provided that the error statistics are sufficiently random that this results in a frame loss ratio (see 1.4.223) of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap when processed according to Clause 120 and then Clause 119. For a complete Physical Layer, the frame loss ratio may be degraded to 6.2×10^{-11} for 64-octet frames with minimum interpacket gap due to additional errors from the electrical interfaces.



200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
DR = PMD FOR SINGLE-MODE FIBER — 500 m

Figure 121-1—200GBASE-DR4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

If the error statistics are not sufficiently random to meet this requirement, then the BER shall be less than that required to give a frame loss ratio of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap.

121.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 200GBASE-DR4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 116.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
PMD:IS_UNITDATA_*i*.indication
PMD:IS_SIGNAL.indication

The 200GBASE-DR4 PMD has four parallel symbol streams, hence $i = 0$ to 3.

In the transmit direction, the PMA continuously sends four parallel symbol streams to the PMD, one per lane, each at a nominal signaling rate of 26.5625 GBd. The PMD then converts these streams of data units into the appropriate signals on the MDI.

In the receive direction, the PMD continuously sends four parallel symbol streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 26.5625 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service interface primitive defined in 116.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_symbol parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_symbol parameters are known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the BER defined in 121.1.1.

121.3 Delay and Skew

121.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 200GBASE-DR4 PMD including 2 m of fiber in one direction shall be no more than 4096 bit times (8 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 116.4 and its references.

121.3.2 Skew constraints

The Skew (relative delay between the lanes) and Skew Variation must be kept within limits so that the information on the lanes can be reassembled by the PCS. Skew and Skew Variation are defined in 116.5 and specified at the points SP1 to SP6 shown in Figure 116-4 and Figure 116-5.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 116.5. The measurements of Skew and Skew Variation are defined in 86.8.3.1.

121.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 121–2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 121–3.

Table 121–2—MDIO/PMD control variable mapping

| MDIO control variable | PMA/PMD register name | Register/bit number | PMD control variable |
|--|-------------------------------|---------------------|--|
| Reset | PMA/PMD control 1 register | 1.0.15 | PMD_reset |
| Global PMD transmit disable | PMD transmit disable register | 1.9.0 | PMD_global_transmit_disable |
| PMD transmit disable 3 to PMD transmit disable 0 | PMD transmit disable register | 1.9.4 to 1.9.1 | PMD_transmit_disable_3 to PMD_transmit_disable_0 |

Table 121–3—MDIO/PMD status variable mapping

| MDIO status variable | PMA/PMD register name | Register/bit number | PMD status variable |
|--|------------------------------------|---------------------|--|
| Fault | PMA/PMD status 1 register | 1.1.7 | PMD_fault |
| Transmit fault | PMA/PMD status 2 register | 1.8.11 | PMD_transmit_fault |
| Receive fault | PMA/PMD status 2 register | 1.8.10 | PMD_receive_fault |
| Global PMD receive signal detect | PMD receive signal detect register | 1.10.0 | PMD_global_signal_detect |
| PMD receive signal detect 3 to PMD receive signal detect 0 | PMD receive signal detect register | 1.10.4 to 1.10.1 | PMD_signal_detect_3 to PMD_signal_detect_0 |

121.5 PMD functional specifications

The 200GBASE-DR4 PMD performs the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

121.5.1 PMD block diagram

The PMD block diagram is shown in Figure 121–2. For purposes of system conformance, the PMD sublayer is standardized at the points described in this subclause. The optical transmit signal is defined at the output end of a single-mode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 121.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 121.11.3). Unless specified otherwise, all receiver measurements and tests defined in 121.8 are made at TP3.

TP1<0:3> and TP4<0:3> are informative reference points that may be useful to implementors for testing components (these test points will not typically be accessible in an implemented system).



The PMD Transmit function shall convert the four symbol streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate optical signals. The four optical signals shall then be delivered to the MDI, which contains four parallel light paths for transmit, according to the transmit optical specifications in this clause. The highest optical power level in each signal stream shall correspond to tx_symbol = three and the lowest shall correspond to tx_symbol = zero.

The PMD Receive function shall convert the four parallel optical signals received from the MDI into separate symbol streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to rx_symbol = three and the lowest shall correspond to rx_symbol = zero.

The PMD global signal detect function shall report the state of SIGNAL_DETECT via the PMD service interface. The SIGNAL_DETECT parameter is signaled continuously, while the PMD:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 116.3.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all four lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 121–4. The PMD receiver is not required to verify whether a compliant 200GBASE-DR4 signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 121–4—SIGNAL_DETECT value definition

| Receive conditions | SIGNAL_DETECT value |
|--|---------------------|
| For any lane; Average optical power at TP3 ≤ -16 dBm | FAIL |
| For all lanes; [(Optical power at TP3 \geq average receive power, each lane (min) Table 121–7) AND (compliant 200GBASE-R signal input)] | OK |
| All other conditions | Unspecified |

As an unavoidable consequence of the requirements for the setting of the SIGNAL_DETECT parameter, implementations must provide adequate margin between the input optical power level at which the SIGNAL_DETECT parameter is set to OK, and the inherent noise level of the PMD including the effects of crosstalk, power supply noise, etc.

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the modulated optical signal.

121.5.5 PMD lane-by-lane signal detect function

Various implementations of the Signal Detect function are permitted by this standard. When the MDIO is implemented, each PMD_signal_detect i , where i represents the lane number in the range 0:3, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 121–4.

121.5.6 PMD reset function

If the MDIO interface is implemented, and if PMD_reset is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

121.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

- a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 121–6.
- b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

121.5.8 PMD lane-by-lane transmit disable function (optional)

The PMD_transmit_disable_*i* (where *i* represents the lane number in the range 0:3) function is optional and allows the optical transmitters in each lane to be selectively disabled.

- a) When a PMD_transmit_disable_*i* variable is set to one, this function shall turn off the optical transmitter associated with that variable so that the transmitter meets the requirements of the average launch power of the OFF transmitter in Table 121–6.
- b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the optical transmitter in each lane.

If the optional PMD_transmit_disable_*i* function is not implemented in MDIO, an alternative method may be provided to independently disable each transmit lane for testing purposes.

121.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to one.

If the MDIO interface is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

121.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the PMD_transmit_fault variable to one.

If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

121.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to one.

If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

121.6 Lane assignments

There are no lane assignments (within a group of transmit or receive lanes) for 200GBASE-DR4. While it is expected that a PMD will map electrical lane *i* to optical lane *i* and vice versa, there is no need to define the physical ordering of the lanes, as the PCS sublayer is capable of receiving the lanes in any arrangement. The positioning of transmit and receive lanes at the MDI is specified in 121.11.3.1.

121.7 PMD to MDI optical specifications for 200GBASE-DR4

The operating range for the 200GBASE-DR4 PMD is defined in Table 121–5. A 200GBASE-DR4 compliant PMD operates on type B1.1, B1.3, or B6 single-mode fibers according to the specifications defined in Table 121–14. A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 200GBASE-DR4 PMD operating at 600 m meets the operating range requirement of 2 m to 500 m).

Table 121–5—200GBASE-DR4 operating range

| PMD type | Required operating range |
|--------------|--------------------------|
| 200GBASE-DR4 | 2 m to 500 m |

121.7.1 200GBASE-DR4 transmitter optical specifications

The 200GBASE-DR4 transmitter shall meet the specifications defined in Table 121–6 per the definitions in 121.8.

Table 121–6—200GBASE-DR4 transmit characteristics

| Description | Value | Unit |
|--|-----------------------|-------|
| Signaling rate, each lane (range) | 26.5625 ± 100 ppm | GBd |
| Modulation format | PAM4 | — |
| Lane wavelength (range) | 1304.5 to 1317.5 | nm |
| Side-mode suppression ratio (SMSR), (min) | 30 | dB |
| Average launch power, each lane (max) | 3 | dBm |
| Average launch power, each lane ^a (min) | –5.1 | dBm |
| Outer Optical Modulation Amplitude (OMA _{outer}), each lane (max) | 2.8 | dBm |
| Outer Optical Modulation Amplitude (OMA _{outer}), each lane (min) ^b | –3 | dBm |
| Launch power in OMA _{outer} minus TDECQ, each lane (min) | –4.4 | dBm |
| Transmitter and dispersion eye closure for PAM4 (TDECQ), each lane (max) | 3.4 | dB |
| Average launch power of OFF transmitter, each lane (max) | –16 | dBm |
| Extinction ratio, each lane (min) | 3.5 | dB |
| RIN _{21.4} OMA (max) | –132 | dB/Hz |
| Optical return loss tolerance (max) | 21.4 | dB |
| Transmitter reflectance ^c (max) | –26 | dB |

^a Average launch power, each lane (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.

^b Even if the TDECQ < 1.4 dB, the OMA_{outer} (min) must exceed this value.

^c Transmitter reflectance is defined looking into the transmitter.

121.7.2 200GBASE-DR4 receive optical specifications

The 200GBASE-DR4 receiver shall meet the specifications defined in Table 121–7 per the definitions in 121.8.

Table 121-7—200GBASE-DR4 receive characteristics

| Description | Value | Unit |
|---|------------------|------|
| Signaling rate, each lane (range) | 26.5625± 100 ppm | GBd |
| Modulation format | PAM4 | — |
| Lane wavelengths (range) | 1304.5 to 1317.5 | nm |
| Damage threshold ^a , each lane | 4 | dBm |
| Average receive power, each lane (max) | 3 | dBm |
| Average receive power, each lane ^b (min) | −8.1 | dBm |
| Receive power (OMA _{outer}), each lane (max) | 2.8 | dBm |
| Receiver reflectance (max) | −26 | dB |
| Receiver sensitivity (OMA _{outer}), each lane ^c (max) | −6.6 | dBm |
| Stressed receiver sensitivity (OMA _{outer}), each lane ^d (max) | −4.1 | dBm |
| Conditions of stressed receiver sensitivity test: ^e | | |
| Stressed eye closure for PAM4 (SECQ), lane under test | 3.4 | dB |
| OMA _{outer} of each aggressor lane | 2.8 | dBm |

^a The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level. The receiver does not have to operate correctly at this input power.

^b Average receive power, each lane (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.

^c Receiver sensitivity (OMA_{outer}), each lane (max) is informative and is defined for a transmitter with SECQ of 0.9 dB.

^d Measured with conformance test signal at TP3 (see 121.8.9) for the BER specified in 121.1.1.

^e These test conditions are for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

121.7.3 200GBASE-DR4 illustrative link power budget

An illustrative power budget and penalties for 200GBASE-DR4 channels are shown in Table 121-8.

Table 121-8—200GBASE-DR4 illustrative link power budget

| Parameter | Value | Unit |
|---|----------------|------|
| Power budget (for max TDECQ) | 6.5 | dB |
| Operating distance | 500 | m |
| Channel insertion loss ^a | 3 | dB |
| Maximum discrete reflectance | See 121.11.2.2 | dB |
| Allocation for penalties ^b (for max TDECQ) | 3.5 | dB |
| Additional insertion loss allowed | 0 | dB |

^a The channel insertion loss is calculated using the maximum distance specified in Table 121-5 and cabled optical fiber attenuation of 0.5 dB/km at 1304.5 nm plus an allocation for connection and splice loss given in 121.11.2.1.

^b Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

121.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

121.8.1 Test patterns for optical parameters

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 121–10 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 121–10 may be used to perform that test. The test patterns used in this clause are shown in Table 121–9.

Table 121–9—Test patterns

| Pattern | Pattern description | Defined in |
|-------------|---------------------------------|--------------|
| Square wave | Square wave (8 threes, 8 zeros) | 120.5.11.2.4 |
| 3 | PRBS31Q | 120.5.11.2.2 |
| 4 | PRBS13Q | 120.5.11.2.1 |
| 5 | Scrambled idle | 119.2.4.9 |
| 6 | SSPRQ | 120.5.11.2.3 |

Table 121–10—Test-pattern definitions and related subclauses

| Parameter | Pattern | Related subclause |
|--|--|-------------------|
| Wavelength | Square wave, 3, 4, 5, 6 or valid 200GBASE-R signal | 121.8.2 |
| Side mode suppression ratio | 3, 5, 6 or valid 200GBASE-R signal | — |
| Average optical power | 3, 5, 6 or valid 200GBASE-R signal | 121.8.3 |
| Outer Optical Modulation Amplitude (OMA _{outer}) | 4 or 6 | 121.8.4 |
| Transmitter and dispersion eye closure for PAM4 (TDECQ) | 6 | 121.8.5 |
| Extinction ratio | 4 or 6 | 121.8.6 |
| RIN _{21.4OMA} | Square wave | 121.8.7 |
| Stressed receiver conformance test signal calibration | 6 | 121.8.9.2 |
| Stressed receiver sensitivity | 3 or 5 | 121.8.9 |

121.8.2 Wavelength

The wavelength of each optical lane shall be within the range given in Table 121–6 if measured per TIA/EIA-455-127-A or IEC 61280-1-3. The lane under test is modulated using the test pattern defined in Table 121–10.

121.8.3 Average optical power

The average optical power of each lane shall be within the limits given in Table 121–6 if measured using the methods given in IEC 61280-1-1. The average optical power is measured using the test pattern defined in Table 121–10, per the test setup in Figure 53-6.

121.8.4 Outer Optical Modulation Amplitude (OMA_{outer})

The OMA_{outer} of each lane shall be within the limits given in Table 121–6. The OMA_{outer} is measured using a test pattern specified for OMA_{outer} in Table 121–10 as the difference between the average optical launch power level P_3 , measured over the central 2 UI of a run of 7 threes, and the average optical launch power level P_0 , measured over the central 2 UI of a run of 6 zeros, as shown in Figure 121–3.




Figure 121–3—Example power levels P_0 and P_3 from PRBS13Q test pattern

121.8.5 Transmitter and dispersion eye closure for PAM4 (TDECQ)

The TDECQ of each lane shall be within the limits given in Table 121–6 if measured using the methods specified in 121.8.5.1, 121.8.5.2, and 121.8.5.3.

TDECQ is a measure of each optical transmitter's vertical eye closure when transmitted through a worst case optical channel (specified in 121.8.5.2), as measured through an optical to electrical converter (O/E) and oscilloscope with the combined frequency response given in 121.8.5.1, and equalized with the reference equalizer (as described in 121.8.5.4). The reference receiver and equalizer may be implemented in software or may be part of an oscilloscope.

Table 121–10 specifies the test patterns to be used for measurement of TDECQ.

121.8.5.1 TDECQ conformance test setup

A block diagram for the TDECQ conformance test is shown in Figure 121–4. Other equivalent measurement implementations may be used with suitable calibration.

Each optical lane is tested individually with all other lanes in operation and all lanes using the same test pattern. There shall be at least 31 UI delay between the test pattern on one lane and the pattern on any other lane, so that the symbols on each lane are not correlated within the PMD. The optical splitter and variable reflector are adjusted so that each transmitter is tested with the optical return loss specified in Table 121–11. The state of polarization of the back reflection is adjusted to create the greatest RIN. Each optical lane is tested with the optical channel described in 121.8.5.2. The combination of the O/E and the oscilloscope has

a fourth-order Bessel-Thomson filter response with a bandwidth of approximately 13.28125 GHz. Compensation may be made for any deviation from an ideal fourth-order Bessel-Thomson response.

The test pattern (specified in Table 121–10) is transmitted repetitively by the optical lane under test and the oscilloscope is set up to capture the complete pattern for TDECQ analysis as described in 121.8.5.3. The clock recovery unit (CRU) has a corner frequency of 4 MHz and a slope of 20 dB/decade. The CRU can be implemented in hardware or software depending on oscilloscope technology.

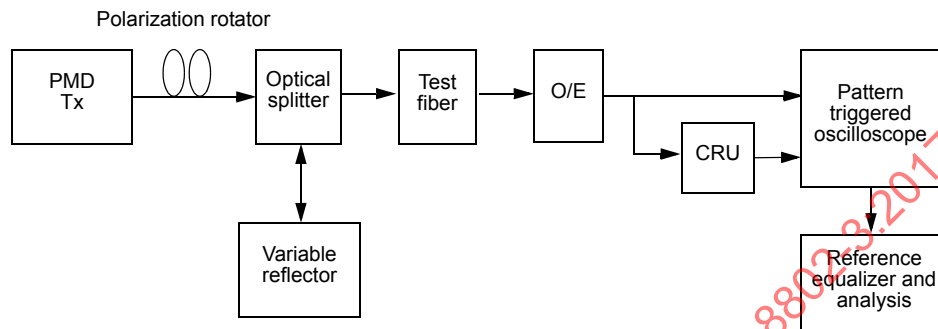


Figure 121–4—TDECQ conformance test block diagram

121.8.5.2 Channel requirements

The transmitter is tested using an optical channel that meets the requirements listed in Table 121–11.

Table 121–11—Transmitter compliance channel specifications

| PMD type | Dispersion ^a (ps/nm) | | Insertion loss ^b | Optical return loss ^c | Max mean DGD |
|--------------|---|---|-----------------------------|----------------------------------|--------------|
| | Minimum | Maximum | | | |
| 200GBASE-DR4 | $0.011625 \times \lambda \times [1 - (1324 / \lambda)^4]$ | $0.011625 \times \lambda \times [1 - (1300 / \lambda)^4]$ | Minimum | 21.4 dB | 0.5 ps |

^a The dispersion is measured for the wavelength of the device under test (λ in nm). The coefficient assumes 500 m for 200GBASE-DR4.

^b There is no intent to stress the sensitivity of the O/E converter associated with the oscilloscope.

^c The optical return loss is applied at TP2.

A 200GBASE-DR4 transmitter is to be compliant with a total dispersion at least as negative as the “minimum dispersion” and at least as positive as the “maximum dispersion” columns specified in Table 121–11 for the wavelength of the device under test. This may be achieved with channels consisting of fibers with lengths chosen to meet the dispersion requirements.

To verify that the fiber has the correct amount of dispersion, the measurement method defined in IEC 60793-1-42 may be used. The measurement is made in the linear power regime of the fiber.

The channel provides an optical return loss specified in Table 121–11. The state of polarization of the back reflection is adjusted to create the greatest RIN.

The mean DGD of the channel is to be less than the value specified in Table 121–11.

121.8.5.3 TDECQ measurement method

The standard deviation of the noise of the O/E and oscilloscope combination, σ_s , is determined with no optical input signal and the same settings as used to capture the histograms described below.

OMA_{outer} is measured according to 121.8.4 on the equalized signal.

The test pattern specified for TDECQ (see Table 121–10) is transmitted repetitively by the optical lane under test and the oscilloscope is set up to capture samples from all symbols in the complete pattern without averaging.

If an equivalent-time sampling oscilloscope is used, the impact of the sampling process and the reference equalizer on transmitter noise must be compensated for, so that the correct magnitude of noise is present at the output of the equalizer.

The captured waveform is processed to find the largest noise that could be combined with the signal by an ideal reference receiver when optimally equalized by a reference equalizer. The optimal equalizer tap coefficients are dependent on the amount of noise that can be added to the signal, so finding the noise that can be added and the optimal equalizer setting is an iterative process. One way of doing this, using estimated PAM4 symbol error ratio as the figure of merit for the equalized signal, is described below.

The reference equalizer (specified in 121.8.5.4) is applied to the waveform. The sum of the equalizer tap coefficients is equal to 1. An eye diagram is formed from the equalized captured waveform.

The average optical power (P_{ave}) of the equalized eye diagram is determined, and the 0 UI and 1 UI crossing points are determined by the average of the eye diagram crossing times, as measured at P_{ave} , as illustrated in Figure 121–5.




Figure 121–5—Illustration of the TDECQ measurement

Two vertical histograms are measured through the eye diagram, nominally centered at 0.45 UI and 0.55 UI. Each of the histogram windows spans all of the modulation levels of the eye diagram, as illustrated in Figure 121–5. The precise time position of the pair of histograms is adjusted to minimize TDECQ while keeping the histograms spaced 0.1 UI apart.

Each histogram window has a width of 0.04 UI. Each histogram window has outer height boundaries which are set beyond the extremes of the eye diagram (so that no further samples would be captured by increasing the vertical separation of the height boundaries).

The sub-eye threshold levels P_{th1} , P_{th2} , and P_{th3} , are determined from the OMA_{outer} and the average optical power of the eye diagram (P_{ave}) as defined in Equation (121-1), Equation (121-2), and Equation (121-3) and illustrated in Figure 121-5.

$$P_{th1} = P_{ave} - \frac{OMA_{outer}}{3} \quad (121-1)$$

$$P_{th2} = P_{ave} \quad (121-2)$$

$$P_{th3} = P_{ave} + \frac{OMA_{outer}}{3} \quad (121-3)$$

Each captured histogram is processed to, in effect, combine the PAM4 waveform with noise, in order to produce an estimate of the partial PAM4 symbol error ratio (SER) for each sub-eye. One way of doing this is described below.

The left and right histograms are each normalized, and can be represented as a series of equally spaced optical power values (y_i) with separation Δy , each with an associated fraction $F(y_i)$, equal to the number of samples captured in that power interval divided by the total number of samples in that histogram. The sum of all $F(y_i)$ for each histogram is equal to 1.

From the left normalized histogram $F(y_i)$, three cumulative probability functions are created, $CF_{L1}(y_i)$, $CF_{L2}(y_i)$, and $CF_{L3}(y_i)$, one around each sub-eye threshold. The right histogram is treated similarly to create three cumulative probability functions $CF_{R1}(y_i)$, $CF_{R2}(y_i)$, and $CF_{R3}(y_i)$. $CF_{L1}(y_i)$ is defined in Equation (121-4).

$$CF_{L1}(y_i) = \begin{cases} \sum_{y = P_{th1}}^{y_i} F(y) & \text{for } y_i \geq P_{th1} \\ \sum_{y = y_i}^{P_{th1}} F(y) & \text{for } y_i < P_{th1} \end{cases} \quad (121-4)$$

Each element of the cumulative probability function, $CF_{L1}(y_i)$, is multiplied by a value $G_{th1}(y_i)$, and then summed to calculate an approximation for SER_{L1} , the partial SER for threshold 1. Each element of the cumulative probability function, $CF_{L2}(y_i)$, is multiplied by a value $G_{th2}(y_i)$, and then summed to calculate an approximation for SER_{L2} . Each element of the cumulative probability function, $CF_{L3}(y_i)$, is multiplied by a value $G_{th3}(y_i)$, and then summed to calculate an approximation for SER_{L3} . The sum of the three partial SERs is the SER associated with the left histogram, SER_L .

Each element of the cumulative probability function, $CF_{R1}(y_i)$, is multiplied by a value $G_{th1}(y_i)$, and then summed to calculate an approximation for SER_{R1} , the partial SER for threshold 1. $CF_{R2}(y_i)$ and $CF_{R3}(y_i)$ are treated similarly to calculate SER_{R2} , and SER_{R3} , the partial SERs for threshold 2 and threshold 3. The sum of the three partial SERs is the SER associated with the right histogram, SER_R .

$G_{th1}(y_i)$ is equivalent to a Gaussian probability density function with an RMS value of σ_G , centered around the sub-eye threshold P_{th1} . $G_{th1}(y_i)$ is given by Equation (121-5) and can be estimated by Equation (121-6).

$$G_{th1}(y_i) = \int_{y_i - \frac{\Delta y}{2}}^{y_i + \frac{\Delta y}{2}} \frac{1}{C_{eq} \sigma_G \sqrt{2\pi}} \times e^{-\left(\frac{y - P_{th1}}{C_{eq} \sigma_G \sqrt{2}}\right)^2} dy \quad (121-5)$$

$$G_{th1}(y_i) = \frac{1}{C_{eq} \sigma_G \sqrt{2\pi}} \times e^{-\left(\frac{y_i - P_{th1}}{C_{eq} \sigma_G \sqrt{2}}\right)^2} \times \Delta y \quad (121-6)$$

$G_{th2}(y_i)$ and $G_{th3}(y_i)$ are similar Gaussian probability density functions with the same RMS value of σ_G , centered around the sub-eye thresholds P_{th2} and P_{th3} respectively. $G_{th2}(y_i)$ and $G_{th3}(y_i)$ are given by Equation (121-7) and Equation (121-8) respectively.

$$G_{th2}(y_i) = \int_{y_i - \frac{\Delta y}{2}}^{y_i + \frac{\Delta y}{2}} \frac{1}{C_{eq} \sigma_G \sqrt{2\pi}} \times e^{-\left(\frac{y - P_{th2}}{C_{eq} \sigma_G \sqrt{2}}\right)^2} dy \quad (121-7)$$

$$G_{th3}(y_i) = \int_{y_i - \frac{\Delta y}{2}}^{y_i + \frac{\Delta y}{2}} \frac{1}{C_{eq} \sigma_G \sqrt{2\pi}} \times e^{-\left(\frac{y - P_{th3}}{C_{eq} \sigma_G \sqrt{2}}\right)^2} dy \quad (121-8)$$

where

C_{eq} is a coefficient which accounts for the reference equalizer noise enhancement

The value of C_{eq} can be calculated from the product of the normalized noise power density spectrum $N(f)$ at the input of the reference equalizer and the normalized frequency response $H_{eq}(f)$ of the reference equalizer, as shown in Equation (121-9).

$$C_{eq} = \sqrt{\int_f N(f) \times |H_{eq}(f)|^2 df} \quad (121-9)$$

where

$N(f)$ is the normalized noise power density spectrum equivalent to white noise filtered by a fourth-order Bessel-Thomson response filter with a bandwidth of 13.28125 GHz.

and

$$\int_f N(f) df = H_{eq}(f=0) = 1 \quad (121-10)$$

The equalizer tap coefficients are iteratively adjusted and SER_L and SER_R calculated until the largest of SER_L and SER_R is minimized. Then, if the larger of SER_L and SER_R is greater than the target SER of 4.8×10^{-4} , the value of σ_G is decreased and the process of equalizer optimization is repeated; If the larger of SER_L and SER_R is lower than the target SER of 4.8×10^{-4} , then the value of σ_G is increased and the process of equalizer optimization is repeated.

When the larger of SER_L and SER_R is equal to the target SER of 4.8×10^{-4} , and the value of σ_G cannot be increased by further optimization of the equalizer tap coefficients, then TDECQ is calculated.

The RMS noise, R , that could be added by a receiver is given by Equation (121–11).

$$R = \sqrt{\sigma_G^2 + \sigma_S^2} \quad (121-11)$$

TDECQ is given by Equation (121–12).

$$TDECQ = 10\log_{10}\left(\frac{OMA_{\text{outer}}}{6} \times \frac{1}{Q_t R}\right) \quad (121-12)$$

where

OMA_{outer} is the Outer Optical Modulation Amplitude as defined in 121.8.4

Q_t is 3.414 consistent with the BER and target symbol error ratio for Gray coded PAM4

121.8.5.4 TDECQ reference equalizer

The reference equalizer for 200GBASE-DR4 is a 5 tap, T spaced, feed-forward equalizer (FFE), where T is the symbol period. The sum of the equalizer tap coefficients is equal to 1.

NOTE—This reference equalizer is part of the TDECQ test and does not imply any particular receiver equalizer implementation.

121.8.6 Extinction ratio

The extinction ratio of each lane shall be within the limits given in Table 121–6 if measured using a test pattern specified for extinction ratio in Table 121–10. The extinction ratio of a PAM4 optical signal is defined as the ratio of the average optical launch power level P_3 , measured over the central 2 UI of a run of 7 threes, and the average optical launch power level P_0 , measured over the central 2 UI of a run of 6 zeros, as shown in Figure 121–3.

121.8.7 Relative intensity noise (RIN_{21.4OMA})

RIN shall be as defined by the measurement methodology of 52.9.6 with the following exceptions:

- a) The optical return loss is 21.4 dB.
- b) Each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm.
- c) The upper –3 dB limit of the measurement apparatus is to be approximately equal to the signaling rate (i.e., 26.6 GHz).
- d) The test pattern is according to Table 121–10.

121.8.8 Receiver sensitivity

Receiver sensitivity, which is defined for an input signal with SECQ of 0.9 dB (e.g., an ideal input signal without overshoot), is informative and compliance is not required. If measured, the test signal should have negligible impairments such as intersymbol interference (ISI), rise/fall times, jitter and RIN. Instead, the normative requirement for receivers is stressed receiver sensitivity.

121.8.9 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 121–7 if measured using the method defined in 121.8.9.1 and 121.8.9.3, with the conformance test signal at TP3 as described in 121.8.9.2, using the test pattern specified for stressed receiver sensitivity in Table 121–10. The BER is required to be met for the lane under test on its own.

Stressed receiver sensitivity is defined with all transmit and receive lanes in operation. Any of the patterns specified for stressed receiver sensitivity in Table 121–10 is sent from the transmit section of the PMD under test. The signal being transmitted is asynchronous to the received signal.

121.8.9.1 Stressed receiver conformance test block diagram

A block diagram for the receiver conformance test is shown in Figure 121–6. The patterns used for the received conformance signal are specified in Table 121–10. The optical test signal is conditioned (stressed) using the stressed receiver methodology defined in 121.8.9.2 and has sinusoidal jitter applied as specified in 121.8.9.4. A suitable test set is needed to characterize the signal used to test the receiver. Stressed receiver conformance test signal verification is described in 121.8.9.3.

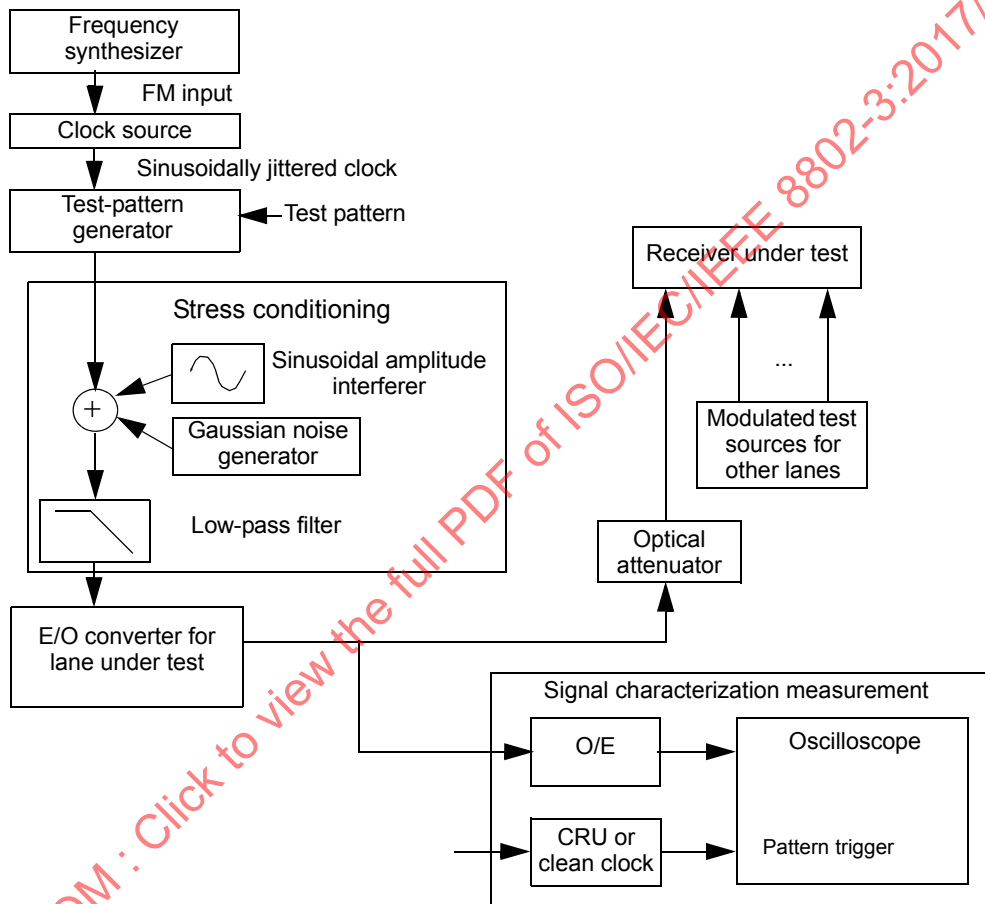


Figure 121–6—Stressed receiver conformance test block diagram

The low-pass filter is used to create ISI. The combination of the low-pass filter and the E/O converter should have a frequency response that results in at least half of the dB value of the stressed eye closure (SECQ) specified in Table 121–7 before the sinusoidal and Gaussian noise terms are added, according to the methods specified in 121.8.9.2. The sinusoidal amplitude interferer causes additional eye closure, but in conjunction with the finite edge rates, also causes some jitter.

The sinusoidally jittered clock represents other forms of jitter and also verifies that the receiver under test can track low-frequency jitter. The sinusoidal amplitude interferer may be set at any frequency between

100 MHz and 2 GHz, although care should be taken to avoid harmonic relationships between the sinusoidal interferer, the sinusoidal jitter, the signaling rate, and the pattern repetition rate. The Gaussian noise generator, the amplitude of the sinusoidal interferer, and the low-pass filter are adjusted so that the SECQ specified in Table 121-7 is met, according to the methods specified in 121.8.9.2.

For improved visibility for calibration, all elements in the signal path (cables, DC blocks, E/O converter, etc.) should have wide and smooth frequency response, and linear phase response, throughout the spectrum of interest. Baseline wander and overshoot and undershoot should be negligible.

121.8.9.2 Stressed receiver conformance test signal characteristics and calibration

The conformance test signal is used to validate that each lane of the PMD receiver meets BER requirements with near worst-case waveforms at TP3.

The primary parameters of the stressed receiver conformance test signal are its stressed eye closure (SECQ), and the sinusoidal jitter applied, as specified in 121.8.9.4. The SECQ of the stressed receiver conformance test signal is measured according to 121.8.5, except that the test fiber is not used. An example stressed receiver conformance test setup is shown in Figure 121-6; however, alternative test setups that generate equivalent stress conditions may be used.

The following steps describe a possible method for setting up and calibrating a stressed receiver conformance test signal when using a stressed receiver conformance test setup as shown in Figure 121-6:

- 1) Set the signaling rate of the test pattern generator to meet the requirements in Table 121-6.
- 2) With the sinusoidal jitter, sinusoidal interferer, and the Gaussian noise generator turned off, set the extinction ratio of the E/O converter to approximately the minimum specified in Table 121-6.
- 3) The required value of SECQ is given in Table 121-7. With the sinusoidal jitter, sinusoidal interferer, and Gaussian noise generator turned off, at least half of the dB value of SECQ should be created by the selection of the appropriate bandwidth for the combination of the low-pass filter and the E/O converter. Any remaining SECQ must be created with a combination of sinusoidal jitter, sinusoidal interference, and Gaussian noise. Sinusoidal jitter is added as specified in Table 121-12.

When calibrating the conformance signal, the sinusoidal jitter frequency should be between 50 MHz and 10 times LB as defined in Table 121-12. Sinusoidal jitter amplitude may be calibrated by measuring the jitter on the oscilloscope, while transmitting the square wave pattern, as defined in Table 121-9, and using a clean clock in place of the CRU to trigger the oscilloscope.

Iterate the adjustments of the sinusoidal interferer, the Gaussian noise generator, and extinction ratio, until the required value of SECQ is met, while also meeting the following conditions: the extinction ratio is approximately the minimum specified in Table 121-6; and sinusoidal jitter is as specified in Table 121-12.

Each receiver lane is conformance tested in turn. The source for the lane under test is adjusted to supply a signal at the input to the receiver under test at the “Stressed receiver sensitivity (OMA_{outer}), each lane (max)” specified in Table 121-7, and the test sources for the other lanes are set to the “ OMA_{outer} of each aggressor lane” specified in Table 121-7.

121.8.9.3 Stressed receiver conformance test signal verification

The SECQ of the stressed receiver conformance test signal is measured according to 121.8.5, except that the test fiber is not used. The clock output from the clock source in Figure 121-6 is modulated with the sinusoidal jitter. To use an oscilloscope to calibrate the final stressed eye jitter that includes the sinusoidal jitter component, a separate clock source (clean clock of Figure 121-6) is required that is synchronized to the source clock, but not modulated with the jitter source.

Care should be taken when characterizing the test signal because excessive noise/jitter in the measurement system would result in an input signal that does not fully stress the receiver under test. Running the receiver tolerance test with a signal that is under-stressed may result in the deployment of non-compliant receivers. The noise/jitter introduced by the O/E, filters, and oscilloscope should be negligible or the results should be corrected for its effects. While the details of test equipment are beyond the scope of this standard, it is recommended that the implementer fully characterize the test equipment and apply appropriate guard bands so that the stressed receiver conformance input signal meets the stress and sinusoidal jitter specified in 121.8.9.2 and 121.8.9.4.

121.8.9.4 Sinusoidal jitter for receiver conformance test

The sinusoidal jitter is used to test receiver jitter tolerance. The amplitude of the applied sinusoidal jitter is dependent on frequency as specified in Table 121–12 and is illustrated in Figure 121–7.

Table 121–12—Applied sinusoidal jitter

| Frequency range | Sinusoidal jitter, peak-to-peak (UI) |
|---|--------------------------------------|
| $f < 40 \text{ kHz}$ | Not specified |
| $40 \text{ kHz} < f \leq 4 \text{ MHz}$ | $2 \times 10^5 / f$ |
| $4 \text{ MHz} < f < 10 \text{ LB}^a$ | 0.05 |

^a LB = loop bandwidth; upper frequency bound for added sine jitter should be at least 10 times the loop bandwidth of the receiver being tested.




Figure 121–7—Illustration of the mask of the sinusoidal component of jitter tolerance

121.9 Safety, installation, environment, and labeling

121.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

121.9.2 Laser safety

200GBASE-DR4 optical transceivers shall conform to Hazard Level 1 laser requirements as defined in IEC 60825-1 and IEC 60825-2, under any condition of operation. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product's laser, safety features, labeling, use, maintenance, and service. This documentation explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.⁶

121.9.3 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

121.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 200GBASE-DR4 PMD over the life of the product while the product operates within the manufacturer's range of environmental, power, and other specifications.

It is recommended that manufacturers indicate in the literature associated with the PHY the operating environmental conditions to facilitate selection, installation, and maintenance.

It is recommended that manufacturers indicate in the literature associated with the components of the optical link, the distance and operating environmental conditions over which the specifications of this clause will be met.

121.9.5 Electromagnetic emission

A system integrating a 200GBASE-DR4 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

121.9.6 Temperature, humidity, and handling

The optical link is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

121.9.7 PMD labeling requirements

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user, with at least the applicable safety warnings and the applicable port type designation (e.g., 200GBASE-DR4).

Labeling requirements for Hazard Level 1 lasers are given in the laser safety standards referenced in 121.9.2.

⁶A host system that fails to meet the manufacturer's requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.

121.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 121–8.




Figure 121–8—Fiber optic cabling model

The channel insertion loss is given in Table 121–13. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, and polarization mode dispersion meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with ANSI/TIA/EIA-526-7/method A-1. The fiber optic cabling model (channel) defined here is the same as a simplex fiber optic link segment. The term channel is used here for consistency with generic cabling standards.

Table 121–13—Fiber optic cabling (channel) characteristics

| Description | 200GBASE-DR4 | Unit |
|--|--------------|-------|
| Operating distance (max) | 500 | m |
| Channel insertion loss ^{a, b} (max) | 3 | dB |
| Channel insertion loss (min) | 0 | dB |
| Positive dispersion ^b (max) | 0.8 | ps/nm |
| Negative dispersion ^b (min) | –0.93 | ps/nm |
| DGD_max ^c | 2.24 | ps |
| Optical return loss (min) | 37 | dB |

^a These channel insertion loss values include cable, connectors, and splices.

^b Over the wavelength range 1304.5 nm to 1317.5 nm.

^c Differential Group Delay (DGD) is the time difference at reception between the fractions of a pulse that were transmitted in the two principal states of polarization of an optical signal. DGD_max is the maximum differential group delay that the system must tolerate.

121.11 Characteristics of the fiber optic cabling (channel)

The 200GBASE-DR4 fiber optic cabling shall meet the specifications defined in Table 121–14. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

121.11.1 Optical fiber cable

The fiber optic cable requirements are satisfied by cables containing IEC 60793-2-50 type B1.1 (dispersion un-shifted single-mode), type B1.3 (low water peak single-mode), or type B6_a (bend insensitive) fibers or the requirements in Table 121–14 where they differ.

Table 121–14—Optical fiber and cable characteristics

| Description | Value | Unit |
|--|---------------------------------|-----------------------|
| Nominal fiber specification wavelength | 1310 | nm |
| Cabled optical fiber attenuation (max) | 0.5 ^a | dB/km |
| Zero dispersion wavelength (λ_0) | $1300 \leq \lambda_0 \leq 1324$ | nm |
| Dispersion slope (max) (S_0) | 0.093 | ps/nm ² km |

^a The 0.5 dB/km attenuation is provided for Outside Plant cable as defined in ANSI/TIA 568-C.3.

121.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 121–8, consists of a mated pair of optical connectors.

121.11.2.1 Connection insertion loss

The maximum link distance is based on an allocation of 2.75 dB total connection and splice loss. For example, this allocation supports five connections with an average insertion loss per connection of 0.5 dB. Connections with different loss characteristics may be used provided the requirements of Table 121–13 are met.

121.11.2.2 Maximum discrete reflectance

The maximum value for each discrete reflectance shall be less than or equal to the value shown in Table 121–15 corresponding to the number of discrete reflectances above –55 dB within the channel. For numbers of discrete reflectances in between two numbers shown in the table, the lower of the two corresponding maximum discrete reflectance values applies.

Table 121–15—Maximum value of each discrete reflectance

| Number of discrete reflectances above –55 dB | Maximum value for each discrete reflectance |
|--|---|
| 1 | –37 dB |
| 2 | –42 dB |
| 4 | –45 dB |
| 6 | –47 dB |
| 8 | –48 dB |
| 10 | –49 dB |

121.11.3 Medium Dependent Interface (MDI)

The 200GBASE-DR4 PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the “fiber optic cabling” (as shown in Figure 121–8). The 200GBASE-DR4 PMD is coupled to the fiber optic cabling through one connector plug into the MDI optical receptacle as shown in Figure 121–10. Example constructions of the MDI include the following:

- PMD with a connectorized fiber pigtail plugged into an adapter
- PMD receptacle

121.11.3.1 Optical lane assignments

The four transmit and four receive optical lanes of 200GBASE-DR4 shall occupy the positions depicted in Figure 121–9 when looking into the MDI receptacle with the connector keyway feature on top. The interface contains eight active lanes within twelve total positions. The transmit optical lanes occupy the left-most four positions. The receive optical lanes occupy the right-most four positions. The four center positions are unused.




Figure 121–9—200GBASE-DR4 optical lane assignments

121.11.3.2 Medium Dependent Interface (MDI) requirements

The MDI shall meet the dimensional specifications of IEC 61754-7-1 interface 7-1-9: *MPO device receptacle, angled interface*. The plug terminating the optical fiber cabling shall meet the dimensional specifications of IEC 61754-7-1 interface 7-1-1: *MPO female plug connector, down-angled interface for 2 to 12 fibres*. The MDI shall optically mate with the plug on the optical fiber cabling. Figure 121–10 shows an MPO female plug connector with down-angled interface, and an MDI as an active device receptacle with angled interface.




Figure 121–10—MPO female plug with down-angled interface and MDI active device receptacle with angled interface

The MDI shall meet the interface performance specifications of IEC 61753-021-2 for performance level D/2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 121.5.1, not at the MDI.

121.12 Protocol implementation conformance statement (PICS) proforma for Clause 121, Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-DR4⁷

121.12.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 121, Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-DR4, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

121.12.2 Identification

121.12.2.1 Implementation identification

| | |
|--|--|
| Supplier ¹ | |
| Contact point for inquiries about the PICS ¹ | |
| Implementation Name(s) and Version(s) ^{1,3} | |
| Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ² | |
| NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model). | |

121.12.2.2 Protocol summary

| | |
|---|---|
| Identification of protocol standard | IEEE Std 802.3bs-2017, Clause 121, Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-DR4 |
| Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS | |
| Have any Exception items been required? No [] Yes [] (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3bs-2017.) | |

| | |
|-------------------|--|
| Date of Statement | |
|-------------------|--|

⁷Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

121.12.3 Major capabilities/options

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--|--------|-------------------|
| *INS | Installation / cable | 121.10 | Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer | O | Yes [] No [] |
| TP1 | Reference point TP1 exposed and available for testing | 121.5.1 | This point may be made available for use by implementers to certify component conformance | O | Yes [] No [] |
| TP4 | Reference point TP4 exposed and available for testing | 121.5.1 | This point may be made available for use by implementers to certify component conformance | O | Yes [] No [] |
| DC | Delay constraints | 121.3.1 | Device conforms to delay constraints | M | Yes [] |
| SC | Skew constraints | 121.3.2 | Device conforms to Skew and Skew Variation constraints | M | Yes [] |
| *MD | MDIO capability | 121.4 | Registers and interface supported | O | Yes [] No [] |

121.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-DR4

121.12.4.1 PMD functional specifications

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--|--------|-------------------|
| F1 | Compatible with 200GBASE-R PCS and PMA | 121.1 | | M | Yes [] |
| F2 | Integration with management functions | 121.1 | | O | Yes [] No [] |
| F3 | Bit error ratio | 121.1.1 | Meets the BER specified in 121.1.1 | M | Yes [] |
| F4 | Transmit function | 121.5.2 | Conveys symbols from PMD service interface to MDI | M | Yes [] |
| F5 | Mapping between optical signal and logical signal for transmitter | 121.5.2 | Highest optical power corresponds to tx_symbol = three | M | Yes [] |
| F6 | Receive function | 121.5.3 | Conveys symbols from MDI to PMD service interface | M | Yes [] |
| F7 | Conversion of four optical signals to four electrical signals | 121.5.3 | For delivery to the PMD service interface | M | Yes [] |
| F8 | Mapping between optical signal and logical signal for receiver | 121.5.3 | Highest optical power corresponds to rx_symbol = three | M | Yes [] |

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|-------------------------------------|-----------|--|--------|------------------------------|
| F9 | Global Signal Detect function | 121.5.4 | Report to the PMD service interface the message PMD:IS_SIGNAL.indication (SIGNAL_DETECT) | M | Yes [] |
| F10 | Global Signal Detect behavior | 121.5.4 | SIGNAL_DETECT is a global indicator of the presence of optical signals on all four lanes | M | Yes [] |
| F11 | Lane-by-lane Signal Detect function | 121.5.5 | Sets PMD_signal_detect_i values on a lane-by-lane basis per requirements of Table 121-4 | MD:O | Yes [] No [] N/A [] |
| F12 | PMD reset function | 121.5.6 | Resets the PMD sublayer | MD:O | Yes [] No [] N/A [] |

121.12.4.2 Management functions

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|--|--------|------------------------------|
| M1 | Management register set | 121.4 | | MD:M | Yes [] N/A [] |
| M2 | Global transmit disable function | 121.5.7 | Disables all of the optical transmitters with the PMD_global_transmit_disable variable | MD:O | Yes [] No [] N/A [] |
| M3 | PMD_lane_by_lane_transmit_disable function | 121.5.8 | Disables the optical transmitter on the lane associated with the PMD_transmit_disable_i variable | MD:O | Yes [] No [] N/A [] |
| M4 | PMD_fault function | 121.5.9 | Sets PMD_fault to one if any local fault is detected | MD:O | Yes [] No [] N/A [] |
| M5 | PMD_transmit_fault function | 121.5.10 | Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane | MD:O | Yes [] No [] N/A [] |
| M6 | PMD_receive_fault function | 121.5.11 | Sets PMD_receive_fault to one if a local fault is detected on any receive lane | MD:O | Yes [] No [] N/A [] |

121.12.4.3 PMD to MDI optical specifications for 200GBASE-DR4

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--------------------------|--------|---------|
| DR1 | Transmitter meets specifications in Table 121-6 | 121.7.1 | Per definitions in 121.8 | M | Yes [] |
| DR2 | Receiver meets specifications in Table 121-7 | 121.7.2 | Per definitions in 121.8 | M | Yes [] |

121.12.4.4 Optical measurement methods

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---|--------|---------|
| OM1 | Measurement cable | 121.8 | 2 m to 5 m in length | M | Yes [] |
| OM2 | Center wavelength | 121.8.2 | Per TIA/EIA-455-127-A or IEC 61280-1-3 under modulated conditions | M | Yes [] |
| OM3 | Average optical power | 121.8.3 | Per IEC 61280-1-1 | M | Yes [] |
| OM4 | OMA measurements | 121.8.4 | Each lane | M | Yes [] |
| OM5 | Transmitter and dispersion eye closure for PAM4 (TDECQ) | 121.8.5 | Each lane | M | Yes [] |
| OM6 | Extinction ratio | 121.8.6 | Each lane | M | Yes [] |
| OM7 | Stressed receiver sensitivity | 121.8.9 | Each lane | M | Yes [] |

121.12.4.5 Environmental specifications

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---------------------------------|-----------|--|--------|---------|
| ES1 | General safety | 121.9.1 | Conforms to IEC 60950-1 | M | Yes [] |
| ES2 | Laser safety—IEC Hazard Level 1 | 121.9.2 | Conforms to Hazard Level 1 laser requirements defined in IEC 60825-1 and IEC 60825-2 | M | Yes [] |
| ES3 | Electromagnetic interference | 121.9.5 | Complies with applicable local and national codes for the limitation of electromagnetic interference | M | Yes [] |

121.12.4.6 Characteristics of the fiber optic cabling and MDI

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|------------------------------|------------|--|--------|--------------------|
| OC1 | Fiber optic cabling | 121.11 | Meets requirements specified in Table 121–13 | INS:M | Yes [] N/A [] |
| OC2 | Maximum discrete reflectance | 121.11.2.2 | Meets requirements specified in Table 121–15 | INS:M | Yes [] N/A [] |
| OC3 | MDI layout | 121.11.3.1 | Optical lane assignments per Figure 121–9 | M | Yes [] |
| OC4 | MDI dimensions | 121.11.3.2 | Per IEC 61754-7-1 interface 7-1-9 | M | Yes [] |
| OC5 | Cabling connector dimensions | 121.11.3.2 | Per IEC 61754-7-1 interface 7-1-1 | INS:M | Yes [] N/A [] |
| OC6 | MDI mating | 121.11.3.2 | MDI optically mates with plug on the cabling | M | Yes [] |

**122. Physical Medium Dependent (PMD) sublayer and medium, type
200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8****122.1 Overview**

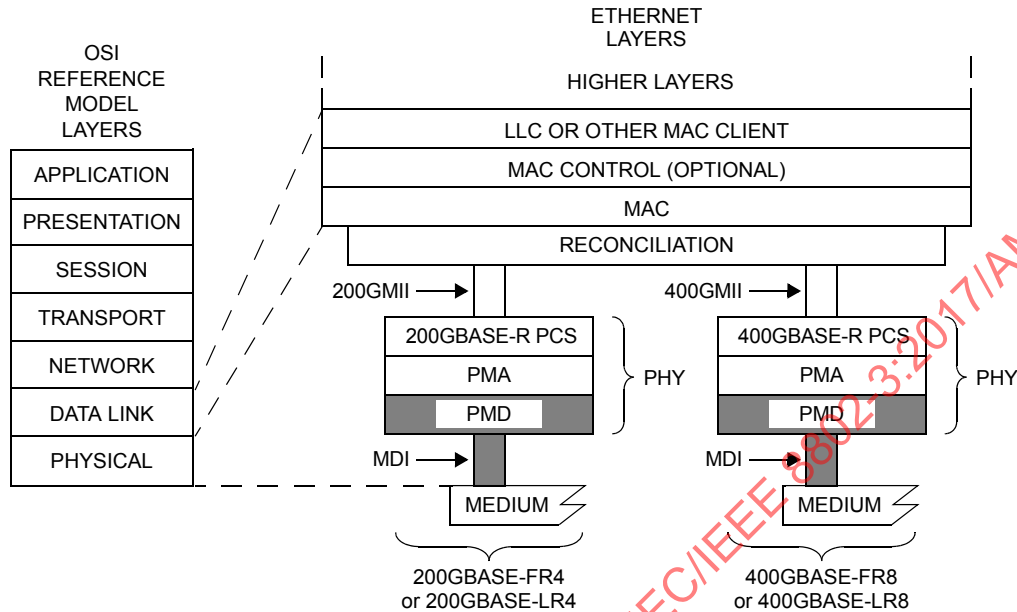
This clause specifies the 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and the 400GBASE-LR8 PMDs together with the single-mode fiber medium. The optical signals generated by these four PMD types are modulated using a 4-level pulse amplitude modulation (PAM4) format. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 122–1, to the medium through the MDI and optionally with the management functions that may be accessible through the management interface defined in Clause 45, or equivalent.

Table 122–1—Physical Layer clauses associated with the 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 PMDs

| Associated clause | 200GBASE-FR4, 200GBASE-LR4 | 400GBASE-FR8, 400GBASE-LR8 |
|--------------------------------|-------------------------------|-------------------------------|
| 117—RS | Required | Required |
| 117—200GMII ^a | Optional | Not applicable |
| 117—400GMII ^a | Not applicable | Optional |
| 118—200GMII Extender | Optional | Not applicable |
| 118—400GMII Extender | Not applicable | Optional |
| 119—PCS for 200GBASE-R | Required | Not applicable |
| 119—PCS for 400GBASE-R | Not applicable | Required |
| 120—PMA for 200GBASE-R | Required | Not applicable |
| 120—PMA for 400GBASE-R | Not applicable | Required |
| 120B—Chip-to-chip 200GAUI-8 | Optional | Not applicable |
| 120B—Chip-to-chip 400GAUI-16 | Not applicable | Optional |
| 120C—Chip-to-module 200GAUI-8 | Optional | Not applicable |
| 120C—Chip-to-module 400GAUI-16 | Not applicable | Optional |
| 120D—Chip-to-chip 200GAUI-4 | Optional | Not applicable |
| 120D—Chip-to-chip 400GAUI-8 | Not applicable | Optional |
| 120E—Chip-to-module 200GAUI-4 | Optional | Not applicable |
| 120E—Chip-to-module 400GAUI-8 | Not applicable | Optional |
| 78—Energy Efficient Ethernet | Optional | Optional |

^a 200GMII and 400GMII are optional interfaces. However, if the appropriate interface is not implemented, a conforming implementation must behave functionally as though the RS and 200GMII or 400GMII were present.

Figure 122–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 200 Gb/s and 400 Gb/s Ethernet are introduced in Clause 116 and the purpose of each PHY sublayer is summarized in 116.2.



400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
 200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 FR = PMD FOR SINGLE-MODE FIBER — 2 km
 LR = PMD FOR SINGLE-MODE FIBER — 10 km

Figure 122–1—200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

200GBASE-FR4 and 200GBASE-LR4 use four lanes, while 400GBASE-FR8 and 400GBASE-LR8 use eight lanes. In this clause, where there are four or eight items (depending on PMD type) such as lanes, the items are numbered from 0 to $n - 1$, and an example item is numbered i . Thus n is 4 or 8.

200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 PHYs with the optional Energy Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

122.1.1 Bit error ratio

The bit error ratio (BER) when processed according to Clause 120 shall be less than 2.4×10^{-4} provided that the error statistics are sufficiently random that this results in a frame loss ratio (see 1.4.223) of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap when processed according to Clause 120 and then Clause 119. For a complete Physical Layer, the frame loss ratio may be degraded to 6.2×10^{-11} for 64-octet frames with minimum interpacket gap due to additional errors from the electrical interfaces.

If the error statistics are not sufficiently random to meet this requirement, then the BER shall be less than that required to give a frame loss ratio of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap.

122.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 PMDs. The service interfaces for these PMDs are described in an abstract manner and do not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 116.3. The PMD service interface primitives are summarized as follows:

```
PMD:IS_UNITDATA_i.request
PMD:IS_UNITDATA_i.indication
PMD:IS_SIGNAL.indication
```

The 200GBASE-FR4 and 200GBASE-LR4 PMDs have four parallel symbol streams, in which case $i = 0$ to 3, and the 400GBASE-FR8 and 400GBASE-LR8 PMDs have eight parallel symbol streams, in which case $i = 0$ to 7.

In the transmit direction, the PMA continuously sends n parallel symbol streams to the PMD, one per lane, each at a nominal signaling rate of 26.5625 GBd. The PMD then converts these streams of data units into the appropriate signals on the MDI.

In the receive direction, the PMD continuously sends n parallel symbol streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 26.5625 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service interface primitive defined in 116.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_symbol parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_symbol parameters are known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the BER defined in 122.1.1.

122.3 Delay and Skew

122.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 200GBASE-FR4 or 200GBASE-LR4 PMD including 2 m of fiber in one direction shall be no more than 4096 bit times (8 pause_quanta or 20.48 ns). The sum of the transmit and receive delays at one end of the link contributed by the 400GBASE-FR8 or 400GBASE-LR8 PMD including 2 m of fiber in one direction shall be no more than 8192 bit times (16 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 116.4 and its references.

122.3.2 Skew constraints

The Skew (relative delay between the lanes) and Skew Variation must be kept within limits so that the information on the lanes can be reassembled by the PCS. Skew and Skew Variation are defined in 116.5 and specified at the points SP1 to SP6 shown in Figure 116-4 and Figure 116-5.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 116.5. The measurements of Skew and Skew Variation are defined in 87.8.2.

122.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 122-2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 122-3.

Table 122-2—MDIO/PMD control variable mapping

| MDIO control variable | PMA/PMD register name | Register/bit number | PMD control variable |
|--|-------------------------------|---------------------|--|
| Reset | PMA/PMD control-1 register | 1.0.15 | PMD_reset |
| Global PMD transmit disable | PMD transmit disable register | 1.9.0 | PMD_global_transmit_disable |
| PMD transmit disable $n-1$ to PMD transmit disable 0 | PMD transmit disable register | 1.9.8 to 1.9.1 | PMD_transmit_disable $n-1$ to PMD_transmit_disable_0 |

Table 122-3—MDIO/PMD status variable mapping

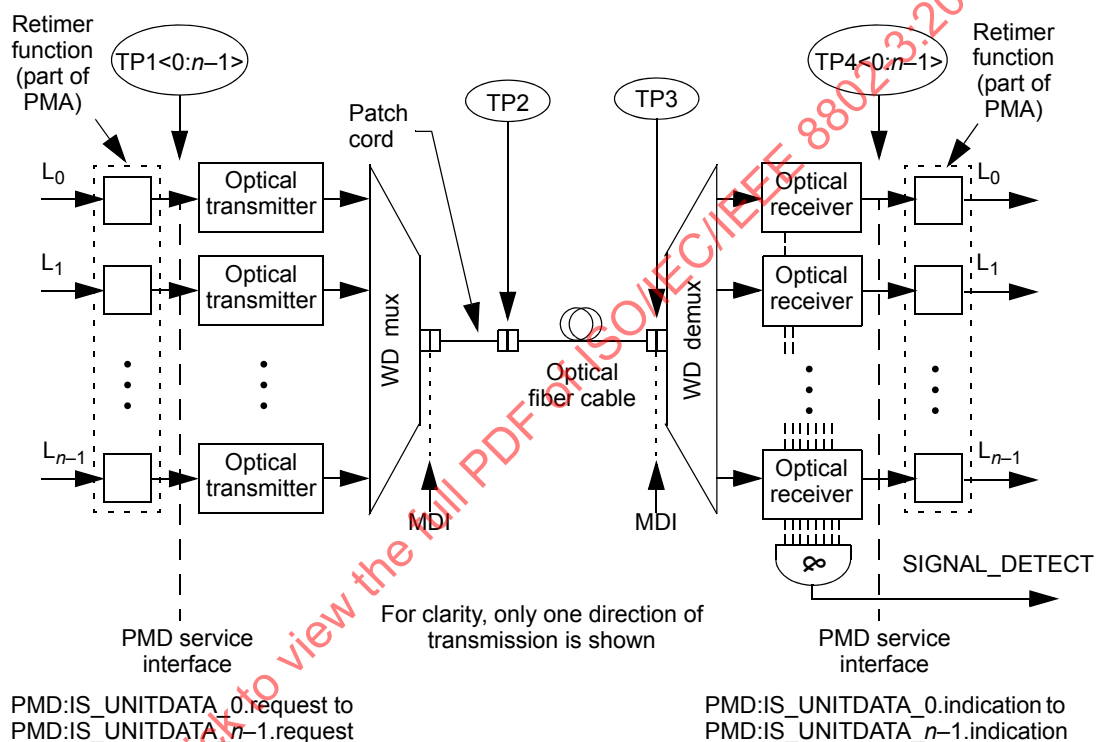
| MDIO status variable | PMA/PMD register name | Register/bit number | PMD status variable |
|--|------------------------------------|---------------------|--|
| Fault | PMA/PMD status 1 register | 1.1.7 | PMD_fault |
| Transmit fault | PMA/PMD status 2 register | 1.8.11 | PMD_transmit_fault |
| Receive fault | PMA/PMD status 2 register | 1.8.10 | PMD_receive_fault |
| Global PMD receive signal detect | PMD receive signal detect register | 1.10.0 | PMD_global_signal_detect |
| PMD receive signal detect $n-1$ to PMD receive signal detect 0 | PMD receive signal detect register | 1.10.8 to 1.10.1 | PMD_signal_detect $n-1$ to PMD_signal_detect_0 |

122.5 PMD functional specifications

The 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 PMDs perform the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

122.5.1 PMD block diagram

The PMD block diagram is shown in Figure 122–2. For purposes of system conformance, the PMD sublayer is standardized at the points described in this subclause. The optical transmit signal is defined at the output end of a single-mode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 122.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 122.11.3). Unless specified otherwise, all receiver measurements and tests defined in 122.8 are made at TP3.



WD = Wavelength division

NOTE—Specification of the retimer function and the electrical implementation of the PMD service interface is beyond the scope of this standard.

Figure 122–2—Block diagram for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 transmit/receive paths

TP1<0:n-1> and TP4<0:n-1> are informative reference points that may be useful to implementers for testing components (these test points will not typically be accessible in an implemented system).

122.5.2 PMD transmit function

The PMD Transmit function shall convert the n symbol streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_n-1.request into n separate optical

signals. The n optical signals shall then be wavelength division multiplexed and delivered to the MDI, all according to the transmit optical specifications in this clause. The highest optical power level in each signal stream shall correspond to $tx_symbol = three$ and the lowest shall correspond to $tx_symbol = zero$.

122.5.3 PMD receive function

The PMD Receive function shall demultiplex the composite optical signal received from the MDI into n separate optical signals. The n optical signals shall then be converted into n symbol streams for delivery to the PMD service interface using the messages `PMD:IS_UNITDATA_0.indication` to `PMD:IS_UNITDATA_n-1.indication`, all according to the receive optical specifications in this clause. The higher optical power level in each signal shall correspond to $rx_symbol = three$ and the lowest shall correspond to $rx_symbol = zero$.

122.5.4 PMD global signal detect function

The PMD global signal detect function shall report the state of `SIGNAL_DETECT` via the PMD service interface. The `SIGNAL_DETECT` parameter is signaled continuously, while the `PMD:IS_SIGNAL.indication` message is generated when a change in the value of `SIGNAL_DETECT` occurs. The `SIGNAL_DETECT` parameter defined in this clause maps to the `SIGNAL_OK` parameter in the inter-sublayer service interface primitives defined in 116.3.

`SIGNAL_DETECT` shall be a global indicator of the presence of optical signals on all n lanes. The value of the `SIGNAL_DETECT` parameter shall be generated according to the conditions defined in Table 122–4. The PMD receiver is not required to verify whether a compliant 200GBASE-R or 400GBASE-R signal is being received. This standard imposes no response time requirements on the generation of the `SIGNAL_DETECT` parameter.

Table 122–4—`SIGNAL_DETECT` value definition

| Receive conditions | <code>SIGNAL_DETECT</code> value |
|--|----------------------------------|
| For any lane; Average optical power at TP3 ≤ -30 dBm | FAIL |
| For all lanes; [(Optical power at TP3 \geq average receive power, each lane (min) in Table 122–11 for 200GBASE-FR4 and 200GBASE-LR4 or Table 122–12 for 400GBASE-FR8 and 400GBASE-LR8) AND (compliant 200GBASE-R or 400GBASE-R signal input)] | OK |
| All other conditions | Unspecified |

As an unavoidable consequence of the requirements for the setting of the `SIGNAL_DETECT` parameter, implementations must provide adequate margin between the input optical power level at which the `SIGNAL_DETECT` parameter is set to OK, and the inherent noise level of the PMD including the effects of crosstalk, power supply noise, etc.

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the `SIGNAL_DETECT` parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the modulated optical signal.

122.5.5 PMD lane-by-lane signal detect function

Various implementations of the Signal Detect function are permitted by this standard. When the MDIO is implemented, each `PMD_signal_detecti`, where i represents the lane number in the range $0:n-1$, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 122-4.

122.5.6 PMD reset function

If the MDIO interface is implemented, and if `PMD_reset` is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

122.5.7 PMD global transmit disable function (optional)

The `PMD_global_transmit_disable` function is optional and allows all of the optical transmitters to be disabled.

- a) When the `PMD_global_transmit_disable` variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 122-9 for 200GBASE-FR4 and 200GBASE-LR4 and Table 122-10 for 400GBASE-FR8 and 400GBASE-LR8.
- b) If a `PMD_fault` is detected, then the PMD may set the `PMD_global_transmit_disable` to one, turning off the optical transmitter in each lane.

122.5.8 PMD lane-by-lane transmit disable function

The `PMD_transmit_disablei` (where i represents the lane number in the range $0:3$ for 200GBASE-FR4 and 200GBASE-LR4 and $0:7$ for 400GBASE-FR8 and 400GBASE-LR8) function allows the optical transmitters in each lane to be selectively disabled.

- a) When a `PMD_transmit_disablei` variable is set to one, this function shall turn off the optical transmitter associated with that variable so that the transmitter meets the requirements of the average launch power of the OFF transmitter in Table 122-9 for 200GBASE-FR4 and 200GBASE-LR4 and Table 122-10 for 400GBASE-FR8 and 400GBASE-LR8.
- b) If a `PMD_fault` is detected, then the PMD may set each `PMD_transmit_disablei` to one, turning off the optical transmitter in each lane.

If the `PMD_transmit_disablei` function is not implemented in MDIO, an alternative method shall be provided to independently disable each transmit lane for testing purposes.

122.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set `PMD_fault` to one.

If the MDIO interface is implemented, `PMD_fault` shall be mapped to the fault bit as specified in 45.2.1.2.3.

122.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the `PMD_transmit_fault` variable to one.

If the MDIO interface is implemented, `PMD_transmit_fault` shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

122.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to one.

If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

122.6 Wavelength-division-multiplexed lane assignments

The wavelength range for each lane of the 200GBASE-FR4 PMD is defined in Table 122-5. The 200GBASE-FR4 center wavelengths are members of the CWDM wavelength grid defined in ITU-T G.694.2 and are spaced at 20 nm.

The wavelength range for each lane of the 200GBASE-LR4 PMD is defined in Table 122-6. The wavelength range for each lane of the 400GBASE-FR8 and 400GBASE-LR8 PMDs is defined in Table 122-7. The 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 center frequencies are members of the frequency grid for 100 GHz spacing and above defined in ITU-T G.694.1 and are spaced at 800 GHz.

Table 122-5—200GBASE-FR4 wavelength-division-multiplexed lane assignments

| Lane | Center wavelength | Wavelength range |
|----------------|-------------------|---------------------|
| L ₀ | 1271 nm | 1264.5 to 1277.5 nm |
| L ₁ | 1291 nm | 1284.5 to 1297.5 nm |
| L ₂ | 1311 nm | 1304.5 to 1317.5 nm |
| L ₃ | 1331 nm | 1324.5 to 1337.5 nm |

Table 122-6—200GBASE-LR4 wavelength-division-multiplexed lane assignments

| Lane | Center frequency | Center wavelength | Wavelength range |
|----------------|------------------|-------------------|-----------------------|
| L ₀ | 231.4 THz | 1295.56 nm | 1294.53 to 1296.59 nm |
| L ₁ | 230.6 THz | 1300.05 nm | 1299.02 to 1301.09 nm |
| L ₂ | 229.8 THz | 1304.58 nm | 1303.54 to 1305.63 nm |
| L ₃ | 229 THz | 1309.14 nm | 1308.09 to 1310.19 nm |

NOTE—There is no requirement to associate a particular electrical lane with a particular optical lane, as the PCS is capable of receiving lanes in any arrangement.

Table 122-7—400GBASE-FR8 and 400GBASE-LR8 wavelength-division-multiplexed lane assignments

| Lane | Center frequency | Center wavelength | Wavelength range |
|----------------|------------------|-------------------|-----------------------|
| L ₀ | 235.4 THz | 1273.54 nm | 1272.55 to 1274.54 nm |
| L ₁ | 234.6 THz | 1277.89 nm | 1276.89 to 1278.89 nm |
| L ₂ | 233.8 THz | 1282.26 nm | 1281.25 to 1283.27 nm |
| L ₃ | 233 THz | 1286.66 nm | 1285.65 to 1287.68 nm |
| L ₄ | 231.4 THz | 1295.56 nm | 1294.53 to 1296.59 nm |
| L ₅ | 230.6 THz | 1300.05 nm | 1299.02 to 1301.09 nm |
| L ₆ | 229.8 THz | 1304.58 nm | 1303.54 to 1305.63 nm |
| L ₇ | 229 THz | 1309.14 nm | 1308.09 to 1310.19 nm |

122.7 PMD to MDI optical specifications for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8

The operating ranges for the 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 PMDs are defined in Table 122-8. A 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, or 400GBASE-LR8 compliant PMD operates on type B1.1, B1.3, or B6 single-mode fibers according to the specifications defined in Table 122-18. A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 400GBASE-FR8 PMD operating at 2.5 km meets the operating range requirement of 2 m to 2 km). The 400GBASE-LR8 PMD interoperates with the 400GBASE-FR8 PMD provided that the channel requirements for 400GBASE-FR8 are met.

Table 122-8—200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 operating ranges

| PMD type | Required operating range |
|-------------------------------|--------------------------|
| 200GBASE-FR4 and 400GBASE-FR8 | 2 m to 2 km |
| 200GBASE-LR4 and 400GBASE-LR8 | 2 m to 10 km |

122.7.1 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 transmitter optical specifications

The 200GBASE-FR4 transmitter shall meet the specifications defined in Table 122-9 per the definitions in 122.8. The 200GBASE-LR4 transmitter shall meet the specifications defined in Table 122-9 per the definitions in 122.8. The 400GBASE-FR8 transmitter shall meet the specifications defined in Table 122-10 per the definitions in 122.8. The 400GBASE-LR8 transmitter shall meet the specifications defined in Table 122-10 per the definitions in 122.8.

Table 122-9—200GBASE-FR4 and 200GBASE-LR4 transmit characteristics

| Description | 200GBASE-FR4 | 200GBASE-LR4 | Unit |
|--|--|--|-------|
| Signaling rate, each lane (range) | 26.5625 ± 100 ppm | | GBd |
| Modulation format | PAM4 | | — |
| Lane wavelengths (range) | 1264.5 to 1277.5 1284.5 to 1297.5 1304.5 to 1317.5 1324.5 to 1337.5 | 1294.53 to 1296.59 1299.02 to 1301.09 1303.54 to 1305.63 1308.09 to 1310.19 | nm |
| Side-mode suppression ratio (SMSR), (min) | 30 | | dB |
| Total average launch power (max) | 10.7 | 11.3 | dBm |
| Average launch power, each lane (max) | 4.7 | 5.3 | dBm |
| Average launch power, each lane ^a (min) | −4.2 | −3.4 | dBm |
| Outer Optical Modulation Amplitude (OMA _{outer}), each lane (max) | 4.5 | 5.1 | dBm |
| Outer Optical Modulation Amplitude (OMA _{outer}), each lane (min) ^b | −1.2 | −0.4 | dBm |
| Difference in launch power between any two lanes (OMA _{outer}) (max) | 4 | | dB |
| Launch power in OMA _{outer} minus TDECQ, each lane (min): for extinction ratio ≥ 4.5 dB for extinction ratio < 4.5 dB | −2.6 | −1.8 | dBm |
| | −2.5 | −1.7 | dBm |
| Transmitter and dispersion eye closure for PAM4 (TDECQ), each lane (max) | 3.3 | 3.4 | dB |
| Average launch power of OFF transmitter, each lane (max) | −30 | | dBm |
| Extinction ratio, each lane (min) | 3.5 | | dB |
| RIN _{16.5} OMA (max) | −132 | — | dB/Hz |
| RIN _{15.1} OMA (max) | — | −132 | dB/Hz |
| Optical return loss tolerance (max) | 16.5 | 15.1 | dB |
| Transmitter reflectance ^c (max) | −26 | | dB |

^a Average launch power, each lane (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.

^b Even if the TDECQ < 1.4 dB for an extinction ratio of ≥ 4.5 dB or TDECQ < 1.3 dB for an extinction ratio of < 4.5 dB, the OMA_{outer} (min) must exceed this value.

^c Transmitter reflectance is defined looking into the transmitter.

Table 122–10—400GBASE-FR8 and 400GBASE-LR8 transmit characteristics

| Description | 400GBASE-FR8 | 400GBASE-LR8 | Unit |
|--|--|--------------|------------|
| Signaling rate, each lane (range) | 26.5625 ± 100 ppm | | GBd |
| Modulation format | PAM4 | | — |
| Lane wavelengths (range) | 1272.55 to 1274.54 1276.89 to 1278.89 1281.25 to 1283.27 1285.65 to 1287.68 1294.53 to 1296.59 1299.02 to 1301.09 1303.54 to 1305.63 1308.09 to 1310.19 | | nm |
| Side-mode suppression ratio (SMSR), (min) | 30 | | dB |
| Total average launch power (max) | 13.2 | | dBm |
| Average launch power, each lane ^a (max) | 5.3 | | dBm |
| Average launch power, each lane ^b (min) | −3.5 | −2.8 | dBm |
| Outer Optical Modulation Amplitude (OMA _{outer}), each lane (max) | 5.5 | 5.7 | dBm |
| Outer Optical Modulation Amplitude (OMA _{outer}), each lane (min) ^c | −0.5 | 0.2 | dBm |
| Difference in launch power between any two lanes (OMA _{outer}) (max) | 4 | | dB |
| Launch power in OMA _{outer} minus TDECQ, each lane (min): for extinction ratio ≥ 4.5 dB for extinction ratio < 4.5 dB | −1.9 −1.8 | −1.2 −1.1 | dBm dBm |
| Transmitter and dispersion eye closure for PAM4 (TDECQ), each lane (max) | 3.1 | 3.3 | dB |
| Average launch power of OFF transmitter, each lane (max) | −30 | | dBm |
| Extinction ratio, each lane (min) | 3.5 | | dB |
| RIN _{16,5} OMA (max) | −132 | — | dB/Hz |
| RIN _{15,1} OMA (max) | — | −132 | dB/Hz |
| Optical return loss tolerance (max) | 16.5 | 15.1 | dB |
| Transmitter reflectance ^d (max) | −26 | | dB |

^a As the total average launch power limit has to be met, not all of the lanes can operate at the maximum average launch power, each lane.

^b Average launch power, each lane (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.

^c Even if the TDECQ < 1.4 dB for an extinction ratio of ≥ 4.5 dB or TDECQ < 1.3 dB for an extinction ratio of < 4.5 dB, the OMA_{outer} (min) must exceed this value.

^d Transmitter reflectance is defined looking into the transmitter.

122.7.2 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 receive optical specifications

The 200GBASE-FR4 receiver shall meet the specifications defined in Table 122–11 per the definitions in 122.8. The 200GBASE-LR4 receiver shall meet the specifications defined in Table 122–11 per the definitions in 122.8. The 400GBASE-FR8 receiver shall meet the specifications defined in Table 122–12 per the definitions in 122.8. The 400GBASE-LR8 receiver shall meet the specifications defined in Table 122–12 per the definitions in 122.8.

Table 122–11—200GBASE-FR4 and 200GBASE-LR4 receive characteristics

| Description | 200GBASE-FR4 | 200GBASE-LR4 | Unit |
|---|--|--|------|
| Signaling rate, each lane (range) | 26.5625 ± 100 ppm | | GBd |
| Modulation format | PAM4 | | — |
| Lane wavelengths (range) | 1264.5 to 1277.5 1284.5 to 1297.5 1304.5 to 1317.5 1324.5 to 1337.5 | 1294.53 to 1296.59 1299.02 to 1301.09 1303.54 to 1305.63 1308.09 to 1310.19 | nm |
| Damage threshold ^a , each lane | 5.7 | 6.3 | dBm |
| Average receive power, each lane (max) | 4.7 | 5.3 | dBm |
| Average receive power, each lane ^b (min) | −8.2 | −9.7 | dBm |
| Receive power (OMA _{outer}), each lane (max) | 4.5 | 5.1 | dBm |
| Difference in receive power between any two lanes (OMA _{outer}) (max) | 4.1 | 4.2 | dB |
| Receiver reflectance (max) | −26 | | dB |
| Receiver sensitivity (OMA _{outer}), each lane ^c (max) | −6 | −7.7 | dBm |
| Stressed receiver sensitivity (OMA _{outer}), each lane ^d (max) | −3.6 | −5.2 | dBm |
| Conditions of stressed receiver sensitivity test: ^e | | | |
| Stressed eye closure for PAM4 (SECQ), lane under test | 3.3 | 3.4 | dB |
| OMA _{outer} of each aggressor lane | 0.5 | −1 | dBm |

^a The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level.

^b Average receive power, each lane (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.

^c Receiver sensitivity (OMA_{outer}), each lane (max) is informative and is defined for a transmitter with SECQ of 0.9 dB.

^d Measured with conformance test signal at TP3 (see 122.8.9) for the BER specified in 122.1.1.

^e These test conditions are for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

Table 122–12—400GBASE-FR8 and 400GBASE-LR8 receive characteristics

| Description | 400GBASE-FR8 | 400GBASE-LR8 | Unit |
|---|--|--------------|------|
| Signaling rate, each lane (range) | 26.5625 ± 100 ppm | | GBd |
| Modulation format | PAM4 | | — |
| Lane wavelengths (range) | 1272.55 to 1274.54 1276.89 to 1278.89 1281.25 to 1283.27 1285.65 to 1287.68 1294.53 to 1296.59 1299.02 to 1301.09 1303.54 to 1305.63 1308.09 to 1310.19 | | nm |
| Damage threshold ^a , each lane | 6.3 | | dBm |
| Average receive power, each lane (max) | 5.3 | | dBm |
| Average receive power, each lane ^b (min) | −7.5 | −9.1 | dBm |
| Receive power (OMA _{outer}), each lane (max) | 5.7 | | dBm |
| Difference in receive power between any two lanes (OMA _{outer}) (max) | 4.1 | 4.5 | dB |
| Receiver reflectance (max) | −26 | | dB |
| Receiver sensitivity (OMA _{outer}), each lane ^c (max) | −5.3 | −7.1 | dBm |
| Stressed receiver sensitivity (OMA _{outer}), each lane ^d (max) | −3.1 | −4.7 | dBm |
| Conditions of stressed receiver sensitivity test: ^e | | | |
| Stressed eye closure for PAM4 (SECQ), lane under test | 3.1 | 3.3 | dB |
| OMA _{outer} of each aggressor lane | 1 | −0.2 | dBm |

^a The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level.

^b Average receive power, each lane (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.

^c Receiver sensitivity (OMA_{outer}), each lane (max) is informative and is defined for a transmitter with SECQ of 0.9 dB.

^d Measured with conformance test signal at TP3 (see 122.8.9) for the BER specified in 122.1.1.

^e These test conditions are for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

122.7.3 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 illustrative link power budgets

Illustrative power budgets and penalties for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 channels are shown in Table 122–13.

Table 122–13—200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 illustrative link power budgets

| Parameter | 200GBASE-FR4 | 400GBASE-FR8 | 200GBASE-LR4 | 400GBASE-LR8 | Unit |
|--|----------------|--------------|----------------|--------------|------|
| Power budget (for maximum TDECQ): for extinction ratio ≥ 4.5 dB | 7.6 | 7.4 | 10.2 | 10.1 | dB |
| for extinction ratio < 4.5 dB | 7.7 | 7.5 | 10.3 | 10.2 | dB |
| Operating distance | 2 | | 10 | | km |
| Channel insertion loss | 4 ^a | | 6.3 | | dB |
| Maximum discrete reflectance | See 122.11.2.2 | | See 122.11.2.2 | | dB |
| Allocation for penalties ^b (for maximum TDECQ): for extinction ratio ≥ 4.5 dB | 3.6 | 3.4 | 3.9 | 3.8 | dB |
| for extinction ratio < 4.5 dB | 3.7 | 3.5 | 4 | 3.9 | dB |
| Additional insertion loss allowed | 0 | | 0 | | dB |

^a The channel insertion loss is calculated using the maximum distance specified in Table 122–8 for 200GBASE-FR4 and 400GBASE-FR8 and fiber attenuation of 0.5 dB/km plus an allocation for connection and splice loss given in 122.11.2.1.

^b Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

122.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

122.8.1 Test patterns for optical parameters

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 122–15 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 122–15 may be used to perform that test. The test patterns used in this clause are shown in Table 122–14.

Table 122–14—Test patterns

| Pattern | Pattern description | Defined in |
|-------------|---------------------------------|--------------|
| Square wave | Square wave (8 threes, 8 zeros) | 120.5.11.2.4 |
| 3 | PRBS31Q | 120.5.11.2.2 |
| 4 | PRBS13Q | 120.5.11.2.1 |
| 5 | Scrambled idle | 119.2.4.9 |
| 6 | SSPRQ | 120.5.11.2.3 |

Table 122–15—Test-pattern definitions and related subclauses

| Parameter | Pattern | Related subclause |
|--|--|-------------------|
| Wavelength | Square wave, 3, 4, 5, 6 or valid 200GBASE-R or 400GBASE-R signal | 122.8.2 |
| Side mode suppression ratio | 3, 5, 6 or valid 200GBASE-R or 400GBASE-R signal | — |
| Average optical power | 3, 5, 6 or valid 200GBASE-R or 400GBASE-R signal | 122.8.3 |
| Outer Optical Modulation Amplitude (OMA _{outer}) | 4 or 6 | 122.8.4 |
| Transmitter and dispersion eye closure for PAM4 (TDECQ) | 6 | 122.8.5 |
| Extinction ratio | 4 or 6 | 122.8.6 |
| RIN _{15.1} OMA and RIN _{16.5} OMA | Square wave | 122.8.7 |
| Stressed receiver conformance test signal calibration | 6 | 122.8.9.2 |
| Stressed receiver sensitivity | 3 or 5 | 122.8.9 |

122.8.2 Wavelength

The wavelength of each optical lane shall be within the ranges given in Table 122–5 for 200GBASE-FR4, in Table 122–6 for 200GBASE-LR4 and Table 122–7 for 400GBASE-FR8 and 400GBASE-LR8, if measured per TIA/EIA-455-127-A or IEC 61280-1-3. The lane under test is modulated using the test pattern defined in Table 122–15.

122.8.3 Average optical power

The average optical power of each lane shall be within the limits given in Table 122–9 for 200GBASE-FR4 and 200GBASE-LR4 and in Table 122–10 for 400GBASE-FR8 and 400GBASE-LR8 if measured using the methods given in IEC 61280-1-1, with the sum of the optical power from all of the lanes not under test below –30 dBm, per the test setup in Figure 53–6.

122.8.4 Outer Optical Modulation Amplitude (OMA_{outer})

The OMA_{outer} of each lane shall be within the limits given in Table 122–9 for 200GBASE-FR4 and 200GBASE-LR4 and in Table 122–10 for 400GBASE-FR8 and 400GBASE-LR8. The OMA_{outer} is measured using a test pattern specified for OMA_{outer} in Table 122–15. It is the difference between the

average optical launch power level P_3 , measured over the central 2 UI of a run of 7 threes, and the average optical launch power level P_0 , measured over the central 2 UI of a run of 6 zeros, as shown in Figure 122–3. For this measurement the sum of the optical power from all of the lanes not under test is below -30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.




Figure 122–3—Power levels P_0 and P_3 from PRBS13Q test pattern

122.8.5 Transmitter and dispersion eye closure for PAM4 (TDECQ)

The TDECQ of each lane shall be within the limits given in Table 122–9 for 200GBASE-FR4 and 200GBASE-LR4 and in Table 122–10 for 400GBASE-FR8 and 400GBASE-LR8 if measured using the methods specified in 122.8.5.1, 122.8.5.2, and 122.8.5.3.

TDECQ is a measure of each optical transmitter's vertical eye closure when transmitted through a worst case optical channel (specified in 122.8.5.2), as measured through an optical to electrical converter (O/E) with a bandwidth equivalent to a reference receiver, and equalized with the reference equalizer (as described in 122.8.5.4). The reference receiver and equalizer may be implemented in software or may be part of the oscilloscope.

Table 122–15 specifies the test patterns to be used for measurement of TDECQ.

122.8.5.1 TDECQ conformance test setup

A block diagram for the TDECQ conformance test is shown in Figure 122–4. Other equivalent measurement implementations may be used with suitable calibration.

Each optical lane is tested individually with all other lanes in operation and all lanes using the same test pattern. There shall be at least 31 UI delay between the test pattern on one lane and the pattern on any other lane, so that the symbols on each lane are not correlated within the PMD. The optical splitter and variable reflector are adjusted so that each transmitter is tested with the optical return loss specified in Table 122–16. The state of polarization of the back reflection is adjusted to create the greatest RIN. The optical filter is used to separate the lane under test from the others. Each optical lane is tested with the optical channel described in 122.8.5.2. The combination of the O/E and the oscilloscope has a fourth-order Bessel-Thomson filter response with a bandwidth of approximately 13.28125 GHz. Compensation may be made for any deviation from an ideal fourth-order Bessel-Thomson response.

The optical filter passband ripple shall be limited to 0.5 dB peak-to-peak and the isolation is chosen such that the ratio of the power in the lane being measured to the sum of the powers of all of the other lanes is greater than 20 dB (see ITU-T G.959.1 Annex B).

The test pattern (specified in Table 122–15) is transmitted repetitively by the optical lane under test and the oscilloscope is set up to capture the complete pattern for TDECQ analysis as described in 122.8.5.3. The clock recovery unit (CRU) has a corner frequency of 4 MHz and a slope of 20 dB/decade. The CRU can be implemented in hardware or software depending on oscilloscope technology.

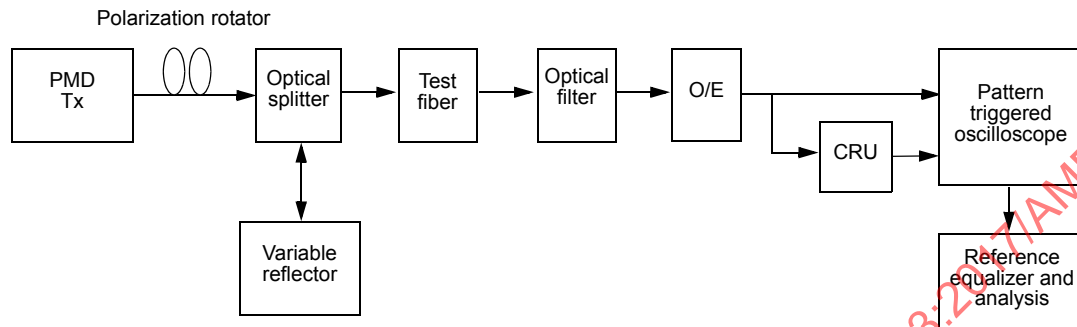


Figure 122–4—TDECQ conformance test block diagram

122.8.5.2 Channel requirements

The transmitter is tested using an optical channel that meets the requirements listed in Table 122–16.

Table 122–16—Transmitter compliance channel specifications

| PMD type | Dispersion ^a (ps/nm) | | Insertion loss ^b | Optical return loss ^c | Max mean DGD |
|------------------------------|---|---|-----------------------------|----------------------------------|--------------|
| | Minimum | Maximum | | | |
| 200GBASE-FR4 or 400GBASE-FR8 | $0.0465 \cdot \lambda \cdot [1 - (1324 / \lambda)^4]$ | $0.0465 \cdot \lambda \cdot [1 - (1300 / \lambda)^4]$ | Minimum | 17.8 dB | 0.8 ps |
| 200GBASE-LR4 or 400GBASE-LR8 | $0.2325 \cdot \lambda \cdot [1 - (1324 / \lambda)^4]$ | $0.2325 \cdot \lambda \cdot [1 - (1300 / \lambda)^4]$ | Minimum | 15.7 dB | 0.8 ps |

^a The dispersion is measured for the wavelength of the device under test (λ in nm). The coefficient assumes 2 km for 200GBASE-FR4 and 400GBASE-FR8 and 10 km for 200GBASE-LR4 and 400GBASE-LR8.

^b There is no intent to stress the sensitivity of the O/E converter associated with the oscilloscope.

^c The optical return loss is applied at TP2.

A 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8 or 400GBASE-LR8 transmitter is to be compliant with a total dispersion at least as negative as the “minimum dispersion” and at least as positive as the “maximum dispersion” columns specified in Table 122–16 for the wavelength of the device under test. This may be achieved with channels consisting of fibers with lengths chosen to meet the dispersion requirements.

To verify that the fiber has the correct amount of dispersion, the measurement method defined in IEC 60793-1-42 may be used. The measurement is made in the linear power regime of the fiber.

The channel provides an optical return loss specified in Table 122–16. The state of polarization of the back reflection is adjusted to create the greatest RIN.

The mean DGD of the channel is to be less than the value specified in Table 122–16.

122.8.5.3 TDECQ measurement method

TDECQ for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 is measured as described in 121.8.5.3 with the exception that the reference equalizer is as specified in 122.8.5.4.

122.8.5.4 TDECQ reference equalizer

The reference equalizer for 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 is a 5 tap, T spaced, feed-forward equalizer (FFE), where T is the symbol period. The sum of the equalizer tap coefficients is equal to 1.

NOTE—This reference equalizer is part of the TDECQ test and does not imply any particular receiver equalizer implementation.

122.8.6 Extinction ratio

The extinction ratio of each lane shall be within the limits given in Table 122–9 for 200GBASE-FR4 and 200GBASE-LR4 and in Table 122–10 for 400GBASE-FR8 and 400GBASE-LR8 if measured using a test pattern specified for extinction ratio in Table 122–15 with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test. The extinction ratio of a PAM4 optical signal is defined as the ratio of the average optical launch power level P_3 , measured over the central 2 UI of a run of 7 threes, and the average optical launch power level P_0 , measured over the central 2 UI of a run of 6 zeros, as shown in Figure 122–3.

122.8.7 Relative intensity noise ($RIN_{16.5\text{OMA}}$ and $RIN_{15.1\text{OMA}}$)

RIN shall be as defined by the measurement methodology of 52.9.6 with the following exceptions:

- The optical return loss is 16.5 dB for 200GBASE-FR4 and 400GBASE-FR8 and 15.1 dB for 200GBASE-LR4 and 400GBASE-LR8.
- Each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.
- The upper –3 dB limit of the measurement apparatus is to be approximately equal to the signaling rate (i.e., 26.6 GHz).
- The test pattern is according to Table 122–15.

122.8.8 Receiver sensitivity

Receiver sensitivity, which is defined for an input signal with SECQ of 0.9 dB (e.g., an ideal input signal without overshoot), is informative and compliance is not required. If measured, the test signal should have negligible impairments such as intersymbol interference (ISI), rise/fall times, jitter and RIN. Instead, the normative requirement for receivers is stressed receiver sensitivity.

122.8.9 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 122–11 for 200GBASE-FR4 and 200GBASE-LR4 and in Table 122–12 for 400GBASE-FR8 and 400GBASE-LR8 if measured using the method defined in 122.8.9.1 and 122.8.9.3, with the conformance test signal at TP3 as described in 122.8.9.2, using the test pattern specified for SRS in Table 122–15. The BER is required to be met for the lane under test on its own.

Stressed receiver sensitivity is defined with all transmit and receive lanes in operation. Any of the patterns specified for SRS in Table 122–15 is sent from the transmit section of the PMD under test. The signal being transmitted is asynchronous to the received signal.

122.8.9.1 Stressed receiver conformance test block diagram

A block diagram for the receiver conformance test is shown in Figure 122–5. The patterns used for the received conformance signal are specified in Table 122–15. The optical test signal is conditioned (stressed) using the stressed receiver methodology defined in 122.8.9.2 and has sinusoidal jitter applied as specified in 121.8.9.4. A suitable test set is needed to characterize the signal used to test the receiver. Stressed receiver conformance test signal verification is described in 122.8.9.3.

The low-pass filter is used to create ISI. The combination of the low-pass filter and the E/O converter should have a frequency response that results in at least half of the dB value of the stressed eye closure (SECQ) specified in Table 122–11 for 200GBASE-FR4 and 200GBASE-LR4 and in Table 122–12 for 400GBASE-FR8 and 400GBASE-LR8 before the sinusoidal and Gaussian noise terms are added, according to the methods specified in 122.8.9.2. The sinusoidal amplitude interferer causes additional eye closure, but in conjunction with the finite edge rates, also causes some jitter.

The sinusoidally jittered clock represents other forms of jitter and also verifies that the receiver under test can track low-frequency jitter. The sinusoidal amplitude interferer may be set at any frequency between 100 MHz and 2 GHz, although care should be taken to avoid harmonic relationships between the sinusoidal interferer, the sinusoidal jitter, the signaling rate, and the pattern repetition rate. The Gaussian noise generator, the amplitude of the sinusoidal interferer, and the low-pass filter are adjusted so that the SECQ specified in Table 122–11 or Table 122–12 is met, according to the methods specified in 122.8.9.2.

For improved visibility for calibration, all elements in the signal path (cables, DC blocks, E/O converter, etc.) should have wide and smooth frequency response, and linear phase response, throughout the spectrum of interest. Baseline wander and overshoot and undershoot should be negligible.

122.8.9.2 Stressed receiver conformance test signal characteristics and calibration

The stressed receiver conformance test signal characteristics and calibration methods are as described in 121.8.9.2 with the following exceptions:

- The SECQ of the stressed receiver conformance test signal is measured according to 122.8.5, except that the test fiber is not used.
- An example stressed receiver conformance test setup is shown in Figure 122–5; however, alternative test setups that generate equivalent stress conditions may be used.
- The signaling rate of the test pattern generator and the extinction ratio of the E/O converter are as given in Table 122–9 for 200GBASE-FR4 and 200GBASE-LR4 and in Table 122–10 for 400GBASE-FR8 and 400GBASE-LR8.
- The required values of the “Stressed receiver sensitivity ($\text{OMA}_{\text{outer}}$), each lane (max)”, “Stressed eye closure for PAM4 (SECQ), lane under test”, and “ $\text{OMA}_{\text{outer}}$ of each aggressor lane” are as given in Table 122–11 for 200GBASE-FR4 and 200GBASE-LR4 and in Table 122–12 for 400GBASE-FR8 and 400GBASE-LR8.

122.8.9.3 Stressed receiver conformance test signal verification

The SECQ of the stressed receiver conformance test signal is measured according to 122.8.5, except that the test fiber is not used. The clock output from the clock source in Figure 122–5 is modulated with the sinusoidal jitter. To use an oscilloscope to calibrate the final stressed eye jitter that includes the sinusoidal jitter component, a separate clock source (clean clock of Figure 122–5) is required that is synchronized to the source clock, but not modulated with the jitter source.

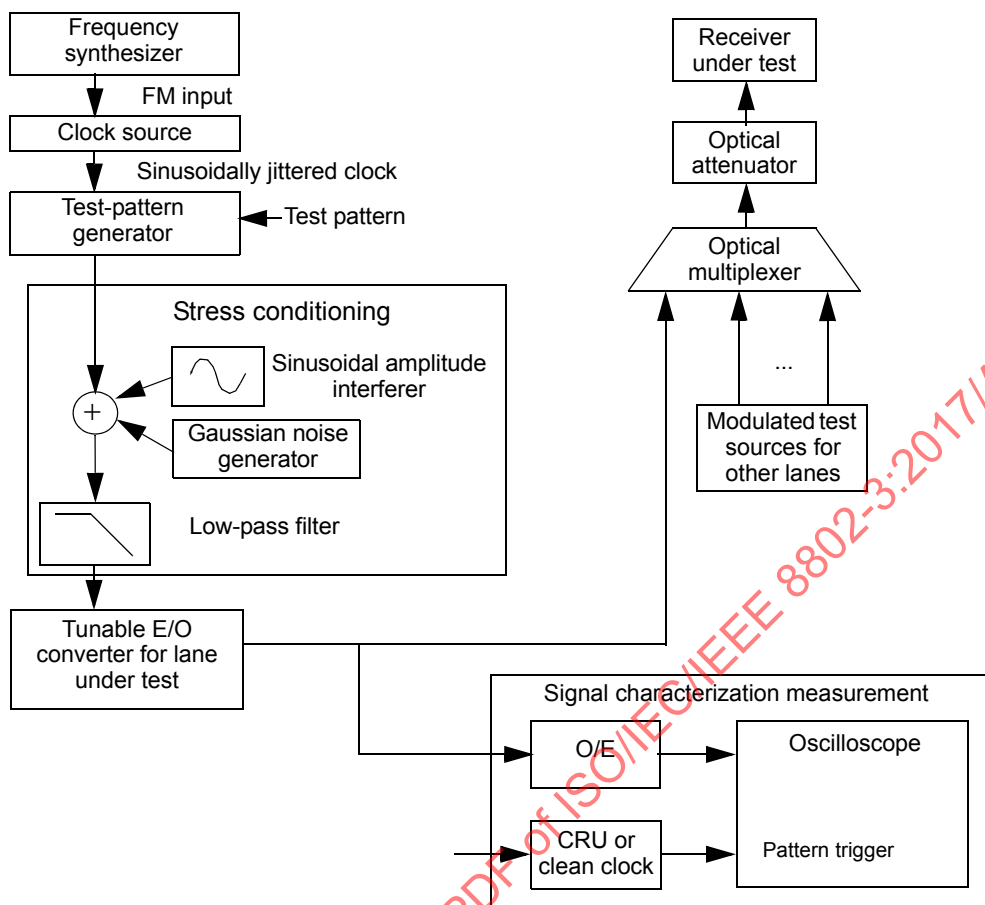


Figure 122-5—Stressed receiver conformance test block diagram

Care should be taken when characterizing the test signal because excessive noise/jitter in the measurement system would result in an input signal that does not fully stress the receiver under test. Running the receiver tolerance test with a signal that is under-stressed may result in the deployment of non-compliant receivers. The noise/jitter introduced by the O/E, filters, and oscilloscope should be negligible or the results should be corrected for its effects. While the details of test equipment are beyond the scope of this standard, it is recommended that the implementer fully characterize the test equipment and apply appropriate guard bands so that the stressed receiver conformance input signal meets the stress and sinusoidal jitter specified in 122.8.9.2 and 121.8.9.4.

122.9 Safety, installation, environment, and labeling

122.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

122.9.2 Laser safety

200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 optical transceivers shall conform to Hazard Level 1 laser requirements as defined in IEC 60825-1 and IEC 60825-2, under any condition of operation. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product's laser, safety features, labeling, use, maintenance, and service. This documentation explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.⁸

122.9.3 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

122.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, or 400GBASE-LR8 PMD over the life of the product while the product operates within the manufacturer's range of environmental, power, and other specifications.

It is recommended that manufacturers indicate in the literature associated with the PHY the operating environmental conditions to facilitate selection, installation, and maintenance.

It is recommended that manufacturers indicate, in the literature associated with the components of the optical link, the distance and operating environmental conditions over which the specifications of this clause will be met.

122.9.5 Electromagnetic emission

A system integrating a 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, or 400GBASE-LR8 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

122.9.6 Temperature, humidity, and handling

The optical link is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

122.9.7 PMD labeling requirements

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user, with at least the applicable safety warnings and the applicable port type designation (e.g., 400GBASE-FR8).

Labeling requirements for Hazard Level 1 lasers are given in the laser safety standards referenced in 122.9.2:

⁸A host system that fails to meet the manufacturer's requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.

122.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 122–6.




Figure 122–6—Fiber optic cabling model

The channel insertion loss is given in Table 122–17. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, and polarization mode dispersion meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with ANSI/TIA/EIA-526-7/method A-1. The fiber optic cabling model (channel) defined here is the same as a simplex fiber optic link segment. The term channel is used here for consistency with generic cabling standards.

Table 122–17—Fiber optic cabling (channel) characteristics

| Description | 200GBASE-FR4 | 400GBASE-FR8 | 200GBASE-LR4 | 400GBASE-LR8 | Unit |
|---|--------------|--------------|------------------|--------------|-------|
| Operating distance (max) | 2 | | 10 | | km |
| Channel insertion loss ^{a,b} (max) | 4 | | 6.3 ^c | | dB |
| Channel insertion loss (min) | 0 | | 0 | | dB |
| Positive dispersion ^b (max) | 6.7 | 1.9 | 9.5 | | ps/nm |
| Negative dispersion ^b (min) | –11.9 | –10.2 | –28.4 | –50.8 | ps/nm |
| DGD_max ^d | 3 | | 8 | | ps |
| Optical return loss (min) | 25 | | 22 | | dB |

^a These channel insertion loss values include cable, connectors, and splices.

^b Over the wavelength range 1264.5 nm to 1337.5 nm for 200GBASE-FR4, 1294.53 nm to 1310.19 nm for 200GBASE-LR4, and 1272.55 nm to 1310.19 nm for 400GBASE-FR8 and 400GBASE-LR8.

^c Using 0.46 dB/km at 1272.55 nm attenuation for optical fiber cables derived from Appendix I of ITU-T G.695 may not support operation at 10 km for 400GBASE-LR8 under worst case conditions.

^d Differential Group Delay (DGD) is the time difference at reception between the fractions of a pulse that were transmitted in the two principal states of polarization of an optical signal. DGD_max is the maximum differential group delay that the system must tolerate.

122.11 Characteristics of the fiber optic cabling (channel)

The 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 fiber optic cabling shall meet the specifications defined in Table 122–17. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

122.11.1 Optical fiber cable

The fiber optic cable requirements are satisfied by cables containing IEC 60793-2-50 type B1.1 (dispersion un-shifted single-mode), type B1.3 (low water peak single-mode), or type B6_a (bend insensitive) fibers or the requirements in Table 122–18 where they differ.

Table 122–18—Optical fiber and cable characteristics

| Description | Value | Unit |
|--|---------------------------------------|-----------------------|
| Nominal fiber specification wavelength | 1310 | nm |
| Cabled optical fiber attenuation (max) | 0.47 ^a or 0.5 ^b | dB/km |
| Zero dispersion wavelength (λ_0) | $1300 \leq \lambda_0 \leq 1324$ | nm |
| Dispersion slope (max) (S_0) | 0.093 | ps/nm ² km |

^a The 0.47 dB/km at 1264.5 nm attenuation for optical fiber cables is derived from Appendix I of ITU-T G.695.

^b The 0.5 dB/km attenuation is provided for Outside Plant cable as defined in ANSI/TIA 568-C.3. Using 0.5 dB/km may not support operation 10 km for 200GBASE-LR4 or 400GBASE-LR8.

122.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 122–6, consists of a mated pair of optical connectors.

122.11.2.1 Connection insertion loss

The maximum link distance for 200GBASE-LR4 and 400GBASE-LR8 is based on an allocation of 2 dB total connection and splice loss. For example, this allocation supports four connections with an average insertion loss per connection of 0.5 dB. The maximum link distance for 200GBASE-FR4 and 400GBASE-FR8 is based on an allocation of 3 dB total connection and splice loss. Connections with different loss characteristics may be used provided the requirements of Table 122–17 are met.

122.11.2.2 Maximum discrete reflectance

The maximum value for each discrete reflectance shall be less than or equal to the value shown in Table 122–19 corresponding to the number of discrete reflectances above –55 dB within the channel. For numbers of discrete reflectances in between two numbers shown in the table, the lower of the two corresponding maximum discrete reflectance values applies.

Table 122–19—Maximum value of each discrete reflectance

| Number of discrete reflectances above –55 dB | Maximum value for each discrete reflectance | |
|--|---|------------------------------|
| | 200GBASE-FR4 or 400GBASE-FR8 | 200GBASE-LR4 or 400GBASE-LR8 |
| 1 | –25 dB | –22 dB |
| 2 | –31 dB | –29 dB |
| 4 | –35 dB | –33 dB |
| 6 | –38 dB | –35 dB |
| 8 | –40 dB | –37 dB |
| 10 | –41 dB | –39 dB |

122.11.3 Medium Dependent Interface (MDI) requirements

The 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, or 400GBASE-LR8 PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the “fiber optic cabling” (as shown in Figure 122–6). Examples of an MDI include the following:

- a) Connectorized fiber pigtail
- b) PMD receptacle

When the MDI is a connector plug and receptacle connection, it shall meet the interface performance specifications of IEC 61753-1-1 and IEC 61753-021-2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 122.5.1, not at the MDI.

122.12 Protocol implementation conformance statement (PICS) proforma for Clause 122, Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8⁹

122.12.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 122, Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

122.12.2 Identification

122.12.2.1 Implementation identification

| | |
|--|--|
| Supplier ¹ | |
| Contact point for inquiries about the PICS ¹ | |
| Implementation Name(s) and Version(s) ^{1,3} | |
| Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ² | |
| NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model). | |

122.12.2.2 Protocol summary

| | |
|---|---|
| Identification of protocol standard | IEEE Std 802.3bs-2017, Clause 122, Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8 |
| Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS | |
| Have any Exception items been required? No <input type="checkbox"/> Yes <input type="checkbox"/> (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3bs-2017.) | |
| Date of Statement | |

⁹Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

122.12.3 Major capabilities/options

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--|--------|-------------------|
| *FR4 | 200GBASE-FR4 PMD | 122.7 | Device supports requirements for 200GBASE-FR4 PHY | O.1 | Yes [] No [] |
| *LR4 | 200GBASE-LR4 PMD | 122.7 | Device supports requirements for 200GBASE-LR4 PHY | O.1 | Yes [] No [] |
| *FR8 | 400GBASE-FR8 PMD | 122.7 | Device supports requirements for 400GBASE-FR8 PHY | O.1 | Yes [] No [] |
| *LR8 | 400GBASE-LR8 PMD | 122.7 | Device supports requirements for 400GBASE-LR8 PHY | O.1 | Yes [] No [] |
| *INS | Installation / cable | 122.10 | Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer | O | Yes [] No [] |
| TP1 | Reference point TP1 exposed and available for testing | 122.5.1 | This point may be made available for use by implementers to certify component conformance | O | Yes [] No [] |
| TP4 | Reference point TP4 exposed and available for testing | 122.5.1 | This point may be made available for use by implementers to certify component conformance | O | Yes [] No [] |
| DC | Delay constraints | 122.5.1 | Device conforms to delay constraints | M | Yes [] |
| SC | Skew constraints | 122.3.2 | Device conforms to Skew and Skew Variation constraints | M | Yes [] |
| *MD | MDIO capability | 122.4 | Registers and interface supported | O | Yes [] No [] |

122.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 200GBASE-FR4, 200GBASE-LR4, 400GBASE-FR8, and 400GBASE-LR8

122.12.4.1 PMD functional specifications

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|---|--------|-------------------|
| F1 | Compatible with 200GBASE-R or 400GBASE-R PCS and PMA | 122.1 | | M | Yes [] |
| F2 | Integration with management functions | 122.1 | | O | Yes [] No [] |
| F3 | Bit error ratio | 122.1.1 | Meets the BER specified in 122.1.1 | M | Yes [] |
| F4 | Transmit function | 122.5.2 | Conveys symbols from PMD service interface to MDI | M | Yes [] |
| F5 | Optical multiplexing and delivery to the MDI | 122.5.2 | Optically multiplexes the four or eight optical signals for delivery to the MDI | M | Yes [] |

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---|--------|------------------------------|
| F6 | Mapping between optical signal and logical signal for transmitter | 122.5.2 | Highest optical power corresponds to tx_symbol = three | M | Yes [] |
| F7 | Receive function | 122.5.3 | Conveys symbols from MDI to PMD service interface | M | Yes [] |
| F8 | Conversion of four or eight optical signals to eight electrical signals | 122.5.3 | For delivery to the PMD service interface | M | Yes [] |
| F9 | Mapping between optical signal and logical signal for receiver | 122.5.3 | Highest optical power corresponds to rx_symbol = three | M | Yes [] |
| F10 | Global Signal Detect function | 122.5.4 | Report to the PMD service interface the message PMD:IS_SIGNAL.indication (SIGNAL_DETECT) | M | Yes [] |
| F11 | Global Signal Detect behavior | 122.5.4 | SIGNAL_DETECT is a global indicator of the presence of optical signals on all four or eight lanes | M | Yes [] |
| F12 | Lane-by-lane Signal Detect function | 122.5.5 | Sets PMD_signal_detect_i values on a lane-by-lane basis per requirements of Table 122-4 | MD:O | Yes [] No [] N/A [] |
| F13 | PMD reset function | 122.5.6 | Resets the PMD sublayer | MD:O | Yes [] No [] N/A [] |

122.12.4.2 Management functions

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|---|--------|------------------------------|
| M1 | Management register set | 122.4 | | MD:M | Yes [] N/A [] |
| M2 | Global transmit disable function | 122.5.7 | Disables all of the optical transmitters with the PMD_global_transmit_disable variable | MD:O | Yes [] No [] N/A [] |
| M3 | PMD_lane_by_lane_transmit_disable function | 122.5.8 | Disables the optical transmitter via MDIO on the lane associated with the PMD_transmit_disable_i variable | MD:O.2 | Yes [] No [] N/A [] |
| M4 | PMD lane-by-lane transmit disable | 122.5.8 | Disables each optical transmitter independently if M3 = No | O.2 | Yes [] N/A [] |
| M5 | PMD_fault function | 122.5.9 | Sets PMD_fault to one if any local fault is detected | MD:O | Yes [] No [] N/A [] |

IEEE Std 802.3bs-2017
 Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 200 Gb/s and 400 Gb/s Operation

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|-----------------------------|-----------|--|--------|------------------------------|
| M6 | PMD_transmit_fault function | 122.5.10 | Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane | MD:O | Yes [] No [] N/A [] |
| M7 | PMD_receive_fault function | 122.5.11 | Sets PMD_receive_fault to one if a local fault is detected on any receive lane | MD:O | Yes [] No [] N/A [] |

122.12.4.3 PMD to MDI optical specifications for 200GBASE-FR4

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--------------------------|--------|--------------------|
| FRF1 | Transmitter meets specifications in Table 122-9 | 122.7.1 | Per definitions in 122.8 | FR4:M | Yes [] N/A [] |
| FRF2 | Receiver meets specifications in Table 122-11 | 122.7.2 | Per definitions in 122.8 | FR4:M | Yes [] N/A [] |

122.12.4.4 PMD to MDI optical specifications for 200GBASE-LR4

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--------------------------|--------|--------------------|
| LRF1 | Transmitter meets specifications in Table 122-9 | 122.7.1 | Per definitions in 122.8 | LR4:M | Yes [] N/A [] |
| LRF2 | Receiver meets specifications in Table 122-11 | 122.7.2 | Per definitions in 122.8 | LR4:M | Yes [] N/A [] |

122.12.4.5 PMD to MDI optical specifications for 400GBASE-FR8

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|--------------------------|--------|--------------------|
| FRE1 | Transmitter meets specifications in Table 122-10 | 122.7.1 | Per definitions in 122.8 | FR8:M | Yes [] N/A [] |
| FRE2 | Receiver meets specifications in Table 122-12 | 122.7.2 | Per definitions in 122.8 | FR8:M | Yes [] N/A [] |

122.12.4.6 PMD to MDI optical specifications for 400GBASE-LR8

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|--------------------------|--------|--------------------|
| LRE1 | Transmitter meets specifications in Table 122-10 | 122.7.1 | Per definitions in 122.8 | LR8:M | Yes [] N/A [] |
| LRE2 | Receiver meets specifications in Table 122-12 | 122.7.2 | Per definitions in 122.8 | LR8:M | Yes [] N/A [] |

122.12.4.7 Optical measurement methods

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---|--------|---------|
| OM1 | Measurement cable | 122.8 | 2 m to 5 m in length | M | Yes [] |
| OM2 | Center wavelength | 122.8.2 | Per TIA/EIA-455-127-A or IEC 61280-1-3 under modulated conditions | M | Yes [] |
| OM3 | Average optical power | 122.8.3 | Per IEC 61280-1-1 | M | Yes [] |
| OM4 | OMA measurements | 122.8.4 | Each lane | M | Yes [] |
| OM5 | Transmitter and dispersion eye closure for PAM4 (TDECQ) | 122.8.5 | Each lane | M | Yes [] |
| OM6 | Extinction ratio | 122.8.6 | Each lane | M | Yes [] |
| OM7 | RIN _{xx} OMA measurement procedure | 122.8.7 | Each lane | M | Yes [] |
| OM8 | Stressed receiver sensitivity | 122.8.9 | Each lane | M | Yes [] |

122.12.4.8 Environmental specifications

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---------------------------------|-----------|--|--------|---------|
| ES1 | General safety | 122.9.1 | Conforms to IEC 60950-1 | M | Yes [] |
| ES2 | Laser safety—IEC Hazard Level 1 | 122.9.2 | Conforms to Hazard Level 1 laser requirements defined in IEC 60825-1 and IEC 60825-2 | M | Yes [] |
| ES3 | Electromagnetic interference | 122.9.5 | Complies with applicable local and national codes for the limitation of electromagnetic interference | M | Yes [] |

122.12.4.9 Characteristics of the fiber optic cabling and MDI

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|------------------------------|------------|--|--------|--------------------|
| OC1 | Fiber optic cabling | 122.11 | Meets requirements specified in Table 122-17 | INS:M | Yes [] N/A [] |
| OC2 | Maximum discrete reflectance | 122.11.2.2 | Meets requirements specified in Table 122-19 | INS:M | Yes [] N/A [] |
| OC3 | MDI requirements | 122.11.3 | Meets IEC 61753-1-1 and IEC 61753-021-2 | INS:M | Yes [] N/A [] |

123. Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-SR16

123.1 Overview

This clause specifies the 400GBASE-SR16 PMD together with the multimode fiber medium. The PMD sublayer provides a point-to-point 400 Gb/s Ethernet link over 16 pairs of multimode fiber, up to at least 100 m. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 123–1, to the medium through the MDI and optionally with the management functions that may be accessible through the management interface defined in Clause 45, or equivalent.

Table 123–1—Physical Layer clauses associated with the 400GBASE-SR16 PMD

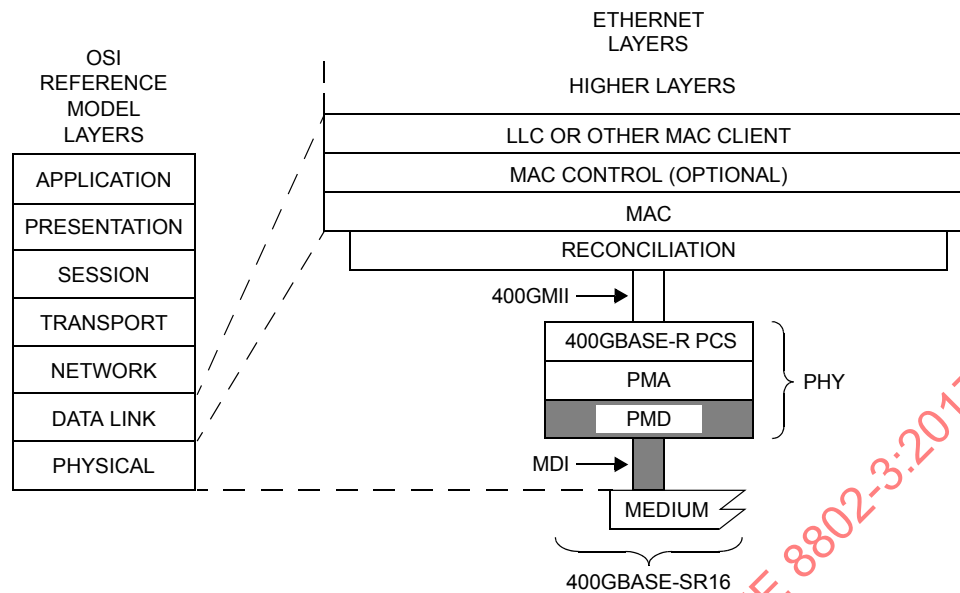
| Associated clause | 400GBASE-SR16 |
|--------------------------------|---------------|
| 117—RS | Required |
| 117—400GMII ^a | Optional |
| 118—400GMII Extender | Optional |
| 119—PCS | Required |
| 120—PMA | Required |
| 120B—Chip-to-chip 400GAUI-16 | Optional |
| 120C—Chip-to-module 400GAUI-16 | Optional |
| 120D—Chip-to-chip 400GAUI-8 | Optional |
| 120E—Chip-to-module 400GAUI-8 | Optional |
| 78—Energy Efficient Ethernet | Optional |

^a The 400GMII is an optional interface. However, if the 400GMII is not implemented, a conforming implementation must behave functionally as though the RS and 400GMII were present.

Figure 123–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 400 Gb/s Ethernet is introduced in Clause 116 and the purpose of each PHY sublayer is summarized in 116.2.

400GBASE-SR16 PHYs with the optional Energy Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

Further relevant information may be found in Clause 1 (terminology and conventions, references, definitions and abbreviations) and Annex A (bibliography, referenced as [B1], [B2], etc.).



400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT

SR = PMD FOR MULTIMODE FIBER

Figure 123–1—400GBASE-SR16 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

123.1.1 Bit error ratio

The bit error ratio (BER) shall be less than 2.4×10^{-4} provided that the error statistics are sufficiently random that this results in a frame loss ratio (see 1.4.223) of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap when processed according to Clause 119. For a complete Physical Layer, the frame loss ratio may be degraded to 6.2×10^{-11} for 64-octet frames with minimum interpacket gap due to additional errors from the electrical interfaces.

If the error statistics are not sufficiently random to meet this requirement, then the BER shall be less than that required to give a frame loss ratio of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap.

123.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 400GBASE-SR16 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 116.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_*i*.request
 PMD:IS_UNITDATA_*i*.indication
 PMD:IS_SIGNAL.indication

The 400GBASE-SR16 PMD has 16 parallel bit streams, hence $i = 0$ to 15.

In the transmit direction, the PMA continuously sends 16 parallel bit streams to the PMD, one per lane, each at a nominal signaling rate of 26.5625 GBd. The PMD converts these streams of bits into appropriate signals on the MDI.

In the receive direction, the PMD continuously sends 16 parallel bit streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 26.5625 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service primitive defined in 116.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_symbol parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_symbol parameters are known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the BER defined in 123.1.1.

123.3 Delay and Skew

123.3.1 Delay constraints

An upper bound to the delay through the PMA and PMD is required for predictable operation of the MAC Control PAUSE operation. The sum of the transmit and receive delays at one end of the link contributed by the 400GBASE-SR16 PMD including 2 m of fiber in one direction shall be no more than 8192 bit times (16 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 116.4 and its references.

123.3.2 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS sublayer. Skew and Skew Variation are defined in 116.5 and specified at the points SP1 to SP6 shown in Figure 116-4 and Figure 116-5.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 116.5. The measurements of Skew and Skew Variation are defined in 86.8.3.1.

123.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 123–2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 123–3.

Table 123–2—MDIO/PMD control variable mapping

| MDIO control variable | PMA/PMD register name | Register/bit number | PMD control variable |
|---|---|---------------------|---|
| Reset | PMA/PMD control 1 register | 1.0.15 | PMD_reset |
| Global PMD transmit disable | PMD transmit disable register | 1.9.0 | PMD_global_transmit_disable |
| PMD transmit disable 15 | PMD transmit disable extension register | 1.27.0 | PMD_transmit_disable_15 |
| PMD transmit disable 14 to PMD transmit disable 0 | PMD transmit disable register | 1.9.15 to 1.9.1 | PMD_transmit_disable_14 to PMD_transmit_disable_0 |

Table 123–3—MDIO/PMD status variable mapping

| MDIO status variable | PMA/PMD register name | Register/bit number | PMD status variable |
|---|--|---------------------|---|
| Fault | PMA/PMD status 1 register | 1.1.7 | PMD_fault |
| Transmit fault | PMA/PMD status 2 register | 1.8.11 | PMD_transmit_fault |
| Receive fault | PMA/PMD status 2 register | 1.8.10 | PMD_receive_fault |
| Global PMD receive signal detect | PMD receive signal detect register | 1.10.0 | PMD_global_signal_detect |
| PMD receive signal detect 15 | PMD receive signal detect extension register | 1.28.0 | PMD_signal_detect_15 |
| PMD receive signal detect 14 to PMD receive signal detect 0 | PMD receive signal detect register | 1.10.15 to 1.10.1 | PMD_signal_detect_14 to PMD_signal_detect_0 |

123.5 PMD functional specifications

The 400GBASE-SR16 PMD performs the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

123.5.1 PMD block diagram

The PMD block diagram is shown in Figure 123–2. For purposes of system conformance, the PMD sublayer is standardized at the points described in this subclause. The optical transmit signal is defined at the output end of a multimode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 123.8 are made at TP2. The optical receive signal is defined at

the output of the fiber optic cabling (TP3) at the MDI (see 123.11.3). Unless specified otherwise, all receiver measurements and tests defined in 123.8 are made at TP3.

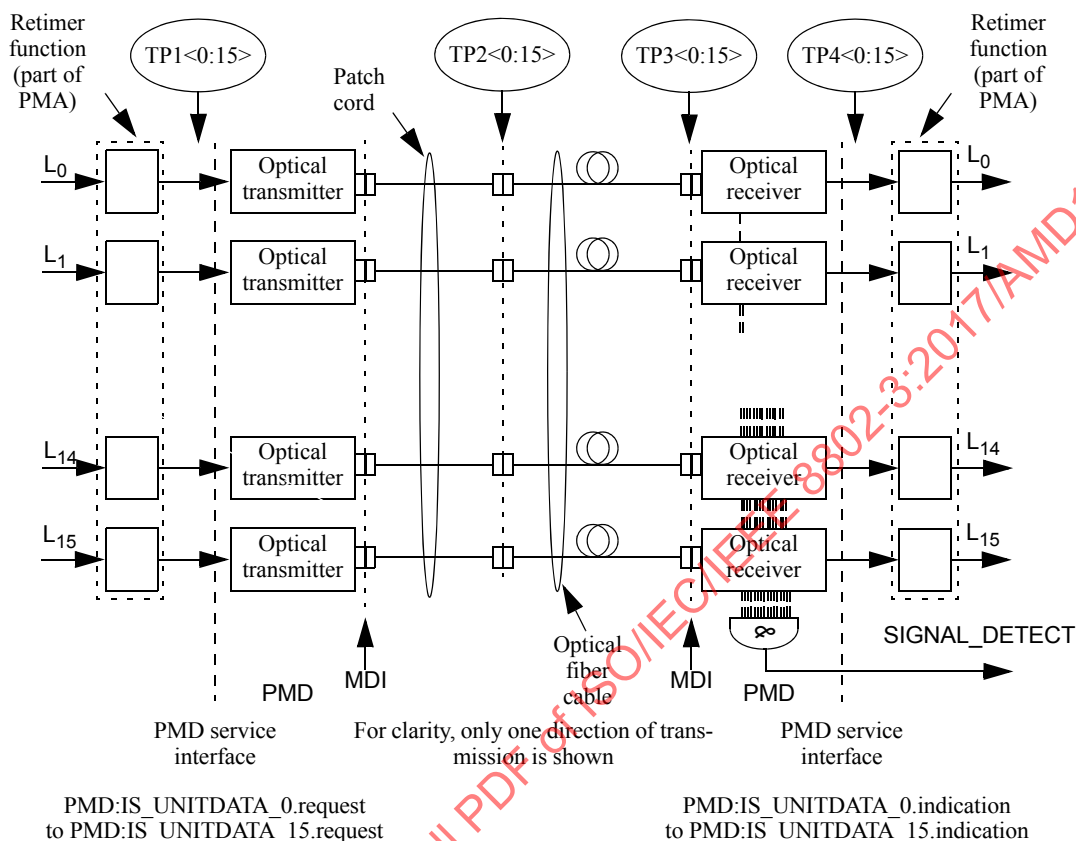


Figure 123-2—Block diagram for 400GBASE-SR16 transmit/receive paths

TP1 and TP4 are informative reference points that may be useful to implementers for testing components (these test points will not typically be accessible in an implemented system).

123.5.2 PMD transmit function

The PMD Transmit function shall convert the 16 bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_15.request into 16 optical signals. The optical signals shall then be delivered to the MDI which contains 16 parallel light paths for transmit, according to the transmit optical specifications in this clause. The higher optical power level in each signal stream shall correspond to tx_symbol = one.

123.5.3 PMD receive function

The PMD Receive function shall convert the 16 optical signals received from the MDI into separate bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_15.indication, all according to the receive optical specifications in this clause. The higher optical power level in the signal stream shall correspond to rx_symbol = one.

123.5.4 PMD global signal detect function

The PMD global signal detect function shall report the state of SIGNAL_DETECT via the PMD service interface. The SIGNAL_DETECT parameter is signaled continuously, while the PMD:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 116.3.

SIGNAL_DETECT shall be an indicator of the presence of the optical signals on all 16 lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 123-4. The PMD receiver is not required to verify whether a compliant 400GBASE-SR16 signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 123-4—SIGNAL_DETECT value definition

| Receive conditions | SIGNAL_DETECT value |
|---|---------------------|
| For any lane; Average optical power at TP3 ≤ -30 dBm | FAIL |
| For all lanes; [(Optical power at TP3 \geq average receive power, each lane (min) in Table 95-7) AND (compliant 400GBASE-SR16 signal input)] | OK |
| All other conditions | Unspecified |

As an unavoidable consequence of the requirements for the setting of the SIGNAL_DETECT parameter, implementations must provide adequate margin between the input optical power level at which the SIGNAL_DETECT parameter is set to OK, and the inherent noise level of the PMD including the effects of crosstalk, power supply noise, etc.

Various implementations of the PMD global signal detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signals and implementations that respond to the average optical power of the modulated optical signals. When the MDIO is implemented, the SIGNAL_DETECT value shall be continuously set in response to the magnitude of the optical signals, according to the requirements of Table 123-4.

123.5.5 PMD lane-by-lane signal detect function

Various implementations of the Signal Detect function are permitted by this standard. When the MDIO is implemented, each PMD_signal_detect_ i , where i represents the lane number in the range 0:15, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 123-4.

123.5.6 PMD reset function

If the MDIO interface is implemented, and if PMD_reset is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

123.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

- a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 95–6.
- b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

123.5.8 PMD lane-by-lane transmit disable function (optional)

The PMD_transmit_disable_*i* (where *i* represents the lane number in the range 0:15) function is optional and allows the optical transmitter in each lane to be selectively disabled.

- a) When a PMD_transmit_disable_*i* variable is set to one, this function shall turn off the optical transmitter associated with that variable so that the transmitter meets the requirements of the average launch power of the OFF transmitter in Table 95–6.
- b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the optical transmitter in each lane.

If the optional PMD_transmit_disable_*i* function is not implemented in MDIO, an alternative method may be provided to independently disable each transmit lane.

123.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to one.

If the MDIO interface is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

123.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on the transmitter, the PMD shall set PMD_transmit_fault to one.

If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

123.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on the receiver, the PMD shall set the PMD_receive_fault variable to one.

If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

123.6 Lane assignments

There are no lane assignments (within a group of transmit or receive lanes) for 400GBASE-SR16. While it is expected that a PMD will map electrical lane *i* to optical lane *i* and vice versa, there is no need to define the physical ordering of the lanes, as the PCS sublayer is capable of receiving the lanes in any arrangement. The positioning of transmit and receive lanes at the MDI is specified in 123.11.3.

123.7 PMD to MDI optical specifications for 400GBASE-SR16

The operating range for the 400GBASE-SR16 PMD is defined in Table 123–5. A 400GBASE-SR16 compliant PMD operates on 50/125 μm multimode fibers, type A1a.2 (OM3), type A1a.3 (OM4), or type A1a.4 (OM5), according to the specifications defined in Table 123–6. A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 400GBASE-SR16 PMD operating at 120 m meets the operating range requirement of 0.5 m to 100 m). The operating range shown in Table 123–5 is the same as 100GBASE-SR4 (See [Clause 95](#)).

Table 123–5—400GBASE-SR16 operating range

| PMD type | Required operating range ^a |
|---------------|---------------------------------------|
| 400GBASE-SR16 | 0.5 m to 70 m for OM3 |
| | 0.5 m to 100 m for OM4 |
| | 0.5 m to 100 m for OM5 |

^a The PCS FEC correction function may not be bypassed for any operating distance.

123.7.1 400GBASE-SR16 transmitter optical specifications

The optical characteristics of each lane of a 400GBASE-SR16 transmitter shall be the same as those of a single lane of 100GBASE-SR4, as specified in [Table 95–6](#) and [95.7.1](#), with the exception that the “signaling rate, each lane” is 26.5625 GBd \pm 100 ppm, TDEC is as specified in 123.8.5, and the BER requirement is as specified in 123.1.1.

123.7.2 400GBASE-SR16 receive optical specifications

The optical characteristics of each lane of a 400GBASE-SR16 receiver shall be the same as those of a single lane of 100GBASE-SR4, as specified in [Table 95–7](#) and [95.7.2](#), with the exception that the “signaling rate, each lane” is 26.5625 GBd \pm 100 ppm, stressed receiver sensitivity is as specified in 123.8.8, and the BER requirement is as specified in 123.1.1.

123.7.3 400GBASE-SR16 illustrative link power budget

The illustrative power budget and penalties for each lane of a 400GBASE-SR16 link are the same as those of a single lane of 100GBASE-SR4, as specified in [Table 95–8](#) and [95.7.3](#).

123.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified. Optical parameters and measurement methods for 400GBASE-SR16 are the same as for 100GBASE-SR4.

123.8.1 Test patterns for optical parameters

The test patterns used in this clause shall be the same as those used for 100GBASE-SR4, as described in [95.8.1](#) and shown in [Table 95–9](#), with the exception that pattern 5 defined in [82.2.11](#) is encoded by Clause 119 PCS for 400GBASE-SR16. [Table 95–10](#) as modified by this subclause (123.8.1) shows the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined.

123.8.2 Center wavelength and spectral width

The center wavelength and RMS spectral width shall be within the range given in Table 95-6 if measured per TIA/EIA-455-127-A or IEC 61280-1-3. The transmitter is modulated using one of the test patterns specified in Table 95-10 as modified by 123.8.1.

123.8.3 Average optical power

The average optical power shall be within the limits given in Table 95-6 if measured using the methods given in IEC 61280-1-1. The average optical power is measured using one of the test patterns specified in Table 95-10 as modified by 123.8.1.

123.8.4 Optical Modulation Amplitude (OMA)

OMA shall be within the limits given in Table 95-6 if measured as defined in 52.9.5 for measurement with a square wave (8 ones, 8 zeros) test pattern or as defined in 68.6.2 (from the variable MeasuredOMA in 68.6.6.2) for measurement with a PRBS9 test pattern. See 123.8.1 for test pattern information.

123.8.5 Transmitter and dispersion eye closure (TDEC)

TDEC is a measure of the optical transmitter's vertical eye closure, as described in 95.8.5. TDEC shall be within the limits given in Table 95-6 if measured using the methods specified in 95.8.5.1 and 95.8.5.2, with the exceptions that the clock recovery unit (CRU) has a corner frequency of 4 MHz and in Equation 95-6, 3.8906R is replaced by 3.4917R, for consistency with the BER of 2.4×10^{-4} given in 123.1.1. Table 95-10 as modified by 123.8.1 specifies the test patterns to be used for measurement of TDEC.

123.8.6 Extinction ratio

The extinction ratio shall be within the limits given in Table 95-6 if measured using the methods specified in IEC 61280-2-2. The extinction ratio is measured using one of the test patterns specified for extinction ratio in Table 95-10 as modified by 123.8.1.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 95-10).

123.8.7 Transmitter optical waveform (transmit eye)

The transmitter optical waveform of a port transmitting the test pattern specified in Table 95-10 as modified by 123.8.1 shall meet the transmitter eye mask coordinates and hit ratio specified in Table 95-6 if measured according to the methods specified in 95.8.7 with the exception that the clock recovery unit's high frequency corner bandwidth is 4 MHz.

123.8.8 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 95-7 if measured using the methods defined in 95.8.8 with the following exceptions:

- The signaling rate of the test pattern generator is set to the rate defined in 123.7.2.
- The clock recovery unit (CRU) has a corner bandwidth of 4 MHz.
- Sinusoidal jitter is added as specified in Table 87-13 instead of Table 95-11.
- When using 95.8.5 to measure the SEC of the stressed receiver conformance test signal, the BER requirement is as specified in 123.1.1 and 3.4917R replaces 3.8906R in Equation 95-6.
- The hit ratio for the stressed receiver eye mask definition is changed from 5×10^{-5} to 2.4×10^{-4} .

123.9 Safety, installation, environment, and labeling

123.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

123.9.2 Laser safety

400GBASE-SR16 optical transceivers shall conform to Hazard Level 1M laser requirements as defined in IEC 60825-1 and IEC 60825-2, under any condition of operation. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product's laser, safety features, labeling, use, maintenance, and service. This documentation explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.¹⁰

123.9.3 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

123.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 400GBASE-SR16 PMD over the life of the product while the product operates within the manufacturer's range of environmental, power, and other specifications.

It is recommended that manufacturers indicate, in the literature associated with the PHY, the operating environmental conditions to facilitate selection, installation, and maintenance.

It is recommended that manufacturers indicate, in the literature associated with the components of the optical link, the distance and operating environmental conditions over which the specifications of this clause will be met.

123.9.5 Electromagnetic emission

A system integrating a 400GBASE-SR16 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

123.9.6 Temperature, humidity, and handling

The optical link is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

¹⁰A host system that fails to meet the manufacturer's requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.

123.9.7 PMD labeling requirements

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user, with at least the applicable safety warnings and the applicable port type designation (e.g. 400GBASE-SR16).

Labeling requirements for Hazard Level 1M lasers are given in the laser safety standards referenced in 123.9.2.

123.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 123–3. The fiber type and length are the same as 100GBASE-SR4 (See [Clause 95](#)).




Figure 123–3—Fiber optic cabling model

The channel insertion loss is given in Table 123–6. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, modal dispersion, reflections and losses of all connectors and splices meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with IEC 61280-4-1:2009. As OM4 and OM5 optical fiber meet the requirements for OM3, a channel compliant to the “OM3” column may use OM4 or OM5 optical fiber, or a combination of OM3, OM4 and OM5. The fiber optic cabling model (channel) defined here is the same as a simplex fiber optic link segment. The term channel is used here for consistency with generic cabling standards.

Table 123–6—Fiber optic cabling (channel) characteristics for 400GBASE-SR16

| Description | OM3 | OM4 | OM5 | Unit |
|---|-----|-----|-----|------|
| Operating distance (max) | 70 | 100 | 100 | m |
| Channel insertion loss ^a (max) | 1.8 | 1.9 | 1.9 | dB |
| Channel insertion loss (min) | 0 | | | dB |

^a These channel insertion loss values include cable loss plus 1.5 dB allocated for connection and splice loss, over the wavelength range 840 nm to 860 nm.

123.11 Characteristics of the fiber optic cabling (channel)

The 400GBASE-SR16 fiber optic cabling shall meet the specifications defined in Table 123–6. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together. The requirements for the optical fiber, connection insertion loss, and maximum discrete reflectance are the same as 100GBASE-SR4 (see [Clause 95](#)).

123.11.1 Optical fiber cable

The fiber contained within the 400GBASE-SR16 fiber optic cabling shall comply with the specifications and parameters of Table 123–7. A variety of multimode cable types may satisfy these requirements, provided the resulting channel also meets the specifications of Table 123–6.

Table 123–7—Optical fiber and cable characteristics

| Description | OM3 ^a | OM4 ^b | OM5 ^c | Unit |
|--|--|------------------|--|--------------------------|
| Nominal core diameter | 50 | | | μm |
| Nominal fiber specification wavelength | 850 | | | nm |
| Effective modal bandwidth (min) ^d | 2000 | 4700 | | MHz.km |
| Cabled optical fiber attenuation (max) | 3.5 | | | dB/km |
| Zero dispersion wavelength (λ ₀) | 1295 ≤ λ ₀ ≤ 1340 | | 1297 ≤ λ ₀ ≤ 1328 | nm |
| Chromatic dispersion slope (max) (S ₀) | 0.105 for 1295 ≤ λ ₀ ≤ 1310 and 0.000375 × (1590 − λ ₀) for 1310 ≤ λ ₀ ≤ 1340 | | − 412/(840(1 − (λ ₀ /840) ⁴)) for 1297 ≤ λ ₀ ≤ 1328 | ps/nm ² km |

^a IEC 60793-2-10 type A1a.2

^b IEC 60793-2-10 type A1a.3

^c IEC 60793-2-10 type A1a.4

^d When measured with the launch conditions specified in Table 95–6

123.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 123–3, consists of a mated pair of optical connectors.

123.11.2.1 Connection insertion loss

The maximum link distance is based on an allocation of 1.5 dB total connection and splice loss. For example, this allocation supports 3 connections with an average insertion loss per connection of 0.5 dB. Connections with lower loss characteristics may be used provided the requirements of Table 123–6 are met. However, the loss of a single connection shall not exceed 0.75 dB.

123.11.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –20 dB.

123.11.3 Medium Dependent Interface (MDI)

The 400GBASE-SR16 PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the “fiber optic cabling” (as shown in Figure 123–3). Examples of an MDI include the following:

- Connectorized fiber pigtail
- PMD receptacle

123.11.3.1 Optical lane assignments

The 16 transmit and 16 receive optical lanes of 400GBASE-SR16 shall occupy the positions depicted in Figure 123-4 when looking into the MDI receptacle with the connector keyway feature on top. The interface contains 32 active lanes. The transmit optical lanes occupy the upper 16 positions. The receive optical lanes occupy the lower 16 positions.




Figure 123-4—400GBASE-SR16 optical lane assignments




Figure 123-5—MPO-16 female two-row plug with flat interface, and an MDI

123.11.3.2 Medium Dependent Interface (MDI) requirements

The MDI adapter or receptacle shall meet the dimensional specifications for designation FOCIS 18 A-k-0 as defined in ANSI/TIA-604-18. The plug terminating the optical fiber cabling shall meet the dimensional specifications of designation FOCIS 18 P-2x16-1-0-2-2 as defined in ANSI/TIA-604-18. The MDI shall optically mate with the plug on the optical fiber cabling. Figure 123-5 shows an MPO-16 female two-row plug with flat interface, and an MDI. The MDI shall meet the interface performance specifications of IEC 61753-1 and IEC 61753-022-2 for performance Class Cm/2m.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 123.5.1, not at the MDI.

123.12 Protocol implementation conformance statement (PICS) proforma for Clause 123, Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-SR16¹¹

123.12.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 123, Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-SR16, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

123.12.2 Identification

123.12.2.1 Implementation identification

| | |
|--|--|
| Supplier ¹ | |
| Contact point for inquiries about the PICS ¹ | |
| Implementation Name(s) and Version(s) ^{1,3} | |
| Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ² | |
| NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model). | |

123.12.2.2 Protocol summary

| | |
|---|--|
| Identification of protocol standard | IEEE Std 802.3bs-2017, Clause 123, Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-SR16 |
| Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS | |
| Have any Exception items been required? No <input type="checkbox"/> Yes <input type="checkbox"/> (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3bs-2017.) | |

| | |
|-------------------|--|
| Date of Statement | |
|-------------------|--|

¹¹Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

123.12.3 Major capabilities/options

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--|--------|-------------------|
| *INS | Installation / cable | 123.10 | Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer | O | Yes [] No [] |
| TP1 | Reference point TP1 exposed and available for testing | 123.5.1 | This point may be made available for use by implementers to certify component conformance | O | Yes [] No [] |
| TP4 | Reference point TP4 exposed and available for testing | 123.5.1 | This point may be made available for use by implementers to certify component conformance | O | Yes [] No [] |
| DC | Delay constraints | 123.3.1 | Device conforms to delay constraints | M | Yes [] |
| SC | Skew constraints | 123.3.2 | Device conforms to Skew and Skew Variation constraints | M | Yes [] |
| *MD | MDIO capability | 123.4 | Registers and interface supported | O | Yes [] No [] |

123.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-SR16**123.12.4.1 PMD functional specifications**

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--|--------|-------------------|
| F1 | Compatible with 400GBASE-R PCS and PMA | 123.1 | | M | Yes [] |
| F2 | Integration with management functions | 123.1 | | O | Yes [] No [] |
| F3 | Bit error ratio | 123.1.1 | Meets the BER specified in 123.1.1 | M | Yes [] |
| F4 | Transmit function | 123.5.2 | Conveys bits from PMD service interface to MDI | M | Yes [] |
| F5 | Mapping between optical signal and logical signal for transmitter | 123.5.2 | Higher optical power is a one | M | Yes [] |
| F6 | Receive function | 123.5.3 | Conveys bits from MDI to PMD service interface | M | Yes [] |
| F7 | Conversion of 16 optical signals to 16 electrical signals | 123.5.3 | For delivery to the PMD service interface | M | Yes [] |
| F8 | Mapping between optical signal and logical signal for receiver | 123.5.3 | Higher optical power is a one | M | Yes [] |

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|-------------------------------------|-----------|--|--------|------------------------------|
| F9 | Global Signal Detect function | 123.5.4 | Report to the PMD service interface the message PMD:IS_SIGNAL.indication (SIGNAL_DETECT) | M | Yes [] |
| F10 | Global Signal Detect behavior | 123.5.4 | SIGNAL_DETECT is a global indicator of the presence of optical signals on all 16 lanes | M | Yes [] |
| F11 | Lane-by-lane Signal Detect function | 123.5.5 | Sets PMD_signal_detect_i values on a lane-by-lane basis per requirements of Table 123-4 | MD:O | Yes [] No [] N/A [] |
| F12 | PMD reset function | 123.5.6 | Resets the PMD sublayer | MD:O | Yes [] No [] N/A [] |

123.12.4.2 Management functions

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|--|--------|------------------------------|
| M1 | Management register set | 123.4 | | MD:M | Yes [] N/A [] |
| M2 | Global transmit disable function | 123.5.7 | Disables all of the optical transmitters with the PMD_global_transmit_disable variable | MD:O | Yes [] No [] N/A [] |
| M3 | PMD_lane_by_lane_transmit_disable function | 123.5.8 | Disables the optical transmitter on the lane associated with the PMD_transmit_disable_i variable | MD:O | Yes [] No [] N/A [] |
| M4 | PMD_fault function | 123.5.9 | Sets PMD_fault to one if any local fault is detected | MD:O | Yes [] No [] N/A [] |
| M5 | PMD_transmit_fault function | 123.5.10 | Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane | MD:O | Yes [] No [] N/A [] |
| M6 | PMD_receive_fault function | 123.5.11 | Sets PMD_receive_fault to one if a local fault is detected on any receive lane | MD:O | Yes [] No [] N/A [] |

123.12.4.3 PMD to MDI optical specifications for 400GBASE-SR16

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--------------------------|--------|---------|
| SR1 | Transmitter meets specifications in 95.7.1 with the exceptions in 123.7.1 | 123.7.1 | Per definitions in 123.8 | M | Yes [] |
| SR2 | Receiver meets specifications in 95.7.2 with the exceptions in 123.7.2 | 123.7.2 | Per definitions in 123.8 | M | Yes [] |

123.12.4.4 Optical measurement methods

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---|--------|---------|
| OM1 | Measurement cable | 123.8 | 2 m to 5 m in length | M | Yes [] |
| OM2 | Center wavelength and spectral width | 123.8.2 | Per TIA/EIA-455-127-A or IEC 61280-1-3 under modulated conditions | M | Yes [] |
| OM3 | Average optical power | 123.8.3 | Per IEC 61280-1-1 | M | Yes [] |
| OM4 | OMA measurements | 123.8.4 | Each lane | M | Yes [] |
| OM5 | Transmitter and dispersion eye closure (TDEC) | 123.8.5 | Each lane | M | Yes [] |
| OM6 | Extinction ratio | 123.8.6 | Per IEC 61280-2-2 | M | Yes [] |
| OM7 | Transmit eye | 123.8.7 | Each lane | M | Yes [] |
| OM8 | Stressed receiver sensitivity | 123.8.8 | | M | Yes [] |

123.12.4.5 Environmental specifications

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|----------------------------------|-----------|--|--------|---------|
| ES1 | General safety | 123.9.1 | Conforms to IEC 60950-1 | M | Yes [] |
| ES2 | Laser safety—IEC Hazard Level 1M | 123.9.2 | Conforms to Hazard Level 1M laser requirements defined in IEC 60825-1 and IEC 60825-2 | M | Yes [] |
| ES3 | Electromagnetic interference | 123.9.5 | Complies with applicable local and national codes for the limitation of electromagnetic interference | M | Yes [] |

123.12.4.6 Characteristics of the fiber optic cabling and MDI

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|-------------------------------|------------|---|--------|--------------------|
| OC1 | Fiber optic cabling | 123.11 | Meets requirements specified in Table 123–6 | INS:M | Yes [] N/A [] |
| OC2 | Optical fiber characteristics | 123.11.1 | Per Table 123–7 | INS:M | Yes [] N/A [] |
| OC3 | Maximum discrete reflectance | 123.11.2.2 | Less than –20 dB | INS:M | Yes [] N/A [] |
| OC4 | MDI layout | 123.11.3.1 | Optical lane assignments per Figure 123–4 | M | Yes [] |
| OC5 | MDI dimensions | 123.11.3.2 | Meets FOCIS 18 A-k-0 as defined in ANSI/TIA-604-18 | M | Yes [] |
| OC6 | Cabling connector dimensions | 123.11.3.2 | Meets FOCIS 18 P-2x16-1-0-2-2 as defined in ANSI/TIA-604-18 | INS:M | Yes [] N/A [] |

IEEE Std 802.3bs-2017
Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters
for 200 Gb/s and 400 Gb/s Operation

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|------------------|------------|--|--------|--------------------|
| OC7 | MDI mating | 123.11.3.2 | MDI optically mates with plug on the cabling | M | Yes [] |
| OC8 | MDI requirements | 123.11.3.2 | Meets IEC 61753-1 and IEC 61753-022-2 | INS:M | Yes [] N/A [] |

IECNORM.COM :: Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/AMD10:2019

124. Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-DR4

124.1 Overview

This clause specifies the 400GBASE-DR4 PMD together with the single-mode fiber medium. The optical signal generated by this PMD type is modulated using a 4-level pulse amplitude modulation (PAM4) format. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 124–1, to the medium through the MDI and optionally with the management functions that may be accessible through the management interface defined in Clause 45, or equivalent.

Table 124–1—Physical Layer clauses associated with the 400GBASE-DR4 PMD

| Associated clause | 400GBASE-DR4 |
|--------------------------------|--------------|
| 117—RS | Required |
| 117—400GMII ^a | Optional |
| 118—400GMII Extender | Optional |
| 119—PCS | Required |
| 120—PMA | Required |
| 120B—Chip-to-chip 400GAUI-16 | Optional |
| 120C—Chip-to-module 400GAUI-16 | Optional |
| 120D—Chip-to-chip 400GAUI-8 | Optional |
| 120E—Chip-to-module 400GAUI-8 | Optional |
| 78—Energy Efficient Ethernet | Optional |

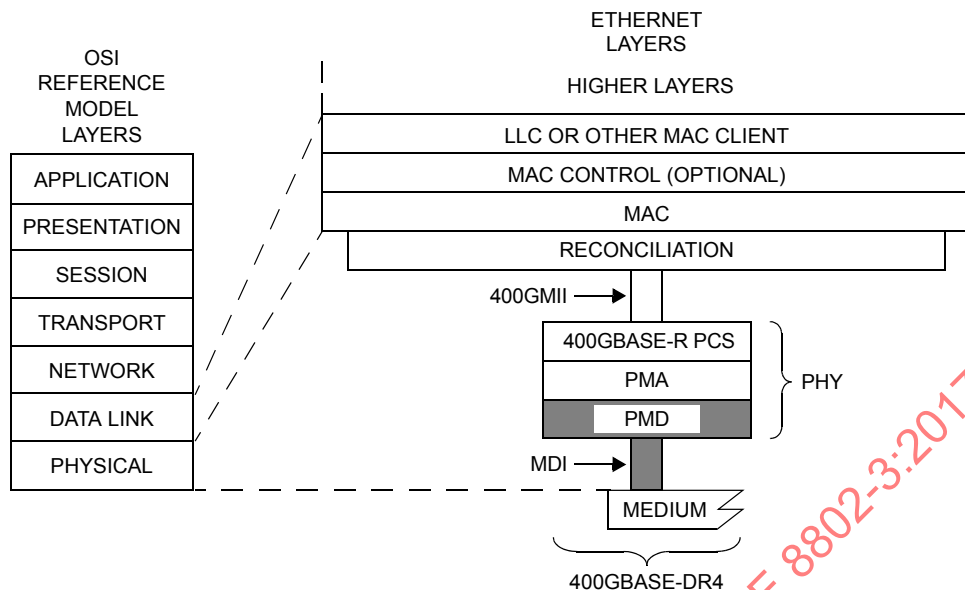
^a The 400GMII is an optional interface. However, if the 400GMII is not implemented, a conforming implementation must behave functionally as though the RS and 400GMII were present.

Figure 124–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 400 Gb/s Ethernet is introduced in Clause 116 and the purpose of each PHY sublayer is summarized in 116.2.

400GBASE-DR4 PHYs with the optional Energy Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

124.1.1 Bit error ratio

The bit error ratio (BER) when processed according to Clause 120 shall be less than 2.4×10^{-4} provided that the error statistics are sufficiently random that this results in a frame loss ratio (see 1.4.223) of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap when processed according to Clause 120 and then Clause 119. For a complete Physical Layer, the frame loss ratio may be degraded to 6.2×10^{-11} for 64-octet frames with minimum interpacket gap due to additional errors from the electrical interfaces.



400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 DR = PMD FOR SINGLE-MODE FIBER — 500 m

Figure 124–1—400GBASE-DR4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

If the error statistics are not sufficiently random to meet this requirement, then the BER shall be less than that required to give a frame loss ratio of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap.

124.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 400GBASE-DR4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 116.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA_ *i*.request
 PMD:IS_UNITDATA_ *i*.indication
 PMD:IS_SIGNAL.indication

The 400GBASE-DR4 PMD has four parallel symbol streams, hence $i = 0$ to 3.

In the transmit direction, the PMA continuously sends four parallel symbol streams to the PMD, one per lane, each at a nominal signaling rate of 53.125 GBd. The PMD then converts these streams of data units into the appropriate signals on the MDI.

In the receive direction, the PMD continuously sends four parallel symbol streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 53.125 GBd. See NOTE at the end of 120.5.2 concerning the transition density of lanes operating at this nominal signaling rate.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service interface primitive defined in 116.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_symbol parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_symbol parameters are known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the BER defined in 124.1.1.

124.3 Delay and Skew

124.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 400GBASE-DR4 PMD including 2 m of fiber in one direction shall be no more than 8192 bit times (16 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 116.4 and its references.

124.3.2 Skew constraints

The Skew (relative delay between the lanes) and Skew Variation must be kept within limits so that the information on the lanes can be reassembled by the PCS. Skew and Skew Variation are defined in 116.5 and specified at the points SP1 to SP6 shown in Figure 116-4 and Figure 116-5.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 116.5. The measurements of Skew and Skew Variation are defined in 86.8.3.1.

124.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO

control variables to PMD control variables shall be as shown in Table 124–2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 124–3.

Table 124–2—MDIO/PMD control variable mapping

| MDIO control variable | PMA/PMD register name | Register/bit number | PMD control variable |
|--|-------------------------------|---------------------|--|
| Reset | PMA/PMD control 1 register | 1.0.15 | PMD_reset |
| Global PMD transmit disable | PMD transmit disable register | 1.9.0 | PMD_global_transmit_disable |
| PMD transmit disable 3 to PMD transmit disable 0 | PMD transmit disable register | 1.9.4 to 1.9.1 | PMD_transmit_disable_3 to PMD_transmit_disable_0 |

Table 124–3—MDIO/PMD status variable mapping

| MDIO status variable | PMA/PMD register name | Register/bit number | PMD status variable |
|--|------------------------------------|---------------------|--|
| Fault | PMA/PMD status 1 register | 1.1.7 | PMD_fault |
| Transmit fault | PMA/PMD status 2 register | 1.8.11 | PMD_transmit_fault |
| Receive fault | PMA/PMD status 2 register | 1.8.10 | PMD_receive_fault |
| Global PMD receive signal detect | PMD receive signal detect register | 1.10.0 | PMD_global_signal_detect |
| PMD receive signal detect 3 to PMD receive signal detect 0 | PMD receive signal detect register | 1.10.4 to 1.10.1 | PMD_signal_detect_3 to PMD_signal_detect_0 |

124.5 PMD functional specifications

The 400GBASE-DR4 PMD performs the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

124.5.1 PMD block diagram

The PMD block diagram is shown in Figure 124–2. For purposes of system conformance, the PMD sublayer is standardized at the points described in this subclause. The optical transmit signal is defined at the output end of a single-mode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 124.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 124.11.3). Unless specified otherwise, all receiver measurements and tests defined in 124.8 are made at TP3.

TP1<0:3> and TP4<0:3> are informative reference points that may be useful to implementors for testing components (these test points will not typically be accessible in an implemented system).



The PMD Transmit function shall convert the four symbol streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate optical signals. The four optical signals shall then be delivered to the MDI, which contains four parallel light paths for transmit, according to the transmit optical specifications in this clause. The highest optical power level in each signal stream shall correspond to tx_symbol = three and the lowest shall correspond to tx_symbol = zero.

The PMD Receive function shall convert the four parallel optical signals received from the MDI into separate symbol streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to rx_symbol = three and the lowest shall correspond to rx_symbol = zero.

The PMD global signal detect function shall report the state of SIGNAL_DETECT via the PMD service interface. The SIGNAL_DETECT parameter is signaled continuously, while the PMD:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 116.3.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all four lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 124-4. The PMD receiver is not required to verify whether a compliant 400GBASE-DR4 signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 124-4—SIGNAL_DETECT value definition

| Receive conditions | SIGNAL_DETECT value |
|--|---------------------|
| For any lane; Average optical power at TP3 ≤ -15 dBm | FAIL |
| For all lanes; [(Optical power at TP3 \geq average receive power, each lane (min) Table 124-7) AND (compliant 400GBASE-R signal input)] | OK |
| All other conditions | Unspecified |

As an unavoidable consequence of the requirements for the setting of the SIGNAL_DETECT parameter, implementations must provide adequate margin between the input optical power level at which the SIGNAL_DETECT parameter is set to OK, and the inherent noise level of the PMD including the effects of crosstalk, power supply noise, etc.

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the modulated optical signal.

124.5.5 PMD lane-by-lane signal detect function

Various implementations of the Signal Detect function are permitted by this standard. When the MDIO is implemented, each PMD_signal_detect i , where i represents the lane number in the range 0:3, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 124-4.

124.5.6 PMD reset function

If the MDIO interface is implemented, and if PMD_reset is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

124.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

- a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 124-6.
- b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

124.5.8 PMD lane-by-lane transmit disable function (optional)

The PMD_transmit_disable_*i* (where *i* represents the lane number in the range 0:3) function is optional and allows the optical transmitters in each lane to be selectively disabled.

- a) When a PMD_transmit_disable_*i* variable is set to one, this function shall turn off the optical transmitter associated with that variable so that the transmitter meets the requirements of the average launch power of the OFF transmitter in Table 124–6.
- b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the optical transmitter in each lane.

If the optional PMD_transmit_disable_*i* function is not implemented in MDIO, an alternative method may be provided to independently disable each transmit lane for testing purposes.

124.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to one.

If the MDIO interface is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

124.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the PMD_transmit_fault variable to one.

If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

124.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to one.

If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

124.6 Lane assignments

There are no lane assignments (within a group of transmit or receive lanes) for 400GBASE-DR4. While it is expected that a PMD will map electrical lane *i* to optical lane *i* and vice versa, there is no need to define the physical ordering of the lanes, as the PCS sublayer is capable of receiving the lanes in any arrangement. The positioning of transmit and receive lanes at the MDI is specified in 124.11.3.1.

124.7 PMD to MDI optical specifications for 400GBASE-DR4

The operating range for the 400GBASE-DR4 PMD is defined in Table 124–5. A 400GBASE-DR4 compliant PMD operates on type B1.1, B1.3, or B6 single-mode fibers according to the specifications defined in Table 124–12. A PMD that exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., a 400GBASE-DR4 PMD operating at 600 m meets the operating range requirement of 2 m to 500 m).

Table 124–5—400GBASE-DR4 operating range

| PMD type | Required operating range |
|--------------|--------------------------|
| 400GBASE-DR4 | 2 m to 500 m |

124.7.1 400GBASE-DR4 transmitter optical specifications

The 400GBASE-DR4 transmitter shall meet the specifications defined in Table 124–6 per the definitions in 124.8.

Table 124–6—400GBASE-DR4 transmit characteristics

| Description | Value | Unit |
|--|----------------------|-------|
| Signaling rate, each lane (range) | 53.125 ± 100 ppm | GBd |
| Modulation format | PAM4 | — |
| Lane wavelength (range) | 1304.5 to 1317.5 | nm |
| Side-mode suppression ratio (SMSR), (min) | 30 | dB |
| Average launch power, each lane (max) | 4 | dBm |
| Average launch power, each lane ^a (min) | –2.9 | dBm |
| Outer Optical Modulation Amplitude (OMA _{outer}), each lane (max) | 4.2 | dBm |
| Outer Optical Modulation Amplitude (OMA _{outer}), each lane (min) ^b | –0.8 | dBm |
| Launch power in OMA _{outer} minus TDECQ, each lane (min) | –2.2 | dBm |
| Transmitter and dispersion eye closure for PAM4 (TDECQ), each lane (max) | 3.4 | dB |
| Average launch power of OFF transmitter, each lane (max) | –15 | dBm |
| Extinction ratio, each lane (min) | 3.5 | dB |
| RIN _{21.4} OMA (max) | –136 | dB/Hz |
| Optical return loss tolerance (max) | 21.4 | dB |
| Transmitter reflectance ^c (max) | –26 | dB |

^a Average launch power, each lane (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.

^b Even if the TDECQ < 1.4 dB, the OMA_{outer} (min) must exceed these values.

^c Transmitter reflectance is defined looking into the transmitter.

124.7.2 400GBASE-DR4 receive optical specifications

The 400GBASE-DR4 receiver shall meet the specifications defined in Table 124–7 per the definitions in 124.8. See NOTE at the end of 120.5.2 concerning the transition density of lanes operating at this nominal signaling rate.

Table 124-7—400GBASE-DR4 receive characteristics

| Description | Value | Unit |
|---|------------------|------|
| Signaling rate, each lane (range) | 53.125 ± 100 ppm | GBd |
| Modulation format | PAM4 | — |
| Lane wavelengths (range) | 1304.5 to 1317.5 | nm |
| Damage threshold ^a , each lane | 5 | dBm |
| Average receive power, each lane (max) | 4 | dBm |
| Average receive power, each lane ^b (min) | −5.9 | dBm |
| Receive power (OMA _{outer}), each lane (max) | 4.2 | dBm |
| Receiver reflectance (max) | −26 | dB |
| Receiver sensitivity (OMA _{outer}), each lane ^c (max) | −4.4 | dBm |
| Stressed receiver sensitivity (OMA _{outer}), each lane ^d (max) | −1.9 | dBm |
| Conditions of stressed receiver sensitivity test: ^e | | |
| Stressed eye closure for PAM4 (SECQ), lane under test | 3.4 | dB |
| OMA _{outer} of each aggressor lane | 4.2 | dBm |

^a The receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level. The receiver does not have to operate correctly at this input power.

^b Average receive power, each lane (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.

^c Receiver sensitivity (OMA_{outer}), each lane (max) is informative and is defined for a transmitter with SECQ of 0.9 dB.

^d Measured with conformance test signal at TP3 (see 124.8.9) for the BER specified in 124.1.1.

^e These test conditions are for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

124.7.3 400GBASE-DR4 illustrative link power budget

An illustrative power budget and penalties for 400GBASE-DR4 channels are shown in Table 124-8.

Table 124-8—400GBASE-DR4 illustrative link power budget

| Parameter | Value | Unit |
|---|----------------|------|
| Power budget (for max TDECQ) | 6.5 | dB |
| Operating distance | 500 | m |
| Channel insertion loss ^a | 3 | dB |
| Maximum discrete reflectance | See 124.11.2.2 | dB |
| Allocation for penalties ^b (for max TDECQ) | 3.5 | dB |
| Additional insertion loss allowed | 0 | dB |

^a The channel insertion loss is calculated using the maximum distance specified in Table 124-5 and cabled optical fiber attenuation of 0.5 dB/km at 1304.5 nm plus an allocation for connection and splice loss given in 124.11.2.1.

^b Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

124.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

124.8.1 Test patterns for optical parameters

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 124–10 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 124–10 may be used to perform that test. The test patterns used in this clause are shown in Table 124–9.

Table 124–9—Test patterns

| Pattern | Pattern description | Defined in |
|-------------|---------------------------------|--------------|
| Square wave | Square wave (8 threes, 8 zeros) | 120.5.11.2.4 |
| 3 | PRBS31Q | 120.5.11.2.2 |
| 4 | PRBS13Q | 120.5.11.2.1 |
| 5 | Scrambled idle | 119.2.4.9 |
| 6 | SSPRQ | 120.5.11.2.3 |

Table 124–10—Test-pattern definitions and related subclauses

| Parameter | Pattern | Related subclause |
|---|--|-------------------|
| Wavelength | Square wave, 3, 4, 5, 6 or valid 400GBASE-R signal | 124.8.2 |
| Side mode suppression ratio | 3, 5, 6 or valid 400GBASE-R signal | — |
| Average optical power | 3, 5, 6 or valid 400GBASE-R signal | 124.8.3 |
| Outer Optical Modulation Amplitude (OMA_{outer}) | 4 or 6 | 124.8.4 |
| Transmitter and dispersion eye closure for PAM4 (TDECQ) | 6 | 124.8.5 |
| Extinction ratio | 4 or 6 | 124.8.6 |
| $RIN_{21.4OMA}$ | Square wave | 124.8.7 |
| Stressed receiver conformance test signal calibration | 6 | 124.8.9 |
| Stressed receiver sensitivity | 3 or 5 | 124.8.9 |

124.8.2 Wavelength

The wavelength of each optical lane shall be within the range given in Table 124–6 if measured per TIA/EIA-455-127-A or IEC 61280-1-3. The lane under test is modulated using the test pattern defined in Table 124–10.

124.8.3 Average optical power

The average optical power of each lane shall be within the limits given in Table 124–6 if measured using the methods given in IEC 61280-1-1. The average optical power is measured using the test pattern defined in Table 124–10, per the test setup in Figure 53-6.

124.8.4 Outer Optical Modulation Amplitude (OMA_{outer})

The OMA_{outer} of each lane shall be within the limits given in Table 124–6. The OMA_{outer} is measured using a test pattern specified for OMA_{outer} in Table 124–10 as the difference between the average optical launch power level P_3 , measured over the central 2 UI of a run of 7 threes, and the average optical launch power level P_0 , measured over the central 2 UI of a run of 6 zeros, as shown in Figure 124–3.




Figure 124–3—Example power levels P_0 and P_3 from PRBS13Q test pattern

124.8.5 Transmitter and dispersion eye closure for PAM4 (TDECQ)

The TDECQ of each lane shall be within the limits given in Table 124–6 if measured using the methods specified in 121.8.5.1, 121.8.5.2, and 121.8.5.3 using a reference equalizer as described in 121.8.5.4, with the following exceptions:

- The signaling rate of the test pattern generator is as given in Table 124–6 and uses the test pattern specified for TDECQ in Table 124–10.
- The combination of the O/E converter and the oscilloscope has a fourth-order Bessel-Thomson filter response with a bandwidth of approximately 26.5625 GHz.
- The normalized noise power density spectrum $N(f)$ is equivalent to white noise filtered by a fourth-order Bessel-Thomson response filter with a bandwidth of 26.5625 GHz.

124.8.6 Extinction ratio

The extinction ratio of each lane shall be within the limits given in Table 124–6 if measured using a test pattern specified for extinction ratio in Table 124–10. The extinction ratio of a PAM4 optical signal is defined as the ratio of the average optical launch power level P_3 , measured over the central 2 UI of a run of 7 threes, and the average optical launch power level P_0 , measured over the central 2 UI of a run of 6 zeros, as shown in Figure 124–3.

124.8.7 Relative intensity noise (RIN_{21.4OMA})

RIN shall be as defined by the measurement methodology of 52.9.6 with the following exceptions:

- a) The optical return loss is 21.4 dB.
- b) Each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm.
- c) The upper –3 dB limit of the measurement apparatus is to be approximately equal to the signaling rate (i.e., 53.2 GHz).
- d) The test pattern is according to Table 124–10.

124.8.8 Receiver sensitivity

Receiver sensitivity, which is defined for an input signal with SECQ of 0.9 dB (e.g., an ideal input signal without overshoot), is informative and compliance is not required. If measured, the test signal should have negligible impairments such as intersymbol interference (ISI), rise/fall times, jitter and RIN. Instead, the normative requirement for receivers is stressed receiver sensitivity.

124.8.9 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 124–7 if measured using the method defined in 121.8.9 with the following exceptions:

- The SECQ of the stressed receiver conformance test signal is measured according to 124.8.5, except that the test fiber is not used.
- The signaling rate of the test pattern generator and the extinction ratio of the E/O converter are as given in Table 124–6 using test patterns specified in Table 124–10.
- The required values of the “Stressed receiver sensitivity (OMA_{outer}), each lane (max)”, “Stressed eye closure for PAM4 (SECQ), lane under test”, and “OMA_{outer} of each aggressor lane” are as given in Table 124–7.

124.9 Safety, installation, environment, and labeling

124.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

124.9.2 Laser safety

400GBASE-DR4 optical transceivers shall conform to Hazard Level 1 laser requirements as defined in IEC 60825-1 and IEC 60825-2, under any condition of operation. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product’s laser, safety features, labeling, use, maintenance, and service. This documentation explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.¹²

¹²A host system that fails to meet the manufacturer’s requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.

124.9.3 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

124.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 400GBASE-DR4 PMD over the life of the product while the product operates within the manufacturer's range of environmental, power, and other specifications.

It is recommended that manufacturers indicate in the literature associated with the PHY the operating environmental conditions to facilitate selection, installation, and maintenance.

It is recommended that manufacturers indicate, in the literature associated with the components of the optical link, the distance and operating environmental conditions over which the specifications of this clause will be met.

124.9.5 Electromagnetic emission

A system integrating a 400GBASE-DR4 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

124.9.6 Temperature, humidity, and handling

The optical link is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

124.9.7 PMD labeling requirements

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user, with at least the applicable safety warnings and the applicable port type designation (e.g., 400GBASE-DR4).

Labeling requirements for Hazard Level 1 lasers are given in the laser safety standards referenced in 124.9.2.

124.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 124-4.




Figure 124-4—Fiber optic cabling model

The channel insertion loss is given in Table 124–11. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, and polarization mode dispersion meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with ANSI/TIA/EIA-526-7/method A-1. The fiber optic cabling model (channel) defined here is the same as a simplex fiber optic link segment. The term channel is used here for consistency with generic cabling standards.

Table 124–11—Fiber optic cabling (channel) characteristics

| Description | 400GBASE-DR4 | Unit |
|---|--------------|-------|
| Operating distance (max) | 500 | m |
| Channel insertion loss ^{a,b} (max) | 3 | dB |
| Channel insertion loss (min) | 0 | dB |
| Positive dispersion ^b (max) | 0.8 | ps/nm |
| Negative dispersion ^b (min) | –0.93 | ps/nm |
| DGD_max ^c | 2.24 | ps |
| Optical return loss (min) | 37 | dB |

^a These channel insertion loss values include cable, connectors, and splices.

^b Over the wavelength range 1304.5 nm to 1317.5 nm

^c Differential Group Delay (DGD) is the time difference at reception between the fractions of a pulse that were transmitted in the two principal states of polarization of an optical signal. DGD_max is the maximum differential group delay that the system must tolerate.

124.11 Characteristics of the fiber optic cabling (channel)

The 400GBASE-DR4 fiber optic cabling shall meet the specifications defined in Table 124–12. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

124.11.1 Optical fiber cable

The fiber optic cable requirements are satisfied by cables containing IEC 60793-2-50 type B1.1 (dispersion un-shifted single-mode), type B1.3 (low water peak single-mode), or type B6_a (bend insensitive) fibers or the requirements in Table 124–12 where they differ.

Table 124–12—Optical fiber and cable characteristics

| Description | Value | Unit |
|--|---------------------------------|-----------------------|
| Nominal fiber specification wavelength | 1310 | nm |
| Cabled optical fiber attenuation (max) | 0.5 ^a | dB/km |
| Zero dispersion wavelength (λ_0) | $1300 \leq \lambda_0 \leq 1324$ | nm |
| Dispersion slope (max) (S_0) | 0.093 | ps/nm ² km |

^a The 0.5 dB/km attenuation is provided for Outside Plant cable as defined in ANSI/TIA 568-C.3.

124.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 124–4, consists of a mated pair of optical connectors.

124.11.2.1 Connection insertion loss

The maximum link distance is based on an allocation of 2.75 dB total connection and splice loss. For example, this allocation supports five connections with an average insertion loss per connection of 0.5 dB. Connections with different loss characteristics may be used provided the requirements of Table 124–11 are met.

124.11.2.2 Maximum discrete reflectance

The maximum value for each discrete reflectance shall be less than or equal to the value shown in Table 124–13 corresponding to the number of discrete reflectances above –55 dB within the channel. For numbers of discrete reflectances in between two numbers shown in the table, the lower of the two corresponding maximum discrete reflectance values applies.

Table 124–13—Maximum value of each discrete reflectance

| Number of discrete reflectances above –55 dB | Maximum value for each discrete reflectance |
|--|---|
| 1 | –37 dB |
| 2 | –42 dB |
| 4 | –45 dB |
| 6 | –47 dB |
| 8 | –48 dB |
| 10 | –49 dB |

124.11.3 Medium Dependent Interface (MDI)

The 400GBASE-DR4 PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the “fiber optic cabling” (as shown in Figure 124–4). The 400GBASE-DR4 PMD is coupled to the fiber optic cabling through one connector plug into the MDI optical receptacle as shown in Figure 124–6. Example constructions of the MDI include the following:

- PMD with a connectorized fiber pigtail plugged into an adapter
- PMD receptacle

124.11.3.1 Optical lane assignments

The four transmit and four receive optical lanes of 400GBASE-DR4 shall occupy the positions depicted in Figure 124–5 when looking into the MDI receptacle with the connector keyway feature on top. The interface contains eight active lanes within twelve total positions. The transmit optical lanes occupy the left-most four positions. The receive optical lanes occupy the right-most four positions. The four center positions are unused.




Figure 124-5—400GBASE-DR4 optical lane assignments

124.11.3.2 Medium Dependent Interface (MDI) requirements

The MDI shall meet the dimensional specifications of IEC 61754-7-1 interface 7-1-9: *MPO device receptacle, angled interface*. The plug terminating the optical fiber cabling shall meet the dimensional specifications of IEC 61754-7-1 interface 7-1-1: *MPO female plug connector, down-angled interface for 2 to 12 fibres*. The MDI shall optically mate with the plug on the optical fiber cabling. Figure 124-6 shows an MPO female plug connector with down-angled interface, and an MDI as an active device receptacle with angled interface.




Figure 124-6—MPO female plug with down-angled interface and MDI active device receptacle with angled interface

The MDI shall meet the interface performance specifications of IEC 61753-021-2 for performance level D/2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 124.5.1, not at the MDI.

124.12 Protocol implementation conformance statement (PICS) proforma for Clause 124, Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-DR4¹³

124.12.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 124, Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-DR4, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

124.12.2 Identification

124.12.2.1 Implementation identification

| | |
|--|--|
| Supplier ¹ | |
| Contact point for inquiries about the PICS ¹ | |
| Implementation Name(s) and Version(s) ^{1,3} | |
| Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ² | |
| NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model). | |

124.12.2.2 Protocol summary

| | |
|---|---|
| Identification of protocol standard | IEEE Std 802.3bs-2017, Clause 124, Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-DR4 |
| Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS | |
| Have any Exception items been required? No [] Yes [] (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3bs-2017.) | |

| | |
|-------------------|--|
| Date of Statement | |
|-------------------|--|

¹³Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

124.12.3 Major capabilities/options

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--|--------|-------------------|
| *INS | Installation / cable | 124.10 | Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer | O | Yes [] No [] |
| TP1 | Reference point TP1 exposed and available for testing | 124.5.1 | This point may be made available for use by implementers to certify component conformance | O | Yes [] No [] |
| TP4 | Reference point TP4 exposed and available for testing | 124.5.1 | This point may be made available for use by implementers to certify component conformance | O | Yes [] No [] |
| DC | Delay constraints | 124.3.1 | Device conforms to delay constraints | M | Yes [] |
| SC | Skew constraints | 124.3.2 | Device conforms to Skew and Skew Variation constraints | M | Yes [] |
| *MD | MDIO capability | 124.4 | Registers and interface supported | O | Yes [] No [] |

124.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 400GBASE-DR4**124.12.4.1 PMD functional specifications**

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--|--------|-------------------|
| F1 | Compatible with 400GBASE-R PCS and PMA | 124.1 | | M | Yes [] |
| F2 | Integration with management functions | 124.1 | | O | Yes [] No [] |
| F3 | Bit error ratio | 124.1.1 | Meets the BER specified in 124.1.1 | M | Yes [] |
| F4 | Transmit function | 124.5.2 | Conveys symbols from PMD service interface to MDI | M | Yes [] |
| F5 | Mapping between optical signal and logical signal for transmitter | 124.5.2 | Highest optical power corresponds to tx_symbol = three | M | Yes [] |
| F6 | Receive function | 124.5.3 | Conveys symbols from MDI to PMD service interface | M | Yes [] |
| F7 | Conversion of four optical signals to four electrical signals | 124.5.3 | For delivery to the PMD service interface | M | Yes [] |
| F8 | Mapping between optical signal and logical signal for receiver | 124.5.3 | Highest optical power corresponds to rx_symbol = three | M | Yes [] |

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|-------------------------------------|-----------|--|--------|------------------------------|
| F9 | Global Signal Detect function | 124.5.4 | Report to the PMD service interface the message PMD:IS_SIGNAL.indication (SIGNAL_DETECT) | M | Yes [] |
| F10 | Global Signal Detect behavior | 124.5.4 | SIGNAL_DETECT is a global indicator of the presence of optical signals on all four lanes | M | Yes [] |
| F11 | Lane-by-lane Signal Detect function | 124.5.5 | Sets PMD_signal_detect_i values on a lane-by-lane basis per requirements of Table 124-4 | MD:O | Yes [] No [] N/A [] |
| F12 | PMD reset function | 124.5.6 | Resets the PMD sublayer | MD:O | Yes [] No [] N/A [] |

124.12.4.2 Management functions

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|--|-----------|--|--------|------------------------------|
| M1 | Management register set | 124.4 | | MD:M | Yes [] N/A [] |
| M2 | Global transmit disable function | 124.5.7 | Disables all of the optical transmitters with the PMD_global_transmit_disable variable | MD:O | Yes [] No [] N/A [] |
| M3 | PMD_lane_by_lane_transmit_disable function | 124.5.8 | Disables the optical transmitter on the lane associated with the PMD_transmit_disable_i variable | MD:O | Yes [] No [] N/A [] |
| M4 | PMD_fault function | 124.5.9 | Sets PMD_fault to one if any local fault is detected | MD:O | Yes [] No [] N/A [] |
| M5 | PMD_transmit_fault function | 124.5.10 | Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane | MD:O | Yes [] No [] N/A [] |
| M6 | PMD_receive_fault function | 124.5.11 | Sets PMD_receive_fault to one if a local fault is detected on any receive lane | MD:O | Yes [] No [] N/A [] |

124.12.4.3 PMD to MDI optical specifications for 400GBASE-DR4

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|--------------------------|--------|---------|
| DR1 | Transmitter meets specifications in Table 124-6 | 124.7.1 | Per definitions in 124.8 | M | Yes [] |
| DR2 | Receiver meets specifications in Table 124-7 | 124.7.2 | Per definitions in 124.8 | M | Yes [] |

124.12.4.4 Optical measurement methods

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---|-----------|---|--------|---------|
| OM1 | Measurement cable | 124.8 | 2 m to 5 m in length | M | Yes [] |
| OM2 | Center wavelength | 124.8.2 | Per TIA/EIA-455-127-A or IEC 61280-1-3 under modulated conditions | M | Yes [] |
| OM3 | Average optical power | 124.8.3 | Per IEC 61280-1-1 | M | Yes [] |
| OM4 | OMA measurements | 124.8.4 | Each lane | M | Yes [] |
| OM5 | Transmitter and dispersion eye closure for PAM4 (TDECQ) | 124.8.5 | Each lane | M | Yes [] |
| OM6 | Extinction ratio | 124.8.6 | Each lane | M | Yes [] |
| OM7 | Stressed receiver sensitivity | 124.8.9 | Each lane | M | Yes [] |

124.12.4.5 Environmental specifications

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|---------------------------------|-----------|--|--------|---------|
| ES1 | General safety | 124.9.1 | Conforms to IEC 60950-1 | M | Yes [] |
| ES2 | Laser safety—IEC Hazard Level 1 | 124.9.2 | Conforms to Hazard Level 1 laser requirements defined in IEC 60825-1 and IEC 60825-2 | M | Yes [] |
| ES3 | Electromagnetic interference | 124.9.5 | Complies with applicable local and national codes for the limitation of electromagnetic interference | M | Yes [] |

124.12.4.6 Characteristics of the fiber optic cabling and MDI

| Item | Feature | Subclause | Value/Comment | Status | Support |
|------|------------------------------|------------|--|--------|--------------------|
| OC1 | Fiber optic cabling | 124.11 | Meets requirements specified in Table 124-11 | INS:M | Yes [] N/A [] |
| OC2 | Maximum discrete reflectance | 124.11.2.2 | Meets requirements specified in Table 124-13 | INS:M | Yes [] N/A [] |
| OC3 | MDI layout | 124.11.3.1 | Optical lane assignments per Figure 124-5 | M | Yes [] |
| OC4 | MDI dimensions | 124.11.3.2 | Per IEC 61754-7-1 interface 7-1-9 | M | Yes [] |
| OC5 | Cabling connector dimensions | 124.11.3.2 | Per IEC 61754-7-1 interface 7-1-1 | INS:M | Yes [] N/A [] |
| OC6 | MDI mating | 124.11.3.2 | MDI optically mates with plug on the cabling | M | Yes [] |