

Second edition
2017-03-01

AMENDMENT 2
2017-11

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

**Part 3:
Standard for Ethernet**

**AMENDMENT 2: Media access control
parameters, physical layers, and
management parameters for 25 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 2: Paramètres de contrôle d'accès au support, couches
physiques et paramètres de gestion pour l'exploitation des interfaces
à 25 Go/s*



Reference number
ISO/IEC/IEEE 8802-3:2017/Amd. 2:2017(E)

© IEEE 2016



COPYRIGHT PROTECTED DOCUMENT

© IEEE 2016

All rights reserved. Unless otherwise specified, 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 address below or ISO's member body in the country of the requester.

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

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

stds.ipr@ieee.org
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.

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.

The main task of ISO/IEC JTC 1 is to prepare International Standards. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

Attention is called to the possibility that implementation of this standard may require the use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. ISO/IEEE is not responsible for identifying essential patents or patent claims for which a license may be required, for conducting inquiries into the legal validity or scope of patents or patent claims or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance or a Patent Statement and Licensing Declaration Form, 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 ISO or the IEEE Standards Association.

ISO/IEC/IEEE 8802-3:2017/Amd.2 was prepared by the LAN/MAN of the IEEE Computer Society (as IEEE STD 802.3by-2016). It was adopted by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 6, *Telecommunications and information exchange between systems*, in parallel with its approval by the ISO/IEC national bodies, under the "fast-track procedure" defined in the Partner Standards Development Organization cooperation agreement between ISO and IEEE. IEEE is responsible for the maintenance of this document with participation and input from ISO/IEC national bodies.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

IEEE 802.3by™-2016
(Amendment to
IEEE Std 802.3™-2015
as amended by
IEEE Std 802.3bw™-2015)

IEEE Standard for Ethernet

Amendment 2: Media Access Control Parameters, Physical Layers, and Management Parameters for 25 Gb/s Operation

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 30 June 2016

IEEE-SA Standards Board

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 88023:2017/Amd 2:2017

Abstract: This amendment to IEEE Std 802.3-2015 adds Physical Layer (PHY) specifications and management parameters for 25 Gb/s operation over twinaxial copper cabling (25GBASE-CR and 25GBASE-CR-S), electrical backplanes (25GBASE-KR and 25GBASE-KR-S), and multimode fiber (25GBASE-SR). This amendment also specifies a 25 Gigabit Attachment Unit Interface (25GAUI) and optional Energy Efficient Ethernet (EEE).

Keywords: 25 Gigabit Ethernet, 25GAUI, 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, 25GBASE-KR-S, 25GBASE-SR, 25GMII, Auto-Negotiation (AN), Backplane Ethernet, Energy Efficient Ethernet (EEE), Ethernet, Forward Error Correction (FEC), IEEE 802[®], IEEE 802.3[™], IEEE 802.3by[™], multimode fiber (MMF), Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, Physical Medium Dependent (PMD) sublayer, Reconciliation Sublayer (RS)

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

Copyright © 2016 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 29 July 2016. 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.

Print: ISBN 978-1-5044-2169-0 STD21020
PDF: ISBN 978-1-5044-2170-6 STDPD21020

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 Notice” or “Important Notices and Disclaimers Concerning IEEE Standards Documents.”

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. Volunteers 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 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-SA Website at <http://ieeexplore.ieee.org/Xplore/home.jsp> or contact IEEE at the address listed previously. For more information about the IEEE-SA or IOWA'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 P802.3by Working Group ballot. Individuals may have not voted, voted for approval, disapproval or abstained on this standard.

David J. Law, *IEEE 802.3 Working Group Chair*

Adam Healey, *IEEE 802.3 Working Group Vice Chair*

Peter 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*

Mark Nowell, *IEEE P802.3by 25 Gb/s Ethernet Task Force Chair*

Matthew Brown, *IEEE P802.3by 25 Gb/s Ethernet Task Force Editor-in-Chief*

John Abbott	Eugene Dai	Brian Holden
David Abramson	Shaoan Dai	Rita Horner
Shadi Abughazaleh	John D'Ambrosia	Bernd Horrmeyer
Faisal Ahmad	Mike Darling	Victor Hou
Dale Amason	Yair Darshan	Liang-wei Huang
J. Michael Andrewartha	Piers J. G. Dawe	Yasuhiro Hyakutake
Oleksandr Babenko	Fred Dawson	Scott Irwin
Kwang-Hyun Baek	Ian Dedic	Kazuhiko Ishibe
Amrik Bains	Christopher Diminico	Hideki Isono
Koussalya Balasubramanian	Thuyen Dinh	Tom Issenhuth
Thananya Baldwin	Curtis Donahue	Kenneth Jackson
Denis Beaudoin	Dan Dove	Andrew Jimenez
Christian Beta	Michael Dudek	Chad Jones
Yakov Belopolsky	Nick Duer	Peter Jones
Michael Bennett	David Dwelley	Antony Joseph
Vipul Bhatt	Frank Effenberger	Manabu Kagami
William Bliss	Hesham Elbakoury	Upen Kareti
Brad Booth	David Estes	Keisuke Kawahara
Martin Bouda	John Ewen	Yasuaki Kawatsu
David Brandt	Josef Faller	Michael Kelsen
Ralf-Peter Braun	Shahar Feldman	Yongbum Kim
Theodore Brillhart	German Feyh	Jonathan King
Paul Brooks	Alan Flatman	Scott Kipp
David Brown	Howard Frazier	Michael Klempa
Thomas Brown	Richard Frosch	Curtis Knittle
Phillip Brownlee	Andrew Gardner	Shigeru Kobayashi
Juan-Carlos Calderon	Mike Gardner	Keisuke Kojima
J. Martin Carroll	Ali Ghiasi	Paul Kolesar
Clark Carty	Joel Goergen	Tom Kolze
Mandeep Chadha	Zhigang Gong	Glen Kramer
David Chalupsky	Steven Gorshe	Hans Lackner
Jacky Chang	James Graba	Brett Lane
Xin Chang	Robert Grow	Jeff Lapak
David Chen	Mark Gustlin	Efstathios Larios
Wheling Cheng	Marek Hajduczenia	Mark Laubach
Ahmad Chini	Bernie Hammond	Greg Le Cheminant
Golam Choudhury	Takehiro Hayashi	Arthur Lee
Keng Hua Chuang	David Hess	David Lewis
Peter Cibula	Yasuo Hidaka	Jon Lewis
Christopher R. Cole	Riu Hirai	Lei Li
Keith Conroy	Thomas Hogenmueller	Mike Peng Li

Shaohua Li	Neven Pischl	Alexander Tan
Thomas Lichtenegger	Rainer Poehmerer	Toshiki Tanaka
Ru Jian Lin	William Powell	Mehmet Tazebay
Robert Lingle	Richard Prodan	Brian Teipen
James liu	Rick Rabinovich	Geoffrey Thompson
Zhenyu Liu	Saifur Rahman	Alan Tipper
William Lo	Adee Ran	Pirooz Toyserkani
Miklos Lukacs	Ram Rao	Nathan Tracy
Kent Lusted	Alon Regev	David Tremblay
Jeffery Maki	Duane Remein	Albert Tretter
James Malkemus	Victor Renteria	Stephen Trowbridge
Yonatan Malkiman	Michael Ressler	Wen-Cheng Tseng
Edwin Mallette	Poldi (Pavlick) Rimboim	Yoshihiro Tsukamoto
Arthur Marris	Martin Rossbach	Mike Tu
Chris Mash	Christopher Roth	Alan Ugolini
Kirsten Matheus	Salvatore Rotolo	Ed Ulrichs
Erdem Matoglu	Hisaya Sakamoto	Sterling A. Vaden
Laurence Matola	Vineet Salunke	Stefano Valle
Brett McClellan	Sam Sambasivan	Paul Vanderlaan
Thomas McDermott	Yasuo Sasaki	Robert Wagner
John McDonough	Fred Schindler	Robert Wang
Richard Mei	Stefan Schneelee	Roy Wang
Richard Mellitz	Peter Scruton	Tongtong Wang
Bryan Moffitt	Alexander Seiger	Xiaofeng Wang
Leo Montreuil	Naoshi Serizawa	Xinyuan Wang
Paul Mooney	Megha Shanbhag	Zhong Feng Wang
Andy Moorwood	Masood Shariff	Markus Weber
Thomas Mueller	Stephen J. Shellhammer	Brian Welch
Ron Muir	Bazhong Shen	Yang Wen
Dale Murray	Mizuki Shirao	Matthias Wendt
Henry Muyschondt	Kapil Shrikhande	Oded Wertheim
Edward Nakamoto	Jeff Slavick	Martin White
Gary Nicholl	Scott Sommers	Natalie Wienckowski
Paul Nikolich	Yoshiaki Sone	Ludwig Winkel
Kevin Noll	Xiaolu Song	Peter Wu
Ronald Nordin	Tom Souvignier	Yu Xu
David Ofelt	Bryan Sparrowhawk	Lennart Yseboodt
Ichiro Ogura	Edward Sprague	Ting-Fa Yu
Tom Palkert	Peter Stassar	Liquan Yuan
Hui Pan	Leonard Stencil	Hayato Yuki
Sujan Pandey	Robert Stone	Garold Yurko
Sesha Panguluri	Steve Swanson	Andrew Zambell
Carlos Pardo	Andre Szczepanek	Jin Zhang
Moon Park	William Szeto	Yan Zhuang
Petar Pepeljigoski	Bharat Tailor	George Zimmerman
Gerald Pepper	Akio Tajima	Helge Zinner
Ruben Perez De Aranda Alonso	Takayuki Tajima	Pavel Zivny
Michael Peters	Tomoo Takahara	Gaoling Zou
John Petrilla	Satoshi Takahashi	
Rick Pimpinella	Kiyoto Takahata	

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

Shadi Abughazaleh	Marco Hernandez	Tomas Pajda
Thomas Alexander	Yasuo Hidaka	Thomas Palkert
Richard Alfvén	Werner Hoelzl	Arumugam Paventhan
Carol Ansley	Rita Horner	Michael Peters
Peter Anslow	Tetsushi Ikegami	Adee Ran
Butch Anton	Noriyuki Ikeuchi	R. K. Rannow
Stefan Aust	Sergiu Iordanescu	Alon Regev
Eric Baden	Scott Irwin	Duane Remein
Amrik Bains	Osamu Ishida	Maximilian Riegel
John Ballingall	Atsushi Ito	Robert Robinson
Saman Behtash	Michael Johas Teener	Benjamin Rolfe
Jacob Ben Ary	Peter Jones	Christopher Roth
Michael Bennett	Adri Jovin	Frank Schewe
Gennaro Boggia	Shinkyō Kaku	Dieter Schickelanz
Brad Booth	Piotr Karocki	Takeshi Shimizu
Ralf-Peter Braun	Stuart Kerry	Kapil Shrikhande
Nancy Bravin	Yongbum Kim	Jeff Slavick
Theodore Brillhart	Jonathan King	David Solomon
Matthew Brown	Scott Kipp	Thomas Starai
Jairo Bustos Heredia	Bruce Kraemer	Peter Stassar
William Byrd	Kumaran Krishnasamy	Robert Stone
Steven B. Carlson	Mark Laubach	Eugene Stoudenmire
Juan Carreon	David J. Law	Walter Struppler
Mandeep Chadha	David Lewis	Mitsutoshi Sugawara
Keng Hua Chuang	Jon Lewis	Junqing Sun
Peter Cibula	Arthur H. Light	Patricia Thaler
Charles Cook	William Lo	James Theodoras
John D'Ambrosia	Kent Lusted	David Thompson
Piers J. G. Dawe	Elvis Maculuba	Geoffrey Thompson
Christopher Diminico	Valerie Maguire	Michael Thompson
Daniel Dove	Jeffery Maki	Sterling A. Vaden
Michael Dudek	Arthur Marris	Dmitri Varsanofiev
Constantin Dumitrescu	Brett McClellan	Balasubramanian Vittal
Sourav Dutta	Thomas Mcdermott	George Vlantis
Liu Fangfang	Richard Mellitz	Roy Wang
German Feyh	Bryan Moffitt	Hung-Yu Wei
Yukihiro Fujimoto	Charles Moorwood	Peter Wu
Eric W. Gray	Michael Newman	Jun Xu
Randall Groves	Nick S. A. Nikjoo	Oren Yuen
Marek Hajduczenia	Mark Nowell	Andrew Zambell
Adam Healey	Satoshi Obara	Zhen Zhou

When the IEEE-SA Standards Board approved this standard on 30 June 2016, it had the following membership:

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

Chuck Adams
Masayuki Ariyoshi
Stephen Dukes
Jianbin Fan
Ronald W. Hotchkiss
J. Travis Griffith

Gary Hoffman
Michael Janezic
Joseph L. Koepfinger*
Hung Ling
Kevin Lu
Gary Robinson
Annette D. Reilly

Mehmet Ulema
Yingli Wen
Howard Wolfman
Don Wright
Yu Yuan
Daidi Zhong

*Member Emeritus

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

Introduction

This introduction is not part of IEEE Std 802.3by™-2016, IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters, Physical Layers, and Management Parameters for 25 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 date of IEEE Std 802.3by-2016 publication, 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 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 include 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.

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.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

Contents

1. Introduction.....	28
1.1 Overview.....	28
1.1.3 Architectural perspectives.....	28
1.1.3.2 Compatibility interfaces.....	28
1.3 Normative references.....	28
1.4 Definitions.....	28
1.5 Abbreviations.....	29
4. Media Access Control.....	30
4.4 Specific implementations.....	30
4.4.2 MAC parameters.....	30
30. Management.....	31
30.3 Layer management for DTEs.....	31
30.3.2 PHY device managed object class.....	31
30.3.2.1 PHY device attributes.....	31
30.3.2.1.2 aPhyType.....	31
30.3.2.1.3 aPhyTypeList.....	31
30.3.2.1.5 aSymbolErrorDuringCarrier.....	31
30.5 Layer management for medium attachment units (MAUs).....	31
30.5.1 MAU managed object class.....	31
30.5.1.1 MAU attributes.....	31
30.5.1.1.2 aMAUType.....	31
30.5.1.1.4 aMediaAvailable.....	32
30.5.1.1.15 aFECAbility.....	32
30.5.1.1.16 aFECMode.....	33
30.5.1.1.17 aFECCorrectedBlocks.....	33
30.5.1.1.18 aFECUnconnectableBlocks.....	34
30.6 Management for link Auto-Negotiation.....	34
30.6.1 Auto-Negotiation managed object class.....	34
30.6.1.1 Auto-Negotiation attributes.....	34
30.6.1.1.5 aAutoNegLocalTechnologyAbility.....	34
45. Management Data Input/Output (MDIO) Interface.....	35
45.2 MDIO Interface Registers.....	35
45.2.1 PMA/PMD registers.....	35
45.2.1.1 PMA/PMD control 1 register (Register 1.0).....	36
45.2.1.1.3 Speed selection (1.0.13, 1.0.6, 1.0.5:2).....	36
45.2.1.2 PMA/PMD status 1 register (Register 1.1).....	36
45.2.1.2.3 Fault (1.1.7).....	36
45.2.1.4 PMA/PMD speed ability (Register 1.4).....	37
45.2.1.4.a 25G capable (1.4.11).....	37
45.2.1.6 PMA/PMD control 2 register (Register 1.7).....	37
45.2.1.7 PMA/PMD status 2 register (Register 1.8).....	38
45.2.1.7.4 Transmit fault (1.8.11).....	38
45.2.1.7.5 Receive fault (1.8.10).....	38
45.2.1.8 PMD transmit disable register (Register 1.9).....	38

45.2.1.10 PMA/PMD extended ability register (Register 1.11)	39
45.2.1.10.aa 25G extended abilities (1.11.12)	39
45.2.1.14 EEE capability (Register 1.16)	39
45.2.1.14.4a 25GBASE-R deep sleep (1.16.2)	39
45.2.1.14b 25G PMA/PMD extended ability register (Register 1.19)	40
45.2.1.14b.1 25GBASE-SR ability (1.19.4)	40
45.2.1.14b.2 25GBASE-CR ability (1.19.3)	40
45.2.1.14b.3 25GBASE-CR-S ability (1.19.2)	40
45.2.1.14b.4 25GBASE-KR ability (1.19.1)	40
45.2.1.14b.5 25GBASE-KR-S ability (1.19.0)	41
45.2.1.80 BASE-R PMD control register (Register 1.150)	41
45.2.1.81 BASE-R PMD status register (Register 1.151)	41
45.2.1.82 BASE-R LP coefficient update, lane 0 register (Register 1.152)	41
45.2.1.83 BASE-R LP status report, lane 0 register (Register 1.153)	41
45.2.1.84 BASE-R LD coefficient update, lane 0 register (Register 1.154)	41
45.2.1.85 BASE-R LD status report, lane 0 register (Register 1.155)	41
45.2.1.94 Single-lane PHY 40GBASE-R FEC corrected blocks counter (Register 1.172, 1.173)	42
45.2.1.95 Single-lane PHY 40GBASE-R FEC uncorrected blocks counter (Register 1.174, 1.175)	42
45.2.1.96 CAUI-4 chip-to-module C2M and 25GAUI C2M recommended CTLE register (Register 1.179)	42
45.2.1.96.1 Recommended CTLE peaking (1.179.4:1)	42
45.2.1.97 25GAUI C2C and lane 0 CAUI-4 chip-to-chip C2C transmitter equalization, receive direction, lane 0 register (Register 1.180)	42
45.2.1.97.1 Request flag (1.180.15)	43
45.2.1.97.2 Post-cursor request (1.180.14:12)	43
45.2.1.97.3 Pre-cursor request (1.180.11:10)	43
45.2.1.97.4 Post-cursor remote setting (1.180.9:7)	43
45.2.1.97.5 Pre-cursor remote setting (1.180.6:5)	43
45.2.1.97.6 Post-cursor local setting (1.180.4:2)	43
45.2.1.97.7 Pre-cursor local setting (1.180.1:0)	43
45.2.1.99 25GAUI C2C and lane 0 CAUI-4 chip-to-chip C2C transmitter equalization, transmit direction, lane 0 register (Register 1.184)	44
45.2.1.99.1 Request flag (1.184.15)	44
45.2.1.99.2 Post-cursor request (1.184.14:12)	44
45.2.1.99.3 Pre-cursor request (1.184.11:10)	44
45.2.1.99.4 Post-cursor remote setting (1.184.9:7)	44
45.2.1.99.5 Pre-cursor remote setting (1.184.6:5)	44
45.2.1.99.6 Post-cursor local setting (1.184.4:2)	45
45.2.1.99.7 Pre-cursor local setting (1.184.1:0)	45
45.2.1.101 RS-FEC control register (Register 1.200)	45
45.2.1.101.a 25G RS-FEC enable (1.200.2)	45
45.2.1.101.1 FEC bypass indication enable (1.200.1)	45
45.2.1.101.2 FEC bypass correction enable (1.200.0)	45
45.2.1.102 RS-FEC status register (Register 1.201)	46
45.2.1.102.1 PCS align status (1.201.15)	46
45.2.1.102.2 RS-FEC align status (1.201.14)	46
45.2.1.102.7 RS-FEC high SER (1.201.2)	46
45.2.1.102.8 FEC bypass indication ability (1.201.1)	46
45.2.1.102.9 FEC bypass correction ability (1.201.0)	46
45.2.1.103 RS-FEC corrected codewords counter (Register 1.202, 1.203)	47
45.2.1.104 RS-FEC uncorrected codewords counter (Register 1.204, 1.205)	47
45.2.1.106 RS-FEC symbol error counter lane 0 (Register 1.210, 1.211)	47

45.2.3 PCS registers.....	47
45.2.3.1 PCS control 1 register (Register 3.0).....	47
45.2.3.2 PCS status 1 register (Register 3.1).....	48
45.2.3.2.7 PCS receive link status (3.1.2).....	48
45.2.3.4 PCS speed ability (Register 3.4).....	48
45.2.3.4.5 25G capable (3.4.4).....	48
45.2.3.6 PCS control 2 register (Register 3.7).....	48
45.2.3.6.1 PCS type selection (3.7.2:0).....	49
45.2.3.7 PCS status 2 register (Register 3.8).....	49
45.2.3.7.3a 25GBASE-R capable (3.8.7).....	49
45.2.3.9 EEE control and capability (Register 3.20).....	49
45.2.3.9.2a 25GBASE-R deep sleep (3.20.11).....	50
45.2.3.9.2b 25GBASE-R fast wake (3.20.10).....	50
45.2.3.13 BASE-R and 10GBASE-T PCS status 1 register (Register 3.32).....	50
45.2.3.13.1 BASE-R and 10GBASE-T receive link status (3.32.12).....	50
45.2.3.13.4 BASE-R and 10GBASE-T PCS high BER (3.32.1).....	50
45.2.3.13.5 BASE-R and 10GBASE-T PCS block lock (3.32.0).....	50
45.2.3.14 BASE-R and 10GBASE-T PCS status 2 register (Register 3.33).....	50
45.2.3.14.1 Latched block lock (3.33.15).....	50
45.2.3.14.2 Latched high BER (3.33.14).....	51
45.2.3.14.3 BER (3.33.13:8).....	51
45.2.3.14.4 Errored blocks (3.33.7:0).....	51
45.2.3.15 10/25GBASE-R PCS test pattern seed A (Registers 3.34 through 3.37).....	51
45.2.3.16 10/25GBASE-R PCS test pattern seed B (Registers 3.38 through 3.41).....	51
45.2.3.17 BASE-R PCS test-pattern control register (Register 3.42).....	51
45.2.7 Auto-Negotiation registers.....	52
45.2.7.12 Backplane Ethernet, BASE-R copper status (Register 7.48).....	52
45.2.7.12.1a RS-FEC negotiated (7.48.7).....	53
45.2.7.12.2 Negotiated Port Type (7.48.1, 7.48.2, 7.48.3, 7.48.5, 7.48.6, 7.48.8, 7.48.9, 7.48.10, 7.48.11, 7.48.12, 7.48.13).....	53
45.2.7.13 EEE advertisement (Register 7.60).....	53
45.2.7.13.a 25GBASE-R EEE supported (7.60.14).....	53
45.2.7.14 EEE link partner ability (Register 7.61).....	54
45.5 Protocol implementation conformance statement (PICS) proforma for Clause 45, Management Data Input/Output (MDIO) interface.....	55
45.5.3 PICS proforma tables for the Management Data Input Output (MDIO) interface.....	55
45.5.3.3 PMA/PMD management functions.....	55
69. Introduction to Ethernet operation over electrical backplanes.....	56
69.1 Overview.....	56
69.1.1 Scope.....	56
69.1.2 Relationship of Backplane Ethernet to the ISO OSI reference model.....	56
69.2 Summary of Backplane Ethernet Sublayers.....	58
69.2.1 Reconciliation sublayer and media independent interfaces.....	58
69.2.3 Physical Layer signaling systems.....	58
69.3 Delay constraints.....	59
69.5 Protocol implementation conformance statement (PICS) proforma.....	59
73. Auto-Negotiation for backplane and copper cable assembly.....	60
73.2 Relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model.....	60
73.3 Functional specifications.....	60
73.6 Link codeword encoding.....	61

73.6.4 Technology Ability Field.....	61
73.6.5 FEC capability	62
73.6.5.1 FEC resolution for 25G PHYs	62
73.6.5.2 FEC resolution for 10 Gb/s per lane PHYs.....	62
73.6.5.3 FEC control variables	63
73.7 Receive function requirements	63
73.7.1 DME page reception	63
73.7.6 Priority Resolution function.....	63
73.6 Management register requirements.....	64
73.10 State diagrams and variable definitions	64
73.10.1 State diagram variables.....	64
74. Forward Error Correction (FEC) sublayer for BASE-R PHYs	66
74.1 Overview.....	66
74.3 Relationship to other sublayers.....	67
74.4 Inter-sublayer interfaces	67
74.4.1a Functional block diagram for 25GBASE-R PHYs	68
74.5 FEC service interface.....	69
74.5.1a 25GBASE-R service primitives	69
74.6 Delay constraints.....	70
74.7 FEC principle of operation	70
74.7.4 Functions within FEC sublayer.....	70
74.7.4.1 Reverse gearbox function	70
74.7.4.1.2 Reverse gearbox function for 25GBASE-R, 40GBASE-R, and 100GBASE-R... 70	
74.7.4.3 FEC transmission bit ordering	71
74.7.4.4 FEC (2112,2080) encoder	72
74.7.4.5 FEC decoder	72
74.7.4.5.1 FEC (2112,2080) decoding.....	73
74.7.4.6 FEC receive bit ordering.....	74
74.7.4.8 FEC rapid block synchronization for EEE (optional).....	75
74.8 FEC MDIO function mapping	76
74.8.1 FEC capability	76
74.9 BASE-R PHY test-pattern mode	76
74.11 Protocol implementation conformance statement (PICS) proforma for Clause 74, Forward Error Correction (FEC) sublayer for BASE-R PHYs	77
74.11.3 Major capabilities/options.....	77
74.11.5 FEC Requirements.....	77
78. Energy-Efficient Ethernet (EEE).....	78
78.1 Overview.....	78
78.1.1 LPI Signaling.....	78
78.1.3 Reconciliation sublayer operation	78
78.1.3.3 PHY LPI operation	78
78.1.3.3.1 PHY LPI transmit operation	78
78.1.4 PHY types optionally supporting EEE	79
78.2 LPI mode timing parameters description.....	79
78.5 Communication link access latency.....	79
78.5.2 25 Gb/s, 40 Gb/s, and 100 Gb/s PHY extension using 25GAUI, XLAUI, or CAUI-n	80

90. Ethernet support for time synchronization protocols.....	81
90.1 Introduction.....	81
105. Introduction to 25 Gb/s networks	82
105.1 Overview.....	82
105.1.1 Scope.....	82
105.1.2 Relationship of 25 Gigabit Ethernet to the ISO OSI reference model.....	82
105.1.3 Nomenclature.....	82
105.2 Physical Layer signaling systems	84
105.3 Summary of 25 Gigabit Ethernet sublayers.....	84
105.3.1 Reconciliation Sublayer (RS) and 25 Gigabit Media Independent Interface (25GMII).....	84
105.3.2 Physical Coding Sublayer (PCS).....	85
105.3.3 Forward Error Correction (FEC) sublayer.....	85
105.3.4 Physical Medium Attachment (PMA) sublayer.....	85
105.3.5 Physical Medium Dependent (PMD) sublayer	85
105.3.6 Auto-Negotiation (AN).....	85
105.3.7 Management interface (MDIO/MDC).....	85
105.3.8 Management.....	85
105.4 Service interface specification method and notation	86
105.4.1 Inter-sublayer service interface.....	86
105.4.2 Instances of the Inter-sublayer service interface.....	86
105.4.3 Semantics of inter-sublayer service interface primitives	87
105.4.3.1 IS_UNITDATA.request.....	87
105.4.3.1.1 Semantics of the service primitive.....	87
105.4.3.1.2 When generated	87
105.4.3.1.3 Effect of receipt	87
105.4.3.2 IS_UNITDATA.indication	87
105.4.3.2.1 Semantics of the service primitive.....	87
105.4.3.2.2 When generated.....	88
105.4.3.2.3 Effect of receipt	88
105.4.3.3 IS_SIGNAL.indication	89
105.4.3.3.1 Semantics of the service primitive.....	89
105.4.3.3.2 When generated	90
105.4.3.3.3 Effect of receipt	90
105.4.3.4 IS_TX_MODE.request.....	90
105.4.3.4.1 Semantics of the service primitive.....	90
105.4.3.4.2 When generated	90
105.4.3.4.3 Effect of receipt	90
105.4.3.5 IS_RX_MODE.request.....	90
105.4.3.5.1 Semantics of the service primitive.....	90
105.4.3.5.2 When generated	90
105.4.3.5.3 Effect of receipt	91
105.4.3.6 IS_RX_LPI_ACTIVE.request	91
105.4.3.6.1 Semantics of the service primitive.....	91
105.4.3.6.2 When generated	91
105.4.3.6.3 Effect of receipt	91
105.4.3.7 IS_ENERGY_DETECT.indication	91
105.4.3.7.1 Semantics of the service primitive.....	91
105.4.3.7.2 When generated	91
105.4.3.7.3 Effect of receipt	91
105.4.3.8 IS_RX_TX_MODE.indication	92
105.4.3.8.1 Semantics of the service primitive.....	92

105.4.3.8.2	When generated	92
105.4.3.8.3	Effect of receipt	92
105.5	Delay constraints.....	92
105.6	State diagrams.....	92
105.7	Protocol implementation conformance statement (PICS) proforma.....	93
106.	Reconciliation Sublayer (RS) and Media Independent Interface (25GMII) for 25 Gb/s operation	94
106.1	Overview.....	94
106.1.1	Summary of major concepts	95
106.1.2	Application.....	95
106.1.3	Rate of operation.....	95
106.1.4	Delay constraints.....	95
106.1.5	Allocation of functions	96
106.1.6	25GMII structure	96
106.1.7	Mapping of 25GMII signals to PLS service primitives	96
106.1.7.1	Mapping of PLS_DATA.request	96
106.1.7.2	Mapping of PLS_DATA.indication.....	96
106.1.7.3	Mapping of PLS_CARRIER.indication	96
106.1.7.4	Mapping of PLS_SIGNAL.indication	96
106.1.7.5	Mapping of PLS_DATA_VALID.indication	97
106.2	25GMII data stream.....	97
106.3	25GMII functional specifications	97
106.4	LPI Assertion and Detection.....	97
106.5	Protocol implementation conformance statement (PICS) proforma for Clause 106 Reconciliation Sublayer (RS) and Media Independent Interface (25GMII) for 25 Gb/s operation.....	98
106.5.1	Introduction.....	98
106.5.2	Identification.....	98
106.5.2.1	Implementation identification.....	98
106.5.2.2	Protocol summary	98
106.5.2.3	Major capabilities/options	99
106.5.3	PICS proforma Tables for Reconciliation Sublayer and 25 Gigabit Media Independent Interface	99
106.5.3.1	General	99
106.5.3.2	Mapping of PLS service primitives	99
106.5.3.3	25GMII signal functional specifications.....	99
107.	Physical Coding Sublayer (PCS) for 64B/66B, type 25GBASE-R	100
107.1	Overview.....	100
107.1.1	Scope.....	100
107.1.2	Relationship of 25GBASE-R to other standards	100
107.1.3	Summary of 25GBASE-R sublayers	100
107.1.3.1	Physical Coding Sublayer (PCS)	100
107.1.4	Inter-sublayer interfaces	100
107.1.4.1	PCS service interface (25GMII)	101
107.1.4.2	Physical Medium Attachment (PMA) service interface	101
107.2	Functions within the PCS	101
107.2.1	Notation conventions	102
107.2.2	Transmission order	102
107.2.3	Test-pattern generator.....	103
107.3	LPI	103
107.3	Delay constraints.....	104

107.4 Support for Auto-Negotiation.....	104
107.5 Protocol implementation conformance statement (PICS) proforma for Clause 107, Physical Coding Sublayer (PCS) for 64B/66B, type 25GBASE-R.....	105
107.5.1 Introduction.....	105
107.5.2 Identification.....	105
107.5.2.1 Implementation identification.....	105
107.5.2.2 Protocol summary.....	105
107.5.3 Major capabilities/options.....	106
107.5.4 25G PCS.....	106
107.5.4.1 Clause 49 functionality.....	106
107.5.4.2 Test-pattern generator.....	106
107.5.4.3 LPI.....	106
107.5.4.4 Delay Constraints.....	107
108. Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs.....	108
108.1 Overview.....	108
108.1.1 Scope.....	108
108.1.2 Position of RS-FEC in the 25GBASE-R PHY sublayers.....	108
108.2 FEC service interface.....	108
108.3 PMA compatibility.....	109
108.4 Delay constraints.....	110
108.5 Functions within the 25GBASE-R RS-FEC sublayer.....	110
108.5.1 Functional block diagram.....	110
108.5.2 Transmit function.....	110
108.5.2.1 Block synchronization.....	110
108.5.2.2 Rate compensation for codeword markers in the transmit direction.....	110
108.5.2.3 64B/66B to 256B/257B transcoder.....	110
108.5.2.4 Codeword marker insertion.....	112
108.5.2.5 Reed-Solomon encoder.....	112
108.5.2.6 Codeword serialization.....	112
108.5.2.7 RS-FEC encoding for rapid codeword lock (EEE deep sleep).....	112
108.5.3 Receive function.....	114
108.5.3.1 Codeword marker lock.....	114
108.5.3.2 Reed-Solomon decoder.....	114
108.5.3.3 Codeword monitor.....	115
108.5.3.4 Codeword marker removal.....	116
108.5.3.5 256B/257B to 64B/66B transcoder.....	116
108.5.3.6 Rate compensation for codeword markers in the receive direction.....	116
108.5.3.7 Rapid codeword lock for EEE deep sleep.....	116
108.5.3.8 Receive bit ordering.....	117
108.5.4 Detailed functions and state diagrams.....	117
108.5.4.1 State diagram conventions.....	117
108.5.4.2 State variables.....	117
108.5.4.3 Functions.....	119
108.5.4.4 Counters.....	119
108.5.4.5 Timers.....	120
108.5.4.6 State diagrams.....	121
108.6 25GBASE-R RS-FEC MDIO function mapping.....	122
108.6.1 FEC_bypass_correction_enable.....	123
108.6.2 FEC_bypass_indication_enable.....	123
108.6.3 25G RS-FEC Enable.....	124
108.6.4 FEC_bypass_correction_ability.....	124
108.6.5 FEC_bypass_indication_ability.....	124

108.6.6 FEC_high_ser	124
108.6.7 FEC_corrected_cw_counter.....	124
108.6.8 FEC_uncorrected_cw_counter.....	124
108.6.9 FEC_symbol_error_counter_0.....	125
108.6.10 align_status	125
108.7 Protocol implementation conformance statement (PICS) proforma for Clause 108, Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs	126
108.7.1 Introduction.....	126
108.7.2 Identification.....	126
108.7.2.1 Implementation identification.....	126
108.7.2.2 Protocol summary	126
108.7.3 Major capabilities/options.....	127
108.7.4 PICS proforma tables for Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs.....	127
108.7.4.1 Transmit function.....	127
108.7.4.2 Receive function	128
108.7.4.3 State diagrams.....	129
108.7.4.4 MDIO function mapping	129
109. Physical Medium Attachment (PMA) sublayer, type 25GBASE-R.....	130
109.1 Overview.....	130
109.1.1 Scope.....	130
109.1.2 Position of the PMA in the 25GBASE-R sublayers.....	130
109.1.3 Summary of functions.....	131
109.1.4 PMA sublayer positioning	132
109.2 PMA service interface	133
109.3 Service interface below PMA	134
109.4 Functions within the PMA	135
109.4.1 Signal drivers	135
109.4.2 PMA local loopback mode.....	135
109.4.3 PMA remote loopback mode.....	135
109.4.4 PMA test patterns.....	135
109.4.4.1 Transmit PRBS31 test-pattern generation	136
109.4.4.2 Receive PRBS31 test-pattern generation.....	136
109.4.4.3 Transmit PRBS31 test-pattern checking.....	136
109.4.4.4 Receive PRBS31 test-pattern checking	137
109.4.4.5 Transmit PRBS9 test-pattern generation	137
109.4.4.6 Receive PRBS9 test-pattern generation.....	137
109.4.4.7 Transmit square wave test-pattern generation	137
109.4.5 Energy Efficient Ethernet for 25GAUI.....	138
109.5 Delay constraints.....	138
109.6 PMA MDIO function mapping.....	138
109.7 Protocol implementation conformance statement (PICS) proforma for Clause 109, Physical Medium Attachment (PMA) sublayer, type 25GBASE-R.....	141
109.7.1 Introduction.....	141
109.7.2 Identification	141
109.7.2.1 Implementation identification.....	141
109.7.2.2 Protocol summary	141
109.7.3 PICS proforma tables for the 25GBASE-R PMA Sublayer	142
109.7.4 Major capabilities/options.....	142
109.7.4.1 PMA functions.....	142
109.7.4.2 PMA characteristics.....	143

110. Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-CR and 25GBASE-CR-S	144
110.1 Overview	144
110.2 PMD service interface	145
110.3 PCS requirements for Auto-Negotiation (AN) service interface	146
110.4 Delay constraints	146
110.5 PMD MDIO function mapping	147
110.6 FEC modes	147
110.7 PMD functional specifications	148
110.7.1 Link block diagram	148
110.7.2 PMD transmit function	149
110.7.3 PMD receive function	150
110.7.4 Global PMD signal detect function	150
110.7.5 Global PMD transmit disable function	150
110.7.6 Loopback mode	151
110.7.7 PMD fault function	151
110.7.8 PMD transmit fault function	151
110.7.9 PMD receive fault function	151
110.7.10 PMD control function	151
110.8 Electrical characteristics	151
110.8.1 Signal levels	151
110.8.2 Signal paths	152
110.8.3 Transmitter characteristics	152
110.8.4 Receiver characteristics	152
110.8.4.1 Receiver input amplitude tolerance	152
110.8.4.2 Receiver interference tolerance test	152
110.8.4.2.1 Test setup	154
110.8.4.2.2 Test channel	154
110.8.4.2.3 Test channel calibration	154
110.8.4.2.4 Pattern generator and noise injection	156
110.8.4.2.5 Test procedure	156
110.8.4.3 Receiver jitter tolerance	156
110.8.4.4 Signaling rate range	157
110.9 Channel characteristics	157
110.10 Cable assembly characteristics	157
110.10.1 Characteristic impedance and reference impedance	158
110.10.2 Cable assembly insertion loss	158
110.10.3 Cable assembly differential return loss	158
110.10.4 Differential to common-mode return loss	159
110.10.5 Differential to common-mode conversion loss	159
110.10.6 Common-mode to common-mode return loss	159
110.10.7 Cable assembly Channel Operating Margin	159
110.10.7.1 Channel signal and crosstalk path calculations	160
110.10.7.1.1 Channel signal path	161
110.10.7.1.2 Channel crosstalk paths	161
110.10.7.2 Signal and crosstalk paths used in calculation of COM	162
110.10.7.2.1 SFP28 to SFP28	162
110.10.7.2.2 QSFP28 to SFP28	162
110.10.7.2.3 SFP28 to QSFP28	162
110.10.7.2.4 QSFP28 to QSFP28	163
110.11 MDI specification	163
110.11.1 Single-lane MDI connectors	163
110.12 Environmental specifications	164

110.13 Protocol implementation conformance statement (PICS) proforma for Clause 110, Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-CR and 25GBASE-CR-S	165
110.13.1 Introduction	165
110.13.2 Identification	165
110.13.2.1 Implementation identification	165
110.13.2.2 Protocol summary	165
110.13.3 Major capabilities/options	166
110.13.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-CR and 25GBASE-CR-S	167
110.13.4.1 PMD functional specifications	167
110.13.4.2 Management functions	168
110.13.4.3 Transmitter specifications	168
110.13.4.4 Receiver specifications	170
110.13.4.5 Cable assembly specifications	171
110.13.4.6 MDI connector specifications	171
110.13.4.7 Environmental specifications	172
111. Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S	173
111.1 Overview	173
111.2 PMD service interface	174
111.3 PCS requirements for Auto-Negotiation (AN) service interface	175
111.4 Delay constraints	175
111.5 PMD MDIO function mapping	175
111.6 FEC modes	176
111.7 PMD functional specifications	177
111.7.1 Link block diagram	177
111.7.2 PMD transmit function	177
111.7.3 PMD receive function	178
111.7.4 Global PMD signal detect function	178
111.7.5 Global PMD transmit disable function	178
111.7.6 Loopback mode	178
111.7.7 PMD fault function	179
111.7.8 PMD transmit fault function	179
111.7.9 PMD receive fault function	179
111.7.10 PMD control function	179
111.8 Electrical characteristics	179
111.8.1 MDI	179
111.8.2 Transmitter characteristics	180
111.8.3 Receiver characteristics	180
111.8.3.1 Receiver interference tolerance	180
111.8.3.2 Receiver jitter tolerance	182
111.9 Channel characteristics	183
111.9.1 25GBASE-KR channel	183
111.9.2 25GBASE-KR-S channel	183
111.10 Environmental specifications	184
111.11 Protocol implementation conformance statement (PICS) proforma for Clause 111, Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S	185
111.11.1 Introduction	185
111.11.2 Identification	185
111.11.2.1 Implementation identification	185

111.11.2.2 Protocol summary	185
111.11.3 Major capabilities/options	186
111.11.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S	187
111.11.4.1 Functional specifications	187
111.11.4.2 Transmitter characteristics	188
111.11.4.3 Receiver characteristics	190
111.11.4.4 Channel characteristics	190
111.11.4.5 Environmental specifications	191
112. Physical Medium Dependent (PMD) sublayer and medium, type 25GBASE-SR	192
112.1 Overview	192
112.1.1 Bit error ratio	193
112.2 Physical Medium Dependent (PMD) service interface	193
112.3 Delay constraints	194
112.4 PMD MDIO function mapping	194
112.5 PMD functional specifications	195
112.5.1 PMD block diagram	195
112.5.2 PMD transmit function	196
112.5.3 PMD receive function	196
112.5.4 PMD global signal detect function	196
112.5.5 PMD reset function	197
112.5.6 PMD global transmit disable function (optional)	197
112.5.7 PMD fault function (optional)	197
112.5.8 PMD transmit fault function (optional)	197
112.5.9 PMD receive fault function (optional)	197
112.6 PMD to MDI optical specifications for 25GBASE-SR	197
112.6.1 25GBASE-SR transmitter optical specifications	198
112.6.2 25GBASE-SR receive optical specifications	198
112.6.3 25GBASE-SR illustrative link power budget	198
112.7 Definition of optical parameters and measurement methods	198
112.7.1 Test patterns for optical parameters	198
112.7.2 Center wavelength and spectral width	199
112.7.3 Average optical power	199
112.7.4 Optical Modulation Amplitude (OMA)	199
112.7.5 Transmitter and dispersion eye closure (TDEC)	199
112.7.6 Extinction ratio	199
112.7.7 Transmitter optical waveform (transmit eye)	199
112.7.8 Stressed receiver sensitivity	199
112.8 Safety, installation, environment, and labeling	199
112.8.1 General safety	199
112.8.2 Laser safety	200
112.8.3 Installation	200
112.8.4 Environment	200
112.8.5 Electromagnetic emission	200
112.8.6 Temperature, humidity, and handling	200
112.8.7 PMD labeling requirements	200
112.9 Fiber optic cabling model	201
112.10 Characteristics of the fiber optic cabling (channel)	201
112.10.1 Optical fiber cable	201
112.10.2 Optical fiber connection	201
112.10.2.1 Connection insertion loss	202
112.10.2.2 Maximum discrete reflectance	202

112.10.3 Medium Dependent Interface (MDI).....	202
112.11 Protocol implementation conformance statement (PICS) proforma for Clause 112, Physical Medium Dependent (PMD) sublayer and medium, type 25GBASE-SR.....	203
112.11.1 Introduction.....	203
112.11.2 Identification.....	203
112.11.2.1 Implementation identification.....	203
112.11.2.2 Protocol summary.....	203
112.11.3 Major capabilities/options.....	204
112.11.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 25GBASE-SR.....	204
112.11.4.1 PMD functional specifications.....	204
112.11.4.2 Management functions.....	205
112.11.4.3 PMD to MDI optical specifications for 25GBASE-SR.....	205
112.11.4.4 Optical measurement methods.....	206
112.11.4.5 Environmental specifications.....	206
112.11.4.6 Characteristics of the fiber optic cabling and MDI.....	206
Annex 4A (normative) Simplified full duplex media access control.....	207
4A.4 Specific implementations.....	207
4A.4.2 MAC parameters.....	207
Annex 31B (normative) MAC Control PAUSE operation.....	208
31B.3 Detailed specification of PAUSE operation.....	208
31B.3.7 Timing considerations for PAUSE operation.....	208
31B.4 Protocol implementation conformance statement (PICS) proforma for MAC Control PAUSE operation.....	209
31B.4.3 Major capabilities/options.....	209
31B.4.6 PAUSE command MAC timing considerations.....	209
Annex 93A (normative) Specification methods for electrical channels.....	210
93A.1 Channel Operating Margin.....	210
93A.1.4 Filters.....	211
93A.2 Test channel calibration using COM.....	211
Annex 109A (normative) Chip-to-chip 25 Gigabit Attachment Unit Interface (25GAUI C2C).....	212
109A.1 Overview.....	212
109A.2 25GAUI C2C compliance point definition.....	213
109A.3 25GAUI C2C electrical characteristics.....	213
109A.3.1 25GAUI C2C transmitter characteristics.....	213
109A.3.2 25GAUI C2C receiver characteristics.....	213
109A.3.3 Optional EEE operation.....	213
109A.4 25GAUI C2C channel characteristics.....	213
109A.5 Protocol implementation conformance statement (PICS) proforma for Annex 109A, Chip-to-chip 25 Gigabit Attachment Unit Interface (25GAUI C2C).....	214
109A.5.1 Introduction.....	214
109A.5.2 Identification.....	214
109A.5.2.1 Implementation identification.....	214
109A.5.2.2 Protocol summary.....	214
109A.5.3 Major capabilities/options.....	215

109A.5.4 PICS proforma tables for chip-to-chip 25 Gigabit Attachment Unit Interface (25GAUI C2C)	215
109A.5.4.1 Transmitter	215
109A.5.4.2 Receiver	216
109A.5.4.3 Channel	216
Annex 109B (normative) Chip-to-module 25 Gigabit Attachment Unit Interface (25GAUI C2M)	217
109B.1 Overview	217
109B.1.1 Bit error ratio	218
109B.2 25GAUI C2M compliance point definitions	218
109B.3 25GAUI C2M electrical characteristics	218
109B.3.1 25GAUI C2M host output characteristics	218
109B.3.2 25GAUI C2M module output characteristics	218
109B.3.2.1 25GAUI C2M module output eye opening	219
109B.3.2.1.1 Eye opening using measurement method A	219
109B.3.2.1.2 Eye opening using measurement method B	219
109B.3.3 25GAUI C2M host input characteristics	219
109B.3.4 25GAUI C2M module input characteristics	219
109B.3.4.1 Module stressed input test using measurement method A	220
109B.3.4.2 Module stressed input test using measurement method B	220
109B.4 25GAUI C2M measurement methodology	220
109B.4.1 Eye width, eye height, and eye closure measurement method B	220
109B.5 Protocol implementation conformance statement (PICS) proforma for Annex 109B, Chip-to-module 25 Gigabit Attachment Unit Interface (25GAUI C2M)	222
109B.5.1 Introduction	222
109B.5.2 Identification	222
109B.5.2.1 Implementation identification	222
109B.5.2.2 Protocol summary	222
109B.5.3 Major capabilities/options	223
109B.5.4 PICS proforma tables for chip-to-module 25 Gigabit Attachment Unit Interface (25GAUI C2M)	223
109B.5.4.1 Host output	223
109B.5.4.2 Module output	224
109B.5.4.3 Host input	224
109B.5.4.4 Module input	225
Annex 109C (informative) 25GBASE-R PMA sublayer partitioning examples	226
Annex 110A (informative) TP0 and TP5 test point parameters and channel characteristics for 25GBASE-CR and 25GBASE-CR-S	230
110A.1 Overview	230
110A.2 Transmitter characteristics at TP0	230
110A.3 Receiver characteristics at TP5	230
110A.4 Transmitter and receiver differential printed circuit board trace loss	230
110A.5 Channel insertion loss	230
110A.6 Channel return loss	232
110A.7 Channel Operating Margin (COM)	232
Annex 110B (normative) Test fixtures for 25GBASE-CR, 25GBASE-CR-S, and 25GAUI C2M	233
110B.1 Test fixtures	233
110B.1.1 SFP28 TP2 or TP3 test fixture	233

110B.1.2 SFP28 Cable assembly test fixture.....	233
110B.1.3 SFP28 Mated test fixtures	234
110B.1.3.1 Mated test fixtures differential insertion loss.....	234
110B.1.3.2 Mated test fixtures differential return loss	234
110B.1.3.3 Mated test fixtures common-mode conversion insertion loss.....	234
110B.1.3.4 Mated test fixtures common-mode return loss.....	234
110B.1.3.5 Mated test fixtures common-mode to differential mode return loss	234
110B.1.3.7 Mated test fixtures integrated near-end crosstalk noise	234
110B.2 Protocol implementation conformance statement (PICS) proforma for Annex 110B, Test fixtures for 25GBASE-CR, 25GBASE-CR-S, and 25GAUI C2M.....	236
110B.2.1 Introduction	236
110B.2.2 Identification	236
110B.2.2.1 Implementation identification	236
110B.2.2.2 Protocol summary	236
110B.2.3 Major capabilities/options.....	237
110B.2.4 PICS proforma tables for test fixtures for 25GBASE-CR, 25GBASE-CR-S, and 25GAUI C2M	237
Annex 110C (informative) Host and cable assembly form factors for 25GBASE-CR and 25GBASE-CR-S PHYs.....	238
110C.1 Overview	238
110C.2 Host form factors.....	239
110C.2.1 SFP28 host form factor	239
110C.2.2 QSFP28 host form factor	239
110C.3 Cable assembly form factors.....	240
110C.3.1 SFP28 to SFP28 cable assembly form factor.....	240
110C.3.2 QSFP28 to QSFP28 cable assembly form factor	241
110C.3.3 QSFP28 to 4×SFP28 cable assembly form factor.....	242

IEEE Standard for Ethernet

Amendment 2: Media Access Control Parameters, Physical Layers, and Management Parameters for 25 Gb/s Operation

IMPORTANT NOTICE: IEEE Standards documents are not intended to ensure safety, health, or environmental protection, or ensure against interference with or from other devices or networks. Implementers 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.

This IEEE document is made available for use subject to important notices and legal disclaimers. These notices and disclaimers appear in all publications containing this document and may be found under the heading “Important Notice” or “Important Notices and Disclaimers Concerning IEEE Documents.” They can also be obtained on request from IEEE or viewed at <http://standards.ieee.org/IPR/disclaimers.html>.

(This amendment is based on IEEE Std 802.3™-2015 as amended by IEEE 802.3bw™-2015.)

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 ~~strikethrough~~ (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.¹

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

¹Notes in text, tables, and figures are given for information only, and do not contain requirements needed to implement the standard.

1. Introduction

1.1 Overview

1.1.3 Architectural perspectives

1.1.3.2 Compatibility interfaces

Insert the following compatibility interface definitions after item h) in the lettered list:

- i) *25 Gigabit Media Independent Interface (25GMII)*. The 25GMII is designed to connect a 25 Gb/s capable MAC to a 25 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 25 Gb/s speeds. The 25GMII is a logical interconnection intended for use as an intra-chip interface. No mechanical connector is specified for use with the 25GMII. The 25GMII is optional.
- j) *25 Gigabit Attachment Unit Interface (25GAUI)*. The 25GAUI is a physical instantiation of the PMA service interface to extend the connection between 25 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 25 Gb/s speeds. The 25GAUI is intended for use as a chip-to-chip or a chip-to-module interface. No mechanical connector is specified for use with the 25GAUI. The 25GAUI is optional.

1.3 Normative references

Insert the following references in alphanumeric order:

SFF-8402, Rev 1.1, September 13, 2014, Specification for SFP+ 1X 28 Gb/s Pluggable Transceiver Solution (SFP28).²

SFF-8432, Rev 5.1, August 8, 2012, Specification for SFP+ Module and Cage.

SFF-8665, Rev 1.9, June 29, 2015, Specification for QSFP+ 28 Gb/s 4X Pluggable Transceiver Solution (QSFP28).

Remove the footnote number from the SFF-8436 reference as shown:

SFF-8436, Rev 4.1, August 24, 2011, Specification for QSFP+ 10 Gbs 4X Pluggable Transceiver.¹⁹

1.4 Definitions

Insert the following new definitions into the list after 1.4.64 10/10G-EPON:

1.4.64a 25GBASE: A family of Physical Layer entities for 25 Gb/s operation. (See IEEE Std 802.3, Clause 105.)

1.4.64b 25GBASE-CR: IEEE 802.3 Physical Layer specification for 25 Gb/s using 25GBASE-R encoding over one lane of twinaxial copper cable. (See IEEE Std 802.3, Clause 110.)

²SFF specifications are available at <ftp://ftp.seagate.com/sff>.

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation

1.4.64c 25GBASE-CR-S: IEEE 802.3 Physical Layer specification equivalent to 25GBASE-CR without support for the RS-FEC sublayer specified in Clause 108. (See IEEE Std 802.3, Clause 110.)

1.4.64d 25GBASE-KR: IEEE 802.3 Physical Layer specification for 25 Gb/s using 25GBASE-R encoding over one lane of an electrical backplane. (See IEEE Std 802.3, Clause 111.)

1.4.64e 25GBASE-KR-S: IEEE 802.3 Physical Layer specification equivalent to 25GBASE-KR without support for the RS-FEC sublayer specified in Clause 108. (See IEEE Std 802.3, Clause 111.)

1.4.64f 25GBASE-R: An IEEE 802.3 physical coding sublayer for one-lane 25 Gb/s operation. (See IEEE Std 802.3, Clause 107.)

1.4.64g 25GBASE-SR: IEEE 802.3 Physical Layer specification for 25 Gb/s using 25GBASE-R encoding over multimode fiber. (See IEEE Std 802.3, Clause 112.)

Insert the following new definitions into the list after 1.4.77 10 Gigabit Sixteen-Bit Interface (XSBI):

1.4.77a 25 Gigabit Attachment Unit Interface (25GAUI): A physical instantiation of the Physical Medium Attachment (PMA) service interface to extend the connection between 25 Gb/s capable PMAs over one lane, used for chip-to-chip or chip-to-module interconnections. (See IEEE Std 802.3, Annex 109A and Annex 109B.)

1.4.77b 25 Gigabit Media Independent Interface (25GMII): The interface between the Reconciliation Sublayer (RS) and the Physical Coding Sublayer (PCS) for 25 Gb/s operation. (See IEEE Std 802.3, Clause 106.)

Change 1.4.107 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, or Clause 82, or Clause 107. (See IEEE Std 802.3, Clause 49, and Clause 82, or Clause 107.)

Change 1.4.134 as follows:

1.4.134 channel: In 10BROAD36, a band of frequencies dedicated to a certain service transmitted on the broadband medium. Otherwise, a defined path along which data in the form of an electrical or optical signal passes. (For 10BROAD36, see IEEE Std 802.3, Clause 11.)

1.5 Abbreviations

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

25GAUI	25 Gigabit Attachment Unit Interface
25GMII	25 Gigabit Media Independent Interface
C2C	chip-to-chip
C2M	chip-to-module

4. Media Access Control

4.4 Specific implementations

4.4.2 MAC parameters

Change Table 4–2 to add 25 Gb/s as follows:

Table 4–2—MAC parameters

Parameters	MAC data rate			
	Up to and including 100 Mb/s	1 Gb/s	10 Gb/s	25 Gb/s, 40 Gb/s, and 100 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	104 bits	not applicable

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

Change Note 4 as follows:

NOTE 4—For 10 Gb/s and 25 Gb/s operation, 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 40 BT (bit times), as measured at the XGMII or 25GMII receive signals at the DTE. This interpacket gap shrinkage may be caused by variable network delays and clock tolerances.

30. Management

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 the following new entry in APPROPRIATE SYNTAX before the entry for 40GBASE-R in 30.3.2.1.2:

APPROPRIATE SYNTAX:

...
25GBASE-R Clause 107 25 Gb/s 64B/66B

30.3.2.1.3 aPhyTypeList

Insert the following new entry in APPROPRIATE SYNTAX before the entry for 40GBASE-R in 30.3.2.1.3:

APPROPRIATE SYNTAX:

...
25GBASE-R Clause 107 25 Gb/s 64B/66B

30.3.2.1.5 aSymbolErrorDuringCarrier

Change the fourth paragraph in BEHAVIOUR DEFINED AS in 30.3.2.1.5 as follows:

BEHAVIOUR DEFINED AS:

For operation at 10 Gb/s, 25 Gb/s, 40 Gb/s, and 100 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 XGMII (see Table 46-4 and Table 81-3), the XLGMII, or the CGMII (see 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 the following new entries in APPROPRIATE SYNTAX before the entry for 40GBASE-R in 30.5.1.1.2:

APPROPRIATE SYNTAX:

...	
25GBASE-R	PCS as specified in Clause 107 with PMA as specified in Clause 109 over undefined PMD
25GBASE-CR	25GBASE-R PCS/PMA over shielded balanced copper cable PMD as specified in Clause 110
25GBASE-CR-S	25GBASE-R PCS/PMA over shielded balanced copper cable PMD as specified in Clause 110 without support for RS-FEC
25GBASE-KR	25GBASE-R PCS/PMA over an electrical backplane PMD as specified in Clause 111
25GBASE-KR-S	25GBASE-R PCS/PMA over an electrical backplane PMD as specified in Clause 111 without support for RS-FEC
25GBASE-SR	25GBASE-R PCS/PMA over multimode fiber PMD as specified in Clause 112
...	

Change BEHAVIOUR DEFINED AS in 30.5.1.1.2 as follows:

BEHAVIOUR DEFINED AS:

Returns a value that identifies the internal MAU type. If an AUI is to be identified to access an external MAU, the type “AUI” is returned. A SET operation to one of the possible enumerations indicated by aMAUTypeList will force the MAU into the new operating mode. If a Clause 22 MII or Clause 35 GMII is present, then this will map to the mode force bits specified in 22.2.4.1. If a Clause 45 MDIO Interface is present, then this will map to the PCS type selection bit(s) in the 10G WIS Control 2 register specified in 45.2.2.6.6, the PCS Control 2 register specified in 45.2.3.6.1, the PMA/PMD type selection bits in the PMA/PMD Control 2 register specified in 45.2.1.6, the PMA/PMD control 1 register specified in 45.2.1.1, the 25G RS-FEC Enable bit in the RS-FEC control register specified in 45.2.1.101.a, and the PCS control 1 register 45.2.3.1. If Clause 28, Clause 37, or Clause 73 Auto-Negotiation is operational, then this will change the advertised ability to the single enumeration specified in the SET operation, and cause an immediate link renegotiation. A change in the MAU type will also be reflected in aPHYType.

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

30.5.1.1.4 aMediaAvailable

Change the first sentence of the eighth paragraph in BEHAVIOUR DEFINED AS in 30.5.1.1.4 as follows:

BEHAVIOUR DEFINED AS:

For 10 Gb/s and 25 Gb/s the enumerations map to value of the link_fault variable within the Link Fault Signaling state diagram (Figure 46–11) as follows: the value OK maps 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”.

30.5.1.1.15 aFECAbility

Change the first sentence of BEHAVIOUR DEFINED AS in 30.5.1.1.15 as follows:

BEHAVIOUR DEFINED AS:

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

30.5.1.1.16 aFECMode

Change APPROPRIATE SYNTAX in 30.5.1.1.16 as follows:

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that meets the requirement of the description below	
unknown	initializing, true state not yet known
disabled	FEC disabled
<u>BASE-R enabled</u>	<u>BASE-R FEC enabled</u>
<u>RS-FEC enabled</u>	<u>RS-FEC enabled</u>
enabled	FEC enabled

Change BEHAVIOUR DEFINED AS in 30.5.1.1.16 as follows:

BEHAVIOUR DEFINED AS:

A read-write value that indicates the mode of operation of the ~~optional~~ FEC sublayer for forward error correction (see ~~65.2, and~~ Clause 74, Clause 91, and Clause 108).

A GET operation returns the current mode of operation of the PHY. A SET operation changes the mode of operation of the PHY to the indicated value. The enumerations “BASE-R enabled” and “RS-FEC enabled” are only used for 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, and 25GBASE-KR-S PHYs where operation in the no-FEC mode maps to the enumerations “disabled”, operation in the BASE-R FEC mode maps to the enumerations “BASE-R enabled”, and operation in the RS-FEC mode maps to the enumerations “RS-FEC enabled” (see 110.6 and 111.6).

When Clause 73 Auto-Negotiation is enabled for a 25GBASE-R PHY, a SET operation is not allowed and a GET operation maps to the variables FEC_enable in Clause 74 and FEC_enable in Clause 108. When Clause 73 Auto-Negotiation is enabled for a non-25GBASE-R PHY supporting Clause 74 FEC a SET operation is not allowed and a GET operation maps to the variable FEC_enable in Clause 74.

If a Clause 45 MDIO Interface is present, then this attribute maps to the FEC control register (see 45.2.8.3) for 1000BASE-PX, to the BASE-R FEC control register (see 45.2.1.93) and the 25G RS-FEC Enable bit in the RS-FEC control register (see 45.2.1.101) for 25GBASE-R, or FEC enable bit in BASE-R FEC control register (see 45.2.1.93).;

30.5.1.1.17 aFECCorrectedBlocks

Change APPROPRIATE SYNTAX in 30.5.1.1.17 as follows:

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.

Change the first sentence of BEHAVIOUR DEFINED AS in 30.5.1.1.17 as follows:

BEHAVIOUR DEFINED AS:

For 1000BASE-PX, 10/25/40/100GBASE-R, 100GBASE-P, 10GBASE-PR, or 10/1GBASE-PRX PHYs, an array of corrected FEC block counters.

30.5.1.1.18 aFECUncorrectableBlocks

Change APPROPRIATE SYNTAX in 30.5.1.1.18 as follows:

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.

Change the first sentence of BEHAVIOUR DEFINED AS in 30.5.1.1.18 as follows:

BEHAVIOUR DEFINED AS:

For 1000BASE-PX, 10/25/40/100GBASE-R, 100GBASE-P, 10GBASE-PR, or 10/1GBASE-PRX PHYs, an array of uncorrectable FEC block counters.

30.6 Management for link Auto-Negotiation

30.6.1 Auto-Negotiation managed object class

30.6.1.1 Auto-Negotiation attributes

30.6.1.1.5 aAutoNegLocalTechnologyAbility

Insert 25GR-S and 25GR after 10GBASE-KRFD, and RS-FEC25G Req and BASE-RFEC25G Req after FEC Requested in APPROPRIATE SYNTAX, in 30.6.1.1.5 as follows:

APPROPRIATE SYNTAX:

25GR-S	25GBASE-CR-S as specified in Clause 110 or 25GBASE-KR-S as specified in Clause 111
25GR	25GBASE-CR as specified in Clause 110 or 25GBASE-KR as specified in Clause 111
RS-FEC25G Req	25G RS-FEC requested as specified in Clause 73 (see 73.6.5) and Clause 108
BASE-RFEC25G Req	25G BASE-R FEC requested as specified in Clause 73 (see 73.6.5) and Clause 74

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 follows:

For devices operating at ~~25~~ 40-Gb/s or higher speeds, the PMA may be instantiated as multiple sublayers (see 83.1.4 and 109.1.4 for how MMD addresses are allocated to multiple PMA sublayers).

Change the identified reserved row in Table 45-3 (as modified by IEEE Std 802.3bw-2015) and insert a new row immediately above the changed row as follows (unchanged rows not shown):

Table 45-3—PMA/PMD registers

Register address	Register name	Subclause
<u>1.19</u>	<u>25G PMA/PMD extended ability</u>	45.2.1.14b
1.19 <u>20</u> through 1.29	Reserved	

Change the rows in Table 45-3 for registers 1.172 through 1.175 and 1.179 through 1.187 (unchanged rows not shown):

Table 45-3—PMA/PMD registers

Register address	Register name	Subclause
1.172 through 1.173	Single-lane PHY 40GBASE-R FEC corrected blocks counter	45.2.1.94
1.174 through 1.175	Single-lane PHY 40GBASE-R FEC uncorrected blocks counter	45.2.1.95
1.179	CAUI-4 chip-to-module C2M and 25GAUI C2M recommended CTLE	45.2.1.96
<u>1.180</u>	<u>25GAUI C2C and lane 0 CAUI-4 C2C transmitter equalization, receive direction</u>	<u>45.2.1.97</u>
1.180 <u>181</u> through 1.183	CAUI-4 chip-to-chip transmitter equalization, receive direction, lane 0 <u>1</u> through lane 3	45.2.1.97 <u>45.2.1.98</u>
<u>1.184</u>	<u>25GAUI C2C and lane 0 CAUI-4 C2C transmitter equalization, transmit direction</u>	<u>45.2.1.99</u>
1.184 <u>185</u> through 1.187	CAUI-4 chip-to-chip transmitter equalization, transmit direction, lane 0 <u>1</u> through lane 3	45.2.1.99 <u>45.2.1.100</u>

45.2.1.1 PMA/PMD control 1 register (Register 1.0)

Change the indicated row of Table 45–4 for 25 Gb/s speed selection 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 x x x = Reserved x 1 x x = Reserved 0 1 1 x = Reserved 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

^aR/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 first sentence of the last paragraph of 45.2.1.1.3 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.

45.2.1.2 PMA/PMD status 1 register (Register 1.1)

45.2.1.2.3 Fault (1.1.7)

Change 45.2.1.2.3 as follows:

Fault is a global PMA/PMD variable. When read as a one, bit 1.1.7 indicates that either (or both) the PMA or the PMD has detected a fault condition on either the transmit or receive paths. When read as a zero, bit 1.1.7 indicates that neither the PMA nor the PMD has detected a fault condition. For 10/25/40/100 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. For 10PASS-TS or 2BASE-TL operations, when read as a one, a fault has been detected and more detailed information is conveyed in 45.2.1.19, 45.2.1.42, 45.2.1.43, and 45.2.1.58.

45.2.1.4 PMA/PMD speed ability (Register 1.4)

Change the row for 1.4.15:10 in Table 45-6 as follows (unchanged rows not shown):

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

Bit(s)	Name	Description	R/W ^a
1.4.15:10 1.4.15:12	Reserved for future speeds	Value always 0	RO
1.4.11	25G capable	1 = PMA/PMD is capable of operating at 25 Gb/s 0 = PMA/PMD is not capable of operating at 25 Gb/s	RO
1.4.10	Reserved	Value always 0	RO

^aRO = Read only

Insert new subclause 45.2.1.4.a before 45.2.1.4.1 as follows:

45.2.1.4.a 25G capable (1.4.11)

When read as a one, bit 1.4.11 indicates that the PMA/PMD is able to operate at a data rate of 25 Gb/s. When read as a zero, bit 1.4.11 indicates that the PMA/PMD is not able to operate at a data rate of 25 Gb/s.

45.2.1.6 PMA/PMD control 2 register (Register 1.7)

Change the indicated row of Table 45-7 (as modified by IEEE Std 802.3bw-2015) for 25G PMA/PMD selection as follows (unchanged rows not shown):

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

Bit(s)	Name	Description	R/W ^a
1.7.5:0	PMA/PMD type selection	5 4 3 2 1 0 1110xx = reserved for future use 111011 = reserved 111010 = 25GBASE-SR PMA/PMD 111001 = 25GBASE-KR or 25GBASE-KR-S PMA/PMD 111000 = 25GBASE-CR or 25GBASE-CR-S PMA/PMD	R/W

^aR/W = Read/Write, RO = Read only

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 between 10GBASE-KX4 and 40GBASE-KR4 in Table 45–9 (unchanged rows not shown):

Table 45–9—Transmit fault description location

PMA/PMD	Description location
25GBASE-KR, 25GBASE-KR-S	111.7.8
25GBASE-CR, 25GBASE-CR-S	110.7.8
25GBASE-SR	112.5.8

45.2.1.7.5 Receive fault (1.8.10)

Insert the following rows between 10GBASE-KX4 and 40GBASE-KR4 in Table 45–10 (unchanged rows not shown):

Table 45–10—Receive fault description location

PMA/PMD	Description location
25GBASE-KR, 25GBASE-KR-S	111.7.9
25GBASE-CR, 25GBASE-CR-S	110.7.9
25GBASE-SR	112.5.9

45.2.1.8 PMD transmit disable register (Register 1.9)

Insert the following rows between 10GBASE-KX4 and 40GBASE-KR4 in Table 45–12 (unchanged rows not shown):

Table 45–12—Transmit disable description location

PMA/PMD	Description location
25GBASE-KR and 25GBASE-KR-S	111.7.5
25GBASE-CR and 25GBASE-CR-S	110.7.5
25GBASE-SR	112.5.6

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

45.2.1.10 PMA/PMD extended ability register (Register 1.11)

Change the row for 1.11.15:12 of Table 45–14 (as modified by IEEE Std 802.3bw-2015) as follows (unchanged rows not shown):

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

Bit(s)	Name	Description	R/W ^a
1.11.15:12 <u>1.11.15:13</u>	Reserved	Value always 0	RO
<u>1.11.12</u>	<u>25G extended abilities</u>	<u>1 = PMA/PMD has 25G extended abilities listed in register 1.19</u> <u>0 = PMA/PMD does not have 25G extended abilities</u>	<u>RO</u>

^aRO = Read only

Insert a new subclause 45.2.1.10.aa before 45.2.1.10.a (as inserted by IEEE Std 802.3bw-2015) as follows:

45.2.1.10.aa 25G extended abilities (1.11.12)

When read as a one, bit 1.11.12 indicates that the PMA/PMD has 25G extended abilities listed in register 1.19. When read as a zero, bit 1.11.12 indicates that the PMA/PMD does not have 25G extended abilities.

45.2.1.14 EEE capability (Register 1.16)

Change the row for 1.16.7:2 of Table 45–17 as follows (unchanged rows not shown):

Table 45–17—EEE capability register bit definitions

Bit(s)	Name	Description	R/W ^a
1.16.7:2 <u>1.16.7:3</u>	Reserved	Value always 0	RO
<u>1.16.2</u>	<u>25GBASE-R deep sleep</u>	<u>1 = EEE deep sleep is supported for 25GBASE-R</u> <u>0 = EEE deep sleep is not supported for 25GBASE-R</u>	<u>RO</u>

^aRO = Read only

Insert new subclause 45.2.1.14.4a before 45.2.1.14.5 as follows:

45.2.1.14.4a 25GBASE-R deep sleep (1.16.2)

If the device supports EEE deep sleep operation for 25GBASE-R, bit 1.16.2 shall be set to a one; otherwise this bit is set to a zero.

Insert 45.2.1.14b and 45.2.1.14b.1 through 45.2.1.14b.5 after 45.2.1.14a as inserted by IEEE Std 802.3bw-2015 as follows:

45.2.1.14b 25G PMA/PMD extended ability register (Register 1.19)

The assignment of bits in the 25G PMA/PMD extended ability register is shown in Table 45–17b.

Table 45–17b—25G PMA/PMD extended ability register bit definitions

Bit(s)	Name	Description	R/W ^a
1.19.15:5	Reserved	Value always 0	RO
1.19.4	25GBASE-SR ability	1 = PMA/PMD is able to perform 25GBASE-SR 0 = PMA/PMD is not able to perform 25GBASE-SR	RO
1.19.3	25GBASE-CR ability	1 = PMA/PMD is able to perform 25GBASE-CR 0 = PMA/PMD is not able to perform 25GBASE-CR	RO
1.19.2	25GBASE-CR-S ability	1 = PMA/PMD is able to perform 25GBASE-CR-S 0 = PMA/PMD is not able to perform 25GBASE-CR-S	RO
1.19.1	25GBASE-KR ability	1 = PMA/PMD is able to perform 25GBASE-KR 0 = PMA/PMD is not able to perform 25GBASE-KR	RO
1.19.0	25GBASE-KR-S ability	1 = PMA/PMD is able to perform 25GBASE-KR-S 0 = PMA/PMD is not able to perform 25GBASE-KR-S	RO

^aRO = Read only

45.2.1.14b.1 25GBASE-SR ability (1.19.4)

When read as a one, bit 1.19.4 indicates that the PMA/PMD is able to operate as a 25GBASE-SR PMA/PMD type. When read as a zero, bit 1.19.4 indicates that the PMA/PMD is not able to operate as a 25GBASE-SR PMA/PMD type.

45.2.1.14b.2 25GBASE-CR ability (1.19.3)

When read as a one, bit 1.19.3 indicates that the PMA/PMD is able to operate as a 25GBASE-CR PMA/PMD type. When read as a zero, bit 1.19.3 indicates that the PMA/PMD is not able to operate as a 25GBASE-CR PMA/PMD type.

45.2.1.14b.3 25GBASE-CR-S ability (1.19.2)

When read as a one, bit 1.19.2 indicates that the PMA/PMD is able to operate as a 25GBASE-CR-S PMA/PMD type. When read as a zero, bit 1.19.2 indicates that the PMA/PMD is not able to operate as a 25GBASE-CR-S PMA/PMD type.

45.2.1.14b.4 25GBASE-KR ability (1.19.1)

When read as a one, bit 1.19.1 indicates that the PMA/PMD is able to operate as a 25GBASE-KR PMA/PMD type. When read as a zero, bit 1.19.1 indicates that the PMA/PMD is not able to operate as a 25GBASE-KR PMA/PMD type.

45.2.1.14b.5 25GBASE-KR-S ability (1.19.0)

When read as a one, bit 1.19.0 indicates that the PMA/PMD is able to operate as a 25GBASE-KR-S PMA/PMD type. When read as a zero, bit 1.19.0 indicates that the PMA/PMD is not able to operate as a 25GBASE-KR-S PMA/PMD type.

45.2.1.80 BASE-R PMD control register (Register 1.150)

Change the first sentence of 45.2.1.80 as follows:

The BASE-R PMD control register is used for 10GBASE-KR and other PHY types using the PMDs described in [Clause 72](#), [Clause 84](#), [Clause 85](#), [Clause 92](#), [Clause 93](#), ~~or Clause 94~~, [Clause 110](#), or [Clause 111](#).

45.2.1.81 BASE-R PMD status register (Register 1.151)

Change the first sentence of 45.2.1.81 as follows:

The BASE-R PMD status register is used for 10GBASE-KR and other PHY types using the PMDs described in [Clause 72](#), [Clause 84](#), [Clause 85](#), [Clause 92](#), [Clause 93](#), ~~or Clause 94~~, [Clause 110](#), or [Clause 111](#).

45.2.1.82 BASE-R LP coefficient update, lane 0 register (Register 1.152)

Change the first sentence of 45.2.1.82 as follows:

The BASE-R LP coefficient update, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in [Clause 72](#), [Clause 84](#), [Clause 85](#), [Clause 92](#), [Clause 93](#), ~~or Clause 94~~, [Clause 110](#), or [Clause 111](#).

45.2.1.83 BASE-R LP status report, lane 0 register (Register 1.153)

Change the first sentence of 45.2.1.83 as follows:

The BASE-R LP status report, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in [Clause 72](#), [Clause 84](#), [Clause 85](#), [Clause 92](#), [Clause 93](#), ~~or Clause 94~~, [Clause 110](#), or [Clause 111](#).

45.2.1.84 BASE-R LD coefficient update, lane 0 register (Register 1.154)

Change the first sentence of 45.2.1.84 as follows:

The BASE-R LD coefficient update, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in [Clause 72](#), [Clause 84](#), [Clause 85](#), [Clause 92](#), [Clause 93](#), ~~or Clause 94~~, [Clause 110](#), or [Clause 111](#).

45.2.1.85 BASE-R LD status report, lane 0 register (Register 1.155)

Change the first sentence of 45.2.1.85 as follows:

The BASE-R LD status report, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in [Clause 72](#), [Clause 84](#), [Clause 85](#), [Clause 92](#), [Clause 93](#), ~~or Clause 94~~, [Clause 110](#), or [Clause 111](#).

Change the subclause title, the first sentence of 45.2.1.94, and the title of Table 45-74 as follows:

45.2.1.94 Single-lane PHY 40GBASE-R FEC corrected blocks counter (Register 1.172, 1.173)

The assignment of bits in the single-lane PHY 40GBASE-R FEC corrected blocks counter register is shown in Table 45-74.

Table 45-74—Single-lane PHY 40GBASE-R FEC corrected blocks counter register bit definitions

Change the subclause title, the first sentence of 45.2.1.95, and the title of Table 45-75 as follows:

45.2.1.95 Single-lane PHY 40GBASE-R FEC uncorrected blocks counter (Register 1.174, 1.175)

The assignment of bits in the single-lane PHY 40GBASE-R FEC uncorrected blocks counter register is shown in Table 45-75.

Table 45-75—Single-lane PHY 40GBASE-R FEC uncorrected blocks counter register bit definitions

Change the subclause title, the first sentence of 45.2.1.96, and the title of Table 45-76 as follows:

45.2.1.96 CAUI-4 chip-to-module C2M and 25GAUI C2M recommended CTLE register (Register 1.179)

The assignment of bits in the CAUI-4 chip-to-module C2M and 25GAUI C2M recommended CTLE register is shown in Table 45-76.

Table 45-76—CAUI-4 chip-to-module C2M and 25GAUI C2M recommended CTLE register bit definitions

45.2.1.96.1 Recommended CTLE peaking (1.179.4:1)

Change the first sentence of 45.2.1.96.1 as follows:

The value of these bits sets the CTLE peaking value recommended by a host that implements the optional CAUI-4 chip-to-module C2M and 25GAUI C2M interface defined in [Annex 83E](#) and [Annex 109B](#), respectively (see [83E.3.1.6](#)).

Change the subclause title, the first sentence of 45.2.1.97, the title of Table 45-77 and the subclauses of 45.2.1.97 as follows:

45.2.1.97 25GAUI C2C and lane 0 CAUI-4 chip-to-chip C2C transmitter equalization, receive direction, lane 0 register (Register 1.180)

The assignment of bits in the 25GAUI C2C and lane 0 CAUI-4 chip-to-chip C2C transmitter equalization, receive direction, lane 0 register is shown in Table 45-77.

Table 45-77—25GAUI C2C and lane 0 CAUI-4 chip-to-chip C2C transmitter equalization, receive direction, lane 0 register bit definitions

45.2.1.97.1 Request flag (1.180.15)

The value of this bit indicates the value of the variable *Request_flag* in the 25GAUI or lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). This indicates whether the 25GAUI or CAUI-4 chip-to-chip C2C device is issuing a request to change the remote transmitter equalization in the 25GAUI or lane 0 CAUI-4 chip-to-chip lane 0 C2C transmitter in the receive direction. If a 25GAUI or lane 0 CAUI-4 receiver in the receive direction is not present in the package, then the value returned for this bit should be zero.

45.2.1.97.2 Post-cursor request (1.180.14:12)

The value of these bits indicates the value of the variable *Requested_eq_c1* in the 25GAUI or lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). 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 25GAUI or lane 0 CAUI-4 chip-to-chip lane 0 C2C transmitter in the receive direction.

45.2.1.97.3 Pre-cursor request (1.180.11:10)

The value of these bits indicates the value of the variable *Requested_eq_cm1* in the 25GAUI or lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). 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 25GAUI or lane 0 CAUI-4 chip-to-chip lane 0 C2C transmitter in the receive direction.

45.2.1.97.4 Post-cursor remote setting (1.180.9:7)

The value of these bits sets the variable *Remote_eq_c1* for the 25GAUI or lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). This is used by a 25GAUI or CAUI-4 receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the post-cursor coefficient $c(1)$ being used in the 25GAUI or lane 0 of the CAUI-4 transmitter in the receive direction (see 83D.3.1.1). It may be used to generate values for the request flag and the request bits. If a 25GAUI or lane 0 CAUI-4 receiver in the receive direction is not present in the package, then these bits have no effect.

45.2.1.97.5 Pre-cursor remote setting (1.180.6:5)

The value of these bits sets the variable *Remote_eq_cm1* for the 25GAUI or lane 0 CAUI-4 receiver in the receive direction (see 83D.3.3.2). This is used by a 25GAUI or CAUI-4 receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the pre-cursor coefficient $c(-1)$ being used in the 25GAUI or lane 0 of the CAUI-4 transmitter in the receive direction (see 83D.3.1.1). It may be used to generate values for the request flag and the request bits. If a 25GAUI or lane 0 CAUI-4 receiver in the receive direction is not present in the package, then these bits have no effect.

45.2.1.97.6 Post-cursor local setting (1.180.4:2)

The value of these bits sets the variable *Local_eq_c1* for the 25GAUI or lane 0 CAUI-4 transmitter in the receive direction (see 83D.3.1.1 and Table 83D-3), which controls the weight of the transmitter equalization post-cursor coefficient $c(1)$. If a 25GAUI or lane 0 CAUI-4 transmitter in the receive direction is not present in the package, then these bits have no effect.

45.2.1.97.7 Pre-cursor local setting (1.180.1:0)

The value of these bits sets the variable *Local_eq_cm1* for the 25GAUI or lane 0 CAUI-4 transmitter in the receive direction (see 83D.3.1.1 and Table 83D-2), which controls the weight of the transmitter equalization precursor coefficient $c(-1)$. If a 25GAUI or lane 0 CAUI-4 transmitter in the receive direction is not present in the package, then these bits have no effect.

Change the subclause title, the first sentence of 45.2.1.99, the title of Table 45-78, and 45.2.1.99.1 through 45.2.1.99.7 as follows:

45.2.1.99 25GAUI C2C and lane 0 CAUI-4 chip-to-chip C2C transmitter equalization, transmit direction, lane 0 register (Register 1.184)

The assignment of bits in the 25GAUI C2C and lane 0 CAUI-4 chip-to-chip C2C transmitter equalization, transmit direction, lane 0 register is shown in Table 45-78.

Table 45-78—25GAUI C2C and lane 0 CAUI-4 chip-to-chip C2C transmitter equalization, transmit direction, lane 0 register bit definitions

45.2.1.99.1 Request flag (1.184.15)

The value of this bit indicates the value of the variable *Request_flag* in the 25GAUI or lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). This indicates whether the 25GAUI or CAUI-4 chip-to-chip C2C device is issuing a request to change the remote transmitter equalization in the 25GAUI or lane 0 CAUI-4 chip-to-chip lane 0 C2C transmitter in the transmit direction. If a 25GAUI or lane 0 CAUI-4 receiver in the transmit direction is not present in the package, then the value returned for this bit should be zero.

45.2.1.99.2 Post-cursor request (1.184.14:12)

The value of these bits indicates the value of the variable *Requested_eq_c1* in the 25GAUI or lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). 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 25GAUI or lane 0 CAUI-4 chip-to-chip lane 0 C2C transmitter in the transmit direction.

45.2.1.99.3 Pre-cursor request (1.184.11:10)

The value of these bits indicates the value of the variable *Requested_eq_cm1* in the 25GAUI or lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). 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 25GAUI or lane 0 CAUI-4 chip-to-chip lane 0 C2C transmitter in the transmit direction.

45.2.1.99.4 Post-cursor remote setting (1.184.9:7)

The value of these bits sets the variable *Remote_eq_c1* for the 25GAUI or lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). This is used by a 25GAUI or CAUI-4 receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the post-cursor coefficient $c(1)$ being used in the 25GAUI or lane 0 of the CAUI-4 or transmitter in the transmit direction (see 83D.3.1.1). It may be used to generate values for the request flag and the request bits. If a 25GAUI or lane 0 CAUI-4 receiver in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.99.5 Pre-cursor remote setting (1.184.6:5)

The value of these bits sets the variable *Remote_eq_cm1* for the 25GAUI or lane 0 CAUI-4 receiver in the transmit direction (see 83D.3.3.2). This is used by a 25GAUI or CAUI-4 receiver that implements the optional transmitter equalization feedback as an indication of the ratio of the pre-cursor coefficient $c(-1)$ being used in the 25GAUI or lane 0 of the CAUI-4 transmitter in the transmit direction (see 83D.3.1.1). It may be used to generate values for the request flag and the request bits. If a 25GAUI or lane 0 CAUI-4 receiver in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.99.6 Post-cursor local setting (1.184.4:2)

The value of these bits sets the variable *Local_eq_c1* for the 25GAUI or lane 0 CAUI-4 transmitter in the transmit direction (see 83D.3.1.1 and Table 83D-3), which controls the weight of the transmitter equalization post-cursor coefficient *c*(1). If a 25GAUI or lane 0 CAUI-4 transmitter in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.99.7 Pre-cursor local setting (1.184.1:0)

The value of these bits sets the variable *Local_eq_cm1* for the 25GAUI or lane 0 CAUI-4 transmitter in the transmit direction (see 83D.3.1.1 and Table 83D-2), which controls the weight of the transmitter equalization precursor coefficient *c*(-1). If a 25GAUI or lane 0 CAUI-4 transmitter in the transmit direction is not present in the package, then these bits have no effect.

45.2.1.101 RS-FEC control register (Register 1.200)

Change row for 1.200.15:2 of Table 45-79 as follows (unchanged rows not shown):

Table 45-79—RS-FEC control register bit definitions

Bit(s)	Name	Description	R/W ^a
1.200.15:2 1.200.15:3	Reserved	Value always 0	RO
<u>1.200.2</u>	<u>25G RS-FEC Enable</u>	<u>1 = The 25GBASE-R Reed-Solomon FEC is enabled</u> <u>0 = The 25GBASE-R Reed-Solomon FEC is disabled</u>	<u>R/W</u>

^aR/W = Read/Write, RO = Read only

Insert new subclause 45.2.1.101.a before 45.2.1.101.1 as follows:

45.2.1.101.a 25G RS-FEC enable (1.200.2)

Bit 1.200.2 enables the 25GBASE-R Reed-Solomon FEC described in Clause 108. When set to a one, this bit enables the 25GBASE-R Reed-Solomon FEC. When set to a zero, bit 1.200.2 disables the 25GBASE-R Reed-Solomon FEC (see 108.6.3).

45.2.1.101.1 FEC bypass indication enable (1.200.1)

Change the last sentence as follows:

Writes to this bit 1.200.1 are ignored and reads return a zero if the RS-FEC does not have the ability to bypass indicating decoding errors to the PCS layer (see 91.5.3.3 and 108.5.3.2).

45.2.1.101.2 FEC bypass correction enable (1.200.0)

Change as follows:

When this bit 1.200.0 is set to one the Reed-Solomon decoder performs error detection without error correction (see 91.5.3.3 and 108.5.3.2). When this bit is set to zero, the decoder also performs error correction. Writes to this bit are ignored and reads return a zero if the RS-FEC does not have the ability to bypass correction.

45.2.1.102 RS-FEC status register (Register 1.201)**45.2.1.102.1 PCS align status (1.201.15)**

Change 45.2.1.102.1 as follows:

Bit 1.201.15 indicates the PCS alignment status of the RS-FEC. For the RS-FEC described in Clause 91, PCS alignment is defined as block lock, alignment marker lock, and deskew of all 20 transmit PCS lanes. For the RS-FEC described in Clause 108, PCS alignment is defined as block lock of the transmit PCS signal. When read as a zero, bit 1.201.15 indicates that the RS-FEC has not obtained PCS alignment. When read as a one, bit 1.201.15 indicates that the RS-FEC has obtained PCS alignment.

~~When read as a one, bit 1.201.15 indicates that the RS-FEC described in Clause 91 has locked and aligned all transmit PCS lanes. When read as a zero, bit 1.201.15 indicates that the RS-FEC has not locked and aligned all transmit PCS lanes. A device that implements the RS-FEC status register but does not implement a separated RS-FEC shall return a one for bit 1.201.15.~~

45.2.1.102.2 RS-FEC align status (1.201.14)

Change 45.2.1.102.2 as follows:

Bit 1.201.14 indicates the PMA alignment status of the RS-FEC. For the RS-FEC described in Clause 91, PMA alignment is defined as alignment marker lock and deskew of all four lanes on the PMA service interface. For the RS-FEC described in Clause 108, PMA alignment is defined as codeword marker lock on the PMA service interface. When read as a zero, bit 1.201.14 indicates that the RS-FEC has not obtained PMA alignment. When read as a one, bit 1.201.14 indicates that the RS-FEC has obtained PMA alignment.

~~When read as a one, bit 1.201.14 indicates that the RS-FEC described in Clause 91 has locked and aligned all receive RS-FEC lanes. When read as a zero, bit 1.201.14 indicates that the RS-FEC has not locked and aligned all receive RS-FEC lanes.~~

45.2.1.102.7 RS-FEC high SER (1.201.2)

Change the first sentence of 45.2.1.102.7 as follows:

When FEC_bypass_indication_enable is set to one, ~~this~~ bit 1.201.2 is set to one if the number of RS-FEC symbol errors in a window of 8192 codewords exceeds the threshold (see 91.5.3.3 and 108.5.3.2) and is set to zero otherwise.

45.2.1.102.8 FEC bypass indication ability (1.201.1)

Change the first sentence of 45.2.1.102.8 as follows:

The Reed-Solomon decoder may have the option to perform error detection without error indication (see 91.5.3.3 and 108.5.3.2) to reduce the delay contributed by the RS-FEC sublayer.

45.2.1.102.9 FEC bypass correction ability (1.201.0)

Change the first sentence of 45.2.1.102.9 as follows:

The Reed-Solomon decoder may have the option to perform error detection without error correction (see 91.5.3.3 and 108.5.3.2) to reduce the delay contributed by the RS-FEC sublayer.

45.2.1.103 RS-FEC corrected codewords counter (Register 1.202, 1.203)

Change the second sentence of 45.2.1.103 as follows:

The assignment of bits in the RS-FEC corrected codewords counter register is shown in Table 45–81. See 91.6.8 and 108.6.7 for a definition of this register.

45.2.1.104 RS-FEC uncorrected codewords counter (Register 1.204, 1.205)

Change the second sentence of 45.2.1.104 as follows:

The assignment of bits in the RS-FEC uncorrected codewords counter register is shown in Table 45–82. See 91.6.9 and 108.6.8 for a definition of this register.

45.2.1.106 RS-FEC symbol error counter lane 0 (Register 1.210, 1.211)

Change the third sentence of 45.2.1.106 as follows:

The assignment of bits in the RS-FEC symbol error counter lane 0 register is shown in Table 45–84. Symbol errors detected in FEC lane 0 are counted and shown in register 1.210.15:0 and 1.211.15:0. See 91.6.11 and 108.6.9 for a definition of this register.

45.2.3 PCS registers

Change the indicated rows of Table 45–119 as follows (unchanged rows not shown):

Table 45–119—PCS registers

Register address	Register name	Subclause
3.34 through 3.37	10/25GBASE-R PCS test pattern seed A	45.2.3.15
3.38 through 3.41	10/25GBASE-R PCS test pattern seed B	45.2.3.16

45.2.3.1 PCS control 1 register (Register 3.0)

Change the indicated row of Table 45–120 for 25 Gb/s speed selection 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 x x x = Reserved 0 1 1 x = Reserved 0 1 0 1 = Reserved 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

^aRO = 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 follows:

When a 10/25/40/100GBASE-R, 10GBASE-W, or 10GBASE-T mode of operation is selected for the PCS using the PCS type selection field (3.7.2:0), this bit is a latching low version of bit 3.32.12.

45.2.3.4 PCS speed ability (Register 3.4)

Change row for 3.4.15:4 of Table 45–122 as follows (unchanged rows not shown):

Table 45–122—PCS speed ability register bit definitions

Bit(s)	Name	Description	R/W ^a
3.4.15:4 3.4.15:5	Reserved for future speeds	Value always 0	RO
3.4.4	25G capable	1 = PCS is capable of operating at 25 Gb/s 0 = PCS is not capable of operating at 25 Gb/s	RO

^aRO = Read only

Insert new subclause 45.2.3.4.5 after 45.2.3.4.4 as follows:

45.2.3.4.5 25G capable (3.4.4)

When read as a one, bit 3.4.4 indicates that the PCS is able to operate at a data rate of 25 Gb/s. When read as a zero, bit 3.4.4 indicates that the PCS is not able to operate at a data rate of 25 Gb/s.

45.2.3.6 PCS control 2 register (Register 3.7)

Change the indicated row of Table 45–123 for 25GBASE-R PCS selection as follows (unchanged rows not shown):

Table 45–123—PCS control 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.7.2:0	PCS type selection	2 1 0 1 1 1 = Select 25GBASE-R PCS type 1 1 0 = reserved 1 0 1 = Select 100GBASE-R PCS type 1 0 0 = Select 40GBASE-R PCS types 0 1 1 = Select 10GBASE-T PCS type 0 1 0 = Select 10GBASE-W PCS type 0 0 1 = Select 10GBASE-X PCS type 0 0 0 = Select 10GBASE-R PCS type	R/W

^aRO = Read only, R/W = Read/Write

45.2.3.6.1 PCS type selection (3.7.2:0)

Change the second sentence of 45.2.3.6.1 as follows:

The PCS type abilities of the PCS are advertised in bits ~~3.8.5:0~~ 3.8.7:0.

45.2.3.7 PCS status 2 register (Register 3.8)

Change the row for 3.8.9:6 of Table 45–124 as follows (unchanged rows not shown):

Table 45–124—PCS status 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.8.9:6 <u>3.8.9:8</u>	Reserved	Value always 0	RO
<u>3.8.7</u>	<u>25GBASE-R capable</u>	<u>1 = PCS is able to support 25GBASE-R PCS type</u> <u>0 = PCS is not able to support 25GBASE-R PCS type</u>	<u>RO</u>
<u>3.8.6</u>	<u>Reserved</u>	<u>Value always 0</u>	<u>RO</u>

^aRO = Read only, LH = Latching high

Insert new subclause 45.2.3.7.3a before 45.2.3.7.4 as follows:

45.2.3.7.3a 25GBASE-R capable (3.8.7)

When read as a one, bit 3.8.7 indicates that the PCS is able to support the 25GBASE-R PCS type. When read as a zero, bit 3.8.7 indicates that the PCS is not able to support the 25GBASE-R PCS type.

45.2.3.9 EEE control and capability (Register 3.20)

Change the row for 3.20.11:10 of Table 45–125 as follows (unchanged rows not shown):

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

Bit(s)	Name	Description	R/W ^a
3.20.11:10	Reserved	Value always 0	RO
<u>3.20.11</u>	<u>25GBASE-R deep sleep</u>	<u>1 = EEE deep sleep is supported for 25GBASE-R</u> <u>0 = EEE deep sleep is not supported for 25GBASE-R</u>	<u>RO</u>
<u>3.20.10</u>	<u>25GBASE-R fast wake</u>	<u>1 = EEE fast wake is supported for 25GBASE-R</u> <u>0 = EEE fast wake is not supported for 25GBASE-R</u>	<u>RO</u>

^aR/W = Read/Write, RO = Read only

Insert new subclauses 45.2.3.9.2a and 45.2.3.9.2b before 45.2.3.9.3 as follows:

45.2.3.9.2a 25GBASE-R deep sleep (3.20.11)

If the device supports EEE deep sleep operation for 25GBASE-R bit 3.20.11 shall be set to a one; otherwise this bit is set to a zero.

45.2.3.9.2b 25GBASE-R fast wake (3.20.10)

If the device supports EEE fast wake operation for 25GBASE-R bit 3.20.10 shall be set to a one; otherwise this bit is set to a zero.

45.2.3.13 BASE-R and 10GBASE-T PCS status 1 register (Register 3.32)

45.2.3.13.1 BASE-R and 10GBASE-T receive link status (3.32.12)

Change last sentence of 45.2.3.13.1 as follows:

This bit is a reflection of the PCS_status variable defined in 49.2.14.1 for 10/25GBASE-R, in 55.3.6.1 for 10GBASE-T and in 82.3.1 for 40/100GBASE-R.

45.2.3.13.4 BASE-R and 10GBASE-T PCS high BER (3.32.1)

Change third sentence of 45.2.3.13.4 as follows:

This bit is a direct reflection of the state of the hi_ber variable in the 64B/66B state diagram and is defined in 49.2.13.2.2 for 10/25GBASE-R and in 82.2.19.2.2 for 40/100GBASE-R.

45.2.3.13.5 BASE-R and 10GBASE-T PCS block lock (3.32.0)

Change third sentence of 45.2.3.13.5 as follows:

This bit is a direct reflection of the state of the block_lock variable in the 64B/66B state diagram and is defined in 49.2.13.2.2 for 10/25GBASE-R and in 82.2.19.2.2 for 40/100GBASE-R.

45.2.3.14 BASE-R and 10GBASE-T PCS status 2 register (Register 3.33)

45.2.3.14.1 Latched block lock (3.33.15)

Change 45.2.3.14.1 as follows:

When read as a one, bit 3.33.15 indicates that the 10/25/40/100GBASE-R or the 10GBASE-T PCS has achieved block lock. When read as a zero, bit 3.33.15 indicates that the 10/25/40/100GBASE-R or the 10GBASE-T PCS has lost block lock.

The latched block lock bit shall be implemented with latching low behavior.

This bit is a latching low version of the 10/25/40/100GBASE-R and 10GBASE-T PCS block lock status bit (3.32.0).

45.2.3.14.2 Latched high BER (3.33.14)

Change 45.2.3.14.2 as follows:

When read as a one, bit 3.33.14 indicates that the 10/25/40/100GBASE-R or the 10GBASE-T PCS has detected a high BER. When read as a zero, bit 3.33.14 indicates that the 10/25/40/100GBASE-R or the 10GBASE-T PCS has not detected a high BER.

The latched high BER bit shall be implemented with latching high behavior.

This bit is a latching high version of the 10/25/40/100GBASE-R and 10GBASE-T PCS high BER status bit (3.32.1).

45.2.3.14.3 BER (3.33.13:8)

Change first sentence of 45.2.3.14.3 as follows:

The BER counter is a six bit count as defined by the ber_count variable in 49.2.14.2 and 82.2.19.2.4 for 10/25/40/100GBASE-R and defined by the lfer_count variable in 55.3.6.2 for 10GBASE-T.

45.2.3.14.4 Errored blocks (3.33.7:0)

Change first sentence of 45.2.3.14.4 as follows:

The errored blocks counter is an eight bit count defined by the errored_block_count counter specified in 49.2.14.2 for 10/25GBASE-R, in 82.3.1 for 40/100GBASE-R and defined by the errored_block_count variable in 55.3.6.2 for 10GBASE-T.

Change the subclause title, the first and second sentences of 45.2.3.15, and the title of Table 45-130 as follows:

45.2.3.15 10/25GBASE-R PCS test pattern seed A (Registers 3.34 through 3.37)

The assignment of bits in the 10/25GBASE-R PCS test pattern seed A registers is shown in Table 45–130. This register is only required when the 10GBASE-R or 25GBASE-R capability is supported.

Table 45–130—10/25GBASE-R PCS test pattern seed A 0-3 register bit definitions

Change the subclause title, the first and second sentences of 45.2.3.16, and the title of Table 45-131 as follows:

45.2.3.16 10/25GBASE-R PCS test pattern seed B (Registers 3.38 through 3.41)

The assignment of bits in the 10/25GBASE-R PCS test pattern seed B registers is shown in Table 45–131. This register is only required when the 10GBASE-R or 25GBASE-R capability is supported.

Table 45–131—10/25GBASE-R PCS test pattern seed B 0-3 register bit definitions

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

Change second sentence from the end of the first paragraph as follows:

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

Change the indicated rows of Table 45-132 as follows (unchanged rows not shown):

Table 45-132—BASE-R PCS test-pattern control register bit definitions

Bit(s)	Name	Description	R/W ^a
3.42.6	Single Lane PHY 40GBASE-R PRBS9 transmit test-pattern enable	1 = Enable PRBS9 test-pattern mode on the transmit path 0 = Disable PRBS9 test-pattern mode on the transmit path	R/W
3.42.5	Single Lane PHY 40GBASE-R PRBS31 receive test-pattern enable	1 = Enable PRBS31 test-pattern mode on the receive path 0 = Disable PRBS31 test-pattern mode on the receive path	R/W
3.42.4	Single Lane PHY 40GBASE-R PRBS31 transmit test-pattern enable	1 = Enable PRBS31 test-pattern mode on the transmit path 0 = Disable PRBS31 test-pattern mode on the transmit path	R/W

^aRO = Read only, R/W = Read/Write

Change titles of subclauses 45.2.3.17.2 to 45.2.3.17.4 as follows:

45.2.3.17.2 Single Lane PHY 40GBASE-R PRBS9 transmit test-pattern enable (3.42.6)

45.2.3.17.3 Single Lane PHY 40GBASE-R PRBS31 receive test-pattern enable (3.42.5)

45.2.3.17.4 Single Lane PHY 40GBASE-R PRBS31 transmit test-pattern enable (3.42.4)

45.2.7 Auto-Negotiation registers

45.2.7.12 Backplane Ethernet, BASE-R copper status (Register 7.48)

Change the row for 7.48.15:12 and 7.48.7 of Table 45-209 as follows (unchanged rows not shown):

Table 45-209—Backplane Ethernet, BASE-R copper status register (Register 7.48) bit definitions

Bit(s)	Name	Description	RO ^a
7.48.15:12 7.48.15:14	Reserved	Value always 0	RO
7.48.13	25GBASE-KR or 25GBASE-CR	1 = PMA/PMD is negotiated to perform 25GBASE-KR or 25GBASE-CR 0 = PMA/PMD is not negotiated to perform 25GBASE-KR or 25GBASE-CR	RO
7.48.12	25GBASE-KR-S or 25GBASE-CR-S	1 = PMA/PMD is negotiated to perform 25GBASE-KR-S or 25GBASE-CR-S 0 = PMA/PMD is not negotiated to perform 25GBASE-KR-S or 25GBASE-CR-S	RO
7.48.7	Reserved RS-FEC negotiated	Value always 0 1 = PMA/PMD is negotiated to perform RS-FEC 0 = PMA/PMD is not negotiated to perform RS-FEC	RO

^aRO = Read only

Insert new subclause 45.2.7.12.1a after 45.2.7.12.1 as follows:

45.2.7.12.1a RS-FEC negotiated (7.48.7)

When the Auto-Negotiation process has completed as indicated by the AN complete bit (7.1.5), bit 7.48.7 indicates that RS-FEC operation has been negotiated. This bit is set only if RS-FEC operation has been negotiated for a BASE-R PHY supporting negotiation of RS-FEC operation.

Change the title of 45.2.7.12.2 and first sentence as follows:

45.2.7.12.2 Negotiated Port Type (7.48.1, 7.48.2, 7.48.3, 7.48.5, 7.48.6, 7.48.8, 7.48.9, 7.48.10, 7.48.11, 7.48.12, 7.48.13)

When the AN process has been completed as indicated by the AN complete bit, these bits in register 7.48 (~~100GBASE-KX, 10GBASE-KX4, 10GBASE-KR, 40GBASE-KR4, 40GBASE-CR4, 100GBASE-CR10, 100GBASE-KP4, 100GBASE-KR4, 100GBASE-CR4~~) indicate the negotiated port type.

45.2.7.13 EEE advertisement (Register 7.60)

Change the row for 7.60.15:14 of Table 45–210 as follows (unchanged rows not shown):

Table 45–210—EEE advertisement register (Register 7.60) bit definitions

Bit(s)	Name	Description	Clause reference; Next Page bit number	R/W ^a
7.60.15:14	Reserved	Value always 0		RO
7.60.14	25GBASE-R EEE	1 = Advertise that the 25GBASE-R PHY has EEE deep sleep capability 0 = Do not advertise that the 25GBASE-R PHY has EEE deep sleep capability	73.7.7.1:U14	R/W

^aR/W = Read/Write, RO = Read only

Insert new subclause 45.2.7.13.a before 45.2.7.13.1 as follows:

45.2.7.13.a 25GBASE-R EEE supported (7.60.14)

If set to a one, bit 7.60.14 indicates that the technologies advertised in bits A9 and A10 of Table 73–4 are capable of supporting EEE deep sleep operation. If set to a zero, EEE deep sleep operation is not supported for 25GBASE-R.

45.2.7.14 EEE link partner ability (Register 7.61)

Change the row for 7.61.15:14 of Table 45–211 as follows (unchanged rows not shown):

Table 45–211—EEE link partner ability (Register 7.61) bit definitions

Bit(s)	Name	Description	Clause reference; Next Page bit number	R/W ^a
7.61.15:14	Reserved	Value always 0		RO
<u>7.61:14</u>	<u>25GBASE-R EEE</u>	<u>1 = Link partner is advertising EEE deep sleep capability for 25GBASE-R</u> <u>0 = Link partner is not advertising EEE deep sleep capability for 25GBASE-R</u>	<u>73.7.7.1:U14</u>	<u>RO</u>

^aRO = Read only

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

**45.5 Protocol implementation conformance statement (PICS) proforma for
 Clause 45, Management Data Input/Output (MDIO) interface³**

45.5.3 PICS proforma tables for the Management Data Input Output (MDIO) interface

45.5.3.3 PMA/PMD management functions

Change the table in 45.5.3.3 as modified by IEEE Std 802.3bw-2015 adding the row for MM129 at the end of the table (unchanged rows not shown):

Item	Feature	Subclause	Value/Comment	Status	Support
<u>MM129</u>	<u>EEE deep sleep capability indicated for each port type</u>	45.2.1.14		<u>EEE:M</u>	Yes <input type="checkbox"/> N/A <input type="checkbox"/>

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

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.

69. Introduction to Ethernet operation over electrical backplanes

69.1 Overview

69.1.1 Scope

Change the second paragraph of 69.1.1 as follows:

Backplane Ethernet supports the IEEE 802.3 full duplex MAC operating at 1000 Mb/s, 10 Gb/s, 25 Gb/s, 40 Gb/s, or 100 Gb/s providing a bit error ratio (BER) better than or equal to 10^{-12} at the MAC/PLS service interface. The following Physical Layers are supported:


- 1000BASE-KX for 1 Gb/s operation over a single lane
- 10GBASE-KX4 for 10 Gb/s operation over four lanes
- 10GBASE-KR for 10 Gb/s operation over a single lane
- 25GBASE-KR and 25GBASE-KR-S for 25 Gb/s operation over a single lane
- 40GBASE-KR4 for 40 Gb/s operation over four lanes
- 100GBASE-KR4 and 100GBASE-KP4 for 100 Gb/s operation over four lanes

69.1.2 Relationship of Backplane Ethernet to the ISO OSI reference model

Change the first paragraph of 69.1.2 as follows and insert Figure 69–1a:

Backplane Ethernet couples the IEEE 802.3 MAC to a family of Physical Layers defined for operation over electrical backplanes. The relationships among Backplane Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in [Figure 69–1](#), [Figure 69–1a](#), and [Figure 69–2](#).

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE
 AN = AUTO-NEGOTIATION
 MDI = MEDIUM DEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 FEC = FORWARD ERROR CORRECTION

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 69–1a—Architectural positioning of 25 Gb/s Backplane Ethernet

Change the lettered list after Figure 69–2 by inserting a new item f), renumbering remaining list items, and changing the newly numbered item i) as follows:

It is important to note that, while this specification defines interfaces in terms of bits, octets, and frames, implementers may choose other data-path widths for implementation convenience. The only exceptions are as follows:

- a) The GMII, which, when implemented at an observable interconnection point, uses an octet-wide data path as specified in Clause 35.
- b) The XGMII, which, when implemented at an observable interconnection point, uses a 4-octet-wide data path as specified in Clause 46.
- c) The management interface, when implemented as the MDIO/MDC (Management Data Input/Output, Management Data Clock) at an observable interconnection point, uses a bit-wide data path as specified in Clause 45.
- d) The 1000BASE-X PMA service interface, when implemented at an observable interconnection point (TBI), uses the 10-bit-wide data path as specified in Clause 36.
- e) The PMA service interface for 10Gb/s serial, when implemented at an observable interconnection point (XSBI), uses the 16-bit-wide data path as specified in Clause 51.
- f) The PMA service interface, which, when physically implemented as 25GAUI (25 Gigabit Attachment Unit Interface) at an observable interconnection port, uses a serial data path as specified in Annex 109A.

- g) The PMA service interface, which, when physically implemented as XLAUI (40 Gb/s Attachment Unit Interface) or as CAUI-4 (100 Gb/s four-lane Attachment Unit Interface) at an observable interconnection port, uses a four-lane data path as specified in Annex 83A or Annex 83D, respectively.
- h) The PMA service interface, which, when physically implemented as CAUI-10 (100 Gb/s ten-lane Attachment Unit Interface) at an observable interconnection port, uses a ten-lane data path as specified in Annex 83A.
- i) The MDIs for 1000BASE-KX, ~~and~~ 10GBASE-KR, 25GBASE-KR, and 25GBASE-KR-S use a serial data path while the MDIs for 10GBASE-KX4, 40GBASE-KR4, 100GBASE-KR4, and 100GBASE-KP4 use a four-lane data path.

69.2 Summary of Backplane Ethernet Sublayers

69.2.1 Reconciliation sublayer and media independent interfaces

Change 69.2.1 as follows:

The [Clause 35](#) RS and GMII, the [Clause 46](#) RS and XGMII, ~~the Clause 106 RS and 25GMII~~, and the [Clause 81](#) RS, XLGMII, and CGMII are employed for the same purpose in Backplane Ethernet, that being the interconnection between the MAC sublayer and the PHY.

69.2.3 Physical Layer signaling systems

Insert the following new paragraph after the third paragraph in 69.2.3:

Backplane Ethernet also specifies 25GBASE-KR and 25GBASE-KR-S. The 25GBASE-KR embodiment employs the PCS defined in Clause 107, the BASE-R FEC defined in Clause 74, the RS-FEC defined in Clause 108, the PMA defined in Clause 109, and the PMD defined in Clause 111, and specifies 25 Gb/s operation over one differential path in each direction. The 25GBASE-KR-S embodiment employs the PCS defined in Clause 107, the BASE-R FEC defined in Clause 74, the PMA defined in Clause 109, and the PMD defined in Clause 111, and specifies 25 Gb/s operation over one differential path in each direction.

Change the last paragraph in 69.2.3 as follows and insert Table 69-1a after Table 69-1:

[Table 69-1](#), [Table 69-1a](#), and [Table 69-2](#) specify the correlation between nomenclature and clauses. A complete implementation conforming to one or more nomenclatures meets the requirements of the corresponding clauses.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

Table 69–1a—Nomenclature and clause correlation for 25 Gb/s Backplane Ethernet Physical Layers

Nomenclature	Clause										
	73	74	78	106		107	108	109	109A	111	
	AN	BASE-R FEC	EEE	RS	25GMII	25GBASE-R PCS	25GBASE-R RS-FEC	25GBASE-R PMA	25GAUI C2C	25GGBASE-KR PMD	25GGBASE-KR-S PMD
25GBASE-KR	M ^a	M	O ^a	M	O	M	M	M	O	M	
25GBASE-KR-S	M	M	O	M	O	M		M	O		M

^aO = Optional, M = Mandatory

69.3 Delay constraints

Insert the following paragraph after the third paragraph of 69.3:

For 25GBASE-KR and 25GBASE-KR-S, normative delay specifications may be found in 106.1.4, 107.3, 109.5, and 111.4, and also referenced in 105.5.

69.5 Protocol implementation conformance statement (PICS) proforma


Change the first paragraph of 69.5 as follows:

The supplier of a protocol implementation that is claimed to conform to any part of IEEE Std 802.3, [Clause 70](#) through [Clause 74](#), [Clause 84](#), [Clause 91](#), [Clause 93](#), [Clause 94](#), [Clause 108](#), [Clause 111](#), and related annexes demonstrates compliance by completing a protocol implementation conformance statement (PICS) proforma.

73. Auto-Negotiation for backplane and copper cable assembly

73.2 Relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model

Change Figure 73–1 as follows:



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE
 AN = AUTO-NEGOTIATION
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 GMII = GIGABIT MEDIA INDEPENDENT INTERFACE
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XGMII = 10 Gb/s MEDIA INDEPENDENT INTERFACE
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 73–1—Location of Auto-Negotiation function within the ISO/IEC OSI reference model

73.3 Functional specifications

Change the third paragraph of 73.3 as follows:

These functions shall comply with the state diagrams from Figure 73–9 through Figure 73–11. The Auto-Negotiation functions shall interact with the technology-dependent PHYs through the Technology-Dependent interface (see 73.9). Technology-Dependent PHYs include 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, 25GBASE-KR, 25GBASE-KR-S, 25GBASE-CR, 25GBASE-CR-S, 40GBASE-KR4, 40GBASE-CR4, 100GBASE-CR10, 100GBASE-KP4, 100GBASE-KR4, and 100GBASE-CR4.

73.6 Link codeword encoding

Replace Figure 73–6 with the following figure and change the last two sentences of the second paragraph of 73.6 as follows:




Figure 73–6—Link codeword Base Page

D[45:21] contains the Technology Ability Field. D[47:44] contains FEC capability (see 73.6.5).

73.6.4 Technology Ability Field

Change Table 73–4 as follows:

Table 73–4—Technology Ability Field encoding

Bit	Technology
A0	1000BASE-KX
A1	10GBASE-KX4
A2	10GBASE-KR
A3	40GBASE-KR4
A4	40GBASE-CR4
A5	100GBASE-CR10
A6	100GBASE-KP4
A7	100GBASE-KR4
A8	100GBASE-CR4
A9	25GBASE-KR-S or 25GBASE-CR-S
A10	25GBASE-KR or 25GBASE-CR
A9 through A24 A11 through A22	Reserved for future technology

Change the third and fourth paragraphs of 73.6.4 adding a new paragraph between them as follows:

For 25 Gb/s operation the same bits are used to advertise backplane and copper cable assembly operation. For other speeds, a PHY for operation over an electrical backplane (e.g., 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, 40GBASE-KR4, 100GBASE-KP4, 100GBASE-KR4) shall not be advertised simultaneously with a PHY for operation over a copper cable assembly (e.g., 40GBASE-CR4, 100GBASE-CR10, 100GBASE-CR4) as the MDI and physical medium are different.

25GBASE-KR-S abilities are a subset of 25GBASE-KR abilities, and likewise 25GBASE-CR-S abilities are a subset of 25GBASE-CR abilities. To allow interoperability between 25GBASE-KR-S and 25GBASE-KR PHY types, and between 25GBASE-CR-S and 25GBASE-CR PHY types, a device that supports 25GBASE-KR or 25GBASE-CR should advertise both A9 and A10 ability bits during auto-negotiation.

The fields A[22:11] A[24:9] are reserved for future use. Reserved fields shall be sent as zero and ignored on receive.

73.6.5 FEC capability

Change the first paragraph of 73.6.5 as follows:

FEC (F2:F3:F0:F1) is encoded in bits D446:D47 of the base link codeword. The four ~~two~~ FEC bits are used as follows:

- a) F0 is 10 Gb/s per lane FEC ability
- b) F1 is 10 Gb/s per lane FEC requested
- c) F2 is 25G RS-FEC requested
- d) F3 is 25G BASE-R FEC requested

Bits F2 and F3 are used for resolving FEC operation for 25G PHYs, while bits F0 and F1 are used for 10 Gb/s per lane operation. Bits F0 and F1 are not used for 25G PHYs.

Insert new subclause 73.6.5.1 before the second paragraph of 73.6.5 as follows:

73.6.5.1 FEC resolution for 25G PHYs

For 25G PHYs if neither PHY requests FEC operation in bits F2 or F3 then FEC is not enabled.

For 25GBASE-KR and 25GBASE-CR PHYs if either PHY requests RS-FEC then RS-FEC operation is enabled, otherwise if either PHY requests BASE-R FEC then BASE-R operation is enabled.

For 25GBASE-KR-S and 25GBASE-CR-S PHYs, if either PHY requests RS-FEC or BASE-R FEC then BASE-R operation is enabled. This is because 25GBASE-KR-S and 25GBASE-CR-S PHYs do not support RS-FEC operation.

Make the second and third paragraphs of 73.6.5 a new subclause 73.6.5.2 and change as follows:

73.6.5.2 FEC resolution for 10 Gb/s per lane PHYs

For 10 Gb/s per lane operation, when the FEC ability bit F0 is set to logical one, it indicates that the PHY has FEC ability (see Clause 74). When the FEC requested F1 bit is set to logical one, it indicates a request to enable FEC on the link.

Since the local device and the link partner may have set the FEC capability bits differently, the priority resolution function is used to enable FEC in the respective PHYs. The FEC function shall be enabled on the link if 10GBASE-KR, 40GBASE-KR4, 40GBASE-CR4, or 100GBASE-CR10 is the HCD technology (see 73.7.6), both devices advertise FEC ability on the F0 bits, and at least one device requests FEC on the F1 bits; otherwise FEC shall not be enabled.

Make fourth and fifth paragraphs of 73.6.5 a new subclause 73.6.5.3 and change as follows:

73.6.5.3 FEC control variables

The variable `an_baser_fec_control` indicates that BASE-R FEC operation has been negotiated. If the value is false, then BASE-R FEC has not been negotiated. If the value is true, then BASE-R FEC has been negotiated. The mapping of this variable to an MDIO bit is defined in Table 73-6.

The variable `an_rs_fec_control` indicates that RS-FEC operation has been negotiated. If the value is false, then RS-FEC has not been negotiated. If the value is true, then RS-FEC has been negotiated.

The mapping of these variables to MDIO register bits is defined in Table 73-6.

If `mr_autoneg_enable` (see 73.10.1) is false, the FEC function is controlled by implementation-dependent means.

73.7 Receive function requirements

73.7.1 DME page reception

Change 73.7.1 as follows:

To be able to detect the DME bits, the receiver should have the capability to receive DME signals sent with the electrical specifications of the PHY (100GBASE-KX, 10GBASE-KX4, 10GBASE-KR, 25GBASE-KR, 25GBASE-KR-S, 25GBASE-CR, 25GBASE-CR-S, 40GBASE-KR4, 40GBASE-CR4, 100GBASE-CR10, 100GBASE-KP4, 100GBASE-KR4, or 100GBASE-CR4). The DME transmit signal level and receive sensitivity are specified in 73.5.1.

73.7.6 Priority Resolution function

Change Table 73-5 to include 25GBASE-KR or 25GBASE-CR and 25GBASE-KR-S or 25GBASE-CR-S and change priority of lower speed PHYs as follows:

Table 73-5—Priority Resolution

Priority	Technology	Capability
1	100GBASE-CR4	100 Gb/s 4 lane, highest priority
2	100GBASE-KR4	100 Gb/s 4 lane
3	100GBASE-KP4	100 Gb/s 4 lane
4	100GBASE-CR10	100 Gb/s 10 lane
5	40GBASE-CR4	40 Gb/s 4 lane
6	40GBASE-KR4	40 Gb/s 4 lane

Table 73–5—Priority Resolution (continued)

Priority	Technology	Capability
<u>7</u>	<u>25GBASE-KR or 25GBASE-CR</u>	<u>25 Gb/s 1 lane</u>
<u>8</u>	<u>25GBASE-KR-S or 25GBASE-CR-S</u>	<u>25 Gb/s 1 lane, short reach</u>
7 <u>9</u>	10GBASE-KR	10 Gb/s 1 lane
8 <u>10</u>	10GBASE-KX4	10 Gb/s 4 lane
9 <u>11</u>	1000BASE-KX	1 Gb/s 1 lane, lowest priority

73.6 Management register requirements

Insert new row for *an_rs_fec_control* at the bottom of Table 73–6 as follows (unchanged rows not shown):

Table 73–6—Backplane Ethernet Auto-Negotiation variable to MDIO register mapping

Variable	Description
<i>an_rs_fec_control</i>	7.48.7 RS-FEC negotiated

73.10 State diagrams and variable definitions

73.10.1 State diagram variables

Change the list of variables to include *25GBASE-CR* and *25GBASE-KR* as follows:

A variable with “[x]” appended to the end of the variable name indicates a variable or set of variables as defined by “x”. “x” may be as follows:

- all; represents all specific technology-dependent PMDs supported in the local device.
- 1GKX; represents the 1000BASE-KX PMD.
- 10GKR; represents the 10GBASE-KR PMD.
- 10GKX4; represents the 10GBASE-KX4 or 10GBASE-CX4 PMD.
- 25GR; represents the 25GBASE-KR, 25GBASE-KR-S, 25GBASE-CR, or 25GBASE-CR-S PMD.
- 40GKR4; represents the 40GBASE-KR4 PMD.
- 40GCR4; represents the 40GBASE-CR4 PMD.
- 100GCR10; represents the 100GBASE-CR10 PMD.
- 100GKP4; represents the 100GBASE-KP4 PMD.
- 100GKR4; represents the 100GBASE-KR4 PMD.
- 100GCR4; represents the 100GBASE-CR4 PMD.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

- HCD; represents the single technology-dependent PMD chosen by Auto-Negotiation as the highest common denominator technology through the Priority Resolution or parallel detection function.
- notHCD; represents all technology-dependent PMDs not chosen by Auto-Negotiation as the highest common denominator technology through the Priority Resolution or parallel detection function.
- PD; represents all of the following that are present: 1000BASE-KX PMD and 10GBASE-KX4 (or 10GBASE-CX4) PMD.

Change `single_link_ready` as follows:

`single_link_ready`

Status indicating that `an_receive_idle = true` and only one of the following indications is being received:

- 1) `link_status_[1GKX] = OK`
- 2) `link_status_[10GKX4] = OK`
- 3) `link_status_[10GKR] = OK`
- 4) `link_status_[25GR] = OK`
- 45) `link_status_[40GKR4] = OK`
- 56) `link_status_[40GCR4] = OK`
- 67) `link_status_[100GCR10] = OK`
- 78) `link_status_[100GKP4] = OK`
- 89) `link_status_[100GKR4] = OK`
- 910) `link_status_[100GCR4] = OK`

Values: `false`; either zero or more than one of the above indications are true or `an_receive_idle = false`.
`true`; Exactly one of the above indications is true and `an_receive_idle = true`.

NOTE—This variable is set by this variable definition; it is not set explicitly in the state diagrams.

74. Forward Error Correction (FEC) sublayer for BASE-R PHYs

74.1 Overview

Change 74.1 as follows:

This clause specifies an ~~optional~~ Forward Error Correction (FEC) sublayer for 10GBASE-R and other BASE-R PHYs. The FEC sublayer can be placed in between the PCS and PMA sublayers of the 10GBASE-R and other BASE-R Physical Layer implementations as shown in [Figure 74-2](#), [Figure 74-2a](#), [Figure 74-3](#), and [Figure 74-4](#). For a PHY with a multi-lane BASE-R PCS, the FEC sublayer is instantiated for each PCS lane and operates autonomously on a per PCS lane basis. The FEC provides coding gain to increase the link budget and BER performance.


The 10GBASE-KR and 40GBASE-KR4 PHYs described in [Clause 72](#), and [Clause 84](#) optionally use the FEC sublayer to increase the performance on a broader set of backplane channels than are defined in [Clause 69](#). The FEC sublayer provides additional margin to account for variations in manufacturing and environmental conditions.

The 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, and 25GBASE-KR-S PHYs described in [Clause 110](#) and [Clause 111](#) are required to implement the FEC sublayer and may use it with links with a BER of 10^{-8} or better.

The 40GBASE-CR4 and 100GBASE-CR10 PHYs described in [Clause 85](#) optionally use the FEC sublayer to improve the BER performance beyond 10^{-12} .

74.3 Relationship to other sublayers

Replace Figure 74–1 with the following figure:



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE PMA = PHYSICAL MEDIUM ATTACHMENT
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE PMD = PHYSICAL MEDIUM DEPENDENT
FEC = FORWARD ERROR CORRECTION XGMII = 10 Gb/s MEDIA INDEPENDENT INTERFACE
MDI = MEDIUM DEPENDENT INTERFACE XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE NOTE 1—OPTIONAL OR CONDITIONAL BASED ON PHY TYPE

Figure 74–1—BASE-R FEC relationship to ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

74.4 Inter-sublayer interfaces

Change 74.4 as follows:

An FEC service interface is provided to allow the FEC sublayer to transfer information to and from the PCS. An abstract service model is used to define the operation of this interface. For 10GBASE-R, the FEC service interface directly maps to the PMA service interface of the PCS defined in Clause 49 and the lower FEC sublayer interface maps to the service interface provided by the serial PMA sublayer defined in Clause 51. For 25GBASE-R, the FEC service interface is an instance of the inter-sublayer service interface defined in 105.4, as is the PMA service interface defined in 109.2. For 40GBASE-R and 100GBASE-R, the FEC service interface is an instance of the inter-sublayer service interface defined in 80.3 as is the PMA service interface defined in 83.2.

For 25GBASE-R the FEC service interface can either connect to the PCS as illustrated in Figure 74–1 or the PMA as illustrated in Figure 109–3 where the FEC and PCS are in separate devices connected by 25GAUI.

For 40GBASE-R and 100GBASE-R the FEC service interface can either connect to the PCS as illustrated in Figure 74–1 or the PMA as illustrated in Figure 83–2 where the FEC and PCS are in separate devices connected by XLAUI/CAUI-n.

This standard defines these interfaces in terms of bits, octets, data-group, data units, and signals; however, implementers may choose other data-path widths and other control mechanisms for implementation convenience, provided that the implementation adheres to the logical model of the service interface.

Insert new subclause 74.4.1a and Figure 74–2a after 74.4.1 as follows:

74.4.1a Functional block diagram for 25GBASE-R PHYs

Figure 74–2a shows the functional block diagram of FEC for 25GBASE-R PHYs and the relationship between the PCS and PMA sublayers.




Figure 74–2a— Functional block diagram for 25GBASE-R PHY

74.5 FEC service interface

Change 74.5 as follows:

The FEC service interface is provided to allow the PCS to transfer information to and from the FEC. The FEC service interface is equivalent to the PMA service interface for 10GBASE-R and an instance of the inter-sublayer service interface defined in 105.4 for 25GBASE-R, and 80.3 for 40GBASE-R and 100GBASE-R. These services are defined in an abstract manner and do not imply any particular implementation. The FEC service interface supports exchange of data units between PCS entities on either side of a link using request and indication primitives. Data units are mapped into FEC blocks by the FEC and passed to the PMA, and vice versa.

The service primitives are defined differently for 10GBASE-R, ~~and for 25GBASE-R, and for 40GBASE-R and 100GBASE-R.~~

Optional physical instantiations of the PMA service interface have been defined (see [Clause 51, Annex 109A, Annex 109B, Annex 83A, Annex 83B, Annex 83D, and Annex 83E](#)). There is XSBI (10 Gigabit Sixteen-Bit Interface) for 10GBASE-R, 25GAUI for 25GBASE-R, XLAUI for 40GBASE-R and CAUI-n for 100GBASE-R. These physical instantiations, with a PMA if required, may also be used for the FEC service interface.

Insert new subclause 74.5.1a after 74.5.1 as follows:

74.5.1a 25GBASE-R service primitives

The FEC service interface for 25GBASE-R is an instance of the inter-sublayer service interface defined in 105.4. The FEC service interface primitives are summarized as follows:

- a) FEC:IS_UNITDATA.request
- b) FEC:IS_UNITDATA.indication
- c) FEC:IS_SIGNAL.indication
- d) FEC:IS_TX_MODE.request(tx_mode)
- e) FEC:IS_RX_MODE.request(rx_mode)
- f) FEC:IS_RX_TX_MODE.indication(rx_tx_mode)
- g) FEC:IS_LPI_ACTIVE.request(rx_lpi_active)
- h) FEC:IS_ENERGY.indication(energy_detect)

Items d), e), f), g), and h) are only required for the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode.

The 25GBASE-R PCS (or PMA) continuously sends bit streams to the FEC via the tx_bit parameter of the FEC:IS_UNITDATA.request at a nominal signaling rate of 25.78125 GBd.

The FEC continuously sends bit streams to the 25GBASE-R PCS (or PMA) via the rx_bit parameter of the FEC:IS_UNITDATA.indication at a nominal signaling rate of 25.78125 GBd.

The SIGNAL_OK parameter of the FEC:IS_SIGNAL.indication primitive can take one of two values: OK or FAIL. A value of OK denotes that the FEC Receive process is successfully delineating valid payload information from the incoming data stream received from the PMA sublayer indicated by the fec_signal_ok variable equal to true, and this payload information is being presented to the PCS (or PMA) via the FEC:IS_UNITDATA.indication primitive. A value of FAIL denotes that errors have been detected by the Receive process indicated by the fec_signal_ok variable equal to false that prevent valid data from being

presented to the PCS, in this case the rx_bit parameter of the FEC:IS_UNITDATA.indication primitive is undefined.

FEC:IS_TX_MODE.request is sourced from the PCS and passed through the FEC to the PMA sublayer. The tx_mode parameter can take the values QUIET, ALERT, or DATA. When tx_mode is QUIET or ALERT, the FEC encoder logic may deactivate functional blocks to conserve energy. When tx_mode is DATA, the FEC encoder logic operates normally.

FEC:IS_RX_MODE.request is sourced from the PCS and passed through the FEC to the PMA sublayer. The rx_mode parameter can take the values QUIET or DATA. When rx_mode is QUIET, the FEC decoder logic may deactivate functional blocks to conserve energy. When rx_mode is DATA, the FEC decoder logic operates normally.

FEC:IS_RX_TX_MODE.indication communicates the rx_tx_mode parameter to a PMA client layer if present. This parameter indicates the value of tx_mode that the PMA sublayer has inferred from the received signal. Without EEE deep sleep capability, the primitive is never generated and the sublayer behaves as if rx_tx_mode = DATA. The parameter rx_tx_mode is assigned one of the following values: DATA, QUIET, or ALERT.

The rx_lpi_active parameter in FEC:IS_LPI_ACTIVE.request is a Boolean variable sent from the PCS that is set to TRUE when LPI mode is active at the receiver and set to FALSE otherwise. When TRUE, rx_lpi_active causes rapid FEC block lock as specified in 74.7.4.8 to be used to quickly determine the start of the FEC block during EEE REFRESH or WAKE.

The energy_detect parameter in FEC:IS_ENERGY.indication is a Boolean variable that indicates to the PCS that energy has been detected at the PMD. This signal is passed up through the FEC from the PMA sublayer to the PCS. It has no effect on FEC operation.

74.6 Delay constraints

Insert a new paragraph after the second paragraph of 74.6 as follows:

The maximum delay contributed by the 25GBASE-R FEC (sum of transmit and receive delays at one end of the link) shall be no more than 6144 BT (or 12 pause quanta or 245.76 ns).

74.7 FEC principle of operation

74.7.4 Functions within FEC sublayer

74.7.4.1 Reverse gearbox function

Change 74.7.4.1.2 as follows:

74.7.4.1.2 Reverse gearbox function for 25GBASE-R, 40GBASE-R, and 100GBASE-R

The reverse gearbox function adapts between the 66-bit width of the 64B/66B blocks and the 1-bit wide lane of the 25GBASE-R, 40GBASE-R, or 100GBASE-R PCS to FEC interface (or PMA to FEC interface). It receives the 1-bit stream from the FEC service interface (or PMA service interface) and converts it back to 66-bit encoded blocks for the FEC Encoder to process. The reverse gearbox function, if implemented, shall operate in the same manner as the lane block sync function defined in 82.2.12.

The reverse gearbox function receives data via the 25GBASE-R FEC:IS_UNITDATA.request primitive (or via the PMA:IS_UNITDATA.request primitive) or via the 40GBASE-R and 100GBASE-R

FEC:IS_UNITDATA_i.request primitive (or via the PMA:IS_UNITDATA_i.request primitive). It obtains lock to the 66-bit blocks in each bit stream using the sync headers and outputs 66-bit blocks to the FEC encoder function (see 74.7.4.4). PCS lane lock is obtained as specified in the PCS lane lock state diagram shown in Figure 82–12.

The internal data-path width from the PCS or PMA is an implementation choice. Depending on the path width, the reverse gearbox function may not be necessary.

74.7.4.3 FEC transmission bit ordering

Change Figure 74–5 as follows:




Figure 74–5—FEC Transmit bit ordering

74.7.4.4 FEC (2112,2080) encoder

Change Figure 74–6 as follows:




Figure 74–6—FEC (2112,2080) encoding

74.7.4.5 FEC decoder

Change third paragraph of 74.7.4.5 as follows:

When the decoder for 10GBASE-R or 25GBASE-R is configured to indicate decoding error, the decoder indicates error to the PCS by means of setting both sync bits to the value 11 in the 1st, 9th, 17th, 25th, and 32nd of the 32 decoded 64B/66B blocks from the corresponding errored FEC block, thus forcing the PCS sublayer to consider this block as invalid.

74.7.4.5.1 FEC (2112,2080) decoding

Change Figure 74–8 and item b) as follows:




Figure 74–8—FEC (2112,2080) decoding

- b) If the variable FEC_Enable_Error_to_PCS is set to 1 to indicate error to PCS layer and the received FEC block has uncorrectable errors then the sync bits for the 1st, 9th, 17th, 25th, and 32nd of the 32 decoded 64B/66B blocks take a value of {SH.0,SH.1} = 11 for the 10GBASE-R and 25GBASE-R PHY. For the 40GBASE-R and 100GBASE-R PHYs, sync bits in all thirty-two 64B/66B decoded 64B/66B blocks take a value of {SH.0,SH.1} = 11. The sync bits for all other 64B/66B blocks take a value as described in item a) above.

74.7.4.6 FEC receive bit ordering

Change Figure 74-9 as follows:




Figure 74-9—FEC Receive bit ordering

74.7.4.8 FEC rapid block synchronization for EEE (optional)

Change the first paragraph of 74.7.4.8 (and split it into multiple paragraphs) as follows:

If the optional EEE capability is supported then during the wake and refresh states the FEC decoder receives one of the two types of deterministic blocks to achieve rapid block synchronization. During these states the reverse gearbox of the remote FEC encoder is receiving unscrambled data from the PCS sublayer via 16-bit FEC_UNITDATA.request primitive.

The ~~A~~ [Clause 49](#) and [Clause 107](#) PCS sublayers encodes /I/ during the wake state and /LI/ during the refresh state, which produces the two types of deterministic FEC blocks.

If the optional EEE deep sleep capability is supported, then a [Clause 82](#) PCS sublayer also encodes /I/ during the wake state and /LI/ during the refresh state, but in addition inserts Rapid Alignment Markers into each of the PCS Lanes (see [82.2.9](#)). This causes the two types of deterministic FEC blocks to have a number of 65-bit words within the deterministic FEC block replaced with Rapid Alignment Markers, thus not matching the two deterministic patterns as shown in [Table 74A–5](#) and [Table 74A–6](#). The locations of the Rapid Alignment Markers within the Rapid FEC block are consistent for a given entry into the wake or refresh states, but the locations can vary for subsequent entries. This modification to the two deterministic patterns needs to be taken into account by the Rapid FEC Lock implementation.

74.8 FEC MDIO function mapping

Change Table 74-1 as follows:

Table 74-42—MDIO/FEC variable mapping

MDIO variable	PMA/PMD register name	Register/bit number	FEC variable
BASE-R FEC ability	BASE-R FEC ability register	1.170.0	FEC_ability
BASE-R FEC Error Indication ability	BASE-R FEC ability register	1.170.1	FEC_Error_Indication_ability
FEC Enable	BASE-R FEC control register	1.171.0	FEC_Enable
FEC Enable Error Indication	BASE-R FEC control register	1.171.1	FEC_Enable_Error_to_PCS
FEC corrected blocks	10GBASE-R Single-lane PHY BASE-R FEC corrected blocks counter register	1.172, 1.173	FEC_corrected_blocks_counter
FEC uncorrected blocks	10GBASE-R Single-lane PHY BASE-R FEC uncorrected blocks counter register	1.174, 1.175	FEC_uncorrected_blocks_counter
FEC corrected blocks, lanes 0 through 19	BASE-R FEC corrected blocks counter register, lanes 0 through 19	1.300 through 1.339	FEC_corrected_blocks_counter_i
FEC uncorrected blocks, lanes 0 through 19	BASE-R FEC uncorrected blocks counter register, lanes 0 through 19	1.700 through 1.739	FEC_uncorrected_blocks_counter_i

74.8.1 FEC capability

Change last paragraph of 74.8.1 as follows:

The FEC capability between the link partners can be negotiated using the Clause 73 Auto-Negotiation as defined in 73.6.5. The FEC function is enabled on the link only if both the link partners advertise they have FEC ability and either one of them requests to enable FEC through the Auto-Negotiation function.

74.9 BASE-R PHY test-pattern mode

Change 74.9 as follows:

The 10GBASE-R PCS and the 25GBASE-R PCS provides test-pattern functionality and the PCS transmit channel and receive channel can each operate in normal mode or test-pattern mode (see 49.2.2). When the 10GBASE-R or 25GBASE-R PHY is configured for test-pattern mode, the FEC function may be disabled by setting the FEC Enable variable to zero, so the test-pattern from the 10GBASE-R PCS can be sent to the PMA service interface, bypassing the FEC Encode and Decode functions.

The Clause 82 and Clause 107 PCS can also operate in test pattern mode (see 82.2.11 and 107.2.3); however, the scrambled idle test pattern does not require bypassing FEC encode and decode.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

74.11 Protocol implementation conformance statement (PICS) proforma for Clause 74, Forward Error Correction (FEC) sublayer for BASE-R PHYs⁴

74.11.3 Major capabilities/options

Change item DC in the table in 74.11.3 as follows:

Item	Feature	Subclause	Value/Comment	Status	Support
DC	FEC Delay Constraints	74.6	Sum of transmit and receive. No more than 12 pause_quanta for 10GBASE-R and 25GBASE-R, 48 pause_quanta for 40GBASE-R, and 240 pause_quanta for 100GBASE-R	M	Yes []

74.11.5 FEC Requirements

Change item FE4 in the table in 74.11.5 as follows:

Item	Feature	Subclause	Value/Comment	Status	Support
FE4	Reverse gearbox function for 25GBASE-R, 40GBASE-R, and 100GBASE-R	74.7.4.1.2	Reverse gearbox function meets the requirements of 82.2.12	O	Yes [] No []

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.

78. Energy-Efficient Ethernet (EEE)

78.1 Overview

Change the third paragraph of 78.1 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. Table 78–1 lists the supported PHYs and interfaces and their associated clauses.

78.1.1 LPI Signaling

Change the fourth and fifth paragraph of 78.1.1 as follows:

The EEE request signals from the PCS control transitions between quiescent and normal operation. The Clause 49 PCS, Clause 107 PCS, and Clause 82 PCS also request transmit alert operation to assist the partner device PMD to detect the end of the quiescent state. Additionally the Clause 49 PCS and Clause 82 these PCS types generate the RX_LPI_ACTIVE signal, which indicates to the Clause 74 BASE-R FEC that it can use rapid block lock because the link partner PCS has bypassed scrambling.

Coding defined in 83.5.11 Clause 83 also allows LPI transmit quiet and alert requests from the PCS to be signaled over the 25GAUI, XLAUI, and CAUI-n interfaces. The 25GAUI, XLAUI, and CAUI-n receive interfaces infer the quiet and alert requests from the data received over the interface and use that to recreate the transmit or receive direction signaling. (See 83.5.11.1.)

78.1.3 Reconciliation sublayer operation

78.1.3.3 PHY LPI operation

78.1.3.3.1 PHY LPI transmit operation

Change the last paragraph of 78.1.3.3.1 as follows:

For PHYs with an operating speed of ~~2540~~ Gb/s or greater that implement the optional EEE capability, two modes of LPI operation may be supported: deep sleep and fast wake. *Deep sleep* refers to the mode for which the transmitter ceases transmission during Low Power Idle (as shown in Figure 78–3) and is equivalent to the only mechanism defined for PHYs with an operating speed of 10 Gb/s or below ~~less than 40 Gb/s~~. Deep sleep support is optional for PHYs with an operating speed of ~~2540~~ Gb/s or greater that implement EEE with the exception of the PHYs noted in Table 78–1 that do not support deep sleep. *Fast wake* refers to the mode for which the transmitter continues to transmit signals during Low Power Idle so that the receiver can resume operation with a shorter wake time (as shown in Figure 78–4). For transmit, other than the PCS encoding LPI, there is no difference between fast wake and normal operation. Fast wake support is mandatory for PHYs with an operating speed of ~~2540~~ Gb/s or greater that implement EEE.

78.1.4 PHY types optionally supporting EEE

Insert new rows into Table 78–1 between 10GBASE-T and XLAUI/CAUI-10 and change the footnote to include 25GAUI as follows (unchanged rows not shown):

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

PHY or interface type	Clause
25GAUI ^a	109A
25GBASE-KR	74, 107, 108, 109, 111
25GBASE-KR-S	74, 107, 109, 111
25GBASE-CR	74, 107, 108, 109, 110
25GBASE-CR-S	74, 107, 109, 110
25GBASE-SR ^b	107, 108, 109, 112

^a25GAUI/XLAUI/CAUI-n shutdown is supported only when deep sleep is enabled for the associated PHY.

78.2 LPI mode timing parameters description

Insert two rows into Table 78–2 between 10GBASE-T and 40GBASE-KR4 as follows (unchanged rows not shown):

Table 78–2—Summary of the key EEE parameters for supported PHYs or interfaces

PHY or interface type	T_s (μs)		T_q (μs)		T_r (μs)	
	Min	Max	Min	Max	Min	Max
25GBASE-KR 25GBASE-CR	4.9	5.1	1 700	1 800	16.9	17.5
25GBASE-KR-S 25GBASE-CR-S	4.9	5.1	1 700	1 800	16.9	17.5

78.5 Communication link access latency

Insert a new paragraph before the last paragraph of 78.5 as follows:

Case-1 of the 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, and 25GBASE-KR-S PHYs applies to PHYs in deep sleep without FEC enabled. Case-2 applies to 25GBASE-R PHYs with the BASE-R FEC in deep sleep. Case-3 applies to 25GBASE-R PHYs with the RS-FEC in deep sleep.

Insert new rows into Table 78–4 between 10GBASE-KR and 40GBASE-KR fast wake and change the last row and footnote to include 25GAUI 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) (μ s)	T_{w_phy} (min) (μ s)	$T_{phy_shrink_tx}$ (max) (μ s)	$T_{phy_shrink_rx}$ (max) (μ s)	$T_{w_sys_rx}$ (min) (μ s)
25GBASE-R fast wake		0.34	0.3	0	0	0.25
25GBASE-CR	Case-1	15.38	12.25	5	7.5	2.88
25GBASE-CR-S	Case-2	17.38	14.25	5	9.5	2.88
25GBASE-KR						
25GBASE-KR-S	Case-3	15.38	12.25	5	7.5	2.88
...						
25GAUI/XLAUI/ CAUI-n ^a		1				

^a $T_{w_sys_tx}$ is increased by 1 μ s for each instance of 25GAUI/XLAUI/CAUI with shutdown enabled on the transmit path. The receiver should negotiate an increase for remote T_{w_sys} for the link partner of 1 μ s for each instance of 25GAUI/XLAUI/CAUI with shutdown enabled on the receive path.

Change 78.5.2 title and text as follows:

78.5.2 25 Gb/s, 40 Gb/s, and 100 Gb/s PHY extension using 25GAUI, XLAUI, or CAUI-n

40 Gb/s PHYs may be extended using 25GAUI, XLAUI, and 100 Gb/s PHYs may be extended using CAUI-10, and CAUI-4 may be used as a physical instantiation of the inter-sublayer service interface to separate functions between devices. The LPI signaling can operate across XLAUI/CAUI-n these interfaces with no change to the PHY timing parameters described in Table 78–4 or the operation of the Data Link Layer Capabilities negotiation described in 78.4.

If PMA Egress AUI Stop Enable (PEASE, see 83.3; MDIO register bit 1.7.8) is asserted for any of the PMA sublayers, the PMA may stop signaling on the XLAUI/CAUI-n AUI in the transmit direction to conserve energy. If PEASE is asserted, the RS defers sending data following deassertion of LPI by an additional time equal to $T_{w_sys_tx} - T_{w_sys_rx}$ as shown in Table 78–4 for each PMA with PEASE asserted (see 81.4.2).

If PMA Ingress AUI Stop Enable (PIASE, see 83.3; MDIO register bit 1.7.9) is asserted for any of the PMA sublayers, the PMA may stop signaling on the XLAUI/CAUI-n AUI in the receive direction to conserve energy. The receiver should negotiate an additional time for the remote T_{w_sys} (equal to $T_{w_sys_tx} - T_{w_sys_rx}$ for the XLAUI/CAUI-n AUI as shown in Table 78–4) for each PMA with PIASE to be asserted before setting the PIASE bits.

90. Ethernet support for time synchronization protocols

90.1 Introduction

Change the second paragraph of 90.1 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) specifications are all compatible with the gRS sublayer defined in 90.5.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

Insert new Clauses 105 to 112 and corresponding new Annexes 109A to 110C as follows:

105. Introduction to 25 Gb/s networks

105.1 Overview

105.1.1 Scope

25 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer, connected through a 25 Gigabit Media Independent Interface (25GMII) to Physical Layer entities such as 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, 25GBASE-KR-S, and 25GBASE-SR.

25 Gb/s Physical Layer entities, such as those specified in Table 105–1, provide a frame loss ratio (see 1.4.223) of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap.

25 Gigabit Ethernet is defined for full duplex operation only.

105.1.2 Relationship of 25 Gigabit Ethernet to the ISO OSI reference model

25 Gigabit Ethernet couples the IEEE 802.3 MAC to a family of 25 Gb/s Physical Layers. The relationships among 25 Gigabit Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in Figure 105–1.

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. The only exceptions are as follows:


- a) The PMA service interface, which, when physically implemented as 25GAUI at an observable interconnection port, uses a single-lane data path as specified in Annex 109A or Annex 109B.
- 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 Media Dependent Interface (MDI) as specified in Clause 110 for 25GBASE-CR and 25GBASE-CR-S, in Clause 111 for 25GBASE-KR and 25GBASE-KR-S, or in Clause 112 for 25GBASE-SR uses a single-lane data path.

105.1.3 Nomenclature

The prefix 25GBASE in the port type (e.g., 25GBASE-R) represents a family of Physical Layer devices operating at a speed of 25 Gb/s.

The term 25GBASE-R refers to a specific family of Physical Layer implementations based upon the 64B/66B data coding method specified in Clause 107. The 25GBASE-R family is composed of 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, 25GBASE-KR-S, and 25GBASE-SR.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE
 AN = AUTO-NEGOTIATION
 FEC = FORWARD ERROR CORRECTION
 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
 NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 105-1—Architectural positioning of 25 Gigabit Ethernet

Physical Layer devices listed in Table 105-1 are defined for operation at 25 Gb/s.

Table 105-1—25 Gb/s PHYs

Name	Description
25GBASE-CR	25 Gb/s PHY using 25GBASE-R encoding over one lane of twinaxial copper cable (see 1.4.407 and Clause 110).
25GBASE-CR-S	25 Gb/s PHY equivalent to 25GBASE-CR without support for the RS-FEC sublayer (see Clause 110).
25GBASE-KR	25 Gb/s PHY using 25GBASE-R encoding over one lane of an electrical backplane (see Clause 111).
25GBASE-KR-S	25 Gb/s PHY equivalent to 25GBASE-KR without support for the RS-FEC sublayer (see Clause 111).
25GBASE-SR	25 Gb/s PHY using 25GBASE-R encoding over a duplex multimode fiber (see Clause 112).

105.2 Physical Layer signaling systems

This standard specifies a family of Physical Layer implementations. The generic term 25 Gigabit Ethernet refers to any use of the 25 Gb/s IEEE 802.3 MAC (the 25 Gigabit Ethernet MAC) coupled with any IEEE 802.3 25GBASE Physical Layer implementation. Table 105–2 specifies the correlation between nomenclature and clauses. Implementations conforming to one or more nomenclatures shall meet the requirements of the corresponding clauses.

Table 105–2—Nomenclature and clause correlation, 25GBASE-R

Nomenclature	Clause														
	73	74	78	106		107	108	109	109A	109B	110		111	112	
	Auto-Negotiation	BASE-R FEC	EEE	RS	25GMII	25GBASE-R PCS	25GBASE-R RS-FEC	PMA	25GAUI C2C	25GAUI C2M	25GBASE-CR PMD	25GBASE-CR-S PMD	25GBASE-KR PMD	25GBASE-KR-S PMD	25GBASE-SR PMD
25GBASE-CR	M ^a	M	O ^a	M	O	M	M	M	O		M				
25GBASE-CR-S	M	M	O	M	O	M		M	O			M			
25GBASE-KR	M	M	O	M	O	M	M	M	O				M		
25GBASE-KR-S	M	M	O	M	O	M		M	O					M	
25GBASE-SR			O	M	O	M	M	M	O	O					M

^aO = Optional, M = Mandatory

105.3 Summary of 25 Gigabit Ethernet sublayers

105.3.1 Reconciliation Sublayer (RS) and 25 Gigabit Media Independent Interface (25GMII)

The 25GMII (Clause 106) provides a logical interconnection between the Media Access Control (MAC) sublayer and Physical Layer entities (PHY). The 25GMII supports 25 Gb/s operation through its 32-bit-wide transmit and receive data paths.

The RS provides a mapping between the signals provided at the 25GMII and the MAC/PLS service definition.

While the 25GMII is an optional interface, it is used extensively in this standard as a basis for functional specification and provides a common service interface for the 25GBASE-R PCS (Clause 107).

105.3.2 Physical Coding Sublayer (PCS)

25GBASE-R PHYs use the PCS specified in Clause 107. The 25GBASE-R PCS performs encoding of data from the 25GMII to 64B/66B code blocks and transfers the encoded data to the PMA and performs decoding of 64B/66B blocks from the PMA and transfers the decoded data to the 25GMII.

105.3.3 Forward Error Correction (FEC) sublayer

An FEC sublayer is available for all 25GBASE-R PHYs. The FEC sublayer can be placed in between the PCS and PMA sublayers or between two PMA sublayers.

The BASE-R FEC (see Clause 74) may be used by some 25GBASE-R PHYs.

The RS-FEC (see Clause 108) may be used by some 25GBASE-R PHYs.

The requirement to implement FEC for each Physical Layer implementation is summarized in 105.2.

105.3.4 Physical Medium Attachment (PMA) sublayer

The PMA sublayer provides a medium-independent means for the PCS to support the use of a range of physical media. The PMA performs the mapping of transmit and receive data streams between the PCS and PMA via the PMA service interface and mapping of transmit and receive data streams between the PMA and PMD via the PMD service interface.

The 25GBASE-R PMA (see Clause 109) 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 25GBASE-R PMA also may provide the observable electrical interface for the 25GAUI chip-to-chip (C2C) or chip-to-module (C2M) (see Annex 109A and Annex 109B).

105.3.5 Physical Medium Dependent (PMD) sublayer

The PMD sublayer is responsible for interfacing to the transmission medium.

The MDI connects the PMD to the medium and is defined in the associated PMD clause.

The 25GBASE-R PMDs and their corresponding media are specified in Clause 110, Clause 111, and Clause 112.

105.3.6 Auto-Negotiation (AN)

AN provides a linked device with the capability to detect the abilities (modes of operation) supported by the device at the other end of the link, determine common abilities, and configure for joint operation.

Clause 73 AN is used by the 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, and 25GBASE-KR-S PHYs.

105.3.7 Management interface (MDIO/MDC)

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

105.3.8 Management

Managed objects, attributes, and actions are defined for all 25 Gigabit Ethernet components. These items are defined in Clause 30.

105.4 Service interface specification method and notation

The service interface specification for 25GBASE-R Physical Layers is as per the definition in 1.2.2 and is based upon the service interface specification in 80.3, but redefined for a single lane.

105.4.1 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.request
IS_UNITDATA.indication
IS_SIGNAL.indication

The IS_UNITDATA.request primitive is used to define the transfer of a stream of data units from a sublayer to the next lower (closer to the medium) sublayer. The IS_UNITDATA.indication primitive is used to define the transfer of a stream 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.

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see 78.1.3.3.1), then the inter-sublayer service interface includes some or all (depending on sublayer, see Figure 105-3) of five additional primitives defined as follows:

IS_TX_MODE.request
IS_RX_MODE.request
IS_ENERGY_DETECT.indication
IS_RX_LPI_ACTIVE.request
IS_RX_TX_MODE.indication

The IS_TX_MODE.request primitive is used to communicate the state of the PCS LPI transmit function to other sublayers in the PHY.

The IS_RX_MODE.request primitive is used to communicate the state of the PCS LPI receive function to other sublayers.

The IS_RX_TX_MODE.indication primitive is used to communicate the state of the rx_tx_mode parameter that reflects the inferred state of the link partner's tx_mode parameter from the PMA to other sublayers.

The IS_RX_LPI_ACTIVE.request primitive is used to communicate to the BASE-R FEC (see Clause 74) that the PCS has detected LPI signaling. This allows the FEC to use rapid block lock; the RS-FEC (see Clause 108) does not use this signal.

The IS_ENERGY_DETECT.indication primitive is used to communicate that the PMD has detected the return of energy on the interface following a period of quiescence.

105.4.2 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.

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation

- b) PMA:—for primitives issued on the interface between the PMA sublayer and the PCS (or the FEC) sublayer called the PMA service interface.
- c) FEC:—for primitives issued on the interface between the FEC sublayer and the PCS (or the PMA) sublayer called the FEC service interface.

Examples of inter-sublayer service interfaces for 25GBASE-R with their corresponding instance names are illustrated in Figure 105–2 and Figure 105–3. 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.request
PMD:IS_UNITDATA.indication
PMD:IS_SIGNAL.indication.
```

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

105.4.3 Semantics of inter-sublayer service interface primitives

The semantics of the inter-sublayer service interface primitives for the 25GBASE-R sublayers are described in 105.4.3.1 through 105.4.3.3.

105.4.3.1 IS_UNITDATA.request

The IS_UNITDATA.request primitive is used to define the transfer of a stream of data units from a sublayer to the next lower sublayer.

105.4.3.1.1 Semantics of the service primitive

```
IS_UNITDATA.request(tx_bit)
```

The data conveyed by IS_UNITDATA.request is a continuous stream of encoded bits. The tx_bit parameter can take one of two values: one or zero.

105.4.3.1.2 When generated

The sublayer continuously sends a bit stream IS_UNITDATA.request(tx_bit) to the next lower sublayer, at a nominal signaling rate of 25.78125 GBd.

105.4.3.1.3 Effect of receipt

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

105.4.3.2 IS_UNITDATA.indication

The IS_UNITDATA.indication primitive is used to define the transfer of a stream of data units from the sublayer to the next higher sublayer.

105.4.3.2.1 Semantics of the service primitive

```
IS_UNITDATA.indication(rx_bit)
```

The data conveyed by IS_UNITDATA.indication is a continuous stream of encoded bits. The rx_bit parameters can take one of two values: one or zero.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation




Figure 105–2—25GBASE-R inter-sublayer service interfaces


105.4.3.2.2 When generated

The sublayer continuously sends a bit stream IS_UNITDATA.indication(rx_bit) to the next higher sublayer, at a nominal signaling rate of 25.78125 GBd.

105.4.3.2.3 Effect of receipt

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

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 105-3—Optional inter-sublayer service interfaces for EEE deep sleep support

105.4.3.3 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.

105.4.3.3.1 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_bit 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.

105.4.3.3.2 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.

105.4.3.3.3 Effect of receipt

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

105.4.3.4 IS_TX_MODE.request

The IS_TX_MODE.request primitive communicates the tx_mode parameter generated by the PCS Transmit Process for EEE capability to invoke the appropriate PMA, FEC, and PMD transmit EEE states. Without EEE deep sleep mode capability, the primitive is never invoked and the sublayers behave as if tx_mode = DATA.

105.4.3.4.1 Semantics of the service primitive

IS_TX_MODE.request(tx_mode)

The tx_mode parameter takes on one of up to three values: DATA, QUIET, or ALERT.

105.4.3.4.2 When generated

This primitive is generated to indicate the low power mode of the transmit path.

105.4.3.4.3 Effect of receipt

The specific effect of receipt of this primitive is defined by the sublayer that receives this primitive. In general, when tx_mode is DATA, the sublayer operates normally and when tx_mode is QUIET, the sublayer may go into a low power mode.

105.4.3.5 IS_RX_MODE.request

The IS_RX_MODE.request primitive communicates the rx_mode parameter generated by the PCS LPI receive function to other sublayers.

Without EEE deep sleep mode capability, the primitive is never invoked and the sublayers behave as if rx_mode = DATA.

105.4.3.5.1 Semantics of the service primitive

IS_RX_MODE.request(rx_mode)

The rx_mode parameter takes on one of two values: DATA or QUIET.

105.4.3.5.2 When generated

This primitive is generated to indicate the state of the PCS LPI receive function.

105.4.3.5.3 Effect of receipt

The specific effect of receipt of this primitive is defined by the sublayer that receives this primitive. In general, when rx_mode is DATA, the sublayer operates normally and when rx_mode is QUIET, the sublayer may go into a low power mode.

105.4.3.6 IS_RX_LPI_ACTIVE.request

The IS_RX_LPI_ACTIVE.request primitive communicates to the FEC that the PCS LPI receive function is active. This primitive may be passed through a PMA sublayer but has no effect on that sublayer. This primitive is only used for a PMA sublayer that is between the PCS and a Clause 74 FEC sublayer; in all other cases the primitive is never invoked and has no effect. Without EEE deep sleep mode capability, the primitive is never invoked and has no effect.

105.4.3.6.1 Semantics of the service primitive

IS_RX_LPI_ACTIVE.request(rx_lpi_active)

The parameter rx_lpi_active is Boolean.

105.4.3.6.2 When generated

This primitive is generated to indicate the state of the PCS LPI receive function. It is FALSE when in the RX_ACTIVE state and TRUE in all other states.

105.4.3.6.3 Effect of receipt

The specific effect of receipt of this primitive is defined by the FEC sublayer that receives this primitive. When rx_lpi_active is true, the FEC sublayer uses rapid block lock to reestablish FEC operation following a period of quiescence.

105.4.3.7 IS_ENERGY_DETECT.indication

The IS_ENERGY_DETECT.indication primitive is used to communicate that the PMD has detected the return of energy on the interface following a period of quiescence. Without EEE deep sleep mode capability, the primitive is never invoked and has no effect.

105.4.3.7.1 Semantics of the service primitive

IS_ENERGY_DETECT.indication(energy_detect)

The parameter energy_detect is Boolean.

105.4.3.7.2 When generated

This primitive is generated by the PMA, reflecting the state of the signal_detect parameter received from the PMD.

105.4.3.7.3 Effect of receipt

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

This parameter is used to indicate that activity has returned on the interface following a period of quiescence.

105.4.3.8 IS_RX_TX_MODE.indication

The IS_RX_TX_MODE.indication primitive communicates the rx_tx_mode parameter. This parameter indicates the value of tx_mode that the PMA sublayer has inferred from the received signal. Without EEE deep sleep capability, the primitive is never generated and the sublayers behave as if rx_tx_mode = DATA.

105.4.3.8.1 Semantics of the service primitive

IS_RX_TX_MODE.indication(rx_tx_mode)

The parameter rx_tx_mode is assigned one of the following values: DATA, QUIET, or ALERT.

105.4.3.8.2 When generated

This primitive is generated whenever there is change in the value of the rx_tx_mode parameter.

105.4.3.8.3 Effect of receipt

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

105.5 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 105–3 contains 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. 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.

Equation (105–1) specifies the calculation of bit time per meter of fiber or electrical cable based upon the parameter n , which represents the ratio of the speed of electromagnetic propagation in the fiber or electrical cable to the speed of light in a vacuum. The value of n should be available from the fiber or electrical cable manufacturer, but if no value is known then a conservative delay estimate can be calculated using a default value of $n = 0.66$. The speed of light in a vacuum is $c = 3 \times 10^8$ m/s. Table 44–3 can be used to convert fiber or electrical cable delay values specified relative to the speed of light or in nanoseconds per meter.

$$\text{Cable delay} = \frac{25 \times 10^9}{nc} \text{ BT/m} \quad (105-1)$$

See 31B.3.7 for PAUSE reaction timing constraints for stations at operating speeds of 25 Gb/s.

105.6 State diagrams

State diagrams take precedence over text.

The conventions of 1.2 are adopted, along with the extensions listed in 21.5.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

Table 105–3—Sublayer delay constraints

Sublayer	Maximum (bit time) ^a	Maximum (pause_quanta) ^b	Maximum (ns)	Notes ^c
25G RS, MAC, and MAC Control	8192	16	327.68	See 106.1.4.
25GBASE-R PCS	3584	7	143.36	See 107.3.
25G BASE-R FEC	6144	12	245.76	See 74.6.
25GBASE-R RS-FEC	24576	48	983.04	See 108.4.
25GBASE-R PMA ^d	4096	8	163.84	See 109.5.
25GBASE-CR PMD	512	1	20.48	See 110.4.
25GBASE-CR-S PMD	512	1	20.48	See 110.4.
25GBASE-KR PMD	512	1	20.48	See 111.4.
25GBASE-KR-S PMD	512	1	20.48	See 111.4.
25GBASE-SR PMD	512	1	20.48	See 112.3.

^a1 bit time (BT) is equal to 40 ps. (See 1.4.117 for the definition of bit time.)

^b1 pause_quantum is equal to 20.48 ns. (See 31B.2 for the definition of pause_quantum.)

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

^dCumulative round-trip delay contributed by up to four PMA stages in a PHY.

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.

105.7 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 45, Clause 73, Clause 74, Clause 106 through Clause 112, 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 25 Gigabit Ethernet PICS conforms to the same notation and conventions used in 2.1.6.

106. Reconciliation Sublayer (RS) and Media Independent Interface (25GMII) for 25 Gb/s operation

106.1 Overview

This clause defines the characteristics for the Reconciliation Sublayer (RS) and the 25 Gigabit Media Independent Interface (25GMII) between the Ethernet MAC and PHY. Figure 106–1 shows the relationship of the RS and 25GMII to the ISO/IEC (IEEE) OSI reference model. The 25 Gb/s RS has identical logical functionality to the 10 Gb/s RS defined in Clause 46. A physical implementation and associated electrical characteristics for the 25GMII are not defined.



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE
 AN = AUTO-NEGOTIATION
 FEC = FORWARD ERROR CORRECTION
 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

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 106–1—RS and 25GMII relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

The 25GMII is an optional logical interface between the MAC sublayer and the Physical Layer.

The RS adapts the bit serial protocols of the MAC to the parallel format of the PCS service interface. Though the 25GMII is an optional interface, it is used extensively in this standard as a basis for specification. The 25 Gb/s Physical Coding Sublayer (PCS) is specified to the 25GMII, so if not implemented, a conforming implementation behaves functionally as if the RS and 25GMII were implemented.

The 25GMII has the following characteristics:

- a) It is capable of supporting 25 Gb/s operation.
- b) Data and delimiters are synchronous to clock reference.
- c) It provides independent 32-bit-wide transmit and receive data paths.
- d) It provides for full duplex operation only.

106.1.1 Summary of major concepts

The following are the major concepts of 25GMII:

- a) The 25GMII is functionally similar to other media independent interfaces that have been defined for other 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 25GMII.
- c) The RS maps the signal set provided at the 25GMII 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 25GMII may also support Low Power Idle (LPI) signaling for PHY types supporting Energy-Efficient Ethernet (EEE) (see Clause 78).

106.1.2 Application

This logical interface is used to provide media independence so that an identical MAC may be used with all 25GBASE PHY types.

106.1.3 Rate of operation

The 25GMII is specified to support 25 Gb/s operation.

106.1.4 Delay constraints

The maximum cumulative MAC Control, MAC, and RS round-trip (sum of transmit and receive) delay shall meet the values specified in Table 106–1. Bit time is defined in 1.4, and pause_quanta is defined in 31B.2.

Table 106–1—Round-trip delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)
MAC, RS, and MAC Control	8192	16

106.1.5 Allocation of functions

The allocation of functions at the 25GMII balances the need for media independence with interface simplicity. The 25GMII maximizes media independence by separating the Data Link and Physical Layers of the OSI seven-layer reference model.

106.1.6 25GMII structure

The 25GMII structure is identical to the XGMII structure specified in 46.1.6.

106.1.7 Mapping of 25GMII signals to PLS service primitives

The Reconciliation Sublayer (RS) shall map the signals provided at the 25GMII 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 25 Gb/s; therefore, PLS service primitives supporting CSMA/CD operation are not mapped through the RS to the 25GMII. The mapping is changed if EEE capability is supported (see 78.3). This behavior and restrictions are the same as described in 22.7, with the details of the signaling described in 106.3. 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 when the PHY is in its low power state.

Mappings for the following primitives are defined for 25 Gb/s operation:

PLS_DATA.request
PLS_DATA.indication
PLS_CARRIER.indication
PLS_SIGNAL.indication
PLS_DATA_VALID.indication

106.1.7.1 Mapping of PLS_DATA.request

The RS maps the primitive PLS_DATA.request to the 25GMII signals TXD<31:0>, TXC<3:0>, and TX_CLK in the same way as for the XGMII as specified in 46.1.7.1.

106.1.7.2 Mapping of PLS_DATA.indication

The RS maps the primitive PLS_DATA.indication to the 25GMII signals RXD<31:0>, RXC<3:0> and RX_CLK in the same way as for the XGMII as specified in 46.1.7.2.

106.1.7.3 Mapping of PLS_CARRIER.indication

25 Gb/s operation supports full duplex operation only. The RS never generates the PLS_CARRIER.indication primitive for PHYs that do not support EEE or Link Interruption.

For PHYs that support EEE capability, CARRIER_STATUS is set in the same way as specified in 46.1.7.3.

106.1.7.4 Mapping of PLS_SIGNAL.indication

25 Gb/s operation supports full duplex operation only. The RS never generates the PLS_SIGNAL.indication primitive.

106.1.7.5 Mapping of PLS_DATA_VALID.indication

The RS maps the primitive PLS_DATA_VALID.indication to the 25GMII signals RXC<3:0> and RXD<31:0> in the same way as for the XGMII as specified in 46.1.7.5.

106.2 25GMII data stream

The 25GMII data stream has the same characteristics as the XGMII data stream described in 46.2.

106.3 25GMII functional specifications

The 25GMII functions identically to the XGMII specified in 46.3 with the exception that the TX_CLK and RX_CLK frequency shall be 390.625 MHz \pm 100 ppm (also one-sixty-fourth of the MAC transmit data rate).

106.4 LPI Assertion and Detection

25 Gb/s PHYs may support Clause 78 Energy-Efficient Ethernet (see Table 105-2).

LPI signaling by the 25 Gb/s RS is identical to that defined for XGMII in 46.4.

The operation of LPI in the PHY requires that the MAC does not send valid data for a time after LPI has been de-asserted as governed by resolved Transmit T_{w_sys} defined in 78.4.2.3.

This wake-up time is enforced by the transmit LPI state diagram using CARRIER_SENSE.indication in an identical manner to that defined in 46.4.

106.5 Protocol implementation conformance statement (PICS) proforma for Clause 106 Reconciliation Sublayer (RS) and Media Independent Interface (25GMII) for 25 Gb/s operation⁵

106.5.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 106, Reconciliation Sublayer (RS) and Media Independent Interface (25GMII) for 25 Gb/s operation, 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](#).

106.5.2 Identification

106.5.2.1 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).	

106.5.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3by-2016, Clause 106, Reconciliation Sublayer (RS) and Media Independent Interface (25GMII) for 25 Gb/s operation
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 802.3by-2016.)	
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.

106.5.2.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PHY	PHY support of 25GMII	106.2, 106.3		O	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
*RS	Reconciliation Sublayer support of 25GMII	106.2, 106.3		O	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
*LPI	Implementation of LPI	106.1.7		O	Yes [<input type="checkbox"/> No [<input type="checkbox"/>

106.5.3 PICS proforma Tables for Reconciliation Sublayer and 25 Gigabit Media Independent Interface

106.5.3.1 General

Item	Feature	Subclause	Value/Comment	Status	Support
G1	PHY support of MAC data rate	106.1.3	Support MAC data rate of 25 Gb/s	PHY:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>
G2	Cumulative MAC Control, MAC and RS round-trip delay	106.1.4	Per Table 106–1	RS:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>

106.5.3.2 Mapping of PLS service primitives

Item	Feature	Subclause	Value/Comment	Status	Support
PL1	Mapping to Clause 6	106.1.7	RS implements mapping to Clause 6 PLS service primitives	RS:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>

106.5.3.3 25GMII signal functional specifications.

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	TX_CLK active edges	106.3	TXD and TXC sampled on both edges of TX_CLK	RS:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>
FS2	TX_CLK frequency	106.3	390.625 MHz ± 100 ppm	RS:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>
FS3	RX_CLK active edges	106.3	RXD and RXC sampled on both edges of RX_CLK	RS:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>
FS4	RX_CLK frequency	106.3	390.625 MHz ± 100 ppm when received data rate is within tolerance	RS:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>

107. Physical Coding Sublayer (PCS) for 64B/66B, type 25GBASE-R

107.1 Overview

107.1.1 Scope

This clause specifies the Physical Coding Sublayer (PCS) that is common to a family of 25 Gb/s Physical Layer implementations known as 25GBASE-R. The 25GBASE-R PCS is a sublayer of the 25 Gb/s PHYs listed in Table 105–1. The term 25GBASE-R is used when referring generally to Physical Layers using the PCS defined in this clause.

107.1.2 Relationship of 25GBASE-R to other standards

Figure 107–1 depicts the relationships among the 25GBASE-R sublayers, the Ethernet MAC and reconciliation layers, and the higher layers.

The 25GBASE-R PCS is identical to the 10GBASE-R PCS specified in Clause 49 with three exceptions:

- 1) It has a single-bit interface to the Physical Medium Attachment (PMA) sublayer rather than the 16-bit interface described in Clause 49.
- 2) It has the ability to generate the scrambled idle test pattern.
- 3) `hi_ber` is asserted if `ber_cnt` reaches 97 in a 2 ms period. This differs from the definition in 49.2.13.3, which defines `hi_ber` as occurring if `ber_cnt` reaches 16 in a 125 μ s period.

107.1.3 Summary of 25GBASE-R sublayers

Figure 107–1 shows the relationship of the 25GBASE-R PCS sublayer with other sublayers to the ISO Open System Interconnection (OSI) reference model.

107.1.3.1 Physical Coding Sublayer (PCS)

The PCS service interface is the 25GMII, which is defined in Clause 106. The 25GMII is identical to the XGMII, which is the service interface of the 10GBASE-R PCS, but operates with 390.625 MHz clocks rather than 156.25 MHz clocks. The 25GMII provides a uniform interface to the Reconciliation Sublayer for all 25 Gb/s PHY implementations.


The 25GBASE-R PCS provides all services required by the 25GMII.

107.1.4 Inter-sublayer interfaces

The upper interface of the PCS may connect to the Reconciliation Sublayer through the 25GMII. The lower interface of the PCS connects to the PMA sublayer to support a PMD. If an FEC sublayer is implemented and the optional physical instantiation of 25GAUI is not implemented directly below the PCS sublayer, then the lower PCS interface connects to the FEC sublayer. The 25GBASE-R PCS has a nominal rate at the PMA (or FEC) service interface of 25.78125 Gb/s, which provides capacity for the MAC data rate of 25 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.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE
 AN = AUTO-NEGOTIATION
 FEC = FORWARD ERROR CORRECTION
 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

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 107–1—25GBASE-R PCS relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

107.1.4.1 PCS service interface (25GMII)

The PCS service interface allows the 25GBASE-R PCS to transfer information to and from a PCS client. The PCS client is the Reconciliation Sublayer. The PCS service interface is defined as the 25GMII in Clause 106.

107.1.4.2 Physical Medium Attachment (PMA) service interface

The PMA (or FEC) service interface for the PCS is described in an abstract manner and does not imply any particular implementation. The PMA (or FEC) Service Interface supports the exchange of encoded data between the PCS and PMA (or FEC) sublayer. The PMA (or FEC) service interface is single-bit and defined in 109.2 and is an instance of the inter-sublayer service interface definition in 105.4.

107.2 Functions within the PCS

The 25GBASE-R PCS shall have all the functionality of the 10GBASE-R PCS specified in Clause 49. In addition, the PCS has the ability to generate the scrambled idle test-pattern generator specified in 107.2.3.

The BER monitor state diagram shown in Figure 49–15 still applies but it shall use a 2 ms timer instead of a 125 μs timer and ber_cnt is tested for a value of 97 rather than 16 in the exit conditions from state BER_BAD_SH. So:

- 1) The definition of “125us_timer” in 49.2.13.2.5 is replaced with “Timer that is triggered every 2 ms +1%, -25%”.
- 2) The definition of “ber_cnt” in 49.2.13.2.4 is replaced with “Count up to a maximum of 97 of the number of invalid sync headers within the current 2 ms period”.
- 3) The definition of “hi_ber” in 49.2.13.2.2 is replaced with “Boolean variable that is asserted true when the ber_cnt reaches 97 indicating a bit error ratio $>10^{-4}$ ”.

The PCS encodes data and control information into 66-bit blocks. The relationship of block bit positions to the 25GMII, PMA, and other PCS constructs is illustrated in Figure 107–2 for transmit and Figure 107–3 for receive.




Figure 107–2—PCS Transmit bit ordering

107.2.1 Notation conventions

Values represented in binary are shown with the first transmitted bit (the LSB) on the left.

Two consecutive 25GMII transfers provide eight characters that are encoded into one 66-bit transmission block.

107.2.2 Transmission order

Block bit transmission order is illustrated in Figure 107–2 and Figure 107–3. Note that these figures show the mapping from 25GMII to 64B/66B block for a block containing eight data characters.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation




Figure 107-3—PCS Receive bit ordering

107.2.3 Test-pattern generator

In addition to those patterns specified in 49.2.8, the PCS shall have the ability to generate a scrambled idle test pattern. This test pattern is suitable for receiver tests and for certain transmitter tests.

When enabled, the scrambled idle test pattern is generated by the scrambler. The input to the scrambler is a control block (block type=0x1E) with all idles as defined in Figure 49-7. Sync headers are added to the scrambled data output.

If a Clause 45 MDIO is implemented, then control of the test-pattern generation is from the BASE-R PCS test-pattern control register (3.42).

107.3 LPI

If the 25GBASE-R PCS is part of a PHY configured for EEE deep sleep operation, the PCS shall follow the state diagrams specified in Figure 49-12 and Figure 49-13.

The LPI functions shall use the timer values in Table 107-1 and Table 107-2 for EEE deep sleep operation.

If the 25GBASE-R PCS is part of a PHY configured for EEE fast wake operation, the PCS shall encode and decode LPI when indicated but the state diagrams specified in Figure 49-12 and Figure 49-13 do not apply. Management functions may use MDIO register bit LPI_FW 3.20.0 to select fast wake operation (see 45.2.3.9.11).

Table 107–1—Transmitter LPI timing parameters

Parameter	Description	Min	Max	Units
T _{SL}	Local Sleep Time from entering the TX_SLEEP state to when tx_mode is set to QUIET	4.9	5.1	μs
T _{QL}	Local Quiet Time from when tx_mode is set to QUIET to entry into the TX_ALERT state	1.7	1.8	ms
T _{WL}	Time spent in the TX_WAKE state	10.9	11.1	μs
T _{IU}	Time spent in the TX_ALERT and TX_SCR_BYPASS states	1.1	1.3	μs

Table 107–2—Receiver LPI timing parameters

Parameter	Description	Min	Max	Units
T _{QR}	The time the receiver waits for energy_detect to be set to TRUE while in the RX_SLEEP and RX_QUIET states before asserting receive fault	2	3	ms
T _{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault (when scr_bypass_enable = FALSE)	—	11.5	μs
T _{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault (when scr_bypass_enable = TRUE)	—	13.7	μs
T _{WTF}	Wake time fault recovery time	—	10	ms

107.3 Delay constraints

The maximum delay contributed by the 25GBASE-R PCS (sum of transmit and receive delays at one end of the link) shall be no more than 3584 BT (7 pause_quanta or 143.36 ns). A description of overall system delay constraints can be found in 105.5.

107.4 Support for Auto-Negotiation

A PCS used with a 25GBASE-KR PMD, 25GBASE-KR-S PMD, 25GBASE-CR PMD, or 25GBASE-CR-S PMD has the same requirements to support auto-negotiation as a PCS used with a 10GBASE-KR PMD. These are specified in 49.2.16.

107.5 Protocol implementation conformance statement (PICS) proforma for Clause 107, Physical Coding Sublayer (PCS) for 64B/66B, type 25GBASE-R⁶

107.5.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 107, Physical Coding Sublayer (PCS) for 64B/66B, type 25GBASE-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](#).

107.5.2 Identification

107.5.2.1 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).	

107.5.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3by-2016, Clause 107, Physical Coding Sublayer (PCS) for 64B/66B, type 25GBASE-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 802.3by-2016.)	
Date of Statement	

⁶Copyright 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.

107.5.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
25GE	25GMII compatibility interface	106	Compatibility interface is supported	O	Yes [] No []
MD	MDIO	45, 49.2.14	Registers and interface supported	O	Yes [] No []
*LPI	Implementation of LPI	107.3		O	Yes [] No []

107.5.4 25G PCS

107.5.4.1 Clause 49 functionality

Item	Feature	Subclause	Value/Comment	Status	Support
PCS1	Supports Clause 49 functionality	107.2		M	Yes [] No []
PCS2	BER monitor runs over a longer window than specified in Clause 49	107.2		M	Yes [] No []

107.5.4.2 Test-pattern generator

Item	Feature	Subclause	Value/Comment	Status	Support
TP1	Scrambled idle test-pattern ability	107.2.3		M	Yes [] No []

107.5.4.3 LPI

Item	Feature	Subclause	Value/Comment	Status	Support
LP1	EEE deep sleep	107.3, 49	PHY configured for deep sleep operation	LPI:O	Yes [] No [] N/A []
LP2	Use the correct timer values	107.3	Uses values from Table 107-1 and Table 107-2	LPI:O	Yes [] No [] N/A []
LP3	EEE fast wake	107.3	Fast wake operation	LPI:M	Yes [] No []

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

107.5.4.4 Delay Constraints

Item	Feature	Subclause	Value/Comment	Status	Support
TIM1	PCS Delay Constraint	107.3	No more than 3584 BT for sum of transmit and receive path delays	M	Yes [] No []

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

108. Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs


108.1 Overview

108.1.1 Scope

This clause specifies a Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs. The specification is closely related to that of the RS-FEC sublayer for 100GBASE-R PHYs, specified in [Clause 91](#). [Annex 91A](#) provides examples of RS-FEC codewords constructed with the method specified in this clause.

108.1.2 Position of RS-FEC in the 25GBASE-R PHY sublayers

Figure 108–1 shows the relationship of the 25GBASE-R RS-FEC sublayer to the ISO/IEC Open System Interconnection (OSI) reference model.



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE
 AN = AUTO-NEGOTIATION
 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
 RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 108–1—25GBASE-R RS-FEC relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

108.2 FEC service interface

This subclause specifies the services provided by the 25GBASE-R RS-FEC sublayer. The service interface is described in an abstract manner and does not imply any particular implementation.

The FEC service interface is provided to allow the PCS to transfer information to and from the 25GBASE-R RS-FEC. The PCS may be connected to the 25GBASE-R RS-FEC using an optional 25GAUI chip-to-chip (C2C) instantiation of the PMA service interface (see Annex 109A), in which case a PMA is the client of the FEC service interface.

The FEC service interface is an instance of the inter-sublayer service interface defined in 105.4. The FEC service interface primitives are summarized as follows:

```
FEC:IS_UNITDATA.request(tx_bit)
FEC:IS_UNITDATA.indication(rx_bit)
FEC:IS_SIGNAL.indication
```

The PCS (or PMA) continuously sends a bit stream to the 25GBASE-R RS-FEC using the FEC:IS_UNITDATA.request(tx_bit) primitive, at a nominal signaling rate of 25.78125 GBd.

The 25GBASE-R RS-FEC continuously sends a bit stream to the PCS (or PMA) using the FEC:IS_UNITDATA.indication(rx_bit) primitive, at a nominal signaling rate of 25.78125 GBd. The actual signaling rate is equal to the underlying PMD signaling rate.

The SIGNAL_OK parameter of the FEC:IS_SIGNAL.indication primitive can take one of two values: OK or FAIL. The value is set to OK when the RS-FEC receive function has identified codeword boundaries as indicated by FEC_align_status equal to true. That value is set to FAIL when the RS-FEC receive function is unable to reliably establish codeword boundaries as indicated by FEC_align_status equal to false.

If the optional EEE deep sleep capability is supported, then the FEC service interface includes four additional primitives as follows:

```
FEC:IS_TX_MODE.request
FEC:IS_RX_MODE.request
FEC:IS_RX_TX_MODE.indication
FEC:IS_ENERGY_DETECT.indication
```

When the tx_mode parameter of the FEC:IS_TX_MODE.request primitive is QUIET or ALERT, the RS-FEC sublayer may disable transmit functional blocks to conserve energy. Otherwise the RS-FEC transmit function operates normally. The value of tx_mode is passed to the client sublayer via the PMA:IS_TX_MODE.request primitive.

When the rx_mode parameter of the FEC:IS_RX_MODE.request primitive is QUIET, the RS-FEC sublayer may disable receive functional blocks to conserve energy. Otherwise the RS-FEC receive function operates normally. The value of rx_mode is passed to the client sublayer via the PMA:IS_RX_MODE.request primitive.

The rx_tx_mode parameter of the FEC:IS_RX_TX_MODE.indication primitive is used to communicate the link partner's value of tx_mode as inferred by the PMA. It is assigned the value that is received via the PMA:IS_RX_TX_MODE.indication primitive.

The energy_detect parameter of the FEC:IS_ENERGY_DETECT.indication primitive is used to communicate that the PMD has detected the return of energy on the interface following a period of quiescence. It is assigned the value that is received via the PMA:IS_ENERGY_DETECT.indication primitive.

108.3 PMA compatibility

The 25GBASE-R RS-FEC sublayer is a client of the 25GBASE-R PMA sublayer defined in Clause 109.

When 25GAUI C2C is used between a device that includes a PCS and a device that includes the RS-FEC, the 25GBASE-R PMA sublayer is the client of the 25GBASE-R RS-FEC sublayer.

108.4 Delay constraints

The maximum delay contributed by the 25GBASE-R RS-FEC sublayer (sum of transmit and receive delays at one end of the link) shall be no more than 24576 bit times (48 pause_quanta or 983.04 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 105.5.

108.5 Functions within the 25GBASE-R RS-FEC sublayer

108.5.1 Functional block diagram

A functional block diagram of the 25GBASE-R RS-FEC sublayer is shown in Figure 108–2.

108.5.2 Transmit function

108.5.2.1 Block synchronization

The RS-FEC transmit function uses the bit stream provided by the FEC:IS_UNITDATA.request primitive. It obtains lock to the 66-bit blocks in the bit stream using the sync headers and outputs 66-bit blocks. Block lock is obtained as specified in the lock state diagram shown in Figure 49–14.

108.5.2.2 Rate compensation for codeword markers in the transmit direction

The RS-FEC transmit process periodically inserts codeword markers into the transcoded block stream (see 108.5.2.4). In order to maintain the same bit rate after codeword marker insertion, the RS-FEC transmit process shall perform the rate compensation function described below, or its functional equivalent:

- a) Decode the PCS blocks received by descrambling (see 49.2.10) and applying the PCS receive process (see 49.2.11) to obtain the 25GMII character stream.
- b) Delete Idle control characters (/I/), Low Power Idle control characters (/LI/), and ordered sets, according to the rules in 49.2.4.7 and 49.2.4.10, to create room as necessary for the periodically occurring codeword markers.
- c) Re-encode the data stream obtained, by applying the PCS transmit process (see 49.2.5) and scrambler (see 49.2.6) to obtain 64B/66B blocks.

If the optional EEE deep sleep capability is supported, when rapid codeword markers are sent, characters are deleted more frequently in item b) as necessary to maintain a constant bit rate.

NOTE—When rapid codeword markers are sent, the PCS LPI transmit state diagram (Figure 49–12) is in the TX_WAKE state, so the data encoded by the PCS comprises one of two control characters, either /I/ or /LI/. Both types of control characters can be deleted for rate compensation. See 78.1.3.1 and 78.1.3.3.1 for more details about LPI wake cycles.

108.5.2.3 64B/66B to 256B/257B transcoder

The transcoder constructs a 257-bit block, tx_scrambled<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.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation




Figure 108-2—Functional block diagram

The construction of the transcoded block shall be according to the process described in 91.5.2.5. For each 257-bit block, bit 0 shall be the first bit transmitted.

108.5.2.4 Codeword marker insertion

In order to support codeword alignment in the receive direction, the 25GBASE-R RS-FEC shall periodically insert codeword markers into the stream of transcoded blocks as the first 257 bits of every 1024th RS-FEC codeword (see 108.5.2.5). The distance between the beginning of successive codeword markers is therefore 20480 257-bit transcoded blocks, equivalent to 81920 64B/66B blocks.

Room for codeword markers is created by the rate compensation for codeword markers in the transmit direction process (see 108.5.2.2) such that the bit rates at the input and the output of the 25GBASE-R RS-FEC sublayer are equal.

The transmitted codeword marker is a 257-bit block, tx_cwm, constructed of four alignment markers followed by a zero bit. Each alignment marker is built from eight octets M_0 , M_1 , M_2 , BIP_3 , M_4 , M_5 , M_6 , and BIP_7 with the bit order shown in Figure 82–9. The BIP_3 field is set to the constant value 0x33 and the BIP_7 field is set to the constant value 0xCC.

tx_cwm is constructed as follows:

- a) tx_cwm<63:0> are set to bits <65:2> of the alignment marker of PCS lane 0 defined in Table 82–2.
- b) tx_cwm<127:64> are set to bits <65:2> of the alignment marker of PCS lane 1 defined in Table 82–3.
- c) tx_cwm<191:128> are set to bits <65:2> of the alignment marker of PCS lane 2 defined in Table 82–3.
- d) tx_cwm<255:192> are set to bits <65:2> of the alignment marker of PCS lane 3 defined in Table 82–3.
- e) tx_cwm<256> is set to 0.

For the optional EEE deep sleep capability, a rapid method of codeword alignment is used following a transition of tx_mode from ALERT to DATA, by transmitting rapid codeword markers. When the value of tx_down_count is nonzero, rapid codeword markers are sent at the beginning of each codeword. Rapid codeword markers are identical to regular codeword markers, with the exception that BIP_3 is set to the value of tx_down_count instead of the constant 0x33 and BIP_7 is set to the bit-wise inversion of tx_down_count instead of the constant 0xCC.

108.5.2.5 Reed-Solomon encoder

The 25GBASE-R RS-FEC sublayer employs the Reed-Solomon code RS(528,514) operating over the Galois Field $GF(2^{10})$ where the symbol size is 10 bits. The encoder described in 91.5.2.7 shall be used.

108.5.2.6 Codeword serialization

Once the data has been Reed-Solomon encoded, it shall be serialized and sent to the PMA using the PMA:IS_UNITDATA.request primitive with the transmit bit ordering illustrated in Figure 108–3.

108.5.2.7 RS-FEC encoding for rapid codeword lock (EEE deep sleep)

If the optional EEE deep sleep capability is supported, the RS-FEC transmit function changes its signaling during LPI refresh and wake periods, as described in this subclause.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation




Figure 108-3—Transmit bit ordering

The transition of the PCS LPI transmit state diagram (Figure 49–12) to TX_WAKE state is detected by the change of the tx_mode parameter of the FEC:IS_TX_MODE.request primitive from ALERT to DATA. Following this transition, the transmit function behaves as follows:

- Set tx_down_count to 40. This causes the codeword marker insertion function (108.5.2.4) to insert a rapid codeword marker in the beginning of each of the following 40 codewords. The value of tx_down_count is encoded in each rapid codeword marker and is decremented after the rapid codeword marker is sent.
- The rapid codeword marker with down_count=1 sets the location of the subsequent regular codeword markers, such that a codeword marker is inserted at the beginning of the 1024th codeword after it. Figure 108–4 depicts the codeword marker locations during this transition.




Figure 108–4—Transition from rapid to normal codeword markers

108.5.3 Receive function

108.5.3.1 Codeword marker lock

The 25GBASE-R RS-FEC shall implement the codeword marker lock process as described in this subclause.

The RS-FEC receive function forms a bit stream by concatenating the bits from the PMA:IS_UNITDATA.indication primitive in the order they are received. This process obtains lock to the codeword markers as shown in FEC synchronization state diagram (Figure 108–6). The status of the codeword marker lock process is reflected by the state variable FEC_align_status.

108.5.3.2 Reed-Solomon decoder

The Reed-Solomon decoder extracts the message symbols from the codeword, corrects them as necessary, and discards the parity symbols. The message symbols correspond to a codeword comprising 20 transcoded blocks rx_scrambled.

The 25GBASE-R RS-FEC sublayer shall be capable of correcting any combination of up to 7 symbol errors in a codeword. The 25GBASE-R RS-FEC sublayer shall also be capable of indicating when a codeword contains errors that were not corrected. The probability that the decoder fails to indicate a codeword with 8 or more symbol errors as uncorrected is expected to be lower than 10^{-6} .

The Reed-Solomon decoder may provide the option to perform error detection without error correction to reduce the delay contributed by the 25GBASE-R RS-FEC sublayer. The presence of this option is indicated by the assertion of the `FEC_bypass_correction_ability` variable (see 108.6.4). When the option is provided, it is enabled by the assertion of the `FEC_bypass_correction_enable` variable (see 108.6.1). This option shall not be used when the 25GBASE-R RS-FEC sublayer is used to form part of a 25GBASE-SR PHY.

NOTE 1—The PHY may rely on the error correction capability of the 25GBASE-R RS-FEC sublayer to achieve its performance objectives. It is recommended that acceptable performance of the underlying link is verified before error correction is bypassed.

The Reed-Solomon decoder shall indicate errors to the PCS sublayer by intentionally corrupting 66-bit block synchronization headers. When the decoder determines that a codeword contains errors (when the bypass correction feature is enabled) or contains errors that were not corrected (when the bypass correction feature is not supported or not enabled), it ensures that, for every other 257-bit block within the codeword starting with the first (1st, 3rd, 5th, etc.), the synchronization header for the first 66-bit block at the output of the 256B/257B to 64B/66B transcoder, `rx_coded_0<1:0>`, is set to 11. In addition, it ensures that `rx_coded_0<1:0>` corresponding to the second 257-bit block and `rx_coded_3<1:0>` corresponding to the last (20th) 257-bit block in the codeword are set to 11.

NOTE 2—Setting `rx_coded_0<1:0>` to 11 as described causes the PCS to assign `R_BLOCK_TYPE=E` to the block and decode its content as `EBLOCK_R` (see 49.2.13.2.1 and 49.2.13.2.3). This will cause all frames 64 bytes and larger that are fully or partially within the codeword to be discarded.

The Reed-Solomon decoder may optionally provide the ability to bypass the error indication feature to reduce the delay contributed by the 25GBASE-R RS-FEC sublayer. The presence of this option is indicated by the assertion of the `FEC_bypass_indication_ability` variable (see 108.6.5). When the option is provided it is enabled by the assertion of the `FEC_bypass_indication_enable` variable (see 108.6.2).

When `FEC_bypass_correction_enable` is asserted, the decoder shall not bypass error indication and the value of `FEC_bypass_indication_enable` has no effect.

When `FEC_bypass_indication_enable` is asserted, additional error monitoring is performed by the 25GBASE-R RS-FEC sublayer to reduce the likelihood that errors in a packet are not detected. The Reed-Solomon decoder counts the number of symbol errors detected in consecutive non-overlapping blocks of 8192 codewords. When the number of symbol errors in a block of 8192 codewords exceeds 417, the Reed-Solomon decoder shall cause synchronization header `rx_coded<1:0>` of each subsequent 66-bit block that is delivered to the PCS to be assigned a value of 00 or 11 for a period of 60 ms to 75 ms.

NOTE 3—This setting of `rx_coded<1:0>` marks the 64B/66B block as bad and inhibits processing of received packets by the PCS. The BER monitor state diagram (Figure 49–15) sets `hi_ber` to true and the PCS `block_lock` is set to false. When AN is supported and enabled, this event causes AN to restart.

For the optional EEE deep sleep capability, the error monitor employed when `FEC_bypass_indication_enable` is asserted shall be disabled when `rx_lpi_active` is true. The next block of 8192 codewords considered by the error monitor begins on the codeword boundary following the transition of `rx_lpi_active` from true to false.

108.5.3.3 Codeword monitor

After codeword marker lock has been achieved, this process continuously checks codeword validity as indicated by the codeword monitor state diagram (Figure 108–7). When three consecutive uncorrected codewords are detected, the codeword monitor shall restart the codeword marker lock process. In such event, it is likely that multiple blocks are marked as bad until codeword marker alignment is found.

NOTE—Marking multiple 64B/66B blocks as bad causes the PCS to assert `hi_ber` and lose its block lock. When AN is supported and enabled, this event causes AN to restart.

108.5.3.4 Codeword marker removal

The first 257 message bits in every 1024th codeword is the vector $rx_cwm\langle 256:0 \rangle$ where bit 0 is the first bit received. The specific codewords that include this vector are indicated by the codeword marker lock function (108.5.3.1).

For the optional EEE deep sleep capability, when FEC_align_status is true and rx_down_count is nonzero, $rx_cwm\langle 256:0 \rangle$ is set to the first 257 message bits in every codeword.

The vector rx_cwm shall be removed prior to transcoding.

108.5.3.5 256B/257B to 64B/66B transcoder

The transcoder extracts a group of four 66-bit blocks, $rx_coded_j\langle 65:0 \rangle$ where $j=0$ to 3, from each 257-bit block $rx_scrambled\langle 256:0 \rangle$. Bit 0 of the 257-bit block is the first bit received.

The transcoder described in 91.5.3.5 shall be used, with the exception that in step f2) $h\langle 3:0 \rangle$ is derived by cross-referencing to $g\langle 3:0 \rangle$ using Figure 49-7 instead of Figure 82-5, to account for the block types used by the 25GBASE-R PCS.

The transcoder creates a stream of 66-bit vectors $rx_coded\langle 65:0 \rangle$ from the groups of four blocks $rx_coded_j\langle 65:0 \rangle$ where $j=0$ to 3. rx_coded_0 is the first block transmitted out of each group.

108.5.3.6 Rate compensation for codeword markers in the receive direction

After the codeword markers have been discarded from transcoding and the stream of $rx_coded\langle 65:0 \rangle$ vectors has been obtained, to compensate for the deleted codeword markers, the RS-FEC receive process shall perform the rate compensation function described below, or its functional equivalent:

- Decode the stream of rx_coded vectors by descrambling (see 49.2.10) and applying the PCS receive process (see 49.2.11) to obtain the 25GMII character stream.
- Insert /I/ or /LI/, according to the rules in 49.2.4.7, to fill in as necessary for any deleted codeword markers or rapid codeword markers.
- Re-encode the data stream obtained, by applying the PCS transmit process (see 49.2.5) and scrambling (see 49.2.6) to obtain 64B/66B blocks $rx_coded_out\langle 65:0 \rangle$.

If $rx_coded\langle 1:0 \rangle$ is either 00 or 11, the process in list item c) shall set $rx_coded_out\langle 1:0 \rangle$ to $rx_coded\langle 1:0 \rangle$ and the process in list item b) shall not insert idle characters at the next block after rx_coded_out .

If the optional EEE deep sleep capability is supported, this function sets rx_lpi_active according to the RX LPI state diagram (Figure 49-13).

108.5.3.7 Rapid codeword lock for EEE deep sleep

If the optional EEE deep sleep capability is supported, the RS-FEC receive function performs rapid codeword lock during LPI refresh and wake periods, as described in this subclause.

When rx_mode (or rx_tx_mode if appropriate) transitions from QUIET to DATA:

- Start timer $hold_off_timer$.
- Enable the RS-FEC rapid codeword lock mechanism, which attempts to detect rapid codeword markers sent by the remote RS-FEC transmit function (see 108.5.2.7). When two rapid codeword markers that are one codeword distance apart are detected (causing entry to state WAKE_GOOD in

the FEC synchronization state diagram, Figure 108–6), the start location of the RS-FEC codeword is set to the start location of the rapid codeword markers. The next codeword marker position is set to 1024 codewords following the rapid codeword marker with `down_count=1`.

NOTE—The rapid codeword lock mechanism is implementation dependent and outside the scope of this standard.

When the start location is found, `FEC_align_status` is set to true and `test_cw` is asserted for each subsequent codeword, enabling aligned codeword decoding. Assuming the rapid codeword lock has determined the correct start of codeword location, the RS-FEC codeword monitor state diagram (Figure 108–7) reaches the `CW_GOOD` state, and then the decoding in item a) of 108.5.3.6 results in one of two deterministic blocks, composed of either `/I/` or `/LI/` control characters.

The RS-FEC sublayer shall hold off asserting `SIGNAL_OK` until one of the following two events occurs:

- 1) The RS-FEC codeword monitor state diagram (Figure 108–7) reaches the `CW_GOOD` state.
- 2) `hold_off_timer_done = true`.

108.5.3.8 Receive bit ordering

The receive bit ordering shall be as illustrated in Figure 108–5.

108.5.4 Detailed functions and state diagrams

108.5.4.1 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. State diagram timers follow the conventions of 14.2.3.2. The notation ++ after a counter or integer variable indicates that its value is to be incremented.

108.5.4.2 State variables

`cw_bad`

Boolean variable that is set to true if the Reed-Solomon decoder (see 108.5.3.2) detects any errors in the current codeword and does not correct them, and is set to false otherwise.

`cwm_counter_done`

Boolean variable that indicates that `cwm_counter` has reached its terminal count.

`cwm_valid`

Boolean variable that is set to true if the received block `rx_cwm<256:0>` (as defined in 108.5.3.4) is a valid codeword marker. Codeword marker validity is tested by comparing bits 23:0 and 55:32 of `rx_cwm<256:0>`, on a nibble-wise basis (12 comparisons), against their respective values in `tx_cwm<256:0>` (as defined in 108.5.2.4). If nine or more nibbles in the candidate block match the corresponding known nibbles in the codeword marker, the candidate block is considered a valid codeword marker.

`FEC_align_status`

Boolean variable that is set to true when the receiver has found the alignment of the codeword marker on the PMA service interface.

`reset`

Boolean variable that controls the resetting of the 25GBASE-R RS-FEC 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 25GBASE-R RS-FEC sublayer into low-power mode.




Figure 108-5—Receive bit ordering

restart_lock

Boolean variable that is set by the codeword monitor process to reset the synchronization process. It is set to true after 3 consecutive uncorrected codewords are received (3_BAD state) and set to false upon entry into the LOSS_OF_ALIGNMENT state.

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 synchronization state diagram has been completed indicating that the next candidate 257-bit block position can be tested.

test_cwm

Boolean variable that is set to true when a candidate block position is available for testing and is set to false according to the FEC synchronization state diagram in Figure 108-6.

test_cw

Boolean variable that is set to true when a new FEC codeword is available for decoding and is set to false according to the codeword monitor state diagram in Figure 108-7.

The following variables are only used for the optional EEE deep sleep capability. If this capability is not supported, lpi_rapid_align, rx_lpi_active, and rcwm_counter_done are set to false, and rx_down_count is set to 0.

lpi_rapid_align

Boolean variable that is set according to the FEC synchronization state diagram in Figure 108-6. Set to false when reset is true.

rx_lpi_active

Boolean variable that is set to true when the rate compensation for codeword markers in the receive direction function (108.5.3.6) infers that the Low Power Idle is being received from the link partner, and is set to false otherwise.

rcwm_counter_done

Boolean variable that indicates that rcwm_counter has reached its terminal count.

rx_down_count

The countdown value of rapid codeword markers inferred by the receiver from rx_cwm. Transmitted rapid codeword markers contain four copies of tx_down_count and four copies of its bitwise-inverse, from which this value may be extracted.

108.5.4.3 Functions

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.

108.5.4.4 Counters

cwm_counter

This counter counts the received codewords that separate the ends of two consecutive codeword markers. An RS-FEC codeword is 5280 bits. The terminal count of this counter is equal to the codeword offset between transmitted codeword markers, 1024.

cw_bad_count

Counts the number of consecutive uncorrected codewords. This counter is set to zero when a codeword is received and cw_bad is false for that codeword.

The following counters are only used for the optional EEE deep sleep capability:

tx_down_count

A counter that is used when transmitting rapid codeword markers and is decremented each time a rapid codeword marker is sent. The counter initial value is set to 40 when the tx_mode parameter of the FEC:IS_TX_MODE.request primitive changes from ALERT to DATA. The terminal count is zero.

rcwm_counter

This counter counts the received codewords that separate the ends of two consecutive rapid codeword markers. An RS-FEC codeword is 5280 bits. The terminal count of this counter is equal to one, which is the codeword offset between transmitted rapid codeword markers.

108.5.4.5 Timers


The following timer is only used for the optional EEE deep sleep capability.

hold_off_timer

This timer is started when rx_mode (or rx_tx_mode if appropriate) transitions from QUIET to DATA. The timer terminal count is 11.5 μ s. The timer is stopped when the terminal count is reached or when the FEC synchronization state diagram enters the 2_GOOD state. When the timer reaches the terminal count, it sets hold_off_timer_done = true.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

108.5.4.6 State diagrams



NOTE—Optional states (inside the dashed box) and transition A are only required to support deep sleep EEE capability.

Figure 108-6—FEC synchronization state diagram




Figure 108-7—Codeword monitor state diagram

108.6 25GBASE-R RS-FEC MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the RS-FEC. If MDIO is implemented, it shall map MDIO control bits to RS-FEC control variables as shown in Table 108-1 and MDIO status bits to RS-FEC status variables as shown in Table 108-2. If a separated PMA (see 45.2.1) is connected to the FEC service interface, the MDIO shall map additional MDIO status bits to additional RS-FEC status variables as shown in Table 108-3.

Table 108–1—MDIO/RS-FEC control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
FEC bypass correction enable	RS-FEC control register	1.200.0	FEC_bypass_correction_enable
FEC bypass indication enable	RS-FEC control register	1.200.1	FEC_bypass_indication_enable
25G RS-FEC enable	RS-FEC control register	1.200.2	FEC_Enable

Table 108–2—MDIO/RS-FEC status variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
FEC bypass correction ability	RS-FEC status register	1.201.0	FEC_bypass_correction_ability
FEC bypass indication ability	RS-FEC status register	1.201.1	FEC_bypass_indication_ability
RS-FEC high SER	RS-FEC status register	1.201.2	FEC_high_ser
RS-FEC align status	RS-FEC status register	1.201.14	FEC_align_status
RS-FEC corrected codewords counter	RS-FEC corrected codewords counter register	1.202, 1.203	FEC_corrected_cw_counter
RS-FEC uncorrected codewords counter	RS-FEC uncorrected codewords counter register	1.204, 1.205	FEC_uncorrected_cw_counter
RS-FEC symbol error counter, lane 0	RS-FEC symbol error counter lane 0 register	1.210, 1.211	FEC_symbol_error_counter_0

Table 108–3—MDIO/RS-FEC status variable mapping for separated PMA

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
PCS align status	RS-FEC status register	1.201.15	align_status

The following subclauses define variables that are not otherwise defined, e.g., for use by state diagrams.

108.6.1 FEC_bypass_correction_enable

When this variable is set to one, the Reed-Solomon decoder performs error detection without error correction (see 108.5.3.2). When this variable is set to zero, the decoder also performs error correction. The default value of the variable is zero. This variable is mapped to the bit defined in 45.2.1.101 (1.200.0).

108.6.2 FEC_bypass_indication_enable

This variable is set to one to bypass the error indication function (see 108.5.3.2) when this ability is supported. When this variable is set to zero, the decoder indicates errors to the PCS sublayer. This variable

has no effect (the decoder does not bypass error indication) if FEC bypass correction enable (1.200.0) is set to one. The default value of this variable is zero. This variable is mapped to the bit defined in 45.2.1.101 (1.200.1).

108.6.3 25G RS-FEC Enable

The RS-FEC sublayer shall have capability to enable or disable the FEC function. An MDIO interface or an equivalent management interface shall be provided to access the variable FEC_Enable for the RS-FEC sublayer. When FEC_Enable variable is set to a one, the RS-FEC performs the transmit function as specified in 108.5.2 and the receive function as specified in 108.5.3. When the variable is set to zero, the transmit and receive functions are disabled, and the RS-FEC sublayer is bypassed, effectively connecting its service interface to the service interface of its underlying sublayer. Implementation of this bypass function should cause a minimal delay in the transmit and receive paths. This variable is mapped to the bit defined in 45.2.1.101.a (1.200.2).

108.6.4 FEC_bypass_correction_ability

The Reed-Solomon decoder may have the option to perform error detection without error correction (see 108.5.3.2) to reduce the delay contributed by the 25GBASE-R RS-FEC sublayer. This variable is set to one to indicate that the decoder has the ability to bypass error correction. The variable is set to zero if this ability is not supported. This variable is mapped to the bit defined in 45.2.1.102 (1.201.0).

108.6.5 FEC_bypass_indication_ability

The Reed-Solomon decoder may have the option to bypass the error indication function (see 108.5.3.2) to reduce the delay contributed by the 25GBASE-R RS-FEC sublayer. This variable is set to one to indicate that the decoder has the ability to bypass error indication. The variable is set to zero if this ability is not supported. This variable is mapped to the bit defined in 45.2.1.102 (1.201.1).

108.6.6 FEC_high_ser

This variable is defined when the FEC_bypass_indication_ability variable is set to one. When FEC_bypass_indication_enable is set to one, this variable is set to one if the number of RS-FEC symbol errors in a window of 8192 codewords exceeds the threshold 417 (see 108.5.3.2) and is set to zero otherwise. This variable is mapped to the bit defined in 45.2.1.102 (1.201.2).

108.6.7 FEC_corrected_cw_counter

A corrected FEC codeword is a codeword that contained errors that were corrected.

FEC_corrected_cw_counter is a 32-bit counter that counts once for each corrected FEC codeword processed when FEC_align_status is true. This variable is mapped to the registers defined in 45.2.1.103 (1.202, 1.203).

108.6.8 FEC_uncorrected_cw_counter

An uncorrected FEC codeword is a codeword that contains errors (when the bypass correction feature is supported and enabled) or contains errors that were not corrected (when the bypass correction feature is not supported or not enabled).

FEC_uncorrected_cw_counter is a 32-bit counter that counts once for each uncorrected FEC codeword processed when FEC_align_status is true. This variable is mapped to the registers defined in 45.2.1.104 (1.204, 1.205).

108.6.9 FEC_symbol_error_counter_0

FEC_symbol_error_counter_0, as defined in 91.6.11, is a 32-bit counter that counts once for each 10-bit symbol corrected in lane 0 when FEC_align_status is true. This variable is mapped to the registers defined in 45.2.1.106 (1.210, 1.211). For the RS-FEC defined in this clause, all symbols are assigned to lane 0 and all symbol errors are counted by this counter.

108.6.10 align_status

This variable is assigned the value of rx_block_lock as defined by the PCS lock state diagram (Figure 49-14) used in the block synchronization function (108.5.2.1). This variable is mapped to the bit defined in 45.2.1.102.1 (1.201.15).

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

108.7 Protocol implementation conformance statement (PICS) proforma for Clause 108, Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs⁷

108.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 108, Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs, 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](#).

108.7.2 Identification

108.7.2.1 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).	

108.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3by-2016, Clause 108, Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs
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 802.3by-2016.)	

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.

108.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*KR	25GBASE-KR		Used to form a complete 25GBASE-KR PHY	O	Yes [] No []
*CR	25GBASE-CR		Used to form a complete 25GBASE-CR PHY	O	Yes [] No []
*SR	25GBASE-SR		Used to form a complete 25GBASE-SR PHY	O	Yes [] No []
DC	Delay constraints	108.4	Conforms to delay constraints specified	M	Yes []
EF	25G RS-FEC enable	108.6.3	Has the capability to enable and disable the 25G RS-FEC function	M	Yes []
*MD	MDIO capability	108.6	Registers and interface supported	O	Yes [] No []
*BEC	Bypass error correction	108.5.3.2	Capability is supported	O	Yes [] No []
*BEI	Bypass error indication	108.5.3.2	Capability is supported	O	Yes [] No []
*AUI	25GAUI C2C	108.3	25GAUI C2C used to connect to a PCS	O	Yes [] No []
*EEE	EEE deep sleep capability		Capability is supported	O	Yes [] No []

108.7.4 PICS proforma tables for Reed-Solomon Forward Error Correction (RS-FEC) sublayer for 25GBASE-R PHYs

108.7.4.1 Transmit function

Item	Feature	Subclause	Value/Comment	Status	Support
TF1	Rate compensation for codeword markers	108.5.2.2	Performs the rate compensation function described	M	Yes []
TF2	64B/66B to 256B/257B transcoder	108.5.2.3	tx_xcoded<256:0> constructed per 91.5.2.5	M	Yes []
TF3	257-bit block transmission order	108.5.2.3	First bit transmitted is bit 0	M	Yes []
TF4	Codeword marker insertion	108.5.2.4	First 257 message bits to be transmitted from every 1024th codeword	M	Yes []
TF5	Reed-Solomon encoder	108.5.2.5	As specified in 91.5.2.7	M	Yes []
TF6	Codeword serialization	108.5.2.6	Transmit bit ordering as in Figure 108-3	M	Yes []

108.7.4.2 Receive function

Item	Feature	Subclause	Value/Comment	Status	Support
RF1	Codeword marker lock	108.5.3.1	Implemented as described	M	Yes []
RF2	Reed-Solomon decoder	108.5.3.2	Capable of correcting any combination of up to 7 symbol errors in a codeword	M	Yes []
RF3	Error correction not bypassed for 25GBASE-SR	108.5.3.2	Error correction is not bypassed	BEC*SR: M	Yes [] N/A []
RF4	Uncorrected error indication	108.5.3.2	Capable of indicating when a codeword contains errors that were not corrected	M	Yes []
RF5	Error indication function	108.5.3.2	Corrupts 66-bit block synchronization headers as specified	M	Yes []
RF6	Error indication when error correction is bypassed	108.5.3.2	Error indication is not bypassed when error correction is bypassed	BEC*BEI: M	Yes [] N/A []
RF7	Error monitoring while error indication is bypassed	108.5.3.2	When the number of symbols errors in a block of 8192 codewords exceeds 417, corrupt 66-bit block synchronization headers for a period of 60 ms to 75 ms	BEI:M	Yes [] N/A []
RF8	Error monitor while rx_lpi_active is true	108.5.3.2	Disabled	BEI*EEE: M	Yes [] N/A []
RF9	Codeword monitor	108.5.3.3	Restarts the codeword marker lock when 3 consecutive uncorrected codewords are detected	M	Yes []
RF10	Alignment marker removal	108.5.3.4	rx_cwm removed prior to transcoding	M	Yes []
RF11	256B/257B to 64B/66B transcoder	108.5.3.5	rx_coded_j<65:0>, j=0 to 3 constructed per 91.5.3.5 with cross-referencing to g<3:0> using Figure 49–7	M	Yes []
RF12	Codeword marker bit rate compensation	108.5.3.6	Performs the rate compensation function described	M	Yes []
RF13	Codeword marker bit rate compensation	108.5.3.6	Content of blocks with invalid block type replaced by eight /E/ characters	M	Yes []
RF14	Codeword marker bit rate compensation	108.5.3.6	Corrupted sync headers passed through and no idle characters inserted after them	M	Yes []
RF15	Receive bit ordering	108.5.3.8	As illustrated in Figure 108–5	M	Yes []

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

108.7.4.3 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SD1	SLIP function	108.5.4.3	Ensure that all possible block positions are evaluated	M	Yes []

108.7.4.4 MDIO function mapping

Item	Feature	Subclause	Value/Comment	Status	Support
MD1	MDIO control and status bit mapping	108.6	Mapped according to Table 108–1 and Table 108–2	MD:M	Yes [] N/A []
MD2	Additional MDIO control and status bit mapping	108.6	Mapped according to Table 108–3	MD*AU I:M	Yes [] N/A []

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

109. Physical Medium Attachment (PMA) sublayer, type 25GBASE-R

109.1 Overview

109.1.1 Scope

This clause specifies the Physical Medium Attachment sublayer (PMA) that is common to a family of 25 Gb/s Physical Layer implementations, known as 25GBASE-R. The PMA allows the PCS (specified in Clause 107) to connect in a media-independent way with a range of physical media. 25GBASE-R can be extended to support any full duplex medium requiring only that the PMD be compliant with the PMA interface.

The PMA service interface is defined in an abstract manner and does not imply any particular implementation. Electrical interfaces connecting PMA sublayers, known as 25GAUI, are defined in Annex 109A and Annex 109B.

109.1.2 Position of the PMA in the 25GBASE-R sublayers

Figure 109–1 shows the relationship of the PMA sublayer (shown shaded) with other sublayers to the ISO Open System Interconnection (OSI) reference model.





Figure 109–1—25GBASE-R PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

109.1.3 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) Provide clock and data recovery.
- b) Provide signal drivers.
- c) Optionally provide local loopback to/from the PMA service interface.
- d) Optionally provide remote loopback to/from the PMD service interface.
- e) Optionally provide test-pattern generation and detection.

The functional block diagram in Figure 109–2 shows the inputs, outputs, test-pattern checking and generation, loopbacks, and Signal Indication Logic (SIL, see 109.2).




inst = PMD, PMA, or FEC, depending on which sublayer is below this PMA
 SIL = Signal Indication Logic

- ^a If 25GAUI immediately above this PMA.
- ^b If 25GAUI or PMD service interface immediately below this PMA.
- ^c Optional.
- ^d Conditional (see 109.4.2)

Figure 109–2—PMA Functional Block Diagram

109.1.4 PMA sublayer positioning

An implementation may use one or more PMA sublayers to provide an interface with a physical electrical interface, 25GAUI, between devices. The number of PMA sublayers required depends on the partitioning of functionality for a particular implementation. An example is illustrated in Figure 109–3. This example illustrates the partitioning that might arise from implementing an FEC sublayer in a device that is separate from the PCS. Additional examples are illustrated in Annex 109C.



25GAUI = 25 GIGABIT ATTACHMENT UNIT INTERFACE MMD = MDIO MANAGEABLE DEVICE
 25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE PCS = PHYSICAL CODING SUBLAYER
 FEC = FORWARD ERROR CORRECTION PMA = PHYSICAL MEDIUM ATTACHMENT
 MAC = MEDIA ACCESS CONTROL PMD = PHYSICAL MEDIUM DEPENDENT
 MDI = MEDIUM DEPENDENT INTERFACE

Figure 109–3—Example 25GBASE-R PMA layering

Management Data Input/Output (MDIO) Manageable Device (MMD) addresses 1, 8, 9, 10, and 11 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 PCS. The example shown in Figure 109–3 could be implemented with four addressable instances: MMD 8 addressing the lowest PMA sublayer (note that this cannot share MMD 1 with the PMD as they are not packaged together in this example), MMD 9 addressing the PMA sublayer above the 25GAUI below the FEC, MMD 10 addressing the PMA sublayer below the 25GAUI above the FEC, and MMD 11 addressing the PMA sublayer closest to the PCS.

The following guidelines apply to the partitioning of PMAs:

- a) The inter-sublayer service interface, defined in 105.4, is used for the PMA, FEC, and PMD service interfaces supporting a flexible architecture with optional FEC and multiple PMA sublayers.

- b) 25GAUI is a 25.78125 GBd serial physical instantiation of the connection between two adjacent PMA sublayers. As a physical instantiation, it defines electrical and timing specification as well as requiring a receive re-timing function.
- c) 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 109–2.
- d) A minimum of one PMA sublayer is required in a PHY.
- e) A maximum of four PMA sublayers are addressable as MDIO MMDs.

109.2 PMA service interface

This subclause specifies the services provided by the 25GBASE-R PMA. The service interface for this PMA is described in an abstract manner and does not imply any particular implementation. The PMA service interface supports the exchange of encoded data. The PMA translates the encoded data to and from signals suitable for the medium.

The PMA service interface is an instance of the inter-sublayer service interface defined in 105.4. The PMA service interface primitives are summarized as follows:

PMA:IS_UNITDATA.request
 PMA:IS_UNITDATA.indication
 PMA:IS_SIGNAL.indication

The PMA continuously sends a bit stream to the PMA client (e.g., PCS) at a nominal signaling rate of 25.78125 GBd. Likewise the PMA client continuously sends a bit stream to the PMA.

In the Tx direction, if data from a PMA:IS_UNITDATA.request primitive is received over a 25GAUI, clock and data are recovered on the lane to recover data one bit at a time. Each recovered bit is sent to the sublayer below using the *inst*:IS_UNITDATA.request primitive.

In the Rx direction, when data is received from the sublayer below the PMA using the *inst*:IS_UNITDATA.request primitive, each received bit is sent to the PMA client using the PMA:IS_UNITDATA.indication primitive.

PMA:IS_SIGNAL.indication(SIGNAL_OK) is generated based on receipt of *inst*:IS_SIGNAL.indication(SIGNAL_OK) from the sublayer below and status of the input signal as determined by the signal detect function (see Figure 109–2). When *inst*:IS_SIGNAL.indication(SIGNAL_OK) has the value FAIL or the signal detect function detects an invalid signal, PMA:IS_SIGNAL.indication(SIGNAL_OK) shall have the value FAIL, otherwise PMA:IS_SIGNAL.indication(SIGNAL_OK) shall have the value OK. The operation of the signal detect function is beyond the scope of this standard.

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see Clause 78, 78.1.3.3.1) then the inter-sublayer service interface includes four additional primitives defined as follows:

PMA:IS_TX_MODE.request
 PMA:IS_RX_MODE.request
 PMA:IS_ENERGY_DETECT.indication
 PMA:IS_RX_TX_MODE.indication

A physically instantiated service interface with the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option (see 78.1.3.3.1) may enter a low power state to conserve energy during periods of low link utilization. The ability to support transition to a low power state in the ingress direction is indicated by register 1.1.9 (PMA Ingress AUI Stop Ability, PIASA). The ability to support transition to a low power state in the egress direction is indicated by register 1.1.8 (PMA Egress AUI Stop Ability, PEASA). Transition to the low power state in the ingress direction is enabled by register 1.7.9 (PMA Ingress AUI Stop Enable, PIASE). Transition to the low power state in the egress direction is enabled by register 1.7.8 (PMA Egress AUI Stop Enable, PEASE). The system shall not assert the enable bit for an interface unless the corresponding ability bit at the other side of the interface is also asserted. If the PIASE bit is TRUE, then the PMA may disable transmitters on the physical instantiation of the ingress AUI when `au_i_tx_mode` is QUIET. If the PEASE bit is TRUE, then the PMA may disable transmitters on the physical instantiation of the egress AUI when `tx_mode` is QUIET.

109.3 Service interface below PMA

Since the architecture supports multiple PMA sublayers, there are several different sublayers that may appear below a PMA, including FEC, the PMD, or another PMA. The variable *inst* represents whichever sublayer appears below the PMA.

The sublayer below the PMA utilizes the inter-sublayer service interface defined in 105.4. The service interface primitives provided to the PMA are summarized as follows:

```
inst:IS_UNITDATA.request(tx_bit)
inst:IS_UNITDATA.indication(rx_bit)
inst:IS_SIGNAL.indication(SIGNAL_OK)
```

Note that electrical and timing specifications of the service interface are defined if the interface is physically instantiated (e.g., 25GAUI), otherwise the service interface is specified abstractly. The service interface below the PMA has an input and output for data transfer and a status input indicating a good signal sent by the sublayer below the PMA (see Figure 109-2).

In the Tx direction, data received via the `PMA:IS_UNITDATA.request` primitive from the PMA client at the PMA service interface is sent to the sublayer below the PMA via the `inst:IS_UNITDATA.indication` primitive.

In the Rx direction, if the bit is received over a physically instantiated interface (e.g., 25GAUI), clock and data are recovered on the lane receiving the bit. The bit is routed through the PMA to an output lane toward the PMA client, sending the bit to the PMA client using the `PMA:IS_UNITDATA.indication` primitive at the PMA service interface.

If the optional Energy Efficient Ethernet (EEE) capability with the deep sleep mode option is supported (see Clause 78, 78.1.3.3.1) then inter-sublayer service interface for the sublayer below the PMA includes four additional primitives defined as follows:

```
inst:IS_TX_MODE.request
inst:IS_RX_MODE.request
inst:IS_ENERGY_DETECT.indication
inst:IS_RX_TX_MODE.indication
```

109.4 Functions within the PMA

109.4.1 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 interface requirements for the 25GAUI chip-to-chip (C2C) are specified in Annex 109A.

The electrical and jitter/timing interface requirements for the 25GAUI chip-to-module (C2M) are specified in Annex 109B.

109.4.2 PMA local loopback mode

PMA local loopback shall be provided by the PMA adjacent to the PMD for 25GBASE-CR, 25GBASE-CR-S, 25GBASE-KR, and 25GBASE-KR-S. PMA local loopback mode is optional for other PMDs or for PMA not adjacent to the PMD. If it is implemented, it shall be as described in this subclause.

The PMA local loopback function involves looping back the input to the output. Each bit received from the PMA:IS_UNITDATA.request(tx_bit) primitive is looped back in the direction of the PCS using the PMA:IS_UNITDATA.indication(rx_bit) primitive.

Ability to perform the local loopback function is indicated by the Local_loopback_ability status variable. 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 the optional Clause 45 MDIO is implemented, the PMA receive process maps the Local_loopback_ability and Local_loopback_enable variables to the registers and bits defined in 109.6.

109.4.3 PMA remote loopback mode

PMA remote loopback mode is optional. If implemented, it shall be as described in this subclause.

When remote loopback is enabled, each bit from the service interface below the PMA via *inst*:IS_UNITDATA.indication is looped back toward the PMD via *inst*:IS_UNITDATA.request. Note that the service interface below the PMA can be provided by the FEC or PMD.

The ability to perform this function is indicated by the Remote_loopback_ability status variable. 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 the optional Clause 45 MDIO is implemented, the PMA receive process maps the Remote_loopback_ability and Remote_loopback_enable variables to the registers and bits defined in 109.6.

109.4.4 PMA test patterns

Where the output of the PMA appears on a physically instantiated interface 25GAUI or the PMD service interface, 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.

109.4.4.1 Transmit PRBS31 test-pattern generation

The ability to generate PRBS31 test patterns in the transmit direction is indicated by the `PRBS31_Tx_generator_ability` status variable. If supported, transmit PRBS31 test-pattern generation is enabled by the `PRBS31_enable` and `PRBS_Tx_gen_enable` control variables. When transmit PRBS31 test-pattern generation is enabled, the PMA shall generate a PRBS31 pattern as defined in 49.2.8 toward the service interface below the PMA via the `inst:IS_UNITDATA.request` primitive.

If the optional Clause 45 MDIO is implemented, the PMA transmit process maps the `PRBS31_Tx_checker_ability`, `PRBS31_Rx_checker_ability`, `PRBS31_enable`, and `PRBS_Tx_gen_enable` variables to the registers and bits defined in 109.6.

To avoid correlated crosstalk between this link and other locally implemented links, it is highly recommended that the PRBS31 patterns generated on each link use independent, random seeds or ensure a minimum offset of 20 000 UI between the PRBS31 sequence on any two links.

109.4.4.2 Receive PRBS31 test-pattern generation

The ability to generate PRBS31 test patterns in the receive direction is indicated by the `PRBS31_Rx_generator_ability` status variable. If supported, receive PRBS31 test-pattern generation is enabled by the `PRBS31_enable` and `PRBS_Rx_gen_enable` control variables. When receive PRBS31 test-pattern generation is enabled, the PMA shall generate a PRBS31 pattern as defined in 49.2.8 toward the PMA client via the `PMA:IS_UNITDATA.indication` primitive. While this test pattern is enabled, the PMA also generates `PMA:IS_SIGNAL.indication(SIGNAL_OK)` indicating a valid signal (`SIGNAL_OK=OK`) toward the PMA client independent of the status at the service interface below the PMA. When receive PRBS31 test-pattern generation is disabled, the PMA returns to normal operation.

If the optional Clause 45 MDIO is implemented, the PMA receive process maps the `PRBS31_Rx_generator_ability`, `PRBS31_enable`, and `PRBS_Rx_gen_enable` variables to the registers and bits defined in 109.6.

To avoid correlated crosstalk between this link and other locally implemented links, it is highly recommended that the PRBS31 patterns generated on each link use independent, random seeds or ensure a minimum offset of 20 000 UI between the PRBS31 sequence on any two links.

109.4.4.3 Transmit PRBS31 test-pattern checking

The ability to check PRBS31 test patterns in the transmit direction is indicated by the `PRBS31_Tx_checker_ability` status variable. If supported, the transmit PRBS31 test-pattern checking is enabled by the `PRBS31_enable` and `PRBS_Tx_check_enable` control variables. When transmit PRBS31 test-pattern checking is enabled, the PMA shall check for the PRBS31 pattern received from the PMA client via the `PMA:IS_UNITDATA.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. The transmit test-pattern error count is indicated by the `Ln0_PRBS_Tx_test_err_counter` status variable. When transmit PRBS31 test-pattern checking is disabled, the PMA returns to normal operation.

If the optional Clause 45 MDIO is implemented, the PMA transmit process maps the `PRBS31_Tx_checker_ability`, `PRBS31_enable`, `PRBS_Tx_check_enable`, and `Ln0_PRBS_Tx_test_err_counter` variables to the registers and bits defined in 109.6.

109.4.4.4 Receive PRBS31 test-pattern checking

The ability to check PRBS31 test patterns in the receive direction is indicated by the PRBS31_Rx_checker_ability status variable. If supported, receive PRBS31 test-pattern checking is enabled by the PRBS31_enable and PRBS_Rx_check_enable control variables. When receive PRBS31 test-pattern checking is enabled, the PMA shall check for the PRBS31 pattern received from the service interface below the PMA via the *inst:IS_UNITDATA.indication* primitive. The receive test-pattern error count is indicated by the Ln0_PRBS_Rx_test_err_counter status variable. While the receive PRBS31 test-pattern mode is enabled, the PMA:IS_SIGNAL.indication(SIGNAL_OK) primitive does not indicate a valid signal (SIGNAL_OK=FAIL). When receive PRBS31 test-pattern checking is disabled, the PMA returns to normal operation.

If the optional Clause 45 MDIO is implemented, the PMA receive process maps the PRBS31_Rx_checker_ability, PRBS31_enable, PRBS_Rx_check_enable, and Ln0_PRBS_Rx_test_err_counter variables to the registers and bits defined in 109.6.

109.4.4.5 Transmit PRBS9 test-pattern generation

The ability to generate the PRBS9 test patterns in the transmit direction is indicated by the PRBS9_Tx_generator_ability status variable. If supported, transmit PRBS9 test-pattern generation is enabled by the PRBS9_enable and PRBS_Tx_gen_enable control variables. When transmit PRBS9 pattern generation is enabled, the PMA shall generate a PRBS9 pattern as defined in Table 68–6 toward the service interface below the PMA via the *inst:IS_UNITDATA.request* primitive. When transmit PRBS9 test-pattern generation is disabled, the PMA returns to normal operation.

If the optional Clause 45 MDIO is implemented, the PMA transmit process maps the PRBS9_Tx_generator_ability, PRBS9_enable, and PRBS_Tx_gen_enable variables to the registers and bits defined in 109.6.

Note that the transmit PRBS9 test pattern is intended to be checked by external test gear and no PRBS9 checking function is provided within the PMA.

109.4.4.6 Receive PRBS9 test-pattern generation

The ability to generate the PRBS9 test patterns in the receive direction is indicated by the PRBS9_Rx_generator_ability status variable. If supported, receive PRBS9 test-pattern generation is enabled by the PRBS9_enable and PRBS_Rx_gen_enable control variables. When receive PRBS9 test-pattern generation is enabled, the PMA shall generate a PRBS9 pattern as defined in Table 68–6 toward the PMA client via the PMA:IS_UNITDATA.indication primitive. The PMA will also generate PMA:IS_SIGNAL.indication(SIGNAL_OK) toward the PMA client independent of the status at the service interface below the PMA. When receive PRBS9 test-pattern generation is disabled, the PMA returns to normal operation.

If the optional Clause 45 MDIO is implemented, the PMA receive process maps the PRBS9_Rx_generator_ability, PRBS9_enable and PRBS_Rx_gen_enable variables to the registers and bits defined in 109.6.

Note that the receive PRBS9 test pattern is intended to be checked by external test gear and no PRBS9 checking function is provided within the PMA.

109.4.4.7 Transmit square wave test-pattern generation

The optional transmit square wave test-pattern generation applies to the PMA transmit direction towards a physically instantiated 25GAUI or towards the PMD service interface.

The ability to generate the square wave test pattern in the transmit direction is indicated by the `Square_wave_ability` status variable. If supported, transmit square wave test-pattern generation is enabled by the `Square_wave_enable_0` control variable. When transmit square wave test-pattern generation is enabled, the PMA shall generate a square wave test pattern composed of 8 ones followed by 8 zeros toward the service interface below the PMA via the `inst:IS_UNITDATA.request` primitive. When transmit square wave test-pattern generation is disabled, the PMA returns to normal operation.

If the optional Clause 45 MDIO is implemented, the PMA transmit process maps the `Square_wave_ability` and `Square_wave_enable_0` variable to the registers and bits defined in 109.6.

109.4.5 Energy Efficient Ethernet for 25GAUI

When the optional Energy Efficient Ethernet (EEE) deep sleep capability is supported and the PMA service interface is physically instantiated as 25GAUI, the additional functions listed in 83.5.11 are required. These functions enable the communication of service interface parameters that are essential to the operation of the EEE deep sleep capability. The timing parameters for EEE operation are shown in Table 83-2.

PMA EEE operation for 25GAUI is specified in 83.5.11 with respect to lane 0 only. Considerations related to multiple lanes do not apply.

109.5 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 109-1. A description of overall system delay constraints and the definitions for `bit-times` and `pause_quanta` can be found in 105.5.

Table 109-1—Delay constraints

Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
4096	8	163.84

109.6 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 109.1.4 for the allocation of MMD addresses to PMA sublayers. Control and status registers for MMD 8, 9, 10, and 11 will 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 109-2. Mapping of MDIO status variables to PMA status variables is shown in Table 109-3. Mapping of MDIO counter to PMA counters is shown in Table 109-4. 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, 10, and 11 as necessary.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

Table 109–2—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
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
PIASE	PMA/PMD control 2	1.7.9	PIASE
PEASE	PMA/PMD control 2	1.7.8	PEASE

Table 109–3—MDIO/PMA status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMA status variable
PMA remote loopback ability	40G/100G PMA/PMD extended ability	1.13.15	Remote_loopback_ability
PMA local loopback ability	PMA/PMD status 2	1.8.0	Local_loopback_ability
PRBS9 Tx generator ability	Test-pattern ability	1.1500.5	PRBS9_Tx_generator_ability
PRBS9 Rx generator ability	Test-pattern ability	1.1500.4	PRBS9_Rx_generator_ability
PRBS31 Tx generator ability	Test-pattern ability	1.1500.3	PRBS31_Tx_generator_ability
PRBS31 Tx checker ability	Test-pattern ability	1.1500.2	PRBS31_Tx_checker_ability
PRBS31 Rx generator ability	Test-pattern ability	1.1500.1	PRBS31_Rx_generator_ability
PRBS31 Rx checker ability	Test-pattern ability	1.1500.0	PRBS31_Rx_checker_ability
Square wave test ability	Test-pattern ability	1.1500.12	Square_wave_ability
PIASA	PMA/PMD status 1	1.1.9	PIASA
PEASA	PMA/PMD status 1	1.1.8	PEASA

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

Table 109–4—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 Rx, lane 0	PRBS Rx pattern testing error counter, lane 0	1.1700	Ln0_PRBS_Rx_test_err_counter

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

109.7 Protocol implementation conformance statement (PICS) proforma for Clause 109, Physical Medium Attachment (PMA) sublayer, type 25GBASE-R⁸

109.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 109, Physical Medium Attachment (PMA) sublayer, type 25GBASE-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](#).

109.7.2 Identification

109.7.2.1 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).	

109.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3by-2016, Clause 109, Physical Medium Attachment (PMA) sublayer, type 25GBASE-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 802.3by-2016.)	

Date of Statement	
-------------------	--

⁸Copyright 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.

109.7.3 PICS proforma tables for the 25GBASE-R PMA Sublayer

109.7.4 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
MD	MDIO	109.6	Registers and interface supported	O	Yes [] No []
*PIA	Physically instantiated 25GAUI above (toward PCS)	109.4.1		O	Yes [] No []
*PIB	Physically instantiated 25GAUI below (toward PMD)	109.4.1		O	Yes [] No []
*PMB	PMD below	109.3		O	Yes [] No []
*LPI	Implementation of LPI with the deep sleep mode option	109.4.5		O	Yes [] No []
*KRCCR	PMA adjacent to the PMD for 25GBASE-KR, 25GBASE-KR-S, 25GBASE-CR, or 25GBASE-CR-S.	109.4.2		O	Yes [] No []
*LBL	Supports PMA local loopback	109.4.2		KRCCR:M !KRCCR:O	Yes [] No [] N/A []
*LBR	Supports PMA remote loopback	109.4.3		O	Yes [] No [] N/A []

109.7.4.1 PMA functions

Item	Feature	Subclause	Value/Comment	Status	Support
PF1	PMA local loopback	109.4.2	Meets the requirements of 109.4.2.	LBL:M	Yes [] No []
PF2	PMA remote loopback	109.4.3	Meets the requirements of 109.4.3.	LBR:M	Yes [] No []
PF3	Transmit PRBS31 generation	109.4.4.1	Meets the requirements of 109.4.4.1.	PMB:O	Yes [] No []
PF4	Receive PRBS31 generation	109.4.4.2	Meets the requirements of 109.4.4.2.	PIA:O	Yes [] No []

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

Item	Feature	Subclause	Value/Comment	Status	Support
PF5	Transmit PRBS31 checking	109.4.4.3	Meets the requirements of 109.4.4.3.	PIA:O	Yes [] No []
PF6	Receive PRBS31 checking	109.4.4.4	Meets the requirements of 109.4.4.4.	PMB:O	Yes [] No []
PF7	Transmit PRBS9 generation	109.4.4.5	Meets the requirements of 109.4.4.5.	PMB:O	Yes [] No []
PF8	Receiver PRBS9 generation	109.4.4.6	Meets the requirements of 109.4.4.6.	PIA:O	Yes [] No []
PF9	Transmit square wave generation	109.4.4.7	Meets the requirements of 109.4.4.7.	PMB:O	Yes [] No []

109.7.4.2 PMA characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
PC1	Cumulative round-trip delay contributed by up to four PMA stages in a PHY.	109.5	No more than 4096 BT or 8 pause quanta	M	Yes [] No []
PC2	Electrical and timing requirements of Annex 109A or Annex 109B as appropriate met by upstream 25GAUI	109.4.1		PIA:M	Yes [] No []
PC3	Electrical and timing requirements of Annex 109A or Annex 109B as appropriate met by downstream 25GAUI	109.4.1		PIB:M	Yes [] No []
PC4	25GAUI deep sleep Rx direction	109.4.5	Meets the requirements of 109.4.5.	LPI*PIA:M LPI*PIB:M	Yes [] No []
PC5	25GAUI deep sleep Tx direction	109.4.5	Meets the requirements of 109.4.5.	LPI*PIA:M LPI*PIB:M	Yes [] No []

**110. Physical Medium Dependent (PMD) sublayer and baseband medium,
type 25GBASE-CR and 25GBASE-CR-S****110.1 Overview**

This clause specifies the 25GBASE-CR PMD, the 25GBASE-CR-S PMD, and the baseband medium. The specifications are closely related to those of 100GBASE-CR4 (Clause 92), but with a single lane instead of four lanes. There are three associated annexes. Annex 110A provides information on parameters with test points that may not be testable in an implemented system, Annex 110B specifies test fixtures, and Annex 110C describes 25GBASE-CR and 25GBASE-CR-S host and cable assembly form factors.

When forming a complete Physical Layer, the PMD shall be connected as illustrated in Figure 110–1, to the appropriate PMA as shown in Table 110–1, to the medium through the MDI, and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 110–1—Physical Layer clauses associated with the 25GBASE-CR and 25GBASE-CR-S PMDs

Associated clause	25GBASE-CR	25GBASE-CR-S
106—RS	Required	Required
106—25GMII ^a	Optional	Optional
107—PCS	Required	Required
74—BASE-R FEC ^b	Required	Required
108—RS-FEC ^b	Required	N/A
109—PMA	Required	Required
109A—25GAUI C2C	Optional	Optional
73—Auto-Negotiation	Required	Required
78—Energy Efficient Ethernet	Optional	Optional

^aThe 25GMII is an optional interface. However, if the 25GMII is not implemented, a conforming implementation must behave functionally as though the RS and 25GMII were present.

^bFEC sublayers can be enabled or disabled according to the FEC mode (see 110.6).

A 25GBASE-CR PHY operates over cable assemblies of types CA-25G-N, CA-25G-S and CA-25G-L (see 110.10). A 25GBASE-CR-S PHY operates over cable assemblies of types CA-25G-N and CA-25G-S. A 25GBASE-CR-S PHY interoperates with a 25GBASE-CR PHY. Table 110–2 summarizes the cable assembly types supported by each of the PHY types.

Table 110–2—Cable assembly types supported by each PHY type

PHY type	CA-25G-N	CA-25G-S	CA-25G-L
25GBASE-CR	Yes	Yes	Yes
25GBASE-CR-S	Yes	Yes	No

When forming a complete 25GBASE-CR or 25GBASE-CR-S Physical Layer, the link BER requirements depend on the FEC mode (see 110.6) according to the following guidelines:

- a) If a PHY operates in the RS-FEC mode, and the RS-FEC decoder does not bypass error correction (see 108.5.3.2), the link is required to operate with a BER of 10^{-5} or better.
- b) If a PHY operates in the BASE-R FEC mode, the link is required to operate with a BER of 10^{-8} or better.
- c) If a PHY operates in the no-FEC mode, or in the RS-FEC mode with error correction bypassed, the link is required to operate with a BER of 10^{-12} or better.

In this context, a link consists of a compliant PMD transmitter, a compliant PMD receiver, and a compliant cable assembly.

For a complete Physical Layer, this specification is considered to be satisfied by a frame loss ratio (see 1.4.223) less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap.

25GBASE-CR and 25GBASE-CR-S PHYs with the optional Energy-Efficient Ethernet (EEE) capability may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization.

Figure 110–1 shows the relationship of the 25GBASE-CR and 25GBASE-CR-S PMD sublayers and MDI to the ISO/IEC Open System Interconnection (OSI) reference model.

110.2 PMD service interface

This subclause specifies the services provided by the 25GBASE-CR and 25GBASE-CR-S PMDs. The service interface for these PMDs is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data. The PMD translates the encoded data to and from signals suitable for the medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 105.4. The PMD service interface primitives are summarized as follows:

```
PMD:IS_UNITDATA.request
PMD:IS_UNITDATA.indication
PMD:IS_SIGNAL.indication
```


The PMA (or the PMD) continuously sends a bit stream to the PMD (or the PMA) at a nominal signaling rate of 25.78125 Gb/s.

The SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive corresponds to the variable Global_PMD_signal_detect as defined in 110.7.4. When Global_PMD_signal_detect is one, SIGNAL_OK shall be assigned the value OK. When Global_PMD_signal_detect is zero, SIGNAL_OK shall be assigned the value FAIL. When SIGNAL_OK is FAIL, the PMD:IS_UNITDATA.indication parameter is undefined.

If the optional EEE deep sleep capability is supported, then the PMD service interface includes two additional primitives as follows:

```
PMD:IS_TX_MODE.request
PMD:IS_RX_MODE.request
```

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE PCS = PHYSICAL CODING SUBLAYER
 AN = AUTO-NEGOTIATION PHY = PHYSICAL LAYER DEVICE
 FEC = FORWARD ERROR CORRECTION PMA = PHYSICAL MEDIUM ATTACHMENT
 LLC = LOGICAL LINK CONTROL PMD = PHYSICAL MEDIUM DEPENDENT
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 110–1—25GBASE-CR and 25GBASE-CR-S relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

110.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 107.4.)

25GBASE-CR and 25GBASE-CR-S PHYs may be extended using a 25GAUI chip-to-chip (C2C) as a physical instantiation of the inter-sublayer service interface between devices. If 25GAUI C2C is instantiated, the AN_LINK(link_status).indication is relayed from the device with the PCS sublayer to the device with the AN sublayer by means at the discretion of the implementer. As examples, the implementer may employ use of pervasive management or employ a dedicated electrical signal to relay the state of link_status as indicated by the PCS sublayer on one device to the AN sublayer on the other device.

110.4 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the PMD and AN shall be no more than 512 bit times (1 pause_quantum or 20.48 ns). It is assumed that the one way delay through the medium is no more than 1500 bit times (60 ns).

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 105.5.

110.5 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the PMD. If MDIO is implemented, it shall map MDIO control bits to PMD control variables as shown in Table 110–3, and MDIO status bits to PMD status variables as shown in Table 110–4.

Table 110–3—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable	1.9.0	Global_PMD_transmit_disable
Restart training	BASE-R PMD control	1.150.0	mr_restart_training
Training enable	BASE-R PMD control	1.150.1	mr_training_enable
Polynomial identifier 0	PMD training pattern lane 0	1.1450.12:11	identifier_0
Seed 0	PMD training pattern lane 0	1.1450.10:0	seed_0

Table 110–4—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect	1.10.0	Global_PMD_signal_detect
Receiver status 0	BASE-R PMD status	1.151.0	rx_trained
Frame lock 0	BASE-R PMD status	1.151.1	frame_lock
Start-up protocol status 0	BASE-R PMD status	1.151.2	training
Training failure 0	BASE-R PMD status	1.151.3	training_failure

110.6 FEC modes

A 25GBASE-CR PHY implements the BASE-R FEC sublayer (Clause 74) and the 25GBASE-R RS-FEC sublayer (Clause 108). A 25GBASE-CR-S PHY implements the BASE-R FEC sublayer (Clause 74). Each FEC sublayer can be either enabled or disabled, according to AN resolution or management control.

Three FEC modes are supported:

- a) When the 25GBASE-R RS-FEC sublayer is enabled, the PHY is defined to operate in the RS-FEC mode.
- b) When the BASE-R FEC sublayer is enabled, the PHY is defined to operate in the BASE-R FEC mode.
- c) When no FEC sublayer is enabled, the PHY is defined to operate in the no-FEC mode.

A 25GBASE-CR PHY can operate in RS-FEC, BASE-R FEC, or no-FEC mode. A 25GBASE-CR-S PHY can operate in either BASE-R FEC or no-FEC mode.

The cable assembly types (CA-25G-N, CA-25G-S, or CA-25G-L, see 110.10) that the PHY supports and the required PMD receiver characteristics (110.8.4) depend on the FEC mode.

The FEC mode is determined using AN (Clause 73) if AN is enabled, or by management control if AN is disabled, and is used in both transmit direction and receive direction. It is recommended to configure the AN FEC advertisement such that only modes that are compatible with the type of the cable assembly attached to the MDI are selected.

110.7 PMD functional specifications

110.7.1 Link block diagram

One direction of a 25GBASE-CR or 25GBASE-CR-S link is shown in Figure 110–2.




Figure 110–2—25GBASE-CR or 25GBASE-CR-S link (one direction is illustrated)

Note that the source lane (SL) signals SL<p> and SL<n> are the positive and negative sides of the transmitter’s differential signal pair and the destination lane (DL) signals DL<p> and DL<n> are the positive and negative sides of the receiver’s differential signal pair.

For purposes of system conformance, the PMD sublayer is standardized at the test points described in this subclause.

The electrical transmit signal is defined at TP2. Unless specified otherwise, all transmitter measurements and tests defined in 110.8.3 are made at TP2 utilizing the test fixture specified in 110B.1.1.

The electrical receive signal is defined at TP3. Unless specified otherwise, all receiver measurements and tests defined in 110.8.4 are performed at TP3 utilizing the test fixture specified in 110B.1.1.

A mated connector pair has been included in both the transmitter and receiver specifications defined in 110.8.3 and 110.8.4. The recommended maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is provided in 92.8.3.6. Annex 110A provides information on parameters associated with test points TP0 and TP5 that may not be testable in an implemented system.

The channel (see 110.9) is defined between the transmitter (TP0) and receiver (TP5) blocks to include the transmitter and receiver differential controlled impedance printed circuit board insertion loss and the cable assembly insertion loss, as illustrated in Figure 110–2. All cable assembly measurements are to be made between TP1 and TP4 as illustrated in Figure 110–2. The cable assembly test fixture of 110B.1.2, or its equivalent, is required for measuring the cable assembly specifications in 110.10 at TP1 and TP4. Two mated connector pairs and the cable assembly test fixture have been included in the cable assembly specifications defined in 110.10. Transmitter and receiver differential controlled impedance printed circuit board insertion losses defined between TP0 and the MDI and between the MDI and TP5, respectively, are provided informatively in 92A.4.

Table 110–5 describes the defined test points illustrated in Figure 110–2.

Table 110–5—Test points

Test points	Description
TP0 to TP5	The channel including the transmitter and receiver differential controlled impedance printed circuit board insertion loss and the cable assembly insertion loss.
TP1 to TP4	All cable assembly measurements are to be made between TP1 and TP4 as illustrated in Figure 110–2. The cable assembly test fixture of 110B.1.2, or its equivalent, is required for measuring the cable assembly specifications in 110.10 at TP1 and TP4.
TP0 to TP2 TP3 to TP5	A mated connector pair has been included in both the transmitter and receiver specifications defined in 110.8.3 and 110.8.4. The recommended maximum insertion loss from TP0 to TP2 or from TP3 to TP5 including the test fixture is provided in 92.8.3.6.
TP2	Unless specified otherwise, all transmitter measurements defined in 110.8.3 are made at TP2 utilizing the test fixture specified in 110B.1.1.
TP3	Unless specified otherwise, all receiver measurements and tests defined in 110.8.4 are made at TP3 utilizing the test fixture specified in 110B.1.1.

110.7.2 PMD transmit function

The PMD transmit function shall convert the bit stream requested by the PMD service interface message PMD:IS_UNITDATA.request(tx_bit) into an electrical signal. The electrical signal shall then be delivered to the MDI, according to the transmit electrical specifications in 110.8.3. A positive differential output voltage (SL<p> minus SL<n>) shall correspond to tx_bit = one.

If the optional EEE deep sleep capability is supported, the following requirements apply. When the PMD service interface message `PMD:IS_TX_MODE.request(tx_mode)` is received with `tx_mode = ALERT`, the PMD transmit function shall transmit a periodic sequence, where each period of the sequence consists of 8 ones followed by 8 zeros, with the transmit equalizer coefficients set to the preset values (see 72.6.10.2.3.1). This sequence is transmitted regardless of the value of `tx_bit` presented by the `PMD:IS_UNITDATA.request` primitive. When `tx_mode` is not set to `ALERT`, the transmit equalizer coefficients are set to the values determined via the start-up protocol (see 110.7.10).

110.7.3 PMD receive function

The PMD receive function shall convert the electrical signal from the MDI into a bit stream for delivery to the PMD service interface using the messages `PMD:IS_UNITDATA.indication(rx_bit)`. A positive differential input voltage ($DL_{<p>} - DL_{<n>}$) shall correspond to `rx_bit = one`.

110.7.4 Global PMD signal detect function

The Global PMD signal detect function is used by the PMD to indicate the successful completion of the start-up protocol by the PMD control function (see 110.7.10). `Global_PMD_signal_detect` is set to zero when the value of the variable `signal_detect` is set to false by the Training state diagram (see Figure 72–5). `Global_PMD_signal_detect` is set to one when the value of `signal_detect` is set to true.

If training is disabled by the management variable `mr_training_enable` (see 110.5), `Global_PMD_signal_detect` shall be set to one.

If the optional EEE deep sleep capability is supported, the following requirements apply. The value of `Global_PMD_signal_detect` is set to zero when the PMD service interface message `PMD:IS_RX_MODE.request(rx_mode)` is initially received with `rx_mode = QUIET`. While `rx_mode` is set to `QUIET`, `Global_PMD_signal_detect` shall be set to one within 500 ns of the application of the `ALERT` pattern defined in 110.7.2, with peak-to-peak differential voltage of 720 mV measured at TP2, to the differential pair at the input of the cable assembly that connects the transmitter to the receiver. While `rx_mode` is set to `QUIET`, `Global_PMD_signal_detect` shall not be set to one when the voltage input to the differential pair of the cable assembly that connects the transmitter to the receiver is less than or equal to 70 mV peak-to-peak differential.

When the MDIO is implemented, this function maps the variables to registers and bits as defined in 110.5.

110.7.5 Global PMD transmit disable function

The Global PMD transmit disable function is mandatory if EEE deep sleep capability is supported and is otherwise optional. When `Global_PMD_transmit_disable` variable is set to one, this function shall turn off the transmitter such that it drives a constant level (i.e., no transitions) and does not exceed the differential peak-to-peak output voltage (max.) with Tx disabled in Table 92–6.

If a PMD fault (110.7.7) is detected, then the PMD may set `Global_PMD_transmit_disable` to one.

Loopback, as defined in 110.7.6, shall not be affected by `Global_PMD_transmit_disable`.

The following additional requirements apply when the optional EEE deep sleep capability is supported:

- a) The Global PMD transmit disable function shall turn off the transmitter as specified in 92.8.3.1 when `tx_mode` transitions to `QUIET` from any other value.
- b) The Global PMD transmit disable function shall turn on the transmitter as specified in 92.8.3.1 when `tx_mode` transitions from `QUIET` to any other value.

110.7.6 Loopback mode

Local loopback mode is provided by the adjacent PMA (see Clause 109) as a test function. When loopback mode is enabled, transmission requests passed to the transmitter are sent directly to the receiver, overriding any signal detected by the receiver on its attached link. Note that loopback mode does not affect the state of the transmitter, which continues to send data (unless disabled).

Control of the loopback function is specified in 109.4.2.

NOTE 1—The signal path that is exercised in the loopback mode is implementation specific, but it is recommended that this signal path encompass as much of the circuitry as is practical. The intention of providing this loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

NOTE 2—Placing a network port into loopback mode can be disruptive to a network.

110.7.7 PMD fault function

PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault, and any other implementation specific fault. If the MDIO interface is implemented, then PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

110.7.8 PMD transmit fault function

The PMD transmit fault function is optional. The faults detected by this function are implementation specific, but the assertion of Global_PMD_transmit_disable is not considered a transmit fault.

If PMD_transmit_fault is set to one, then Global_PMD_transmit_disable should also be set to one.

If the MDIO interface is implemented, then PMD_transmit_fault shall be mapped to the Transmit fault bit as specified in 45.2.1.7.4.

110.7.9 PMD receive fault function

The PMD receive fault function is optional. The faults detected by this function are implementation specific. A fault is indicated by setting the variable PMD_receive_fault to one.

If the MDIO interface is implemented, then PMD_receive_fault shall be mapped to the Receive fault bit as specified in 45.2.1.7.5.

110.7.10 PMD control function

25GBASE-CR and 25GBASE-CR-S PMDs shall use the same control function as lane 0 of 100GBASE-CR4, as defined in 92.7.12.

The variables seed_0 and polynomial_0 control the training pattern. It is recommended that implementations with multiple PMDs use distinct values of polynomial_0 for PMDs that are coupled by crosstalk.

110.8 Electrical characteristics

110.8.1 Signal levels

The MDI (see 110.11) is a low-swing AC-coupled differential interface. AC-coupling within the plug connectors, as defined in 110.11.1, allows for interoperability between components operating from different supply voltages.

110.8.2 Signal paths

The MDI transmit and receive paths are point-to-point connections. Each path comprises two complementary signals, which form a balanced differential pair. There is one differential path in each direction for a total of two pairs, or four connections.

110.8.3 Transmitter characteristics

Transmitter electrical characteristics at TP2 for 25GBASE-CR and 25GBASE-CR-S PHYs shall be the same as those of a single lane of 100GBASE-CR4, as summarized in Table 92–6 and detailed in 92.8.3.1 through 92.8.3.9, with the exception that the value of linear fit pulse peak (min.) is $0.49 \times v_f$. The transmitter specifications at TP0 are provided informatively in 110A.2.

110.8.4 Receiver characteristics

Receiver electrical characteristics are specified at TP3. The receiver shall meet the return loss requirements specified in 92.8.4.2 and 92.8.4.3. In addition, the requirements in 110.8.4.1, 110.8.4.2, 110.8.4.3 and 110.8.4.4 apply.

The receiver specifications at TP5 are provided informatively in 110A.3.

110.8.4.1 Receiver input amplitude tolerance

When a 25GBASE-CR or 25GBASE-CR-S PMD receiver is connected to a compliant transmitter whose peak-to-peak differential output voltage, as defined by 92.8.3.1 using preset equalizer coefficients, is 1200 mV, using a compliant cable assembly with the minimum insertion loss specified in 110.10.2, the following BER requirements apply.

- a) When the RS-FEC mode is used, the PMD receiver shall operate at a BER better than 10^{-5} .
- b) When the BASE-R FEC mode is used, the PMD receiver shall operate at a BER better than 10^{-8} .
- c) When the no-FEC mode is used, the PMD receiver shall operate at a BER better than 10^{-12} .

The receiver is allowed to control the transmitter equalizer coefficients, using the protocol defined in 92.7.12 or an equivalent process, to meet these requirements.

110.8.4.2 Receiver interference tolerance test

Receiver interference tolerance is measured according to the requirements listed in 110.8.4.2.1 through 110.8.4.2.5.

A 25GBASE-CR PHY shall comply with the receiver interference tolerance test requirements for the RS-FEC, BASE-R FEC, and no-FEC modes. A 25GBASE-CR-S PHY shall comply with the receiver interference tolerance test requirements for the BASE-R FEC and no-FEC modes.

Two tests are defined for each FEC mode. Test 1 includes a low-loss channel, identical for all FEC modes. Test 2 includes a high-loss channel, which is different for each FEC mode. The cable assembly used in the test channel (see 110.8.4.2.2) shall meet the cable assembly Channel Operating Margin (COM) as specified in 110.10.7, with CA-25G-L COM parameter values being used for RS-FEC mode test, CA-25G-S COM parameter values being used for BASE-R FEC mode test, and CA-25G-N COM parameter values being used for no-FEC mode test.

The test requirements for the RS-FEC mode are provided in Table 110–6. The test requirements for the BASE-R FEC mode are provided in Table 110–7. The test requirements for the no-FEC mode are provided in Table 110–8.

Table 110-6—25GBASE-CR interference tolerance parameters, RS-FEC mode

Parameter	Test 1 (low loss)		Test 2 (high loss)		Units
	Min	Max	Min	Max	
Test pattern	Scrambled idle encoded by RS-FEC				
RS-FEC symbol error ratio required ^a	$< 10^{-4}$				
Test channel insertion loss at 12.8906 GHz ^b	14.3	14.8	29.44	29.94	dB
Cable assembly insertion loss at 12.8906 GHz	8	10	20.48	22.48	dB
COM		3		3	dB
b_{\max} used in COM calculation	1				
DER_0 used in COM calculation	10^{-5}				

^aThe RS-FEC symbol error ratio is measured using the RS-FEC symbol error counter (see 108.6.9).

^bInsertion loss between the two test references (see Figure 110-3b).

Table 110-7—25GBASE-CR and 25GBASE-CR-S interference tolerance parameters, BASE-R FEC mode

Parameter	Test 1 (low loss)		Test 2 (high loss)		Units
	Min	Max	Min	Max	
Test pattern	Scrambled idle encoded by BASE-R FEC				
BASE-R FEC corrected block ratio required ^{a,b}	$< 2.1 \times 10^{-5}$				
Test channel insertion loss at 12.8906 GHz ^c	14.3	14.8	23.44	23.94	dB
Cable assembly insertion loss at 12.8906 GHz	8	10	14.48	16.48	dB
COM		3		3	dB
b_{\max} used in COM calculation	0.5				
DER_0 used in COM calculation	10^{-8}				

^aThe BASE-R FEC corrected block ratio is measured using the FEC corrected blocks counter (see 74.8.4.1).

^bThe FEC uncorrected blocks counter (see 74.8.4.2) is required to indicate zero errors during the test unless the test duration is such that the uncorrected block ratio can be verified to be less than 4.7×10^{-10} .

^cInsertion loss between the two test references (see Figure 110-3b).

Table 110–8—25GBASE-CR and 25GBASE-CR-S interference tolerance parameters, no-FEC mode

Parameter	Test 1 (low loss)		Test 2 (high loss)		Units
	Min	Max	Min	Max	
Test pattern	Scrambled idle or PRBS31				
Bit error ratio required ^a	< 10 ⁻¹²				
Test channel insertion loss at 12.8906 GHz ^b	14.3	14.8	22.48	22.98	dB
Cable assembly insertion loss at 12.9806 GHz	8	10	13.5	15.5	dB
COM		3		2.2	dB
b_{\max} used in COM calculation	0.35				
DER_0 used in COM calculation	10 ⁻¹²				

^aThe bit error ratio is measured using the PCS errored blocks counter (see 49.2.14.2) or the PMA PRBS31 error counter (see 109.4.4.4) as appropriate.

^bInsertion loss between the two test references (see Figure 110–3b).

110.8.4.2.1 Test setup

The interference tolerance test is performed with the setup shown in Figure 110–3a. The requirements of this subclause are verified at the test references in Figure 110–3a and Figure 110–3b. The cable assembly unused single-ended paths are terminated in 50 Ω to provide 100 Ω differential termination.

NOTE—The cable assembly provides AC-coupling as specified in 110.11.

110.8.4.2.2 Test channel

The test channel (as depicted in Figure 110–3a) consists of the following:

- A cable assembly meeting the requirements of 110.10 and the insertion loss specified for the test being performed.
- A cable assembly test fixture (see 110B.1.2 and 92.11.2).
- A frequency-dependent attenuator.


NOTE—The frequency-dependent attenuator represents a Tx host channel and may be implemented with PCB traces and test cables.

110.8.4.2.3 Test channel calibration

The scattering parameters of the test channel are measured at the test references as illustrated in Figure 110–3b using the cable assembly test fixtures specified in 110B.1.

The insertion loss at 12.8906 GHz of the signal path between the test references in Figure 110–3b is within the limits in Table 110–6, Table 110–7, or Table 110–8, as appropriate for the test being performed.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation



NOTE—The MDI of the host under test is not included in the test channel.

Figure 110-3a—Interference tolerance test setup




Figure 110-3b—Test channel calibration

The COM is calculated using the method and parameters of 110.10.7 with the following considerations:

- The channel signal path is $SCHS_p = \text{cascade}(S^{CTSP}, S^{HOSP})$, where $\text{cascade}()$ is defined in 93A.1.2.1, S^{HOSP} is defined in 110.10.7.1.1, and S^{CTSP} is the measured channel between the test references in Figure 110-3b.
- The COM parameters are as modified by Table 110-6, Table 110-7, or Table 110-8, as appropriate for the test being performed.
- COM is calculated using both Test 1 and Test 2 device package model transmission line lengths listed in Table 110-11 on the receiver side. The value of COM is taken as the lower of the two calculated values.
- The augmented signal path in 93A.1.2 is replaced by S_p determined from Equation (110-1) (effectively omitting the transmitter device package model $S^{(p)}$). The filtered voltage transfer function $H^{(k)}(f)$ calculated in Equation (93A-19) uses the filter $H_t(f)$ defined by Equation (93A-46), where β is 2 and T_r is the 20% to 80% transition time at the Tx test reference. T_r is measured using the method in 86A.5.3.3, with the exception that the observation filter bandwidth is 33 GHz instead of 12 GHz. T_r is measured with the transmit equalizer turned off (i.e., coefficients set to the preset values, see 72.6.10.2.3.1).
- Even-odd jitter, effective bounded uncorrelated jitter and effective random jitter without noise injection (see 110.8.4.2.4) are measured at the Tx test reference and comply with the specification in 92.8.3.8.1 and 92.8.3.8.2. In the calculation of COM, A_{DD} is set to half of the value of $EBUJ$ and σ_{RJ} is set to the value of ERJ , replacing the values in Table 110-11. It is recommended to adjust the

pattern generator jitter such that the effective bounded uncorrelated jitter and the effective total uncorrelated jitter (see 92.8.3.8.2) are as close as practical to their limits in Table 92–6.

- f) The SNR_{TX} value that results in the required COM value for the test is calculated. The injected noise (see 110.8.4.2.4) is set such that the SNDR, as measured at the Tx test reference using the procedure in 92.8.3.7, matches the calculated SNR_{TX} value.

$$S_p = \text{cascade}(SCHS_p, S^{(rp)}) \quad (110-1)$$

where

$\text{cascade}()$	is defined in 93A.1.2.1
$SCHS_p$	is defined in item a) above
$S^{(rp)}$	is defined in 93A.1.2.4

110.8.4.2.4 Pattern generator and noise injection

The pattern generator transmits data to the device under test. At the start of transmitter training, the pattern generator output amplitude shall be 800 mV peak-to-peak differential when measured on an alternating one-zero pattern. The output amplitude, measured on an alternating one-zero pattern, is not permitted to exceed 800 mV peak-to-peak differential during transmitter training. The output waveform of the pattern generator shall comply with 110.8.1.

Broadband noise is added to the signal before the Tx test reference, with noise level set according to step f in 110.8.4.2.3.

110.8.4.2.5 Test procedure

The pattern generator is first configured to transmit the training pattern defined in 110.7.10. During this initialization period, the device under test (DUT) configures the pattern generator transmit equalizer to the coefficient settings it would select using the protocol described in 72.6.10 and the receiver is tuned using its optimization method. The coefficient settings may be communicated via the start-up protocol or by other means.

After the pattern generator equalizer has been configured and the receiver tuned, the pattern generator is set to generate the test pattern specified in Table 110–6, Table 110–7, or Table 110–8, as appropriate for the test being performed. During the test, the transmitter in the device under test transmits the same pattern type specified for the test, with equalization turned off (preset condition).

For 25GBASE-CR and 25GBASE-CR-S PHYs, the receiver under test meets the error requirements specified for the tests in Table 110–7 and Table 110–8. For a 25GBASE-CR PHY, the receiver under test meets the error requirement specified for the test in Table 110–6.

110.8.4.3 Receiver jitter tolerance

Jitter tolerance in RS-FEC mode is measured with a channel meeting the insertion loss of Test 2 and the RS-FEC symbol error ratio requirement specified in Table 110–6. Jitter tolerance in BASE-R FEC mode is measured with a channel meeting the insertion loss of Test 2 and the corrected block ratio requirement specified in Table 110–7. Jitter tolerance in no-FEC mode is measured with a channel meeting the insertion loss of Test 2 and the bit error ratio requirement specified in Table 110–8.

Receiver jitter tolerance is verified for each pair of jitter frequency and peak-to-peak amplitude values listed in Table 110–9. The test setup and procedure of 110.8.4.2 are used, with the exception that no noise is injected (i.e., step f in 110.8.4.2.3 is not performed). In addition, jitter with the specified frequency is applied

to the transmitter and the jitter amplitude is adjusted to obtain the specified peak-to-peak jitter for that frequency at the Tx test reference.

For a 25GBASE-CR-S PHY, the receiver under test shall meet the error requirements specified for the tests in Table 110-7 and Table 110-8, for each case listed in Table 110-9. For a 25GBASE-CR PHY, the receiver under test shall also meet the error requirement specified for the test in Table 110-6 for each case listed in Table 110-9.

Table 110-9—Receiver jitter tolerance parameters

Parameter	Case A values	Case B values	Units
Jitter frequency	190	940	kHz
Peak-to-peak jitter amplitude	5	1	UI

110.8.4.4 Signaling rate range

25GBASE-CR and 25GBASE-CR-S PHYs shall comply with the receiver requirements of 110.8.4.2 and 110.8.4.3 for any signaling rate in the range $25.78125 \text{ GBd} \pm 100 \text{ ppm}$. The corresponding unit interval is approximately 38.787879 ps.

110.9 Channel characteristics

The channel is defined between TP0 and TP5 to include the transmitter and receiver differential controlled impedance printed circuit board and the cable assembly as illustrated in Figure 110-2. The channel insertion loss, return loss, COM, and the transmitter and receiver differential controlled impedance printed circuit board parameters are provided informatively in 110A.4 through 110A.7.

Channel definitions apply for links between two PHYs that can each be either 25GBASE-CR or 25GBASE-CR-S.

110.10 Cable assembly characteristics

Cable assemblies defined in this subclause contain insulated conductors terminated in a connector at each end for use as link segments between MDIs. Cable assemblies are primarily intended as point-to-point links between 25GBASE-CR or 25GBASE-CR-S PHYs using controlled impedance cables. Since 25GBASE-CR and 25GBASE-CR-S PHYs have two specified MDI connectors, single-lane (SFP28, specified in 110.11.1) and multi-lane (QSFP28, specified in 92.12), there are three possible combinations of the connectors at each end. The possible cable assembly types are described in Annex 110C.

All cable assembly measurements are to be made between TP1 and TP4 with cable assembly test fixtures as specified in Annex 110B. These cable assembly specifications are based upon twinaxial cable characteristics, but other cable types are acceptable if the specifications of this subclause are met.

Three cable assembly types are specified, with different COM requirements:

- a) Cable assembly long (CA-25G-L): Cable assembly that supports links between two PHYs that operate in RS-FEC mode with error correction enabled on both receivers, with achievable cable length of at least 5 m.

- b) Cable assembly short (CA-25G-S): Cable assembly that supports links between two PHYs that operate in BASE-R FEC mode, with achievable cable length of at least 3 m.
- c) Cable assembly no-FEC (CA-25G-N): Cable assembly that supports links between two PHYs that operate in no-FEC mode, with achievable cable length of at least 3 m.

NOTE—It may be possible to construct compliant cable assemblies longer than indicated. Length of a cable assembly does not imply compliance to specifications.

Table 110–10 provides a summary of the cable assembly characteristics for CA-25G-L, CA-25G-S, and CA-25G-N, and references to the subclauses addressing each parameter. Due to the improved error correction capability in RS-FEC mode, the specified maximum insertion loss for CA-25G-L is larger than that of CA-25G-S. Likewise, due to the error correction capability provided in BASE-R FEC mode, the specified maximum insertion loss for CA-25G-S is larger than that of CA-25G-N. All other parameters except for some of the input parameters for the COM calculation are identically specified.

Table 110–10—Cable assembly characteristics summary

Description	Reference	CA-25G-L	CA-25G-S	CA-25G-N	Unit
Maximum insertion loss at 12.8906 GHz	110.10.2	22.48	16.48	15.5	dB
Minimum insertion loss at 12.8906 GHz	110.10.2		8		dB
Minimum differential return loss at 12.8906 GHz	110.10.3		6		dB
Differential to common-mode return loss	110.10.4		Equation (92–28)		dB
Differential to common-mode conversion loss	110.10.5		Equation (92–29)		dB
Common-mode to common-mode return loss	110.10.6		Equation (92–30)		dB
COM	110.10.7		See Table 110–11		dB

110.10.1 Characteristic impedance and reference impedance

The nominal differential characteristic impedance of the cable assembly is 100 Ω. The differential reference impedance for cable assembly specifications shall be 100 Ω.

110.10.2 Cable assembly insertion loss

The measured insertion loss of the CA-25G-L, CA-25G-S, and CA-25G-N cable assembly shall be greater than or equal to the minimum cable assembly insertion loss given in Equation (92–26) and illustrated in Figure 92–12.

The measured insertion loss at 12.8906 GHz of the CA-25G-L cable assembly shall be less than or equal to 22.48 dB. The measured insertion loss at 12.8906 GHz of the CA-25G-S cable assembly shall be less than or equal to 16.48 dB. The measured insertion loss at 12.8906 GHz of the CA-25G-N cable assembly shall be less than or equal to 15.5 dB.

110.10.3 Cable assembly differential return loss

The cable assembly differential return loss shall meet the requirements of 92.10.3.

110.10.4 Differential to common-mode return loss

The cable assembly differential to common-mode return loss shall meet the requirements of 92.10.4.

110.10.5 Differential to common-mode conversion loss

The cable assembly differential to common-mode conversion loss shall meet the requirements of 92.10.5.

110.10.6 Common-mode to common-mode return loss

The cable assembly common-mode to common-mode return loss shall meet the requirements of 92.10.6.

110.10.7 Cable assembly Channel Operating Margin

The cable assembly Channel Operating Margin (COM) for each lane is derived from measurements of the cable assembly signal, near-end crosstalk and far-end crosstalk paths. COM is computed using the path calculations defined in 110.10.7.1 and the procedure in 93A.1, where T_r is 8 ps and β is 2 for $H_r(f)$ as used in Equation (93A-19). The specific paths used depend on cable assembly form factor (see Annex 110C), as described in 110.10.7.2.

COM parameter values for the three cable assembly types, CA-25G-L, CA-25G-S, and CA-25G-N, are provided in Table 110-11.

Table 110-11—COM parameter values

Parameter	Symbol	CA-25G-N	CA-25G-S	CA-25G-L ^a	Units
Signaling rate	f_b		25.78125		GBd
Maximum start frequency	f_{min}		0.05		GHz
Maximum frequency step ^b	Δf		0.01		GHz
Device package model					
Single-ended device capacitance	C_d		2.5×10^{-4}		nF
Transmission line length, Test 1	z_p		12		mm
Transmission line length, Test 2	z_p		30		mm
Single-ended package capacitance at package-to-board interface	C_p		1.8×10^{-4}		nF
Single-ended reference resistance	R_0		50		Ω
Single-ended termination resistance	R_d		55		Ω
Receiver 3 dB bandwidth	f_r		$0.75 \times f_b$		GHz
Transmitter equalizer, minimum cursor coefficient	$c(0)$		0.62		—
Transmitter equalizer, pre-cursor coefficient	$c(-1)$				—
Minimum value			-0.18		
Maximum value			0		
Step size			0.02		
Transmitter equalizer, post-cursor coefficient	$c(1)$				—
Minimum value			-0.38		
Maximum value			0		
Step size			0.02		

Table 110–11—COM parameter values (continued)

Parameter	Symbol	CA-25G-N	CA-25G-S	CA-25G-L ^a	Units
Continuous time filter, DC gain	g_{DC}				
Minimum value		–16	–13	–13	dB
Maximum value		0	0	0	dB
Step size		1	1	1	dB
Continuous time filter, zero frequency	f_z	$f_b / 4$			GHz
Continuous time filter, pole frequencies	f_{p1} f_{p2}	$f_b / 4$ f_b			GHz
Transmitter differential peak output voltage					
Victim	A_v		0.4		V
Far-end aggressor	A_{fe}		0.6		V
Near-end aggressor	A_{ne}		0.6		V
Number of signal levels	L	2			—
Level separation mismatch ratio	R_{LM}	1			—
Transmitter signal-to-noise ratio	SNR_{TX}	29	29	29	dB
Number of samples per unit interval	M	32			—
Decision feedback equalizer (DFE) length	N_b	14			—
Normalized DFE coefficient magnitude limit, for $n = 1$ to N_b	$b_{max}(n)$	0.35	0.5	1	—
Random jitter, RMS	σ_{RJ}	0.01			UI
Dual-Dirac jitter, peak	A_{DD}	0.05			UI
One-sided noise spectral density	η_0	5.2×10^{-8}			V ² /GHz
Target detector error ratio	DER_0	10^{-12}	10^{-8}	10^{-5}	—
Channel Operating Margin (min.)	COM	3 ^c	3	3	dB

^aThe parameters for CA-25G-L are the same as those for 100GBASE-CR4 (Table 93–8), except for g_{DC} , A_{fe} , and SNR_{TX} .

^bFor cable lengths greater than 4 m, a frequency step (Δf) no larger than 5 MHz is recommended.

^cFor CA-25G-N cable assemblies with insertion loss at 12.8906 GHz greater than 12 dB, the minimum COM is relaxed to 2.2 dB.

Test 1 and Test 2 differ in the value of the device package model transmission line length z_p . COM for any channel within the cable assembly shall be greater than or equal to the Channel Operating Margin (min.) value specified in Table 110–11 for both Test 1 and Test 2.

110.10.7.1 Channel signal and crosstalk path calculations

The channel paths between TP0 and TP5 used for calculation of the cable assembly COM consist of measured cable assembly signal and crosstalk paths, representative transmitter PCB signal paths, and representative receiver PCB signal paths.

The transmitter and receiver PCB signal paths are calculated using the method defined in 93A.1.2.3. The scattering parameters for a PCB are defined by Equation (93A–13), Equation (93A–14), and the parameter values given in Table 92–12. The PCB trace length parameter z_p has different value for each specific signal path, as specified in 110.10.7.1.1 and 110.10.7.1.2.

The channel path calculations use the function `cascade()` defined in 93A.1.2.1.

110.10.7.1.1 Channel signal path

The scattering parameters of the channel signal path from TP0 to TP5 are calculated using Equation (110–2). The transmitter and receiver PCB signal paths are both denoted as $S^{(HOSP)}$ and are calculated from Equation (93A–13) and Equation (93A–14) using $z_p = 151$ mm in length, representing an insertion loss of 6.26 dB at 12.8906 GHz on each PCB.

$$SCHS_p^{(k)} = \text{cascade}(\text{cascade}(S^{(HOSP)}, S^{(CASP)}), S^{(HOSP)}) \quad (110-2)$$

where

- $SCHS_p^{(k)}$ is the channel signal path
- $S^{(HOSP)}$ is the host (transmitter or receiver) PCB signal path
- $S^{(CASP)}$ is the cable assembly signal path (TP1 to TP4)
- k is equal to zero

110.10.7.1.2 Channel crosstalk paths

The MDI is the significant contributor to crosstalk and is included in and characterized by the cable assembly crosstalk measurements. Crosstalk includes a near-end path where the aggressor is the PMD transmitter, and for specific form factors, near-end and alien far-end crosstalk paths where the aggressors are other PMD transmitters that are connected to the same cable assembly.

For the channel crosstalk paths, the receiver PCB model is $S^{(HOSP)}$ as defined in 110.10.7.1.1. The aggressor transmitter host PCB model is denoted as $S^{(HOTxSP)}$ and is calculated from Equation (93A–13) and Equation (93A–14) using $z_p = 72$ mm in length, representing an insertion loss of 3 dB at 12.8906 GHz. The transmitter host PCB insertion loss is lower than that used for the signal path, to allow for a reasonable worst-case crosstalk in the COM calculation.

The scattering parameters of the channel near-end crosstalk paths are calculated using Equation (110–3). The scattering parameters of the channel alien far-end crosstalk paths are calculated using Equation (110–4).

$$SCHNXT_p^{(k)} = \text{cascade}(\text{cascade}(S^{(HOTxSP)}, S^{(CANXTk)}), S^{(HOSP)}) \quad (110-3)$$

where

- $SCHNXT_p^{(k)}$ is the near-end crosstalk path
- $S^{(HOSP)}$ is the host receiver PCB signal path defined in 110.10.7.1.1
- $S^{(HOTxSP)}$ is the aggressor transmitter PCB signal path
- $S^{(CANXTk)}$ is the cable assembly near-end crosstalk path k (TP1 to TP4)
- k is the index of the near-end crosstalk path

$$SCHAFXT_p^{(k)} = \text{cascade}(\text{cascade}(S^{(HOTxSP)}, S^{(CAFXTk)}), S^{(HOSP)}) \quad (110-4)$$

where

$SCHAFXT_p^{(k)}$	is the alien far-end crosstalk path
$S^{(HOSP)}$	is the host receiver PCB signal path defined in 110.10.7.1.1
$S^{(HOTxSP)}$	is the aggressor transmitter PCB signal path
$S^{(CAFXTk)}$	is the cable assembly far-end crosstalk path k (TP1 to TP4)
k	is the index of the alien far-end crosstalk path

110.10.7.2 Signal and crosstalk paths used in calculation of COM

Cable assemblies have several form factors, as described in Annex 110C. The choice of signal and crosstalk paths for calculation of COM is specific to each cable assembly form factor, as specified in 110.10.7.2.1 through 110.10.7.2.4.

110.10.7.2.1 SFP28 to SFP28

The SFP28 to SFP28 channel structure includes the signal path, one near-end crosstalk path and no alien far-end crosstalk. The signal and near-end crosstalk paths are used in calculation of COM.

The signal path is calculated using Equation (110-2).

The near-end crosstalk path is calculated using Equation (110-3).

110.10.7.2.2 QSFP28 to SFP28

The QSFP28 to SFP28 channel structure includes the signal path, three alien far-end and one near-end crosstalk path. These five paths are used in calculation of COM. Crosstalk from transmitters on other SFP28 connectors is assumed to be insignificant.

The signal path is calculated using Equation (110-2).

The near-end crosstalk path is calculated using Equation (110-3), with k equal to 1.

The three alien far-end crosstalk paths are calculated using Equation (110-4), with k values from 1 to 3.

110.10.7.2.3 SFP28 to QSFP28

The SFP28 to QSFP28 channel structure includes the signal path, three alien far-end and four near-end crosstalk paths. These eight paths are used in calculation of COM.

The signal path is calculated using Equation (110-2).

The near-end crosstalk paths are calculated using Equation (110-3), with k values from 1 to 4.

The three alien far-end crosstalk paths are calculated using Equation (110-4), with k values from 1 to 3.

110.10.7.2.4 QSFP28 to QSFP28

The QSFP28 to QSFP28 channel structure includes the same paths defined for the SFP28 to QSFP28 channel, and COM is calculated in the same way, as defined in 110.10.7.2.3.

110.11 MDI specification

This subclause defines the 25GBASE-CR Media Dependent Interface (MDI), used for both 25GBASE-CR and 25GBASE-CR-S Physical Layers. The MDI couples the PMD (110.7 and 110.8) to the cable assembly (110.10).

The mechanical interface between the PMD and the cable assembly may be either a mated pair of connectors meeting the requirements of 110.11.1 (single-lane MDI) or a mated pair of connectors meeting the requirements of 92.12.1.1 (multi-lane MDI). The plug connector is used on the cable assembly and the receptacle is used on the PMD.

For the multi-lane MDI, each of the paired transmit and receive lanes (SL1, DL1), (SL2, DL2), (SL3, DL3) or (SL4, DL4) may be used for the transmit and receive connections (SL and DL) of a 25GBASE-CR or 25GBASE-CR-S PHY.

For 25GBASE-CR plug connectors, the receive lanes are AC-coupled; the AC-coupling shall be within the plug connectors. It should be noted that there may be various methods for AC-coupling in actual implementations. The low-frequency 3 dB cutoff of the AC-coupling shall be less than 50 kHz. It is recommended that the value of the coupling capacitors be 100 nF. The capacitor limits the inrush charge and baseline wander.

110.11.1 Single-lane MDI connectors

The single-lane MDI uses the SFP+ 28 Gb/s 1X Pluggable (SFP28) plug and receptacle as defined in SFF-8402 and SFF-8432.

The cable assembly connector shall be the SFP28 plug as illustrated in Figure 110–4. The PMD connector shall be the SFP28 receptacle with the mechanical mating interface as illustrated in Figure 110–5. These connectors shall have data signal and signal ground contact assignments as specified in Table 110–12 and electrical performance consistent with the signal quality and electrical requirements of 110.8, 110.9, and 110.10.




Figure 110–4—SFP28 cable assembly plug

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation




Figure 110-5—SFP28 example MDI board receptacle

Table 110-12—Lane to MDI connector contact mapping

Rx lane	MDI connector contact	Tx lane	MDI connector contact
signal gnd	S11	signal gnd	S17
DL<n>	S12	SL<p>	S18
DL<p>	S13	SL<n>	S19
signal gnd	S14	signal gnd	S20

110.12 Environmental specifications

All equipment subject to this clause shall conform to the applicable requirements of 14.7.

110.13 Protocol implementation conformance statement (PICS) proforma for Clause 110, Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-CR and 25GBASE-CR-S⁹

110.13.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 110, Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-CR and 25GBASE-CR-S, 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](#).

110.13.2 Identification

110.13.2.1 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).	

110.13.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3by-2016, Clause 110, Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-CR and 25GBASE-CR-S
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 802.3by-2016.)	

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.

110.13.3 Major capabilities/options

Item ^a	Feature	Subclause	Value/Comment	Status	Support
*CR	25GBASE-CR	110.1	Supports requirements of 25GBASE-CR PHY	O	Yes [] No []
CR-S	25GBASE-CR-S	110.1	Supports requirements of 25GBASE-CR-S PHY	M	Yes []
25GMII	25GMII	110.1	Interface is supported	O	Yes [] No []
PCS	25GBASE-R PCS	110.1		M	Yes []
BR-FEC	BASE-R FEC	110.1	Implemented	M	Yes []
RS-FEC	25GBASE-R RS-FEC	110.1	Implemented	CR:M	Yes [] N/A []
PMA	25GBASE-R PMA	110.1		M	Yes []
AUI	25GAUI C2C	110.1	Interface is supported	O	Yes [] No []
AN	Auto-negotiation	110.1	Device implements Auto-Negotiation	M	Yes []
DC	Delay constraints	110.4	Device conforms to delay constraints specified	M	Yes []
*MD	MDIO capability	110.5	Registers and interface supported	O	Yes [] No []
*EEE	EEE deep sleep capability	110.1	Capability is supported	O	Yes [] No []
*GTD	Global PMD transmit disable function	110.7.5	Function is supported	EEE: M	Yes [] No []
MDIS	Single-lane MDI connector	110.11	Device uses SFP28 as MDI connector	O:2	Yes [] N/A []
MDIQ	Quad-lane MDI connector	110.11	Device uses QSFP28 as MDI connector	O:2	Yes [] N/A []
*CBL	Cable assembly	110.10	Items marked with CBL include cable assembly specifications not applicable to a PHY manufacturer	O	Yes [] No []
CAN	Cable assembly – no FEC	110.10	Cable assembly meets CA-25G-N specifications	CBL: O.1	Yes [] No []
CAS	Cable assembly – short	110.10	Cable assembly meets CA-25G-S specifications	CBL: O.1	Yes [] No []
CAL	Cable assembly – long	110.10	Cable assembly meets CA-25G-L specifications	CBL: O.1	Yes [] No []
FFSS	SFP28 to SFP28	110.10.7.2.1	Cable assembly form factor is SFP28 to SFP28	CBL: O.2	Yes [] No []

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation

Item ^a	Feature	Subclause	Value/Comment	Status	Support
FFQS	QSFP28 to 4×SFP28	110.10.7.2.2 and 110.10.7.2.3	Cable assembly form factor is QSFP28 to 4×SFP28	CBL: O.2	Yes [] No []
FFQQ	QSFP28 to QSFP28	110.10.7.2.4	Cable assembly form factor is QSFP28 to 4×SFP28	CBL: O.2	Yes [] No []

^aA “*” preceding an “Item” identifier indicates there are other PICS that depend on whether or not this item is supported.

110.13.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-CR and 25GBASE-CR-S

110.13.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
PF1	Transmit function	110.7.2	Converts a logical bit streams from the PMD service interface into an electrical signal and delivers it to the MDI	M	Yes []
PF2	Transmitter signal	110.7.2	A positive differential voltage corresponds to tx_bit = one	M	Yes []
PF3	ALERT signal	110.7.2	Transmit a periodic sequence, where each period of the sequence consists of 8 ones followed by 8 zeros, when tx_mode is set to ALERT	EEE:M	Yes [] N/A []
PF4	Receive function	110.7.2	Converts an electrical signal from the MDI into a logical bit stream and delivers it to the PMD service interface	M	Yes []
PF5	Receiver signal	110.7.2	A positive differential voltage corresponds to rx_bit = one	M	Yes []
PF6	When training disabled by management	110.7.2	Global_PMD_signal_detect set to one	M	Yes []
PF7	Global_PMD_signal_detect asserted, rx_mode = QUIET	110.7.4	Set to one within 500 ns following the application of the signal defined in 110.7.4 to the input of the channel	EEE:M	Yes []
PF8	Global_PMD_signal_detect not asserted, rx_mode = QUIET	110.7.4	Not set to one when the signal applied to the input of the channel is less than or equal to 70 mV peak-to-peak differential	EEE:M	Yes []
PF9	Global_PMD_transmit_disable	110.7.5	Disables all transmitters by forcing a constant output level	GTD:M	Yes [] N/A []
PF10	Global_PMD_transmit_disable effect on loopback	110.7.5	No effect	GTD:M	Yes [] N/A []
PF11	Global PMD transmit disable function, tx_mode transition to QUIET	110.7.5	Turn off transmitter when tx_mode transitions to QUIET from any other value	EEE:M	Yes [] N/A []

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation

Item	Feature	Subclause	Value/Comment	Status	Support
PF12	Global PMD transmit disable function, tx_mode transition from QUIET	110.7.5	Turn on transmitter when tx_mode transitions from QUIET to any other value	EEE:M	Yes [] N/A []
PF13	PMD_fault variable mapping to MDIO	110.7.7	Mapped to the fault bit as specified in 45.2.1.2.3	MD:M	Yes [] N/A []
PF14	PMD_transmit_fault variable mapping to MDIO	110.7.8	Mapped to Transmit fault bit as specified in 45.2.1.7.4	MD:M	Yes [] N/A []
PF15	PMD_receive_fault variable mapping to MDIO	110.7.9	Mapped to Receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [] N/A []
PF16	PMD control function	110.7.10	Uses the same control function as lane 0 of 100GBASE-CR4, as defined in 92.7.12	M	Yes []

110.13.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	PMD_fault function	110.7.7	Mapped to the fault bit as specified in 45.2.1.2.3	MD:M	Yes [] N/A []
MF2	PMD_transmit_fault function	110.7.8	Mapped to the PMD_transmit_fault bit as specified in 45.2.1.7.4	MD:M	Yes [] N/A []
MF3	PMD_receive_fault function	110.7.9	Contributes to the PMA/PMD receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [] N/A []

110.13.4.3 Transmitter specifications

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Test fixture return loss	92.11.1.1	Meets equation constraints	M	Yes []
TC2	Test fixture insertion loss	92.11.1.2	Meets equation constraints	M	Yes []
TC3	Signaling rate	92.8.3.9	25.78125 GBd ± 100 ppm	M	Yes []
TC4	Peak-to-peak differential output voltage	92.8.3.1	Less than or equal to 1200 mV regardless of transmit equalizer setting	M	Yes []
TC5	Peak-to-peak differential output voltage, transmitter disabled	92.8.3.1	Less than or equal to 35 mV	M	Yes []
TC6	DC common-mode output voltage	92.8.3.1	Between 0 V and 1.9 V with respect to signal ground	M	Yes []
TC7	AC common-mode output voltage	92.8.3.1	Less than or equal to 30 mV RMS with respect to signal ground	M	Yes []

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation

Item	Feature	Subclause	Value/Comment	Status	Support
TC8	The peak-to-peak differential output voltage	92.8.3.1	Less than 30 mV within 500 ns of the transmitter being disabled.	EEE:M	Yes []
TC9	The peak-to-peak differential output voltage	92.8.3.1	Greater than 720 mV within 500 ns of the transmitter being enabled	EEE:M	Yes []
TC10	The peak-to-peak differential output voltage	92.8.3.1	Meets the requirements of 92.8.3 within 1 μ s of the transmitter being enabled	EEE:M	Yes []
TC11	DC common-mode output voltage while the transmitter is disabled.	92.8.3.1	Maintained to within ± 150 mV of the value for the enabled transmitter	EEE:M	Yes []
TC12	Common-mode output voltage requirements	92.8.3.1	Met regardless of the transmit equalizer setting	M	Yes []
TC13	Differential output return loss (min)	92.8.3.2	Meets equation constraints	M	Yes []
TC14	Reference impedance for differential return loss measurements	92.8.3.2	100 Ω	M	Yes []
TC15	Common-mode to differential mode output return loss	92.8.3.3	Meets equation constraints	M	Yes []
TC16	Steady-state voltage, v_f	92.8.3.5.2	0.34 V min, 0.6 V max	M	Yes []
TC17	Linear fit pulse peak (min)	110.8.3	$0.49 \times v_f$	M	Yes []
TC18	Coefficient initialization	92.8.3.5.3	Satisfies the requirements of 92.8.3.5.3.	M	Yes []
TC19	Normalized coefficient step size for “increment”	92.8.3.5.4	Between 0.0083 and 0.05	M	Yes []
TC20	Normalized coefficient step size for “decrement”	92.8.3.5.4	Between -0.05 and -0.0083	M	Yes []
TC21	Maximum post-cursor equalization ratio	92.8.3.5.5	Greater than or equal to 4	M	Yes []
TC22	Maximum pre-cursor equalization ratio	92.8.3.5.5	Greater than or equal to 1.54	M	Yes []
TC23	Transmitter output SNDR	92.8.3.7	Greater than or equal to 26 dB	M	Yes []
TC24	Even-odd jitter	92.8.3.8.1	Less than or equal to 0.035 UI regardless of the transmit equalization setting	M	Yes []
TC25	Effective bounded uncorrelated jitter	92.8.3.8.2	Less than or equal to 0.1 UI peak-to-peak regardless of the transmit equalization setting	M	Yes []
TC26	Effective total uncorrelated jitter	92.8.3.8.2	Less than or equal to 0.18 UI RMS regardless of the transmit equalization setting	M	Yes []

110.13.4.4 Receiver specifications

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Input amplitude tolerance, RS-FEC mode	110.8.4.1	PMD BER better than 10^{-5} under described conditions	CR:M	Yes []
RC2	Input amplitude tolerance, BASE-R FEC mode	110.8.4.1	PMD BER better than 10^{-8} under described conditions	M	Yes []
RC3	Input amplitude tolerance, no-FEC mode	110.8.4.1	PMD BER better than 10^{-12} under described conditions	M	Yes []
RC4	Differential input return loss	92.8.4.2	Meets equation constraints	M	Yes []
RC5	Reference impedance for differential return loss measurements	92.8.4.2	100 Ω	M	Yes []
RC6	Common-mode input return loss	92.8.4.3	Meets equation constraints	M	Yes []
RC7	Interference tolerance, RS-FEC mode	110.8.4.2	Satisfy requirements summarized in Table 110–6	CR:M	Yes [] N/A []
RC8	Interference tolerance, BASE-R FEC and no-FEC modes	110.8.4.2	Satisfy requirements summarized in Table 110–7 and Table 110–8	M	Yes []
RC9	Interference tolerance	110.8.4.2	Cable assembly used meets COM requirements	M	Yes []
RC10	Interference tolerance	110.8.4.2.4	Pattern generator amplitude and waveform	M	Yes []
RC11	Jitter tolerance, RS-FEC mode	110.8.4.3	Meets the error requirement specified in Table 110–6 for each case listed in Table 110–9	CR:M	Yes [] N/A []
RC12	Jitter tolerance, BASE-R FEC and no-FEC modes	110.8.4.3	Meets the error requirement specified in Table 110–7 and Table 110–8 for each case listed in Table 110–9	M	Yes []
RC13	Signaling rate range	110.8.4.4	25.78125 GBd \pm 100 ppm	M	Yes []

110.13.4.5 Cable assembly specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CA1	Differential reference impedance	92.10.1	100 Ω	CBL:M	Yes [] N/A []
CA2	Minimum insertion loss at 12.8906 GHz	110.10.2	Per Equation (92–26)	CBL:M	Yes [] N/A []
CA3	Maximum insertion loss at 12.8906 GHz	110.10.2	Less than or equal to 15.5 dB	CAN:M	Yes [] N/A []
CA4	Maximum insertion loss at 12.8906 GHz	110.10.2	Less than or equal to 16.48 dB	CAS:M	Yes [] N/A []
CA5	Maximum insertion loss at 12.8906 GHz	110.10.2	Less than or equal to 22.48 dB	CAL:M	Yes [] N/A []
CA6	Differential return loss	110.10.3	Per Equation (92–27)	CBL:M	Yes [] N/A []
CA7	Differential to common-mode input and output return loss	110.10.4	Per Equation (92–28)	CBL:M	Yes [] N/A []
CA8	Differential to common-mode conversion loss	110.10.5	Per Equation (92–29)	CBL:M	Yes [] N/A []
CA9	Common-mode to common-mode return loss	110.10.6	Per Equation (92–30)	CBL:M	Yes [] N/A []
CA10	Cable assembly Channel Operating Margin (COM)	110.10.7	Greater than or equal to the minimum value (in Table 110–11), both Test 1 and Test 2, for all channels within the cable assembly	CBL:M	Yes [] N/A []
CA11	AC-coupling	110.11	Within the plug connector, on the receive lane, 3 dB cutoff frequency less than 50 kHz	CBL:M	Yes [] N/A []

110.13.4.6 MDI connector specifications

Item	Feature	Subclause	Value/Comment	Status	Support
MDC1	Single-lane MDI connector	110.11.1	SFP28 receptacle.	MDIS:M	Yes [] N/A []
MDC2	Quad-lane MDI connector	92.12	QSFP28 receptacle.	MDIQ:M	Yes [] N/A []

110.13.4.7 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Environmental specifications	110.12	Conform to applicable requirements of 14.7	M	Yes []

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

111. Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S

111.1 Overview

This clause specifies the 25GBASE-KR PMD, the 25GBASE-KR-S PMD, and the baseband medium. The specifications are closely related to those of 100GBASE-KR4 (Clause 93) but with a single lane instead of four lanes. This clause also makes use of Annex 93A, Annex 93B, and Annex 93C.

When forming a complete Physical Layer, the PMD shall be connected as illustrated in Figure 111-1, to the appropriate PMA as shown in Table 111-1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 111-1—Physical Layer clauses associated with the 25GBASE-KR and 25GBASE-KR-S PMDs

Associated clause	25GBASE-KR	25GBASE-KR-S
106—RS	Required	Required
106—25GMII ^a	Optional	Optional
107—PCS	Required	Required
74—BASE-R FEC ^b	Required	Required
108—RS-FEC ^b	Required	N/A
109—PMA	Required	Required
109A—25GAUI C2C	Optional	Optional
73—Auto-Negotiation	Required	Required
78—Energy Efficient Ethernet	Optional	Optional

^aThe 25GMII is an optional interface. However, if the 25GMII is not implemented, a conforming implementation must behave functionally as though the RS and 25GMII were present.

^bFEC sublayers can be enabled or disabled according to the FEC mode (see 111.6).

A 25GBASE-KR PHY operates over a channel meeting the requirements of 111.9.1 or 111.9.2. A 25GBASE-KR-S PHY operates over a channel meeting the requirements of 111.9.2. A 25GBASE-KR PHY interoperates with a 25GBASE-KR-S PHY.

When forming a complete 25GBASE-KR or 25GBASE-KR-S Physical Layer, the link BER requirements depend on the FEC mode (see 111.6) according to the following guidelines.


- a) If a PHY operates in the RS-FEC mode, and the RS-FEC decoder does not bypass error correction (see 108.5.3.2), the link is required to operate with a BER of 10^{-5} or better.
- b) If a PHY operates in the BASE-R FEC mode, the link is required to operate with a BER of 10^{-8} or better.
- c) If a PHY operates in the no-FEC mode, or in the RS-FEC mode with error correction bypassed, the link is required to operate with a BER of 10^{-12} or better.

In this context, a link consists of a compliant PMD transmitter, a compliant PMD receiver, and a channel meeting the requirements of 111.9.1 or 111.9.2.

For a complete Physical Layer, this specification is considered to be satisfied by a frame loss ratio (see 1.4.223) of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap.

25GBASE-KR and 25GBASE-KR-S PHYs with the optional Energy-Efficient Ethernet (EEE) capability may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization.

Figure 111–1 shows the relationship of the PMD and MDI to the ISO/IEC Open System Interconnection (OSI) reference model.



25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE
AN = AUTO-NEGOTIATION
FEC = FORWARD ERROR CORRECTION
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
NOTE 1—CONDITIONAL BASED ON PHY TYPE

Figure 111–1—25GBASE-KR and 25GBASE-KR-S relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

111.2 PMD service interface

This subclause specifies the services provided by the 25GBASE-KR and 25GBASE-KR-S PMDs. The service interface for these PMDs is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data. The PMD translates the encoded data to and from signals suitable for the medium.

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation

The PMD service interface is an instance of the inter-sublayer service interface defined in 105.4. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA.request
PMD:IS_UNITDATA.indication
PMD:IS_SIGNAL.indication

The PMA (or the PMD) continuously sends a bit stream to the PMD (or the PMA) at a nominal signaling rate of 25.78125 GBd.

The SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive corresponds to the variable Global_PMD_signal_detect as defined in 111.7.4. When Global_PMD_signal_detect is one, SIGNAL_OK shall be assigned the value OK. When Global_PMD_signal_detect is zero, SIGNAL_OK shall be assigned the value FAIL. When SIGNAL_OK is FAIL, the PMD:IS_UNITDATA.indication parameter is undefined.

If the optional EEE deep sleep capability is supported, then the PMD service interface includes two additional primitives as follows:

PMD:IS_TX_MODE.request
PMD:IS_RX_MODE.request

111.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 107.4.)

25GBASE-KR and 25GBASE-KR-S PHYs may be extended using a 25GAUI chip-to-chip (C2C) as a physical instantiation of the inter-sublayer service interface between devices. If 25GAUI C2C is instantiated, the AN_LINK(link_status).indication is relayed from the device with the PCS sublayer to the device with the AN sublayer by means at the discretion of the implementer. As examples, the implementer may employ use of pervasive management or employ a dedicated electrical signal to relay the state of link_status as indicated by the PCS sublayer on one device to the AN sublayer on the other device.

111.4 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the PMD, AN, and the medium in one direction shall be no more than 512 bit times (1 pause_quantum or 20.48 ns). It is assumed that the one way delay through the medium is no more than 200 bit times (8 ns).

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 105.5.

111.5 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the PMD. If MDIO is implemented, it shall map MDIO control bits to PMD control variables as shown in Table 111-2, and MDIO status bits to PMD status variables as shown in Table 111-3.

Table 111–2—25GBASE-KR MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable	1.9.0	Global_PMD_transmit_disable
Restart training	BASE-R PMD control	1.150.0	mr_restart_training
Training enable	BASE-R PMD control	1.150.1	mr_training_enable
Polynomial identifier 0	PMD training pattern lane 0	1.1450.12:11	identifier_0
Seed 0	PMD training pattern lane 0	1.1450.10:0	seed_0

Table 111–3—25GBASE-KR MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect	1.10.0	Global_PMD_signal_detect
Receiver status 0	BASE-R PMD status	1.151.0	rx_trained
Frame lock 0	BASE-R PMD status	1.151.1	frame_lock
Start-up protocol status 0	BASE-R PMD status	1.151.2	training
Training failure 0	BASE-R PMD status	1.151.3	training_failure

111.6 FEC modes

A 25GBASE-KR PHY implements the BASE-R FEC sublayer (Clause 74) and the 25GBASE-R RS-FEC sublayer (Clause 108). A 25GBASE-KR-S PHY implements the BASE-R FEC sublayer (Clause 74). Each FEC sublayer can be either enabled or disabled, according to AN resolution or management control.

Three FEC modes are supported:

- a) When the 25GBASE-R RS-FEC sublayer is enabled, the PHY is defined to operate in the RS-FEC mode.
- b) When the BASE-R FEC sublayer is enabled, the PHY is defined to operate in the BASE-R FEC mode.
- c) When no FEC sublayer is enabled, the PHY is defined to operate in the no-FEC mode.

A 25GBASE-KR PHY can operate in RS-FEC, BASE-R FEC, or no-FEC mode. A 25GBASE-KR-S PHY can operate in either BASE-R FEC or no-FEC mode.

The channel specification (111.9) that the PHY supports and the required PMD receiver characteristics (111.8.3) depend on the FEC mode.

The FEC mode is determined using AN (Clause 73) and is used in both transmit direction and receive direction. It is recommended to configure the AN FEC advertisement such that only modes that are compatible with the specific channel are selected.

111.7 PMD functional specifications

111.7.1 Link block diagram

One direction of a 25GBASE-KR or 25GBASE-KR-S link is shown in Figure 111–2.




Figure 111–2—25GBASE-KR or 25GBASE-KR-S link (one direction is illustrated)

111.7.2 PMD transmit function

The PMD transmit function shall convert the bit stream requested by the PMD service interface message `PMD:IS_UNITDATA.request(tx_bit)` into an electrical signal. The electrical signal shall then be delivered to the MDI, according to the transmit electrical specifications in 111.8.2. A positive differential output voltage ($SL<p>$ minus $SL<n>$) shall correspond to $tx_bit = one$.

If the optional EEE deep sleep capability is supported, the following requirements apply. When the PMD service interface message `PMD:IS_TX_MODE.request(tx_mode)` is received with $tx_mode = ALERT$, the PMD transmit function shall transmit a periodic sequence, where each period of the sequence consists of 8 ones followed by 8 zeros, with the transmit equalizer coefficients set to the preset values (see 72.6.10.2.3.1). This sequence is transmitted regardless of the value of tx_bit presented by the `PMD:IS_UNITDATA.request` primitive. When tx_mode is not set to `ALERT`, the transmit equalizer coefficients are set to the values determined via the start-up protocol (see 111.7.10).

111.7.3 PMD receive function

The PMD receive function shall convert the electrical signal from the MDI into a bit stream for delivery to the PMD service interface using the message PMD:IS_UNITDATA.indication(rx_bit). A positive differential input voltage (DL<p> minus DL<n>) shall correspond to rx_bit = one.

111.7.4 Global PMD signal detect function

The Global PMD signal detect function is used by the PMD to indicate the successful completion of the start-up protocol by the PMD control function (see 111.7.10). Global_PMD_signal_detect is set to zero when the value of the variable signal_detect is set to false by the Training state diagram (see Figure 72-5). Global_PMD_signal_detect is set to one when the value of signal_detect is set to true.

If training is disabled by the management variable mr_training_enable (see 111.5), Global_PMD_signal_detect shall be set to one.

If the optional EEE deep sleep capability is supported, the following requirements apply. The value of Global_PMD_signal_detect is set to zero when the PMD service interface message PMD:IS_RX_MODE.request(rx_mode) is initially received with rx_mode = QUIET. While rx_mode is set to QUIET, Global_PMD_signal_detect shall be set to one within 500 ns of the application of the ALERT pattern defined in 111.7.2, with peak-to-peak differential voltage of 720 mV as measured at TP0a, to the differential pair at the input of the channel that connects the transmitter to the receiver. While rx_mode is set to QUIET, Global_PMD_signal_detect shall not be set to one when the voltage applied to the input of the differential pair of the channel that connects the transmitter to the receiver is less than or equal to 60 mV peak-to-peak differential.

When the MDIO is implemented, this function maps the variables to registers and bits as defined in 111.5.

111.7.5 Global PMD transmit disable function

The Global PMD transmit disable function is mandatory if the EEE deep sleep capability is supported and is otherwise optional. When this function is supported, it shall meet the following requirements:

- a) When Global_PMD_transmit_disable variable is set to one, this function shall turn off the transmitter such that it drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage in Table 93-4.
- b) If a PMD fault (111.7.7) is detected, then the PMD may set Global_PMD_transmit_disable to one.
- c) Loopback, as defined in 111.7.6, shall not be affected by Global_PMD_transmit_disable.
- d) The following additional requirements apply when the optional EEE deep sleep capability is supported. The Global PMD transmit disable function shall turn off the transmitter as specified in 93.8.1.3 when tx_mode transitions to QUIET from any other value. The Global PMD transmit disable function shall turn on the transmitter as specified in 93.8.1.3 when tx_mode transitions from QUIET to any other value.

111.7.6 Loopback mode

Local loopback mode is provided by the adjacent PMA (see Clause 109) as a test function. When loopback mode is enabled, transmission requests passed to the transmitter are sent directly to the receiver, overriding any signal detected on its attached link. Note that loopback mode does not affect the state of the transmitter, which continues to send data (unless disabled).

Control of the loopback function is specified in 109.4.2.

NOTE 1—The signal path that is exercised in the loopback mode is implementation specific, but it is recommended that this signal path encompass as much of the circuitry as is practical. The intention of providing this loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

NOTE 2—Placing a network port into loopback mode can be disruptive to a network.

111.7.7 PMD fault function

PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault, and any other implementation specific fault. If the MDIO is implemented, PMD_fault shall be mapped to the fault bit as specified in 45.2.1.2.3.

111.7.8 PMD transmit fault function

The PMD transmit fault function is optional. The faults detected by this function are implementation specific, but the assertion of Global_PMD_transmit_disable is not considered a transmit fault.

If PMD_transmit_fault is set to one, then Global_PMD_transmit_disable should also be set to one.

If the MDIO interface is implemented, then PMD_transmit_fault shall be mapped to the Transmit fault bit as specified in 45.2.1.7.4.

111.7.9 PMD receive fault function

The PMD receive fault function is optional. The faults detected by this function are implementation specific. A fault is indicated by setting the variable PMD_receive_fault to one.

If the MDIO interface is implemented, then PMD_receive_fault shall be mapped to the Receive fault bit specified in 45.2.1.7.5.

111.7.10 PMD control function

25GBASE-KR and 25GBASE-KR-S PMDs shall use the same control function as lane 0 of 100GBASE-CR4, as defined in 92.7.12.

The variables seed_0 and polynomial_0 control the training pattern choice. It is recommended that implementations with multiple PMDs use distinct values of polynomial_0 for PMDs that are coupled by crosstalk.

111.8 Electrical characteristics

111.8.1 MDI

The MDI for 25GBASE-KR and 25GBASE-KR-S PHYs is an implementation-dependent direct electrical connection between the PMD and the medium. The MDI comprises two differential pairs, one for the transmit function and one for the receive function, marked by TP0 and TP5 in Figure 111-2.

Transmitter and receiver characteristics are defined at TP0a and TP5a, which are connected to the MDI through the test fixtures described in 93.8.1.1 and 93.8.2.1.

111.8.2 Transmitter characteristics

Transmitter electrical characteristics at TP0a for 25GBASE-KR and 25GBASE-KR-S shall be the same as those of a single lane of 100GBASE-KR4, as summarized in Table 93–4 and detailed in 93.8.1.1 through 93.8.1.7, with the exception that the value of linear fit pulse peak (min.) is $0.75 \times v_f$.

111.8.3 Receiver characteristics

Receiver electrical characteristics are specified at TP5a. The receiver shall meet the return loss specified in 93.8.2.2 using the test fixture specified in 93.8.2.1. In addition, the requirements in 111.8.3.1 and 111.8.3.2 apply.

111.8.3.1 Receiver interference tolerance

The receiver interference tolerance test setup and method are as specified in 93.8.2.3, for a single lane, with the following considerations:

- a) The test requirements in this subclause replace the test requirements in Table 93–6. The test requirements for RS-FEC mode are provided in Table 111–4. The test requirements for BASE-R FEC mode are provided in Table 111–5. The test requirements for no-FEC mode are provided in Table 111–6.
- b) COM is calculated using both Test 1 and Test 2 device package model transmission line lengths listed in Table 111–8 on the receiver side. The value of COM is taken as the lower of the two calculated values.
- c) The transmitter device package model $S^{(p)}$ is omitted from Equation (93A–3) in the calculation of COM. The filtered voltage transfer function $H^{(k)}(f)$ calculated in Equation (93A–19) uses the filter $H_r(f)$ defined by Equation (93A–46), where β is 2, T_r is calculated as $T_r = 1.09 \times T_{rm} - 4.32$ ps, and T_{rm} is the measured 20% to 80% transition time of the signal at TP0a. T_{rm} is measured using the method in 86A.5.3.3, with the exception that the observation filter bandwidth is 33 GHz instead of 12 GHz. T_{rm} is measured with the transmit equalizer turned off (i.e., coefficients set to the preset values, see 72.6.10.2.3.1).

It is recommended to adjust the test transmitter jitter such that the effective bounded uncorrelated jitter and the effective total uncorrelated jitter are as close as practical to their limits in Table 93–4.

A 25GBASE-KR PHY shall comply with the receiver interference tolerance test requirements for RS-FEC mode, BASE-R FEC mode and no-FEC mode. A 25GBASE-KR-S PHY shall comply with the receiver interference tolerance test requirements for BASE-R FEC mode and no-FEC mode.

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

Table 111-4—25GBASE-KR interference tolerance parameters, RS-FEC mode

Parameter	Test 1 (low loss)		Test 2 (high loss)		Units
	Min	Max	Min	Max	
Insertion loss at 12.89 GHz ^a	30	30.5	35	35.5	dB
COM	—	3	—	3	dB
Test pattern	Scrambled idle encoded by RS-FEC				
RSS_DFE4 ^b	0.05	—	0.05	—	
RS-FEC symbol error ratio required ^c	< 10 ⁻⁴				
<i>b</i> _{max} used in COM calculation	1				
<i>DER</i> ₀ used in COM calculation	10 ⁻⁵				

^aMeasured between TPt and TP5 (see Figure 93C-4).

^bThe parameter RSS_DFE4 is a figure of merit for the test channel that is defined in 93A.2.

^cThe RS-FEC symbol error ratio is measured using the RS-FEC symbol error counter (see 108.6.9).

Table 111-5—25GBASE-KR and 25GBASE-KR-S interference tolerance parameters, BASE-R FEC mode

Parameter	Test 1 (low loss)		Test 2 (high loss)		Units
	Min	Max	Min	Max	
Insertion loss at 12.89 GHz ^a	16	16.5	30	30.5	dB
COM	—	3	—	3	dB
Test pattern	Scrambled idle encoded by BASE-R FEC				
RSS_DFE4 ^b	0.05	—	0.05	—	
BASE-R FEC corrected block ratio required ^{c, d}	< 2.1 × 10 ⁻⁵				
<i>b</i> _{max} used in COM calculation	0.5				
<i>DER</i> ₀ used in COM calculation	10 ⁻⁸				

^aMeasured between TPt and TP5 (see Figure 93C-4).

^bThe parameter RSS_DFE4 is a figure of merit for the test channel that is defined in 93A.2.

^cThe BASE-R FEC corrected block ratio is measured using the FEC corrected blocks counter (see 74.8.4.1).

^dThe FEC uncorrected blocks counter (see 74.8.4.2) is required to indicate zero errors during the test unless the test duration is such that the uncorrected block ratio can be verified to be less than 4.7 × 10⁻¹⁰.

Table 111–6—25GBASE-KR and 25GBASE-KR-S interference tolerance parameters, no-FEC mode

Parameter	Test 1 (low loss)		Test 2 (high loss)		Units
	Min	Max	Min	Max	
Insertion loss at 12.89 GHz ^a	16	16.5	30	30.5	dB
COM	—	3	—	3	dB
Test pattern	Scrambled idle or PRBS31				
RSS_DFE4 ^b	0.05	—	0.05	—	
Bit error ratio required ^c	< 10 ⁻¹²				
<i>b</i> _{max} used in COM calculation	0.35				
<i>DER</i> ₀ used in COM calculation	10 ⁻¹²				

^aMeasured between TPt and TP5 (see Figure 93C–4).

^bThe parameter RSS_DFE4 is a figure of merit for the test channel that is defined in 93A.2.

^cThe bit error ratio is measured using the PCS errored blocks counter (see 49.2.14.2) or the PMA PRBS31 error counter (see 109.4.4.4) as appropriate.

111.8.3.2 Receiver jitter tolerance

Jitter tolerance in the RS-FEC mode is measured with a channel meeting the insertion loss of Test 2 and the RS-FEC symbol error ratio requirement specified in Table 111–4. Jitter tolerance in the BASE-R FEC mode is measured with a channel meeting the insertion loss of Test 2 and the corrected block ratio requirement specified in Table 111–5. Jitter tolerance in the no-FEC mode is measured with a channel meeting the insertion loss of Test 2 and the bit error ratio requirement specified in Table 111–6.

Receiver jitter tolerance is verified for each pair of jitter frequency and peak-to-peak amplitude values listed in Table 111–7. The test setup shown in Figure 93–12, or its equivalent, is used. The synthesizer frequency is set to the specified jitter frequency and the synthesizer output amplitude is adjusted until the specified peak-to-peak jitter amplitude for that frequency is measured at TP0a. The test procedure is the same as the one described in 111.8.3.1, with the exception that no broadband noise is added.

For 25GBASE-KR and 25GBASE-KR-S PHYs, the receiver under test shall meet the error requirement specified for the tests in Table 111–5 and Table 111–6, for each case listed in Table 111–7. For a 25GBASE-KR PHY, the receiver under test shall also meet the error requirement specified for the test in Table 111–4 for each case listed in Table 111–7.

Table 111–7—Receiver jitter tolerance parameters

Parameter	Case A values	Case B values	Units
Jitter frequency	190	940	kHz
Peak-to-peak jitter amplitude	5	1	UI

111.9 Channel characteristics

Channel characteristics are defined by Channel Operating Margin (COM), computed using the procedure in 93A.1, where T_r is 12 ps and β is 2 for $H_i(f)$ as used in Equation (93A-19). The parameters used for calculation of COM are different for channels used to connect two PHYs of type 25GBASE-KR and for channels used to connect two PHYs where one or both are of type 25GBASE-KR-S.

Channels used to connect PHYs that operate in the no-FEC mode of operation or in the RS-FEC mode with error correction bypassed shall meet the COM requirement with DER_0 set to 10^{-12} and b_{max} set to 0.35.

All channels shall have AC-coupling as specified in 93.9.4 and are recommended to meet the insertion loss limits in 93.9.2 and the return loss limits in 93.9.3.

111.9.1 25GBASE-KR channel

A 25GBASE-KR channel is used as a link connecting two 25GBASE-KR PHYs. COM parameter values for this channel are provided in Table 111-8.

COM is calculated for two test cases, Test 1 and Test 2, which differ in the value of the device package model transmission line length z_p . COM shall be greater than or equal to 3 dB for each test.

111.9.2 25GBASE-KR-S channel

A 25GBASE-KR-S channel is used as a link between a pair of 25GBASE-KR-S PHYs, between a pair of 25GBASE-KR PHYs, or between a 25GBASE-KR-S PHY and a 25GBASE-KR PHY. COM parameter values for this channel are provided in Table 111-8.

COM is calculated for two test cases, Test 1 and Test 2, which differ in the value of the device package model transmission line length z_p . COM shall be greater than or equal to 3 dB for each test.

Table 111-8—COM parameter values

Parameter	Symbol	25GBASE-KR ^a	25GBASE-KR-S	Units
Signaling rate	f_b	25.78125		GBd
Maximum start frequency	f_{min}	0.05		GHz
Maximum frequency step	Δf	0.01		GHz
Device package model				
Single-ended device capacitance	C_d	2.5×10^{-4}		nF
Transmission line length, Test 1	z_p	12		mm
Transmission line length, Test 2	z_p	30		mm
Single-ended package capacitance at package-to-board interface	C_p	1.8×10^{-4}		nF
Single-ended reference resistance	R_0	50		Ω
Single-ended termination resistance	R_d	55		Ω
Receiver 3 dB bandwidth	f_r	$0.75 \times f_b$		GHz
Transmitter equalizer, minimum cursor coefficient	$c(0)$	0.62		—

Table 111–8—COM parameter values (continued)

Parameter	Symbol	25GBASE-KR ^a	25GBASE-KR-S	Units
Transmitter equalizer, pre-cursor coefficient	$c(-1)$			—
Minimum value		−0.18		
Maximum value		0		
Step size		0.02		
Transmitter equalizer, post-cursor coefficient	$c(1)$			—
Minimum value		−0.38		
Maximum value		0		
Step size		0.02		
Continuous time filter, DC gain	g_{DC}			
Minimum value		−13		dB
Maximum value		0		dB
Step size		1		dB
Continuous time filter, zero frequency	f_z		$f_b / 4$	GHz
Continuous time filter, pole frequencies	f_{p1} f_{p2}		$f_b / 4$ f_b	GHz
Transmitter differential peak output voltage				
Victim	A_v	0.43		V
Far-end aggressor	A_{fe}	0.63		V
Near-end aggressor	A_{ne}	0.63		V
Number of signal levels	L	2		—
Level separation mismatch ratio	R_{LM}	1		—
Transmitter signal-to-noise ratio	SNR_{TX}	27		dB
Number of samples per unit interval	M	32		—
Decision feedback equalizer (DFE) length	N_b	14		—
Normalized DFE coefficient magnitude limit, for $n = 1$ to N_b	$b_{\max}(n)$	1	0.5	—
Random jitter, RMS	σ_{RJ}	0.01		UI
Dual-Dirac jitter, peak	A_{DD}	0.05		UI
One-sided noise spectral density	η_0	5.2×10^{-8}		V ² /GHz
Target detector error ratio	DER_0	10^{-5}	10^{-8}	—

^aThe parameters for the 25GBASE-KR channel are the same as those for 100GBASE-KR4 (Table 93–8), except for g_{DC} , A_v , A_{fe} , and A_{ne} .

111.10 Environmental specifications

All equipment subject to this clause shall conform to the applicable requirements of 93.10.

111.11 Protocol implementation conformance statement (PICS) proforma for Clause 111, Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S¹⁰

111.11.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 111, Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S, 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](#).

111.11.2 Identification

111.11.2.1 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).	

111.11.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3by-2016, Clause 111, Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S
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 802.3by-2016.)	

Date of Statement	
-------------------	--

¹⁰Copyright 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.

111.11.3 Major capabilities/options

Item ^a	Feature	Subclause	Value/Comment	Status	Support
*KR	25GBASE-KR	111.1	Supports requirements of 25GBASE-KR PHY	O	Yes [] No []
KR-S	25GBASE-KR-S	111.1	Supports requirements of 25GBASE-KR-S PHY	M	Yes []
25GMII	25GMII	111.1	Interface is supported	O	Yes [] No []
PCS	25GBASE-R PCS	111.1		M	Yes []
BR-FEC	BASE-R FEC	111.1	Implemented	M	Yes []
RS-FEC	25GBASE-R RS-FEC	111.1	Implemented	KR:M	Yes [] N/A []
PMA	25GBASE-R PMA	111.1		M	Yes []
25GAUI	25GAUI	111.1	Interface is supported	O	Yes [] No []
AN	Auto-negotiation	111.1		M	Yes []
DC	Delay constraints	111.4	Conforms to delay constraints specified	M	Yes []
*MD	MDIO capability	111.5	Registers and interface supported	O	Yes [] No []
*EEE	EEE capability	111.1	Capability is supported	O	Yes [] No []
*GTD	Global PMD transmit disable function	111.7.5	Function is supported	EEE:M	Yes [] No []
*CHNL	Channel	111.9	Channel specifications not applicable to a PHY manufacturer.	O	Yes [] No []

^aA “*” preceding an “Item” identifier indicates there are other PICS that depend on whether or not this item is supported.

111.11.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S

111.11.4.1 Functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	PMD transmit function	111.7.2	Converts a logical bit stream from the PMD service interface into an electrical signal and delivers it to the MDI	M	Yes []
FS2	Mapping of logical signals to electrical signals	111.7.2	Positive differential output voltage corresponds to tx_bit = one	M	Yes []
FS3	ALERT signal	111.7.2	Transmit a periodic sequence, where each period of the sequence consists of 8 ones followed by 8 zeros, when tx_mode is set to ALERT	EEE:M	Yes [] N/A []
FS4	PMD receive function	111.7.3	Converts an electrical signal from the MDI into a logical bit stream and delivers it to the PMD service interface	M	Yes []
FS5	Mapping of electrical signals to logical signals	111.7.3	Positive differential input voltage corresponds to rx_bit = one	M	Yes []
FS6	SIGNAL_OK mapping	111.2	Set to OK when Global_PMD_signal_detect is one and set to FAIL when Global_PMD_signal_detect is zero	M	Yes []
FS7	Global_PMD_signal_detect	111.7.4	Set to 1 when training disabled by variable mr_training_enable	M	Yes []
FS8	Global_PMD_signal_detect when rx_mode = QUIET	111.7.4	Set to one within 500 ns following the application of the signal defined in 111.7.2 to the input of the channel	EEE:M	Yes [] N/A []
FS9	Global_PMD_signal_detect when rx_mode = QUIET	111.7.4	Not set to one when the signal applied to the input of the channel is less than or equal to 60 mV peak-to-peak differential	EEE:M	Yes [] N/A []
FS10	Global_PMD_transmit_disable	111.7.5	When set to one, transmitter is turned off	GTD:M	Yes [] N/A []
FS11	Global PMD transmit disable function affect on loopback	111.7.5	No effect	GTD:M	Yes [] N/A []
FS12	Global PMD transmit disable function, tx_mode transition to QUIET	111.7.5	Turns off transmitter when tx_mode transitions to QUIET from any other value	EEE:M	Yes [] N/A []

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation

Item	Feature	Subclause	Value/Comment	Status	Support
FS13	Global PMD transmit disable function, tx_mode transition from QUIET	111.7.5	Turns on transmitter when tx_mode transitions from QUIET to any other value	EEE:M	Yes [] N/A []
FS14	PMD_fault variable mapping to MDIO	111.7.7	Mapped to the fault bit as specified in 45.2.1.2.3	MD:M	Yes [] N/A []
FS15	PMD_transmit_fault variable mapping to MDIO	111.7.8	Mapped to Transmit fault bit as specified in 45.2.1.7.4	MD:M	Yes [] N/A []
FS16	PMD_receive_fault variable mapping to MDIO	111.7.9	Mapped to Receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [] N/A []
FS17	PMD control function	111.7.10	Defined in 92.7.12	M	Yes []

111.11.4.2 Transmitter characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Test fixture insertion loss	93.8.1.1	Between 1.2 dB and 1.6 dB at 12.89 GHz	M	Yes []
TC2	Test fixture insertion loss deviation	93.8.1.1	Magnitude less than 0.1 dB	M	Yes []
TC3	Test fixture differential return loss	93.8.1.1	Meets equation constraints	M	Yes []
TC4	Test fixture common-mode return loss	93.8.1.1	Greater than or equal to 10 dB from 0.05 to 13 GHz	M	Yes []
TC5	Signaling rate	93.8.1.2	25.78125 GBd ± 100 ppm	M	Yes []
TC6	Peak-to-peak differential output voltage	93.8.1.3	Less than or equal to 1200 mV regardless of transmit equalizer setting	M	Yes []
TC7	Peak-to-peak differential output voltage, transmitter disabled	93.8.1.3	Less than or equal to 30 mV	M	Yes []
TC8	DC common-mode output voltage	93.8.1.3	Between 0 V and 1.9 V with respect to signal ground	M	Yes []
TC9	AC common-mode output voltage	93.8.1.3	Less than or equal to 12 mV RMS with respect to signal ground	M	Yes []
TC10	Common-mode output voltage requirements	93.8.1.3	Met regardless of the transmit equalizer setting	M	Yes []
TC11	Transmitter disable timing	93.8.1.3	Peak-to-peak differential output voltage less than 30 mV within 500 ns of the transmitter being disabled	EEE:M	Yes [] N/A []

IEEE Std 802.3by-2016
IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
Physical Layers, and Management Parameters for 25 Gb/s Operation

Item	Feature	Subclause	Value/Comment	Status	Support
TC12	Transmitter enable timing	93.8.1.3	Peak-to-peak differential output voltage greater than 720 mV within 500 ns of the transmitter being enabled and meet all requirements of 93.8.1 within 1 μ s	EEE:M	Yes [] N/A []
TC13	Common-mode output voltage, transmitter disabled	93.8.1.3	Maintained to within \pm 150 mV of the value for the enabled transmitter	EEE:M	Yes [] N/A []
TC14	Differential input return loss	93.8.1.4	Meets equation constraints	M	Yes []
TC15	Reference impedance for differential return loss measurements	93.8.1.4	100 Ω	M	Yes []
TC16	Common-mode output return loss	93.8.1.4	Meets equation constraints	M	Yes []
TC17	Reference impedance for common-mode return loss measurements	93.8.1.4	25 Ω	M	Yes []
TC18	Steady-state voltage, v_f	93.8.1.5.2	Greater than or equal to 0.4 V and less than or equal to 0.6 V after the transmit equalizer coefficients have been set to the “preset” values	M	Yes []
TC19	Linear fit pulse peak	111.8.2	Greater than $0.75 \times v_f$ after the transmit equalizer coefficients have been set to the “preset” values	M	Yes []
TC20	Coefficient initialization	93.8.1.5.3	Satisfies the requirements of 93.8.1.5.3.	M	Yes []
TC21	Normalized coefficient step size for “increment”	93.8.1.5.4	Between 0.0083 and 0.05	M	Yes []
TC22	Normalized coefficient step size for “decrement”	93.8.1.5.4	Between -0.05 and -0.0083	M	Yes []
TC23	Maximum post-cursor equalization ratio	93.8.1.5.5	Greater than or equal to 4	M	Yes []
TC24	Maximum pre-cursor equalization ratio	93.8.1.5.5	Greater than or equal to 1.54	M	Yes []
TC25	Transmitter output noise and distortion	93.8.1.6	SNDR greater than or equal to 27 dB	M	Yes []
TC26	Even-odd jitter	93.8.1.7	Less than or equal to 0.035 UI regardless of the transmit equalization setting	M	Yes []
TC27	Effective bounded uncorrelated jitter	93.8.1.7	Less than or equal to 0.1 UI peak-to-peak regardless of the transmit equalization setting	M	Yes []
TC28	Effective total uncorrelated jitter	93.8.1.7	Less than or equal to 0.18 UI peak-to-peak regardless of the transmit equalization setting	M	Yes []

111.11.4.3 Receiver characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Test fixture insertion loss	93.8.2.1	Between 1.2 dB and 1.6 dB at 12.89 GHz	M	Yes []
RC2	Test fixture insertion loss deviation	93.8.2.1	Magnitude less than 0.1 dB	M	Yes []
RC3	Test fixture differential return loss	93.8.2.1	Meets equation constraints	M	Yes []
RC4	Test fixture common-mode return loss	93.8.2.1	Greater than or equal to 10 dB from 0.05 GHz to 13 GHz	M	Yes []
RC5	Differential input return loss	93.8.2.2	Meets equation constraints	M	Yes []
RC6	Reference impedance for differential return loss measurements	93.8.2.2	100 Ω	M	Yes []
RC7	Differential to common-mode return loss	93.8.2.2	Meets equation constraints	M	Yes []
RC8	Interference tolerance, RS-FEC mode	111.8.3.1	Satisfy requirements summarized in Table 111-4	KR:M	Yes [] N/A []
RC9	Interference tolerance, BASE-R FEC and no-FEC modes	111.8.3.1	Satisfy requirements summarized in Table 111-5 and Table 111-6	M	Yes []
RC10	Jitter tolerance, RS-FEC mode	111.8.3.2	Meets the error requirement specified in Table 111-4 for each case listed in Table 111-7	KR:M	Yes [] N/A []
RC11	Jitter tolerance, BASE-R FEC and no-FEC modes	111.8.3.2	Meets the error requirement specified in Table 111-5 and Table 111-6 for each case listed in Table 111-7	M	Yes []

111.11.4.4 Channel characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	COM when PHYs use no-FEC mode or RS-FEC mode with error correction bypassed	111.9	Meet requirements with $DER_{\sigma}=10^{-12}$ and $b_{max}=0.35$	CHNL:M	Yes [] N/A []
CC2	AC-coupling	111.9	As specified in 93.9.4	CHNL:M	Yes [] N/A []
CC3	Channel Operating Margin (COM)	111.9.1	Greater than or equal to 3 dB for each test case	CHNL *KR:M	Yes [] N/A []
CC4	Channel Operating Margin (COM)	111.9.2	Greater than or equal to 3 dB for each test case	CHNL *!KR:M	Yes [] N/A []

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation

111.11.4.5 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	Environmental specifications	111.10	Conform to applicable requirements of 93.10	M	Yes []

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd 2:2017

112. Physical Medium Dependent (PMD) sublayer and medium, type 25GBASE-SR**112.1 Overview**

This clause specifies the 25GBASE-SR PMD together with the multimode fiber medium. The PMD sublayer provides a point-to-point 25 Gb/s Ethernet link over a pair of multimode fibers, 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 112–1, 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 112–1—Physical Layer clauses associated with the 25GBASE-SR PMD

Associated clause	25GBASE-SR
106—RS	Required
106—25GMII ^a	Optional
107—PCS for 25GBASE-R	Required
108—RS-FEC ^b	Required
109—PMA for 25GBASE-R	Required
109A—25GAUI C2C	Optional
109B—25GAUI C2M	Optional
78—Energy Efficient Ethernet	Optional

^aThe 25GMII is an optional interface. However, if the 25GMII is not implemented, a conforming implementation must behave functionally as though the RS and 25GMII were present.


^bThe option to bypass the Clause 108 RS-FEC correction function is not supported.

Figure 112–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 25 Gigabit Ethernet is introduced in Clause 105 and the purpose of each PHY sublayer is summarized in 105.2.

25GBASE-SR 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.).

IEEE Std 802.3by-2016
 IEEE Standard for Ethernet—Amendment 2: Media Access Control Parameters,
 Physical Layers, and Management Parameters for 25 Gb/s Operation



25GMII = 25 GIGABIT 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
 RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION
 SR = PMD FOR MULTIMODE FIBER

Figure 112-1—25GBASE-SR PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

112.1.1 Bit error ratio

The bit error ratio (BER) shall be less than 5×10^{-5} provided that the error statistics are sufficiently random that this results in a frame loss ratio (see 1.4.223) of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap when processed according to Clause 108.

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 6.2×10^{-10} for 64-octet frames with minimum interpacket gap when processed according to Clause 108.

112.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 25GBASE-SR 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 105.3. The PMD service interface primitives are summarized as follows:

PMD:IS_UNITDATA.request
PMD:IS_UNITDATA.indication
PMD:IS_SIGNAL.indication

In the transmit direction, the PMA continuously sends a bit stream to the PMD, at a nominal signaling rate of 25.78125 GBd. The PMD converts this stream of bits into appropriate signals on the MDI.

In the receive direction, the PMD continuously sends a bit stream to the PMA corresponding to the signals received from the MDI, at a nominal signaling rate of 25.78125 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 105.4.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_bit parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit 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 112.1.1.

112.3 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 25GBASE-SR PMD including 2 m of fiber in one direction shall be no more than 512 bit times (1 pause_quantum or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quantum can be found in 105.5 and its references.

112.4 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 112–2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 112–3.

Table 112–2—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