

INTERNATIONAL
STANDARD

ISO/IEC/
IEEE
8802-3

Second edition
2017-03-01

AMENDMENT 7
2017-11

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

**Part 3:
Standard for Ethernet**

**AMENDMENT 7: Media access control
parameters, physical layers, and
management parameters for 2.5 Gb/s
and 5 Gb/s operation, types 2.5GBASE-T
and 5GBASE-T**

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

Partie 3: Norme pour Ethernet

*AMENDEMENT 7: Paramètres de commandes d'accès aux
supports, couches physiques et paramètres de gestion en vue d'un
fonctionnement à 2,5 Gb/s et à 5 Gb/s, de types 2.5GBASE-T et
5GBASE-T*



Reference number
ISO/IEC/IEEE 8802-3:2017/Amd.7:2017(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.7 was prepared by the LAN/MAN of the IEEE Computer Society (as IEEE STD 802.3bz-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/AMD7:2017

IEEE Std 802.3bz™-2016

(Amendment to
IEEE Std 802.3™-2015
as amended by
IEEE Std 802.3bw™-2015,
IEEE Std 802.3by™-2016,
IEEE Std 802.3bq™-2016,
IEEE Std 802.3bp™-2016,
IEEE Std 802.3br™-2016, and
IEEE Std 802.3bn™-2016)

IEEE Standard for Ethernet

Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters for 2.5 Gb/s and 5 Gb/s Operation, Types 2.5GBASE-T and 5GBASE-T

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 22 September 2016
IEEE-SA Standards Board

STANDARDSISO.COM :: Click to view the full PDF of ISO/IEC JTC1/SC29/SC32/IEEE 802-3:2017/AMD7:2017

Abstract: Ethernet Media Access Control (MAC) parameters, Physical Layer specifications, and management objects for the transfer of Ethernet format frames at 2.5 Gb/s and 5 Gb/s over balanced twisted-pair transmission media used in structured cabling are defined in this amendment to IEEE Std 802.3-2015.

Keywords: 2.5G/5GBASE-T, amendment, Auto-Negotiation, Ethernet, IEEE 802.3™, IEEE 802.3bz™, Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, structured cabling, XGMII

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 18 October 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 978-1-5044-2370-0 STD21142
PDF: ISBN 978-1-5044-2371-7 STDPD21142

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://standards.ieee.org> 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 P802.3bz Working Group ballot. Individuals may have not voted, voted for approval, disapproval, or abstained on this amendment.

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

Adam Healey, *IEEE 802.3 Working Group Vice-Chair*

Pete Anslow, *IEEE 802.3 Working Group Secretary*

Steven B. Carlson, *IEEE 802.3 Working Group Executive Secretary*

Valerie Maguire, *IEEE 802.3 Working Group Treasurer*

David Chalupsky, *IEEE P802.3bz 2.5G/5GBASE-T Task Force Chair*

George Zimmerman, *IEEE P802.3bz 2.5G/5GBASE-T Task Force Editor-in-Chief*

Eric Baden	Shahar Feldman	Miklos Lukacs
Amrik Bains	German Feyh	Kent Lusted
Thananya Baldwin	Alan Flatman	Jeffery Maki
Denis Beaudoin	Matthias Fritsche	David Malicoat
Christian Beia	Richard Frosch	James Malkemus
Michael Bennett	Andrew Gardner	Yonatan Malkiman
Vipul Bhatt	Ali Ghiasi	Arthur Marris
William Bliss	Joel Goergen	Takeo Masuda
Brad Booth	Steven Gorshe	Kirsten Matheus
Martin Bouda	James Graba	Erdem Matoglu
Ralf-Peter Braun	Robert Grow	Laurence Matola
Theodore Brillhart	Mark Gustlin	Brett McClellan
Paul Brooks	Marek Hajduczenia	Thomas McDermott
Matthew Brown	Bernie Hammond	John McDonough
Jairo Bustos Heredia	Takehiro Hayashi	Richard Mei
Adrian Butter	Yasuo Hidaka	Richard Mellitz
Juan-Carlos Calderon	Brian Holden	Bryan Moffitt
J. Martin Carroll	Rita Horner	Ardeshir Mohammadian
Clark Carty	Bernd Horrmeyer	Leo Montreuil
Craig Chabot	Victor Hou	Paul Mooney
Geoffrey Chacon Simon	Yasuhiro Hyakutake	Dale Murray
Mandeep Chadha	Hideki Isono	Henry Muysshondt
Jacky Chang	Tom Issenhuth	Edward Nakamoto
Xin Chang	Kenneth Jackson	Gary Nicholl
Ahmad Chini	Andrew Jimenez	Paul Nikolich
Keng Hua Chuang	Chad Jones	Kevin Noll
Peter Cibula	Peter Jones	Mark Nowell
Christopher R. Cole	Manabu Kagami	David Ofelt
Shaoan Dai	Upen Kareti	Tom Palkert
John D. Ambrosia	Keisuke Kawahara	Hui Pan
Yan Darshan	Yasuaki Kawatsu	Sesha Panguluri
Piers Dawe	Michael Kelsen	Carlos Pardo
Wael Diab	Yong Kim	Petar Pepeljugoski
Eric DiBiao	Jonathan King	Gerald Pepper
Daniel Dillow	Scott Kipp	Ruben Perez De Aranda
Chris Diminico	Michael Klempa	Alonso
Thuyen Dinh	Shigeru Kobayashi	Michael Peters
Curtis Donahue	Keisuke Kojima	Rick Pimpinella
Dan Dove	Paul Kolesar	William Powell
Mike Dudek	Tom Kolze	Richard Prodan
Nick Duer	Glen Kramer	Rick Rabinovich
David Dwelley	Hans Lackner	Adee Ran
Frank Effenberger	Jeffrey Lapak	Alon Regev
		Duane Remein
		Victor Renteria

Martin Rossbach
 Christopher Roth
 Salvatore Rotolo
 Vineet Salunke
 Sam Sambasivan
 Edward Sayre
 Dieter Schicketanz
 Fred Schindler
 Peter Scruton
 Hossein Sedarat
 Naoshi Serizawa
 Megha Shanbhag
 Stephen Shellhammer
 Ramin Shirani
 Tom Skaar
 Jeff Slavick
 Scott Sommers
 Tom Souvignier
 Edward Sprague
 Peter Stassar

Robert Stone
 Junqing Sun
 Steve Swanson
 Andre Szczepanek
 William Szeto
 Bharat Tailor
 Akio Tajima
 Takayuki Tajima
 Satoshi Takahashi
 Kohichi Tamura
 Alexander Tan
 Mehmet Tazebay
 Brian Teipen
 Geoffrey Thompson
 Pirooz Tooyserkani
 Albert Tretter
 Stephen Trowbridge
 Yoshihiro Tsukamoto
 Mike Tu

Sterling A. Vaden
 Stefano Valle
 Paul Vanderlaan
 Robert Wagner
 Roy Wang
 Xinyuan Wang
 Matthias Wendt
 Oded Wertheim
 Martin White
 Natalie Wienckowski
 Ludwig Winkel
 Peter Wu
 Yu Xu
 Jun Yi
 Lennart Yseboodt
 Ting-Fa Yu
 Hayato Yuki
 Andrew Zambell
 Yan Zhuang
 Helge Zinner

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

Iwan Adhicandra
 Thomas Alexander
 Dale Amason
 Pete Anslow
 Lee Armstrong
 Eric Baden
 Gennaro Boggia
 Christian Boiger
 Ralf-Peter Braun
 Nancy Bravin
 Theodore Brillhart
 Jairo Bustos Heredia
 William Byrd
 Steven B. Carlson
 Juan Carreon
 Clark Carty
 Mandeep Chadha
 David Chalupsky
 Keith Chow
 Peter Cibula
 Charles Cook
 Rodney Cummings
 Shaoan Dai
 John D'Ambrosia
 Janos Farkas
 German Feyh
 Matthias Fritsche
 Yukihiko Fujimoto
 Devon Gayle
 Joel Goergen
 Zhigang Gong
 James Graba
 Randall Groves
 Marek Hajduczenia
 Jerome Henry
 Marco Hernandez
 David Hess
 Werner Hoelzl

Rita Horner
 C. Huntley
 Noriyuki Ikeuchi
 Sergiu Iordanescu
 Atsushi Ito
 Raj Jain
 Peter Jones
 Adri Jovin
 Piotr Karocki
 Stuart Kerry
 Yongbum Kim
 Scott Kipp
 Geoff Ladwig
 Mark Laubach
 David J. Law
 Jon Lewis
 Arthur H. Light
 William Lo
 Michael Lynch
 Elvis Maculuba
 Valerie Maguire
 Arthur Marris
 Michael Maytum
 Mick McCarthy
 Brett McClellan
 Thomas McDermott
 Richard Mellitz
 John Messenger
 Bryan Moffitt
 Michael Newman
 Charles Ngethe
 Mark Nowell
 Satoshi Obara
 Thomas Palkert
 Bansi Patel
 Arumugam Paventhan
 Ruben Perez De Aranda Alonso
 Michael Peters

Rick Pimpinella
 William Powell
 Adee Ran
 Alon Regev
 Maximilian Riegel
 Robert Robinson
 Benjamin Rolfe
 Frank Schewe
 Dieter Schicketanz
 Hossein Sedarat
 Omer Sella
 Takeshi Shimizu
 Thomas Starai
 Walter Struppler
 Mitsutoshi Sugawara
 Junqing Sun
 Patrik Sundstrom
 James Theodoras
 Geoffrey Thompson
 Michael Thompson
 Mark-Rene Uchida
 Alexander Umnov
 Sterling Vaden
 Dmitri Varsanofiev
 Prabodh Varshney
 George Vlantis
 Lisa Ward
 Stephen Webb
 Hung-Yu Wei
 Ludwig Winkel
 Peter Wu
 Dayin Xu
 Jun Xu
 Oren Yuen
 Zhen Zhou
 George Zimmerman

When the IEEE-SA Standards Board approved this standard on 22 September 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
J. Travis Griffith
Gary Hoffman

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

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/AMD7:2017

Introduction

This introduction is not part of IEEE Std 802.3bz-2016, IEEE Standard for Ethernet—Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters for 2.5 Gb/s and 5 Gb/s Operation, Types 2.5GBASE-T and 5GBASE-T.

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 802.3bz-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 Std 802.3 format frames between stations in a subscriber access network. Clause 68 specifies a 10 Gb/s Physical Layer specification. Clause 69 through Clause 74 and associated annexes specify Ethernet operation over electrical backplanes at speeds of 1000 Mb/s and 10 Gb/s.

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

IEEE Std 802.3bw-2015

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

IEEE Std 802.3by-2016

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

IEEE Std 802.3bq-2016

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

IEEE Std 802.3bp-2016

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

IEEE Std 802.3br-2016

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

IEEE Std 802.3bn-2016

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

IEEE Std 802.3bz™-2016

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

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/Amd7:2017

Contents

1. Introduction.....	22
1.1 Overview.....	22
1.1.3 Architectural perspectives.....	22
1.1.3.2 Compatibility interfaces.....	22
1.3 Normative references.....	22
1.4 Definitions.....	22
1.5 Abbreviations.....	24
4. Media Access Control.....	25
4.4 Specific implementations.....	25
4.4.2 MAC parameters.....	25
28. Physical Layer link signaling for Auto-Negotiation on twisted pair.....	26
28.3 State diagrams and variable definitions.....	26
28.3.1 State diagram variables.....	26
30. Management.....	27
30.3 Layer management for DTEs.....	27
30.3.2 PHY device managed object class.....	27
30.3.2.1 PHY device attributes.....	27
30.3.2.1.2 aPhyType.....	27
30.3.2.1.3 aPhyTypeList.....	27
30.5 Layer management for medium attachment units (MAUs).....	27
30.5.1 MAU managed object class.....	27
30.5.1.1 MAU attributes.....	27
30.5.1.1.2 aMAUType.....	27
30.5.1.1.4 aMediaAvailable.....	27
30.5.1.1.24 aLDFastRetrainCount.....	28
30.5.1.1.25 aLPFastRetrainCount.....	28
30.6 Management for link Auto-Negotiation.....	28
30.6.1 Auto-Negotiation managed object class.....	28
30.6.1.1 Auto-Negotiation attributes.....	28
30.6.1.1.5 aAutoNegLocalTechnologyAbility.....	28
45. Management Data Input/Output (MDIO) Interface.....	29
45.1 Overview.....	29
45.2 MDIO Interface Registers.....	29
45.2.1 PMA/PMD registers.....	29
45.2.1.1 PMA/PMD control 1 register (Register 1.0).....	29
45.2.1.1.3 Speed selection (1.0.13, 1.0.6, 1.0.5:2).....	30
45.2.1.4 PMA/PMD speed ability (Register 1.4).....	30
45.2.1.4.aa 5G capable (1.4.14).....	30
45.2.1.4.ab 2.5G capable (1.4.13).....	30
45.2.1.6 PMA/PMD control 2 register (Register 1.7).....	31
45.2.1.7 PMA/PMD status 2 register (Register 1.8).....	31
45.2.1.7.4 Transmit fault (1.8.11).....	31

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

45.2.1.7.5	Receive fault (1.8.10).....	31
45.2.1.8	PMD transmit disable register (Register 1.9)	32
45.2.1.10	PMA/PMD extended ability register (Register 1.11)	32
45.2.1.10.aaa	2.5G/5G extended abilities (1.11.14).....	32
45.2.1.14c	2.5G/5G PMA/PMD extended ability register (Register 1.21)	32
45.2.1.14c.1	5GBASE-T ability (1.21.1).....	32
45.2.1.14c.2	2.5GBASE-T ability (1.21.0).....	33
45.2.1.62	MultiGBASE-T status (Register 1.129).....	33
45.2.1.62.1	LP information valid (1.129.0)	33
45.2.1.64	MultiGBASE-T TX power backoff and PHY short reach setting (Register 1.131).....	33
45.2.1.64.1	MultiGBASE-T TX power backoff settings (1.131.15:10)	33
45.2.1.65	MultiGBASE-T test mode register (Register 1.132)	34
45.2.1.65.1	Test mode control (1.132.15:13).....	34
45.2.1.65.2	Transmitter test frequencies (1.132.12:10)	34
45.2.1.74	RX signal power channel A register (Register 1.141)	34
45.2.1.75	RX signal power channel B register (Register 1.142)	34
45.2.1.76	RX signal power channel C register (Register 1.143)	34
45.2.1.77	RX signal power channel D register (Register 1.144)	35
45.2.1.78	MultiGBASE-T skew delay register (Registers 1.145 and 1.146)	35
45.2.1.79	MultiGBASE-T fast retrain status and control register (Register 1.147)	35
45.2.1.79.1	LP fast retrain count (1.147.15:11)	35
45.2.1.79.2	LD fast retrain count (1.147.10:6)	35
45.2.1.79.5	Fast retrain signal type (1.147.2:1)	36
45.2.1.79.6	Fast retrain enable (1.147.0)	36
45.2.3	PCS registers.....	36
45.2.3.1	PCS control 1 register (Register 3.0)	36
45.2.3.1.2	Loopback (3.0.14).....	36
45.2.3.4	PCS speed ability register (Register 3.4)	37
45.2.3.4.6	2.5G capable (3.4.6).....	37
45.2.3.4.7	5G Capable (3.4.7).....	37
45.2.3.6	PCS control 2 register (Register 3.7)	38
45.2.3.7	PCS status 2 register (Register 3.8)	38
45.2.3.7.1a	5GBASE-T capable (3.8.13).....	38
45.2.3.7.1b	2.5GBASE-T capable (3.8.12).....	38
45.2.3.9a	EEE control and capability 2 register (Register 3.21)	39
45.2.3.9a.2	5GBASE-T EEE supported (3.21.1).....	39
45.2.3.9a.3	2.5GBASE-T EEE supported (3.21.0).....	39
45.2.3.13	BASE-R and MultiGBASE-T PCS status 1 register (Register 3.32)	39
45.2.3.13.1	BASE-R and MultiGBASE-T receive link status (3.32.12)	39
45.2.3.13.4	BASE-R and MultiGBASE-T PCS high BER (3.32.1)	39
45.2.3.13.5	BASE-R and MultiGBASE-T block lock (3.32.0)	40
45.2.3.14	BASE-R and MultiGBASE-T PCS status 2 register (Register 3.33)	40
45.2.3.14.3	BER (3.33.13:8).....	40
45.2.3.14.4	Errored blocks (3.33.7:0)	40
45.2.7	Auto-Negotiation registers.....	41
45.2.7.10	MultiGBASE-T AN control 1 register (Register 7.32).....	41
45.2.7.10.4ca	5GBASE-T capability (7.32.8)	41
45.2.7.10.4cb	2.5GBASE-T capability (7.32.7)	41
45.2.7.10.4cc	5GBASE-T Fast retrain ability (7.32.6).....	42
45.2.7.10.4cd	2.5GBASE-T Fast retrain ability (7.32.5).....	42
45.2.7.11	MultiGBASE-T AN status 1 register (Register 7.33).....	42
45.2.7.11.1	MASTER-SLAVE configuration fault (7.33.15)	42
45.2.7.11.2	MASTER-SLAVE configuration resolution (7.33.14).....	43

45.2.7.11.7ba	Link partner 5GBASE-T capability (7.33.6)	43
45.2.7.11.7bb	Link partner 2.5GBASE-T capability (7.33.5)	43
45.2.7.11.7bc	5GBASE-T Fast retrain ability (7.33.4)	43
45.2.7.11.7bd	2.5GBASE-T Fast retrain ability (7.33.3)	43
45.2.7.13	EEE advertisement 1 (Register 7.60)	43
45.2.7.14	EEE link partner ability 1 (Register 7.61)	44
45.2.7.14aa	EEE advertisement 2 (Register 7.62)	45
45.2.7.14aa.1	5GBASE-T EEE (7.62.1)	45
45.2.7.14aa.2	2.5GBASE-T EEE (7.62.0)	45
45.2.7.14ab	EEE link partner ability 2 (Register 7.63)	45
45.2.7.14a	MultiGBASE-T AN control 2 (Register 7.64)	46
45.2.7.14a.a	2.5GBASE-T THP Bypass Request	46
45.2.7.14a.b	5GBASE-T THP Bypass Request	47
45.2.7.14b	MultiGBASE-T AN status 2 (Register 7.65)	47
45.2.7.14b.a	2.5GBASE-T Link Partner THP Bypass Request	47
45.2.7.14b.b	5GBASE-T Link Partner THP Bypass Request	47
45.5	Protocol implementation conformance statement (PICS) proforma for Clause 45, Management Data Input/Output (MDIO) interface	48
45.5.3	PICS proforma tables for the Management Data Input Output (MDIO) interface	48
45.5.3.2	PMA/PMD MMD options	48
45.5.3.3	PMA/PMD management functions	48
45.5.3.9	Auto-Negotiation management functions	49
46.	Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)	50
46.1	Overview	50
46.1.1	Summary of major concepts	50
46.1.2	Application	50
46.1.3	Rate of operation	50
46.3.1.1	TX_CLK (10 Gb/s transmit clock)	51
46.3.2.1	RX_CLK (receive clock)	51
46.5	XGMII electrical characteristics	51
46.6	Protocol implementation conformance statement (PICS) proforma for Clause 46, Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)	52
46.6.3	PICS proforma Tables for Reconciliation Sublayer and 10 Gigabit Media Independent Interface	52
46.6.3.1	General	52
46.6.3.6	XGMII signal functional specifications	52
78.	Energy-Efficient Ethernet (EEE)	53
78.1	Overview	53
78.1.4	PHY types optionally supporting EEE	53
78.2	LPI mode timing parameters description	53
78.3	Capabilities Negotiation	53
78.5	Communication link access latency	54
125.	Introduction to 2.5 Gb/s and 5 Gb/s networks	55
125.1	Overview	55
125.1.1	Scope	55
125.1.2	Relationship of 2.5 Gigabit and 5 Gigabit Ethernet to the ISO OSI reference model	55
125.1.3	Nomenclature	55

125.1.4	Physical Layer signaling systems	56
125.2	Summary of 2.5 Gigabit and 5 Gigabit Ethernet sublayers	57
125.2.1	Reconciliation Sublayer (RS) and Media Independent Interface	57
125.2.2	Physical Coding Sublayer (PCS)	57
125.2.3	Physical Medium Attachment sublayer (PMA).....	57
125.2.4	Auto-Negotiation, type BASE-T	58
125.2.5	Management interface (MDIO/MDC).....	58
125.2.6	Management.....	58
125.3	Delay Constraints.....	58
126.	Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, and baseband medium, types 2.5GBASE-T and 5GBASE-T	59
126.1	Overview.....	59
126.1.1	Nomenclature.....	59
126.1.2	Relationship of 2.5GBASE-T and 5GBASE-T to other standards.....	59
126.1.3	Operation of 2.5GBASE-T and 5GBASE-T.....	60
126.1.3.1	Summary of Physical Coding Sublayer (PCS)	64
126.1.3.2	Summary of Physical Medium Attachment (PMA) sublayer.....	64
126.1.3.3	Summary of EEE capability	65
126.1.4	Signaling.....	66
126.1.5	Interfaces.....	66
126.1.6	Conventions in this clause	67
126.2	2.5GBASE-T and 5GBASE-T service primitives and interfaces	67
126.2.1	Technology Dependent Interface.....	67
126.2.1.1	PMA_LINK.request.....	67
126.2.1.1.1	Semantics of the primitive.....	67
126.2.1.1.2	When generated	68
126.2.1.1.3	Effect of receipt	68
126.2.1.2	PMA_LINK.indication	68
126.2.1.2.1	Semantics of the primitive.....	68
126.2.1.2.2	When generated	68
126.2.1.2.3	Effect of receipt	68
126.2.2	PMA service interface	68
126.2.2.1	PMA_TXMODE.indication.....	70
126.2.2.1.1	Semantics of the primitive.....	70
126.2.2.1.2	When generated	70
126.2.2.1.3	Effect of receipt	70
126.2.2.2	PMA_CONFIG.indication.....	70
126.2.2.2.1	Semantics of the primitive.....	70
126.2.2.2.2	When generated	70
126.2.2.2.3	Effect of receipt	71
126.2.2.3	PMA_UNITDATA.request.....	71
126.2.2.3.1	Semantics of the primitive.....	71
126.2.2.3.2	When generated	71
126.2.2.3.3	Effect of receipt	71
126.2.2.4	PMA_UNITDATA.indication	71
126.2.2.4.1	Semantics of the primitive.....	72
126.2.2.4.2	When generated	72
126.2.2.4.3	Effect of receipt	72
126.2.2.5	PMA_SCRSTATUS.request.....	72
126.2.2.5.1	Semantics of the primitive.....	72
126.2.2.5.2	When generated	72
126.2.2.5.3	Effect of receipt	72

126.2.2.6	PMA_PCSSTATUS.request	72
126.2.2.6.1	Semantics of the primitive	72
126.2.2.6.2	When generated	73
126.2.2.6.3	Effect of receipt	73
126.2.2.7	PMA_RXSTATUS.indication	73
126.2.2.7.1	Semantics of the primitive	73
126.2.2.7.2	When generated	73
126.2.2.7.3	Effect of receipt	73
126.2.2.8	PMA_REMRXSTATUS.request	73
126.2.2.8.1	Semantics of the primitive	73
126.2.2.8.2	When generated	74
126.2.2.8.3	Effect of receipt	74
126.2.2.9	PMA_ALERTDETECT.indication	74
126.2.2.9.1	Semantics of the primitive	74
126.2.2.9.2	When generated	74
126.2.2.9.3	Effect of receipt	74
126.2.2.10	PCS_RX_LPI_STATUS.request	74
126.2.2.10.1	Semantics of the primitive	74
126.2.2.10.2	When generated	75
126.2.2.10.3	Effect of receipt	75
126.2.2.11	PMA_PCSDATAMODE.indication	75
126.2.2.11.1	Semantics of the primitive	75
126.2.2.11.2	When generated	75
126.2.2.11.3	Effect of receipt	75
126.2.2.12	PMA_FR_ACTIVE.indication	75
126.2.2.12.1	Semantics of the primitive	75
126.2.2.12.2	When generated	76
126.2.2.12.3	Effect of receipt	76
126.3	Physical Coding Sublayer (PCS)	76
126.3.1	PCS service interface (XGMII)	76
126.3.2	PCS functions	76
126.3.2.1	PCS Reset function	77
126.3.2.2	PCS Transmit function	78
126.3.2.2.1	Use of blocks	79
126.3.2.2.2	65B-LDPC transmission code	79
126.3.2.2.3	Notation conventions	79
126.3.2.2.4	Transmission order	79
126.3.2.2.5	Block structure	79
126.3.2.2.6	Control codes	82
126.3.2.2.7	Ordered sets	82
126.3.2.2.8	Idle (/I/)	82
126.3.2.2.9	LPI (/LI/)	84
126.3.2.2.10	Start (/S/)	84
126.3.2.2.11	Terminate (/T/)	84
126.3.2.2.12	ordered set (/O/)	84
126.3.2.2.13	Error (/E/)	85
126.3.2.2.14	Transmit process	85
126.3.2.2.15	PCS Scrambler	85
126.3.2.2.16	LDPC framing and LDPC encoder	85
126.3.2.2.17	Substitution for zero-bit fill	86
126.3.2.2.18	PAM16 bit mapping	86
126.3.2.2.19	EEE capability	87
126.3.2.3	PCS Receive function	88
126.3.2.3.1	Frame and block synchronization	89

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

126.3.2.3.2	PCS descrambler.....	90
126.3.2.3.3	Invalid blocks.....	90
126.3.3	Test-pattern generators	91
126.3.4	PMA training side-stream scrambler polynomials	91
126.3.4.1	Generation of bits San, Sbn, Scn, Sdn	92
126.3.4.2	Generation of 4D symbols TAn, TBn, TCn, TDn	92
126.3.4.3	PMA training mode descrambler polynomials	93
126.3.5	LPI signaling.....	93
126.3.5.1	LPI Synchronization	94
126.3.5.2	Quiet period signaling.....	95
126.3.5.3	Refresh period signaling	95
126.3.6	Detailed functions and state diagrams	95
126.3.6.1	State diagram conventions	95
126.3.6.2	State diagram parameters.....	96
126.3.6.2.1	Constants.....	96
126.3.6.2.2	Variables	96
126.3.6.2.3	Timers	98
126.3.6.2.4	Functions.....	99
126.3.6.2.5	Counters	100
126.3.6.3	State diagrams.....	101
126.3.7	PCS management.....	101
126.3.7.1	Status.....	101
126.3.7.2	Counters	102
126.3.7.3	Loopback	109
126.4	Physical Medium Attachment (PMA) sublayer.....	109
126.4.1	PMA functional specifications.....	109
126.4.2	PMA functions	110
126.4.2.1	PMA Reset function.....	110
126.4.2.2	PMA Transmit function.....	110
126.4.2.2.1	Alert signal.....	110
126.4.2.2.2	Link failure signal.....	112
126.4.2.3	PMA transmit disable function	112
126.4.2.3.1	Global PMA transmit disable function	112
126.4.2.3.2	PMA pair by pair transmit disable function.....	112
126.4.2.3.3	PMA MDIO function mapping.....	112
126.4.2.4	PMA Receive function.....	113
126.4.2.5	PHY Control function	114
126.4.2.5.1	Infofield notation	115
126.4.2.5.2	Start of Frame Delimiter	115
126.4.2.5.3	Current transmitter settings.....	115
126.4.2.5.4	Next transmitter settings	116
126.4.2.5.5	Requested transmitter settings	116
126.4.2.5.6	Message field	116
126.4.2.5.7	SNR_margin	117
126.4.2.5.8	Transition counter.....	118
126.4.2.5.9	Coefficient exchange handshake	118
126.4.2.5.10	Ability fields	118
126.4.2.5.11	Reserved fields.....	118
126.4.2.5.12	Vendor-specific field	118
126.4.2.5.13	Coefficient field	119
126.4.2.5.14	CRC16	119
126.4.2.5.15	Startup sequence	119
126.4.2.5.16	Fast retrain function.....	122
126.4.2.6	Link Monitor function	123

126.4.2.7	Refresh Monitor function.....	123
126.4.2.8	Clock Recovery function	123
126.4.3	MDI.....	123
126.4.3.1	MDI signals transmitted by the PHY.....	123
126.4.3.2	Signals received at the MDI.....	125
126.4.4	Automatic MDI/MDI-X configuration	125
126.4.5	State variables	126
126.4.5.1	State diagram variables.....	126
126.4.5.2	Timers	129
126.4.5.3	Functions.....	130
126.4.5.4	Counters	130
126.4.6	State diagrams.....	131
126.4.6.1	PHY Control state diagram.....	131
126.4.6.2	Transition counter state diagrams	132
126.4.6.3	Link Monitor state diagram	134
126.4.6.4	EEE Refresh monitor state diagram.....	135
126.4.6.5	Fast retrain state diagram	136
126.5	PMA electrical specifications	136
126.5.1	Isolation requirement	136
126.5.2	Test modes	136
126.5.2.1	Test fixtures	139
126.5.3	Transmitter electrical specifications	140
126.5.3.1	Maximum output droop	140
126.5.3.2	Transmitter nonlinear distortion	141
126.5.3.3	Transmitter timing jitter.....	141
126.5.3.4	Transmitter power spectral density (PSD) and power level	142
126.5.3.5	Transmit clock frequency	143
126.5.4	Receiver electrical specifications.....	143
126.5.4.1	Receiver differential input signals	143
126.5.4.2	Receiver frequency tolerance.....	143
126.5.4.3	Rejection of External EM Fields	144
126.5.4.4	Alien crosstalk noise rejection.....	144
126.6	Management interfaces	145
126.6.1	Support for Auto-Negotiation.....	145
126.6.1.1	2.5GBASE-T and 5GBASE-T use of registers during Auto-Negotiation	145
126.6.1.2	2.5GBASE-T and 5GBASE-T Auto-Negotiation page use.....	146
126.6.1.3	Sending Next Pages	148
126.6.2	MASTER-SLAVE configuration resolution	148
126.7	Link segment characteristics.....	150
126.7.1	Cabling system characteristics.....	151
126.7.2	Link segment transmission parameters.....	151
126.7.2.1	Insertion loss	152
126.7.2.2	Differential characteristic impedance	152
126.7.2.3	Return loss	152
126.7.2.4	Coupling parameters between duplex channels comprising one link segment	153
126.7.2.4.1	Differential near-end crosstalk.....	153
126.7.2.4.2	Multiple disturber near-end crosstalk (MDNEXT) loss	153
126.7.2.4.3	Multiple disturber power sum near-end crosstalk (PSNEXT) loss.....	154
126.7.2.4.4	Attenuation to crosstalk ratio, far-end (ACRF)	154
126.7.2.4.5	Multiple disturber attenuation to crosstalk ratio, far-end (MDACRF).....	155
126.7.2.4.6	Multiple disturber power sum attenuation to crosstalk ratio, far-end (PS ACRF).....	156
126.7.2.5	Maximum link delay	156
126.7.2.6	Link delay skew	156

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

126.7.3	Coupling parameters between link segments.....	156
126.7.3.1	Alien crosstalk limited signal-to-noise ratio criteria.....	156
126.8	MDI specification	162
126.8.1	MDI connectors	162
126.8.2	MDI electrical specifications	163
126.8.2.1	MDI FEXT.....	163
126.8.2.2	MDI return loss.....	163
126.8.2.3	MDI impedance balance	163
126.8.2.4	MDI fault tolerance.....	164
126.9	Environmental specifications.....	164
126.9.1	General safety	164
126.9.2	Network safety	165
126.9.3	Installation and maintenance guidelines	165
126.9.4	Telephone voltages	165
126.9.5	Electromagnetic compatibility	165
126.9.6	Temperature and humidity	166
126.10	PHY labeling.....	166
126.11	Delay constraints.....	166
126.12	Protocol implementation conformance statement (PICS) proforma for Clause 126— Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, and baseband medium, types 2.5GBASE-T and 5GBASE-T	167
126.12.1	Identification.....	167
126.12.1.1	Implementation identification.....	167
126.12.1.2	Protocol summary	167
126.12.2	Major capabilities/options.....	168
126.12.3	Physical Coding Sublayer (PCS)	168
126.12.3.1	PCS Transmit functions	168
126.12.3.2	PCS Receive functions.....	170
126.12.3.3	Other PCS functions	170
126.12.4	Physical Medium Attachment (PMA)	170
126.12.5	PMA Electrical Specifications.....	173
126.12.6	PMA Management Interface.....	175
126.12.7	Characteristics of the link segment	176
126.12.8	MDI requirements.....	176
126.12.9	General safety and environmental requirements	177
126.12.10	Timing requirements.....	177
Annex 28B (normative)	IEEE 802.3 Selector Base Page definition	178
28B.3	Priority resolution	178
Annex 28C (normative)	Next Page Message Code field definitions.....	179
28C.11	Message code 9—MultiGBASE-T and 1000BASE-T technology message code.....	179
Annex 28D (normative)	Description of extensions to Clause 28 and associated annexes	180
28D.9	Extensions required for Clause 126 (2.5G/5GBASE-T).....	180
Annex 31B (normative)	MAC Control PAUSE operation	181
31B.3	Detailed specification of PAUSE operation.....	181
31B.3.7	Timing considerations for PAUSE operation	181
31B.4	Protocol implementation conformance statement (PICS) proforma for MAC Control	

IEEE Std 802.3bz-2016
Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
for 2.5 Gb/s and 5 Gb/s Operation

PAUSE operation	182
31B.4.3 Major capabilities/options.....	182
31B.4.6 PAUSE command MAC timing considerations	183

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/AMD7:2017

IEEE Standard for Ethernet

Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters for 2.5 Gb/s and 5 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>.

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.¹

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 ~~strike through~~ (to remove old material) and underscore (to add new material). **Delete** removes existing material. **Insert** adds new material without disturbing the existing material. Deletions and insertions may require renumbering. If so, renumbering instructions are given in the editing instruction. **Replace** is used to make changes in figures or equations by removing the existing figure or equation and replacing it with a new one. Editing instructions, change markings, and this NOTE will not be carried over into future editions because the changes will be incorporated into the base standard.

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

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

Change item f) XGMII as follows:

- f) *10 Gigabit Media Independent Interface (XGMII)*. The XGMII is designed to connect a 2.5 Gb/s, 5 Gb/s, or 10 Gb/s capable MAC to a ~~10 Gb/s~~ PHY of the same rate. While conformance with implementation of this interface is not necessary to ensure communication, it allows maximum flexibility in intermixing PHYs and DTEs at 2.5 Gb/s, 5 Gb/s, and 10 Gb/s speeds. The XGMII is intended for use as a chip-to-chip interface. No mechanical connector is specified for use with the XGMII. The XGMII is optional.

1.3 Normative references

Insert the following reference in alphanumeric order:

TIA TSB-5021, Guidelines for the Assessment and Mitigation of Installed Cabling to Support 2.5GBASE-T and 5GBASE-T.

1.4 Definitions

Insert the definitions for 2.5GBASE-T after the definition for 1.4.64 10/10G-EPON and before the definition for 25GBASE, inserted by IEEE Std 802.3by-2016:

1.4.64aa 2.5GBASE-T: IEEE 802.3 Physical Layer specification for a 2.5 Gb/s LAN using four pairs of Category 5e/Class D balanced copper cabling. (See IEEE Std 802.3, Clause 126.)

Insert the definitions for 5GBASE-T after the definition for 1.4.72a 40GBASE-T, inserted by IEEE Std 802.3bq-2016:

1.4.72b 5GBASE-T: IEEE 802.3 Physical Layer specification for a 5 Gb/s LAN using four pairs of Category 5e/Class D balanced copper cabling. (See IEEE Std 802.3, Clause 126.)

Change the definition for 10 Gigabit Media Independent Interface (XGMII) as follows:

1.4.76 10 Gigabit Media Independent Interface (XGMII): The interface between the Reconciliation Sublayer (RS) and the Physical Coding Sublayer (PCS) for 2.5 Gb/s, 5 Gb/s, and 10 Gb/s operation. (See IEEE Std 802.3, Clause 46.)

Insert a new definition for Category 5e balanced cabling, after Category 5 balanced cabling as follows:

1.4.127a Category 5e balanced cabling: Balanced 100 Ω cables and associated connecting hardware whose transmission characteristics are specified up to 100 MHz per ISO/IEC 11801:2002 and ANSI/TIA-568-B.2-2001. (See IEEE Std 802.3, [Clause 14](#), [Clause 25](#), [Clause 40](#), [Clause 33](#), and Clause 126.)

Change the definition for Category 6 balanced cabling, as follows:

1.4.128 Category 6 balanced cabling: Balanced 100 Ω cables and associated connecting hardware whose transmission characteristics are specified up to 250 MHz per ISO/IEC 11801:2002 and ANSI/TIA-568-C.2-2009. (See IEEE Std 802.3, [Clause 14](#), [Clause 25](#), [Clause 40](#), [Clause 55](#), [Clause 33](#), and [Clause 126](#).) (~~i.e., cabling components meet the performance specified in ISO/IEC 11801:2002 and ANSI/TIA-568-C.2). In addition to the requirements outlined in ISO/IEC 11801:1995 and ANSI/TIA-568-C.2, IEEE 802.3 Clause 14, Clause 23, Clause 25, Clause 40, and Clause 55 specify additional requirements for this cabling when used with 10BASE-T, 100BASE-T, and 10GBASE-T.~~ Balanced 100 Ω cables and associated connecting hardware whose transmission characteristics are specified up to 100 MHz per ISO/IEC 11801:2002 and ANSI/TIA-568-B.2-2001.

Change the definition for Category 6A balanced cabling, as follows:

1.4.129 Category 6A balanced cabling: Balanced 100 Ω cables and associated connecting hardware whose transmission characteristics are specified up to 500 MHz (i.e., cabling components meet the performance specified in ISO/IEC 11801:2002 Amendment 2 and ANSI/TIA-568-C.2). In addition to the requirements outlined in ISO/IEC 11801:2002 Amendment 2 and ANSI/TIA-568-C.2, IEEE 802.3 [Clause 14](#), [Clause 23](#), [Clause 25](#), [Clause 40](#), ~~and Clause 55~~, and [Clause 126](#) specify additional requirements for this cabling when used with 10BASE-T, 100BASE-T, 2.5GBASE-T, 5GBASE-T, and 10GBASE-T.

Change the definition for Category 7 balanced cabling, as follows:

1.4.130 Category 7 balanced cabling: Balanced 100 Ω cables and associated connecting hardware whose transmission characteristics are specified up to 600 MHz (i.e., cabling components meet the performance specified in ISO/IEC 11801:2002). In addition to the requirements outlined in ISO/IEC 11801:2002, IEEE 802.3 [Clause 14](#), [Clause 23](#), [Clause 25](#), [Clause 40](#), ~~and Clause 55~~, and [Clause 126](#) specify additional requirements for this cabling when used with 10BASE-T, 100BASE-T, 2.5GBASE-T, 5GBASE-T, and 10GBASE-T.

Change the definition for Category 7A balanced cabling, as follows:

1.4.131 Category 7A balanced cabling: Balanced 100 Ω cables and associated connecting hardware whose transmission characteristics are specified up to 1000 MHz (i.e., cabling components meet the performance specified in ISO/IEC 11801:2002 Amendment 2). In addition to the requirements outlined in ISO/IEC 11801:2002 Amendment 2, IEEE 802.3 [Clause 14](#), [Clause 23](#), [Clause 25](#), [Clause 40](#), ~~and Clause 55~~, and [Clause 126](#) specify additional requirements for this cabling when used with 10BASE-T, 100BASE-T, 2.5GBASE-T, 5GBASE-T, and 10GBASE-T.

Change the definition for Category 8 balanced cabling (as inserted by IEEE Std 802.3bq-2016) as follows:

1.4.131a Category 8 balanced cabling: Balanced 100 Ω cables and associated connecting hardware whose transmission characteristics are specified up to 2000 MHz (i.e., cabling components that meet the Category 8.1 or Category 8.2 requirements specified in ISO/IEC 11801-1 or Category 8 specified in ANSI/TIA-568-C.2-1). In addition to the requirements outlined in ISO/IEC 11801-1 and ANSI/TIA-568-C.2-1, IEEE Std 802.3 [Clause 14](#), [Clause 23](#), [Clause 25](#), [Clause 40](#), [Clause 55](#), ~~and Clause 113~~, and [Clause 126](#) specify additional requirements for this cabling when used with 10BASE-T, ~~100BASE-T~~ 100BASE-T4, 100BASE-TX, 1000BASE-T, 2.5GBASE-T, 5GBASE-T, 10GBASE-T, 25GBASE-T, and 40GBASE-T.

Change the definition for Infocfield (as inserted by IEEE Std 802.3bq-2016) as follows:

1.4.237a Infocfield: A sixteen octet frame transmitted at regular intervals containing messages for startup operation by certain PHYs. (See IEEE Std 802.3, [Clause 55](#), ~~and Clause 113~~, and [Clause 126](#).)

Change the definition for MultiGBASE-T (as inserted by IEEE Std 802.3bq-2016) as follows:

1.4.277a MultiGBASE-T: PHYs that belong to the set of specific BASE-T Ethernet PHYs at speeds in excess of 1000 Mb/s, including 2.5GBASE-T, 5GBASE-T, 10GBASE-T, 25GBASE-T, and 40GBASE-T. [See IEEE Std 802.3, Clause 126 (for both 2.5GBASE-T and 5GBASE-T), Clause 55 (10GBASE-T), and Clause 113 (for both 25GBASE-T and 40GBASE-T).]

1.5 Abbreviations

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

ALSNR alien limited signal-to-noise ratio

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd7:2017

4. Media Access Control

4.4 Specific implementations

4.4.2 MAC parameters

Change Table 4-2 as follows, inserting a new column between 1 Gb/s and 10 Gb/s, and deleting rightmost column (25 Gb/s, 40Gb/s, and 100 Gb/s) (as modified by IEEE Std 802.3by-2016):

Table 4-2—MAC parameters

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

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

Change NOTE 4 (as modified by IEEE Std 802.3by-2016) as follows:

NOTE 4—For 2.5 Gb/s, 5 Gb/s, 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.

28. Physical Layer link signaling for Auto-Negotiation on twisted pair

28.3 State diagrams and variable definitions

28.3.1 State diagram variables

Insert new rows for 2p5GigT and 5GigT in the first list in 28.3.1 (as modified by IEEE Std 802.3bq-2016) above the row for 10GigT as follows:

2p5GigT; represents that the 2.5GBASE-T PMA is the signal source.

5GigT; represents that the 5GBASE-T PMA is the signal source.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/AMD7:2017

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 entries in "APPROPRIATE SYNTAX" after 1000BASE-T1 (inserted by IEEE Std 802.3bp-2016):

2.5GBASE-T	Clause 126 2.5 Gb/s PAM16
5GBASE-T	Clause 126 5 Gb/s PAM16

30.3.2.1.3 aPhyTypeList

Insert the following new entries in "APPROPRIATE SYNTAX" after 1000BASE-T1 (inserted by IEEE Std 802.3bp-2016):

2.5GBASE-T	Clause 126 2.5 Gb/s PAM16
5GBASE-T	Clause 126 5 Gb/s PAM16

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" after 1000BASE-T1 (inserted by IEEE Std 802.3bp-2016):

2.5GBASE-T	Four-pair twisted-pair balanced copper cabling PHY as specified in Clause 126
5GBASE-T	Four-pair twisted-pair balanced copper cabling PHY as specified in Clause 126

30.5.1.1.4 aMediaAvailable

Change the eighth paragraph of 30.5.1.1.4 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bq-2016) as follows:

For 2.5 Gb/s, 5 Gb/s, 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 values OK and Link Interruption map to the enumeration “available”, the value Local Fault maps to the enumeration “not available” and the value Remote Fault maps to the enumeration “remote fault”.

30.5.1.1.24 aLDFastRetrainCount

Change 30.5.1.1.24 aLDFastRetrainCount (as modified by IEEE Std 802.3bq-2016) to include 2.5GBASE-T and 5GBASE-T [as part of the Energy-Efficient Ethernet package (optional)] as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized nonresettable counter. This counter has a maximum increment rate of 1000 counts per second

BEHAVIOUR DEFINED AS:

A count of the number of fast retrains initiated by the local device. The counter can be derived from fr_tx_counter (see 55.4.5.4, and 113.4.5.4, and 126.4.5.4). If a Clause 45 MDIO Interface to the PMA/PMD is present, then this attribute can be derived from the LD fast retrain count register (see 45.2.1.79.2).;

30.5.1.1.25 aLPFastRetrainCount

Change 30.5.1.1.25 aLPFastRetrainCount (as modified by IEEE Std 802.3bq-2016) to include 2.5GBASE-T and 5GBASE-T [as part of the Energy-Efficient Ethernet package (optional)] as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized nonresettable counter. This counter has a maximum increment rate of 1000 counts per second

BEHAVIOUR DEFINED AS:

A count of the number of fast retrains initiated by the link partner. The counter can be derived from fr_rx_counter (see 55.4.5.4, and 113.4.5.4, and 126.4.5.4). If a Clause 45 MDIO Interface to the PMA/PMD is present, then this attribute can be derived from the LP fast retrain count register (see 45.2.1.79.1).;

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 the following new entries in "APPROPRIATE SYNTAX" after 1000BASE-T1 (inserted by IEEE Std 802.3bp-2016):

2.5GBASE-T	2.5GBASE-T PHY as specified in Clause 126
5GBASE-T	5GBASE-T PHY as specified in Clause 126

45. Management Data Input/Output (MDIO) Interface

45.1 Overview

Change third paragraph of 45.1 as follows:

This extension to the MDIO interface is applicable to the following:

- Implementations that operate at speeds of 2.5~~10~~ Gb/s and above.

45.2 MDIO Interface Registers

45.2.1 PMA/PMD registers

Insert rows for 1.20 and 1.21 and change the reserved row for 1.20 through 1.29 in Table 45-3, (as modified by IEEE Std 802.3by-2016) as follows (unchanged rows not shown):

Table 45-3—PMA/PMD registers

Register address	Register name	Subclause
<u>1.20</u>	<u>Reserved</u>	
<u>1.21</u>	<u>2.5G/5G PMA/PMD extended ability</u>	<u>45.2.1.14c</u>
1.20 through 1.29	Reserved	

45.2.1.1 PMA/PMD control 1 register (Register 1.0)

Change Table 45-4 (as modified by IEEE Std 802.3by-2016) as follows (unchanged rows not shown):

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

Bit	Name	Description	R/W
1.0.5:2	Speed selection	5 4 3 2 1 x x x = Reserved 0 1 1 x = Reserved 0 1 1 1 = 5 Gb/s 0 1 1 0 = 2.5 Gb/s 0 1 0 1 = Reserved 0 1 0 0 = 25 Gb/s 0 0 1 1 = 100 Gb/s 0 0 1 0 = 40 Gb/s 0 0 0 1 = 10PASS-TS/2BASE-TL 0 0 0 0 = 10 Gb/s	R/W

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 modified by IEEE Std 802.3by-2016) as follows:

When bits 5 through 2 are set to 0010 the use of a 40G PMA/PMD is selected; when set to 0011 the use of a 100G PMA/PMD is selected; when set to 0100 the use of a 25G PMA/PMD is selected; when set to 0110 the use of a 2.5G PMA/PMD is selected; when set to 0111 the use of a 5G PMA/PMD is selected.

45.2.1.4 PMA/PMD speed ability (Register 1.4)

Change the Reserved row for 1.4.15:12 in Table 45–6 (as modified by IEEE Std 802.3by-2016) and insert rows for 1.4.14, 1.4.13, and 1.4.12 as follows (unchanged rows not shown):

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

Bit(s)	Name	Description	R/W ^a
1.4.15:12	Reserved for future speeds	Value always 0	RO
<u>1.4.14</u>	<u>5G capable</u>	<u>1 = PMA/PMD is capable of operating at 5 Gb/s</u> <u>0 = PMA/PMD is not capable of operating as 5 Gb/s</u>	<u>RO</u>
<u>1.4.13</u>	<u>2.5G capable</u>	<u>1 = PMA/PMD is capable of operating at 2.5 Gb/s</u> <u>0 = PMA/PMD is not capable of operating as 2.5 Gb/s</u>	<u>RO</u>
<u>1.4.12</u>	<u>Reserved for future speeds</u>	<u>Value always 0</u>	<u>RO</u>

^aRO = Read only

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

45.2.1.4.aa 5G capable (1.4.14)

When read as a one, bit 1.4.14 indicates that the PMA/PMD is able to operate at a data rate of 5 Gb/s. When read as a zero, bit 1.4.14 indicates that the PMA/PMD is not able to operate at a data rate of 5 Gb/s.

45.2.1.4.ab 2.5G capable (1.4.13)

When read as a one, bit 1.4.13 indicates that the PMA/PMD is able to operate at a data rate of 2.5 Gb/s. When read as a zero, bit 1.4.13 indicates that the PMA/PMD is not able to operate at a data rate of 2.5 Gb/s.

45.2.1.6 PMA/PMD control 2 register (Register 1.7)

Change the description for bit 1.7.5:0 in Table 45–7 (as modified by IEEE Std 802.3bw-2015, IEEE Std 802.3bp-2016, IEEE Std 802.3by-2016, and IEEE Std 802.3bq-2016) as follows and adjust reserved rows as appropriate (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 <u>1 1 0 0 0 1 = 5GBASE-T PMA</u> <u>1 1 0 0 0 0 = 2.5GBASE-T PMA</u>	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 description location for 2.5GBASE-T and 5GBASE-T before the row for 10GBASE-KR in Table 45–9 as follows (unchanged rows not shown):

Table 45–9—Transmit fault description location

PMA/PMD	Description location
2.5GBASE-T, 5GBASE-T	126.4.2.2

45.2.1.7.5 Receive fault (1.8.10)

Insert description location for 2.5GBASE-T and 5GBASE-T above the row for 10GBASE-KR in Table 45–10 as follows (unchanged rows not shown):

Table 45–10—Receive fault description location

PMA/PMD	Description location
2.5GBASE-T, 5GBASE-T	126.4.2.4

45.2.1.8 PMD transmit disable register (Register 1.9)

Insert description location for 2.5GBASE-T and 5GBASE-T after the row for 10GBASE-KR in Table 45–12 as follows (unchanged rows not shown):

Table 45–12—Transmit disable description location

PMA/PMD	Description location
2.5GBASE-T and 5GBASE-T	126.4.2.3

45.2.1.10 PMA/PMD extended ability register (Register 1.11)

Change the Reserved row for 1.11.15:13 in Table 45-14 (as modified by IEEE Std 802.3by-2016) and insert rows for 1.11.14 and 1.11.13 as follows (unchanged rows not shown):

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

Bit(s)	Name	Description	R/W ^a
1.11.15:13	Reserved	Value always zero	RO
1.11.14	2.5G/5G extended abilities	1 = PMA/PMD has 2.5G/5G extended abilities listed in register 1.21 0 = PMA/PMD does not have 2.5G/5G extended abilities	RO
1.11.13	Reserved	Value always zero	RO

^aRO = Read only

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

45.2.1.10.aaa 2.5G/5G extended abilities (1.11.14)

When read as a one, bit 1.11.14 indicates that the PMA/PMD has 2.5G/5G extended abilities listed in register 1.21. When read as a zero, bit 1.11.14 indicates that the PMA/PMD does not have 2.5G/5G extended abilities.

Insert 45.2.1.14c after 45.2.1.14b (as inserted by IEEE Std 802.3by-2016) as follows:

45.2.1.14c 2.5G/5G PMA/PMD extended ability register (Register 1.21)

The assignment of bits in the 2.5G/5G PMA/PMD extended ability register is shown in Table 45–17c. All of the bits in the PMA/PMD extended ability register are read only; a write to the PMA/PMD extended ability register shall have no effect.

45.2.1.14c.1 5GBASE-T ability (1.21.1)

When read as a one, bit 1.21.1 indicates that the PMA/PMD is able to operate as a 5GBASE-T PMA type. When read as a zero, bit 1.21.1 indicates that the PMA is not able to operate as a 5GBASE-T PMA type.

Table 45–17c—2.5G/5G PMA/PMD extended ability register bit definitions

Bit	Name	Description	R/W ^a
1.21.15:2	Reserved	Value always 0	RO
1.21.1	5GBASE-T ability	1 = PMA/PMD is able to perform 5GBASE-T 0 = PMA/PMD is not able to perform 5GBASE-T	RO
1.21.0	2.5GBASE-T ability	1 = PMA/PMD is able to perform 2.5GBASE-T 0 = PMA/PMD is not able to perform 2.5GBASE-T	RO

^aRO = Read only

45.2.1.14c.2 2.5GBASE-T ability (1.21.0)

When read as a one, bit 1.21.0 indicates that the PMA/PMD is able to operate as a 2.5GBASE-T PMA type. When read as a zero, bit 1.21.0 indicates that the PMA/PMD is not able to operate as a 2.5GBASE-T PMA type.

45.2.1.62 MultiGBASE-T status (Register 1.129)

45.2.1.62.1 LP information valid (1.129.0)

Change the text of 45.2.1.62.1 (as modified by IEEE Std 802.3bq-2016) to add references 2.5G/5GBASE-T as follows:

When read as a one, bit 1.129.0 indicates that the startup protocol defined in 55.4.2.5 (for 10GBASE-T), or 113.4.2.5 (for 25G/40GBASE-T), or 126.4.2.5 (for 2.5G/5GBASE-T) has been completed, and that the contents of bits 1.130.11:0, 1.131.15:10, 1.145.14:8, 1.146.14:8, and 1.146.6:0, which are established during the startup protocol, are valid. When read as a zero, bit 1.129.0 indicates that the startup process has not been completed, and that the contents of these bits that are established during the startup protocol are invalid. A PMA in the MultiGBASE-T set shall return a value of zero in bit 1.129.1 if PMA link_status=FAIL.

45.2.1.64 MultiGBASE-T TX power backoff and PHY short reach setting (Register 1.131)

45.2.1.64.1 MultiGBASE-T TX power backoff settings (1.131.15:10)

Change the text of 45.2.1.64.1 (as modified by IEEE Std 802.3bq-2016) to reference 2.5GBASE-T and 5GBASE-T as follows:

The MultiGBASE-T TX power backoff settings reflects the TX power backoff selected during the startup negotiation process. The 10GBASE-T startup negotiation process and all TX power backoff settings are defined in 55.4.2.5 and 55.4.5.1. The 25GBASE-T and 40GBASE-T startup negotiation process and all TX power backoff settings are defined in 113.4.2.5 and 113.4.5.1. For 2.5GBASE-T and 5GBASE-T, startup negotiation process and all TX power backoff settings are defined in 126.4.2.5 and 126.4.5.1. If LP information valid bit, 1.129.0, is set to one then bits 1.131.15:13 indicates the TX power backoff setting of the link partner and bits 1.131.12:10 indicates the TX power backoff setting of the local device.

45.2.1.65 MultiGBASE-T test mode register (Register 1.132)**45.2.1.65.1 Test mode control (1.132.15:13)**

Change text of 45.2.1.65.1 (as modified by IEEE Std 802.3bq-2016) to include references to 2.5G/5GBASE-T and Clause 126 as follows:

Transmitter test mode operations defined by bits 1.132.15:13, are described for 10GBASE-T in 55.5.2 and Table 55–12, for 2.5G/5GBASE-T in 126.5.2 and Table 126–13, and for 25G/40GBASE-T in 113.5.2 and Table 113–17. The default value for bits 1.132.15:13 is zero.

45.2.1.65.2 Transmitter test frequencies (1.132.12:10)

Change text of 45.2.1.65.2 (as modified by IEEE Std 802.3bq-2016) to include references to 2.5G/5GBASE-T and Clause 126 as follows:

When test mode 4 is selected by setting bits 1.132.15:13 to one, zero, zero respectively, bits 1.132.12:10 select the transmit test frequency as follows in Table 45–57. Detailed use and operation of these transmitter test frequencies is described in 55.5.2, and 113.5.2, and 126.5.2.

45.2.1.74 RX signal power channel A register (Register 1.141)

Change the text of 45.2.1.74 (as modified by IEEE Std 802.3bq-2016) to add reference to 2.5G/5GBASE-T specifications as follows:

The RX signal power channel A register is read only and contains the receive signal power measured at the MDI during training as described in 55.4.3.1, and 113.4.3.1, and 126.4.6.1. The RX signal power should reflect the power measured when the device transitions out of the state PMA_Training_Init_M or MA_Training_Init_S (as appropriate, see 55.4.6.1, and 113.4.6.1, and 126.4.6.1), when the link partner is transmitting with PBO_tx = 4 (8 dB power backoff). It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –20 dBm to 5.5 dBm. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.75 RX signal power channel B register (Register 1.142)

Change the text of 45.2.1.75 (as modified by IEEE Std 802.3bq-2016) to add reference to 2.5G/5GBASE-T specifications as follows:

The RX signal power channel B register is read only and contains the receive signal power measured at the MDI during training as described in 55.4.3.1, and 113.4.3.1, and 126.4.6.1. The RX signal power should reflect the power measured when the device transitions out of the state PMA_Training_Init_M or MA_Training_Init_S (as appropriate, see 55.4.6.1, and 113.4.6.1, and 126.4.6.1), when the link partner is transmitting with PBO_tx = 4 (8 dB power backoff). It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of –20 dBm to 5.5 dBm. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.76 RX signal power channel C register (Register 1.143)

Change the text of 45.2.1.76 (as modified by IEEE Std 802.3bq-2016) to add reference to 2.5G/5GBASE-T specifications as follows:

The RX signal power channel C register is read only and contains the receive signal power measured at the MDI during training as described in 55.4.3.1, and 113.4.3.1, and 126.4.6.1. The RX signal power should reflect the power measured when the device transitions out of the state PMA_Training_Init_M or

MA_Training_Init_S (as appropriate, see ~~55.4.6.1, and 113.4.6.1, and 126.4.6.1~~), when the link partner is transmitting with PBO_tx = 4 (8 dB power backoff). It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of -20 dBm to 5.5 dBm. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.77 RX signal power channel D register (Register 1.144)

Change the text of 45.2.1.77 (as modified by IEEE Std 802.3bq-2016) to add reference to 2.5G/5GBASE-T specifications as follows:

The RX signal power channel D register is read only and contains the receive signal power measured at the MDI during training as described in ~~55.4.3.1, and 113.4.3.1, and 126.4.6.1~~. The RX signal power should reflect the power measured when the device transitions out of the state PMA_Training_Init_M or MA_Training_Init_S (as appropriate, see ~~55.4.3.1, and 113.4.6.1, and 126.4.6.1~~), when the link partner is transmitting with PBO_tx = 4 (8 dB power backoff). It is reported in units of 0.1 dB to an accuracy of 0.5 dB within the range of -20 dBm to 5.5 dBm. The number is in offset two's complement notation, with 0.0 dB represented by 0x8000. Implementation of this register is optional.

45.2.1.78 MultiGBASE-T skew delay register (Registers 1.145 and 1.146)

Change the text of 45.2.1.78 (as modified by IEEE Std 802.3bq-2016) to reference 2.5G/5GBASE-T as follows:

The skew delay register reports the current skew delay on each of the pair with respect to physical pair A (see Table 45-58). It is reported with resolution equal to one symbol period (see ~~55.1.3, and 113.1.2, and 126.1.3~~) of the PHY (e.g., 1.25 ns for 10GBASE-T) to an accuracy of two symbol periods (e.g., 2.5 ns for 10GBASE-T). The number reported is in two's complement notation with positive values representing delay and negative values representing advance with respect to physical pair A. If the delay exceeds the maximum amount that can be represented by the range (-64 symbols to +63 symbols), the field displays the maximum respective value. The value shall be updated at least once per second.

45.2.1.79 MultiGBASE-T fast retrain status and control register (Register 1.147)

45.2.1.79.1 LP fast retrain count (1.147.15:11)

Change the text of 45.2.1.79.1 (as modified by IEEE Std 802.3bq-2016) to include 2.5GBASE-T and 5GBASE-T as follows:

These bits map to fr_rx_counter as defined in ~~55.4.5.4 for 10GBASE-T, and 113.4.5.4 for 25GBASE-T and 40GBASE-T, and 126.4.5.4 for 2.5GBASE-T and 5GBASE-T~~. The counter is a 5-bit count of the number of fast retrains requested by the link partner. These bits shall be reset to all zeros when read or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.1.79.2 LD fast retrain count (1.147.10:6)

Change the text of 45.2.1.79.2 (as modified by IEEE Std 802.3bq-2016) to include 2.5GBASE-T and 5GBASE-T as follows:

These bits map to fr_tx_counter as defined in ~~55.4.5.4 for 10GBASE-T, and 113.4.5.4 for 25GBASE-T and 40GBASE-T, and 126.4.5.4 for 2.5GBASE-T and 5GBASE-T~~. The counter is a 5-bit count of the number of fast retrains requested by the local device. These bits shall be reset to all zeros when read or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.1.79.5 Fast retrain signal type (1.147.2:1)

Change the text of 45.2.1.79.5 (as modified by IEEE Std 802.3bq-2016) to add reference to 126.3.6.2.2 as follows:

For PHYs that support fast retrain, these bits map to fr_sigtype as defined in 55.3.6.2.2, ~~and 113.3.6.2.2~~, and 126.3.6.2.2. When Fast retrain signal type is set to 00, the PMA sends IDLE characters on the receive path during fast retrain. When Fast retrain signal type is set to 01, the PMA sends Local Fault on the receive path during fast retrain. When Fast retrain signal type is set to 10, the PMA sends Link Interruption on the receive path during fast retrain.

45.2.1.79.6 Fast retrain enable (1.147.0)

Change the text of 45.2.1.79.6 (as modified by IEEE Std 802.3bq-2016) to add reference to 126.4.5.1 as follows:

For PHYs that support fast retrain, ~~this bit 1.147.0~~ controls fr_enable as defined in 55.4.5.1, ~~and 113.4.5.1~~, and 126.4.5.1. When PMA reset is executed, this bit is set to one.

NOTE—Setting this bit to zero while a link is up will cause the PHY to stop supporting fast retrain, and the link will drop if the link partner initiates a fast retrain.

45.2.3 PCS registers

45.2.3.1 PCS control 1 register (Register 3.0)

Change the description for bit 3.0.5:2 in Table 45–120 (as modified by IEEE Std 802.3by-2016) as follows, and adjust reserved rows as appropriate (unchanged rows not shown):

Table 45–120—PCS control 1 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.0.5:2	Speed selection	5 4 3 2 1 1 x x = Reserved 1 0 1 x = Reserved 1 0 0 1 = Reserved 1 0 0 0 = 5 Gb/s 0 1 1 1 = 2.5 Gb/s 0 1 1 0 = Reserved	R/W

^aR/W = Read/Write

Change the text of 45.2.3.1.2 (as modified by IEEE Std 802.3bq-2016) to include 2.5GBASE-T and 5GBASE-T as follows:

45.2.3.1.2 Loopback (3.0.14)

When any MultiGBASE-T or the 10GBASE-R mode of operation is selected for the PCS using the PCS type selection field (3.7.1:0), the PCS shall be placed in a loopback mode of operation when bit 3.0.14 is set to a one. When bit 3.0.14 is set to a one, the 10GBASE-R, 10GBASE-T₁ or any PCS in the MultiGBASE-T set shall accept data on the transmit path and return it on the receive path. The speed of the loopback is selected by the PCS control 1 (Register 3.0) defined in 45.2.3.1. The specific behavior of the 10GBASE-R PCS during loopback is specified in 49.2. The specific behavior for the 10GBASE-T PCS during loopback is

specified in 55.3.6.3. The specific behavior for the 25GBASE-T and 40GBASE-T PCS during loopback is specified in 113.3.7.3. The specific behavior for the 2.5GBASE-T or 5GBASE-T PCS during loopback is specified in 126.3.7.3. For all other port types, the PCS loopback functionality is not applicable and writes to this bit shall be ignored and reads from this bit shall return a value of zero.

45.2.3.4 PCS speed ability register (Register 3.4)

Change Table 45–122 (as modified by IEEE Std 802.3by-2016) as follows (unchanged rows not shown):

Table 45–122—PCS speed ability register bit definitions

Bit(s)	Name	Description	R/W ^a
3.4.15:58	Reserved for future speeds	Value always 0	RO
3.4.7	5G capable	1 = PCS is capable of operating at 5 Gb/s 0 = PCS is not capable of operating at 5 Gb/s	RO
3.4.6	2.5G capable	1 = PCS is capable of operating at 2.5 Gb/s 0 = PCS is not capable of operating at 2.5 Gb/s	RO
3.4.5	Reserved	Value always 0	RO

^aRO = Read only

Insert 45.2.3.4.6 and 45.2.3.4.7 after 45.2.3.4.5 (as inserted by IEEE Std 802.3by-2016) as follows:

45.2.3.4.6 2.5G capable (3.4.6)

When read as a one, bit 3.4.6 indicates that the PCS is able to operate at a data rate of 2.5 Gb/s. When read as a zero, bit 3.4.6 indicates that the PCS is not able to operate at a data rate of 2.5 Gb/s.

45.2.3.4.7 5G Capable (3.4.7)

When read as a one, bit 3.4.7 indicates that the PCS is able to operate at a data rate of 5 Gb/s. When read as a zero, bit 3.4.7 indicates that the PCS is not able to operate at a data rate of 5 Gb/s.

45.2.3.6 PCS control 2 register (Register 3.7)

Change the description for bit 3.7.3:0 in Table 45–123 (as modified by IEEE Std 802.3bq-2016) as follows (unchanged rows not shown):

Table 45–123—PCS control 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.7.3:0	PCS type selection	3 2 1 0 1 0 1 x=reserved 1 0 1 1= Select 5GBASE-T PCS type 1 0 1 0= Select 2.5GBASE-T PCS type	R/W

^aR/W = Read/Write

45.2.3.7 PCS status 2 register (Register 3.8)

Change 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
3.8.13:12	Reserved 5GBASE-T capable	Value always 01 = PCS is able to support 2.5GBASE-T PCS type 0 = PCS is not able to support 2.5GBASE-T PCS type	RO
3.8.12	2.5GBASE-T capable	1 = PCS is able to support 5GBASE-T PCS type 0 = PCS is not able to support 5GBASE-T PCS type	RO

^aRO = Read only, LH = Latching high

Insert 45.2.3.7.1a and 45.2.3.7.1b after 45.2.3.7.1 as follows:

45.2.3.7.1a 5GBASE-T capable (3.8.13)

When read as a one, bit 3.8.13 indicates that the PCS is able to support the 5GBASE-T PCS type. When read as a zero, bit 3.8.13 indicates that the PCS is not able to support the 5GBASE-T PCS type.

45.2.3.7.1b 2.5GBASE-T capable (3.8.12)

When read as a one, bit 3.8.12 indicates that the PCS is able to support the 2.5GBASE-T PCS type. When read as a zero, bit 3.8.12 indicates that the PCS is not able to support the 2.5GBASE-T PCS type.

45.2.3.9a EEE control and capability 2 register (Register 3.21)

Change Table 45–125a (as inserted by IEEE Std 802.3bq-2016) as follows:

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

Bit(s)	Name	Description	R/W ^a
3.21.1	5GBASE-T EEE	1 = EEE is supported for 5GBASE-T 0 = EEE is not supported for 5GBASE-T	RO
3.21.4:0	Reserved 2.5GBASE-T EEE	Value always 0 1 = EEE is supported for 2.5GBASE-T 0 = EEE is not supported for 2.5GBASE-T	RO

^a RO = Read only

Insert 45.2.3.9a.2 and 45.2.3.9a.3 after 45.2.3.9a.1 (as inserted by IEEE Std 802.3bq-2016) as follows:

45.2.3.9a.2 5GBASE-T EEE supported (3.21.1)

If the device supports EEE operation for 5GBASE-T as defined in 126.1.3.3, bit 3.21.1 shall be set to one.

45.2.3.9a.3 2.5GBASE-T EEE supported (3.21.0)

If the device supports EEE operation for 2.5GBASE-T as defined in 126.1.3.3, bit 3.21.0 shall be set to one.

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

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

Change the text of 45.2.3.13.1 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bq-2016) to reference 2.5GBASE-T and 5GBASE-T as follows:

When read as a one, bit 3.32.12 indicates that the PCS is in a fully operational state. When read as a zero, bit 3.32.12 indicates that the PCS is not fully operational. This bit is a reflection of the PCS_status variable defined in 49.2.14.1 for 10/25GBASE-R, in 126.3.7.1 for 2.5GBASE-T and 5GBASE-T, in 55.3.6.155.3.7.1 for 10GBASE-T, in 113.3.7.1 for 25GBASE-T and 40GBASE-T, and in 82.3.1 for 40/100GBASE-R.

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

Change the text of the second paragraph of 45.2.3.13.4 (as modified by IEEE Std 802.3bq-2016) to reference 2.5GBASE-T and 5GBASE-T as follows:

For any member of the MultiGBASE-T set, when read as a one, bit 3.32.1 indicates that the 64B/65B receiver is detecting a BER of $\geq 10^{-4}$. When read as a zero, bit 3.32.1 indicates that the 64B/65B receiver is detecting a BER of $< 10^{-4}$. This bit is a direct reflection of the state of the hi_lfer variable in the MultiGBASE-T 64B/65B state diagrams, defined in 126.3.6.2.2 for 2.5GBASE-T and 5GBASE-T, in 55.3.6.1 for 10GBASE-T, and in 113.3.6.2.2 for 25GBASE-T and 40GBASE-T.

45.2.3.13.5 BASE-R and MultiGBASE-T block lock (3.32.0)

Change the text of 45.2.3.13.5 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bq-2016) to reference 2.5GBASE-T and 5GBASE-T as follows:

When read as a one, bit 3.32.0 indicates that the 64B/66B receiver for BASE-R or the 64B/65B receiver for a member of the MultiGBASE-T set has block lock. When read as a zero, bit 3.32.0 indicates that the 64B/66B receiver for BASE-R or the 64B/65B receiver for a member of the MultiGBASE-T set has not achieved block lock. 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. For both the 2.5GBASE-T and 5GBASE-T PCS, the block_lock variable in the 64B/65B state diagram is defined in 126.3.6.2.2. For the 10GBASE-T PCS the block_lock variable in the 64B/65B state diagram is defined in 55.3.2.3. For both the 25GBASE-T and 40GBASE-T PCS the block_lock variable in the 64B/65B state diagram is defined in 113.3.6.2.2. For a multi-lane PCS, this bit indicates that the receiver has both block lock and alignment for all lanes and is identical to 3.50.12 (see 45.2.3.25.1).

45.2.3.14 BASE-R and MultiGBASE-T PCS status 2 register (Register 3.33)**45.2.3.14.3 BER (3.33.13:8)**

Change the text of 45.2.3.14.3 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bq-2016) to include 2.5GBASE-T and 5GBASE-T 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 counter in 126.3.6.2 in 2.5GBASE-T and 5GBASE-T, 55.3.6.2 for 10GBASE-T, and in 113.3.6.2.2 for 25GBASE-T and 40GBASE-T. These bits shall be reset to all zeros when the BASE-R and MultiGBASE-T PCS status 2 register is read by the management function or upon execution of the PCS reset. If the BER high order counter, 3.44 (see 45.2.3.23) is not implemented then these bits shall be held at all ones in the case of overflow.

45.2.3.14.4 Errored blocks (3.33.7:0)

Change the text of 45.2.3.14.4 (as modified by IEEE Std 802.3by-2016 and IEEE Std 802.3bq-2016) to include 2.5GBASE-T and 5GBASE-T 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 counter errored_block_count in 126.3.6.2 in 2.5GBASE-T and 5GBASE-T, 55.3.6.2 for 10GBASE-T and in 113.3.6.2 for 25GBASE-T and 40GBASE-T. These bits shall be reset to all zeros when the errored blocks count is read by the management function or upon execution of the PCS reset. If the Errored blocks high order counter, 3.45 (see 45.2.3.24) is not implemented then these bits shall be held at all ones in the case of overflow.

45.2.7 Auto-Negotiation registers

Change the names of registers 7.60 and 7.61 (EEE advertisement 1 and EEE LP ability 1), and the reserved row (7.62:63) in Table 45–200 (as modified by IEEE Std 802.3bq-2016, IEEE Std 802.3bw-2015, and IEEE Std 802.3bp-2016) to add EEE advertisement 2 register 7.62, and insert row for register 7.63, EEE LP ability 2, after 7.62 as follows (unchanged rows not shown):

Table 45–200—Auto-Negotiation MMD registers

Register address	Register name	Subclause
7.60	EEE advertisement 1	45.2.7.13
7.61	EEE LP link partner ability 1	45.2.7.14
7.62:63	Reserved	45.2.7.14aa
7.63	EEE LP link partner ability 2	45.2.7.14ab

45.2.7.10 MultiGBASE-T AN control 1 register (Register 7.32)

Change the reserved row in Table 45–207 (as modified by IEEE Std 802.3bq-2016) and insert rows for bits 7.32.8, 7.32.7, 7.32.6, and 7.32.5 above the reserved row as follows (unchanged rows not shown):

Table 45–207—MultiGBASE-T AN control 1 register

Bit(s)	Name	Description	R/W ^a
7.32.8	5GBASE-T ability	1 = Advertise PHY as 5GBASE-T capable 0 = Do not advertise PHY as 5GBASE-T capable	R/W
7.32.7	2.5GBASE-T ability	1 = Advertise PHY as 2.5GBASE-T capable 0 = Do not advertise PHY as 2.5GBASE-T capable	R/W
7.32.6	5GBASE-T Fast retrain ability	1 = Advertise PHY as 5GBASE-T fast retrain capable 0 = Do not advertise PHY as 5GBASE-T fast retrain capable	R/W
7.32.5	2.5GBASE-T Fast retrain ability	1 = Advertise PHY as 2.5GBASE-T fast retrain capable 0 = Do not advertise PHY as 2.5GBASE-T fast retrain capable	R/W
7.32.8:4	Reserved	Value always 0	RO

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

Insert 45.2.7.10.4ca through 45.2.7.10.4cd after 45.2.7.10.4c (as inserted by IEEE Std 802.3bq-2016) as follows:

45.2.7.10.4ca 5GBASE-T capability (7.32.8)

Bit 7.32.8 is used to select whether or not Auto-Negotiation advertises the ability to operate as a 5GBASE-T PHY. If bit 7.32.8 is set to one the PHY shall advertise 5GBASE-T PHY capability. If bit 7.32.8 is set to zero the PHY shall not advertise 5GBASE-T PHY capability.

45.2.7.10.4cb 2.5GBASE-T capability (7.32.7)

Bit 7.32.7 is used to select whether or not Auto-Negotiation advertises the ability to operate as a 2.5GBASE-T PHY. If bit 7.32.7 is set to one the PHY shall advertise 2.5GBASE-T PHY capability. If bit 7.32.7 is set to zero the PHY shall not advertise 2.5GBASE-T PHY capability.

45.2.7.10.4cc 5GBASE-T Fast retrain ability (7.32.6)

Bit 7.32.6 is used to select whether or not the 5GBASE-T PHY advertises the ability to support 5GBASE-T fast retrain. Fast retrain ability is exchanged during link training, see 126.4.2.5.10. If bit 7.32.6 is set to one, the PHY shall advertise fast retrain ability. If bit 7.32.6 is set to zero, the PHY shall not advertise fast retrain ability.

45.2.7.10.4cd 2.5GBASE-T Fast retrain ability (7.32.5)

Bit 7.32.5 is used to select whether or not the 2.5GBASE-T PHY advertises the ability to support 2.5GBASE-T fast retrain. Fast retrain ability is exchanged during link training, see 126.4.2.5.10. If bit 7.32.5 is set to one, the PHY shall advertise fast retrain ability. If bit 7.32.5 is set to zero, the PHY shall not advertise fast retrain ability.

45.2.7.11 MultiGBASE-T AN status 1 register (Register 7.33)

Change the reserved row in Table 45–208 (as modified by IEEE Std 802.3bz-2016) and insert three new rows below it as follows (unchanged rows not shown):

Table 45–208—MultiGBASE-T AN status 1 register

Bit(s)	Name	Description	R/W ^a
7.33.6-3	Reserved <u>Link partner 5GBASE-T capability</u>	Value always 0 <u>1 = Link partner is able to operate as 5GBASE-T</u> <u>0 = Link partner is not able to operate as 5GBASE-T</u>	RO
7.33.5	<u>Link partner 2.5GBASE-T capability</u>	<u>1 = Link partner is able to operate as 2.5GBASE-T</u> <u>0 = Link partner is not able to operate as 2.5GBASE-T</u>	RO
7.33.4	<u>5GBASE-T Fast retrain ability</u>	<u>1 = Link partner is capable of 5GBASE-T fast retrain</u> <u>0 = Link partner is not capable of 5GBASE-T fast retrain</u>	RO
7.33.3	<u>2.5GBASE-T Fast retrain ability</u>	<u>1 = Link partner is capable of 2.5GBASE-T fast retrain</u> <u>0 = Link partner is not capable of 2.5GBASE-T fast retrain</u>	RO

^aRO = Read only, SC = Self-clearing, LH = Latching high

45.2.7.11.1 MASTER-SLAVE configuration fault (7.33.15)

Change the text of 45.2.7.11.1 to include 2.5GBASE-T and 5GBASE-T as follows:

MASTER-SLAVE configuration fault bit 7.33.15 shall be set in the event that determination of the MASTER-SLAVE cannot be successfully concluded. MASTER-SLAVE configuration fault, as well as the criteria and method of fault detection, is PHY specific. Additional information regarding the resolution and selection of MASTER-SLAVE and bit 7.33.15 for 10GBASE-T is contained in 55.6.2. Additional information regarding the resolution and selection of MASTER-SLAVE and bit 7.33.15 for 25GBASE-T and 40GBASE-T is contained in 113.6.2. Additional information regarding the resolution and selection of MASTER-SLAVE and bit 7.33.15 for 2.5GBASE-T and 5GBASE-T is contained in 126.6.2. The MASTER-SLAVE configuration fault bit 7.33.15 shall be cleared each time the MultiGBASE-T status register 7.33 is read via the management interface and shall be cleared by a PMA reset of a PHY in the MultiGBASE-T set. Bit 7.33.15 shall self clear upon Auto-Negotiation enable.

45.2.7.11.2 MASTER-SLAVE configuration resolution (7.33.14)

Change the text of 45.2.7.11.2 to include 2.5GBASE-T and 5GBASE-T as follows:

Bit 7.33.14 is determined by the ~~10GBASE-T~~ MASTER-SLAVE configuration resolution function, described in ~~55.6.2~~ or the ~~25GBASE-T/40GBASE-T MASTER-SLAVE configuration resolution function~~ described in ~~113.6.2~~. These are described in 126.6.2, 55.6.2, and 113.6.2, for 2.5G/5GBASE-T, 10GBASE-T, and 25G/40GBASE-T, respectively. If the MASTER-SLAVE configuration resolution bit 7.33.14 is set to one and the Auto-Negotiation complete bit 7.1.5 is set and MASTER-SLAVE configuration fault bit 7.33.15 in the MultiGBASE-T status register is zero, then MASTER mode of operation has been selected. If the MASTER-SLAVE configuration resolution bit 7.33.14 is set to zero and the Auto-Negotiation complete bit 7.1.5 is set and MASTER-SLAVE configuration fault bit 7.33.15 in the MultiGBASE-T status register is zero, then SLAVE mode of operation has been selected. In all other cases, neither SLAVE mode nor MASTER mode has been selected.

Insert the following new subclauses (45.2.7.11.7ba, 45.2.7.11.7bb, 45.2.7.11.7bc, and 45.2.7.11.7bd) after 45.2.7.11.7b, inserted by IEEE Std 802.3bq-2016:

45.2.7.11.7ba Link partner 5GBASE-T capability (7.33.6)

Bit 7.33.6 is only valid when page received bit 7.1.6 is set to one. When read as a one, bit 7.33.6 indicates that the link partner is able to operate as 5GBASE-T. When read as a zero, bit 7.33.6 indicates that the link partner is not able to operate as 5GBASE-T.

45.2.7.11.7bb Link partner 2.5GBASE-T capability (7.33.5)

Bit 7.33.5 is only valid when page received bit 7.1.5 is set to one. When read as a one, bit 7.33.5 indicates that the link partner is able to operate as 2.5GBASE-T. When read as a zero, bit 7.33.5 indicates that the link partner is not able to operate as 2.5GBASE-T.

45.2.7.11.7bc 5GBASE-T Fast retrain ability (7.33.4)

When read as a one, bit 7.33.4 indicates that the link partner has the ability to support the 5GBASE-T fast retrain capability as specified in 126.4.2.5.16. When read as a zero, bit 7.33.4 indicates that the PHY lacks the ability to support the 5GBASE-T fast retrain capability. This bit is only valid after link is established.

45.2.7.11.7bd 2.5GBASE-T Fast retrain ability (7.33.3)

When read as a one, bit 7.33.3 indicates that the link partner has the ability to support the 2.5GBASE-T fast retrain capability as specified in 126.4.2.5.16. When read as a zero, bit 7.33.3 indicates that the PHY lacks the ability to support the 2.5GBASE-T fast retrain capability. This bit is only valid after link is established.

Change the title of 45.2.7.13 and the text of the first paragraph of 45.2.7.13 (as modified by IEEE Std 802.3bq-2016) to rename the register EEE advertisement 1 as follows (unchanged paragraphs not shown):

45.2.7.13 EEE advertisement 1 (Register 7.60)

This register defines the EEE advertisement for several device types. Devices that use Clause 28 Auto-Negotiation send EEE advertisement in the Unformatted Next Page following a EEE technology message code as defined in 28C.12 as part of the 10GBASE-T and 1000BASE-T technology message code as defined in 28C.11. Devices that use Clause 73 Auto-Negotiation send EEE advertisement in the unformatted code field of Message Next Page with EEE technology message code as defined in 73A.4. For 25GBASE-T and 40GBASE-T, the EEE advertisement is exchanged in the Infocfield during training as defined in 113.4.2.5.10.

For 2.5GBASE-T and 5GBASE-T, the EEE advertisement is exchanged in the Infofield during training as defined in 126.4.2.5.10. The assignment of bits in the EEE advertisement 1 register and the correspondence with the bits in the Next Page messages or in the training Infofield are shown in Table 45–210.

Change the title of Table 45–210 as follows:

Table 45–210—EEE advertisement 1 register (Register 7.60) bit definitions

Change the title and text of 45.2.7.14 (as modified by IEEE Std 802.3bq-2016) to change the name of register 7.61 as follows:

45.2.7.14 EEE link partner ability 1 (Register 7.61)

All of the bits in the EEE LP ability 1 register are read-only. A write to the EEE LP ability 1 register shall have no effect. Except for 10GBASE-T, members of the MultiGBASE-T PHY set exchange the EEE ability in the Infofield during link training. For these PHYs, the EEE LP ability register is updated after link is established. For all other PHYs, when the AN process has been completed, this register shall reflect the contents of the link partner's EEE advertisement 1 register. The assignment of bits in the EEE link partner ability 1 register and the correspondence with the bits in the Next Page messages are shown in Table 45–211.

Change the title of Table 45–211 as follows (unchanged rows not shown):

Table 45–211—EEE link partner ability 1 (Register 7.61) bit definitions

Change the second paragraph of 45.2.7.14 as follows:

The definitions for the contents of the EEE LP ability register are given by the definitions for the contents on the link partner's EEE advertisement 1 register, 7.60 (see 45.2.7.13).

Insert a new subclause (45.2.7.14aa including Table 45–211aa) after 45.2.7.14 and before 45.2.7.14a (inserted by IEEE Std 802.3bz-2016) as follows:

45.2.7.14aa EEE advertisement 2 (Register 7.62)

EEE advertisement 2 register is a continuation of EEE advertisement 1 register. The assignment of bits in the EEE advertisement 2 register are shown in Table 45–211aa.

Table 45–211aa—EEE advertisement 2 register (Register 7.62) bit definitions

Bit(s)	Name	Description	Clause reference; Next Page bit number	R/W ^a
7.62.15:2	Reserved	Value always 0		RO
7.62.1	5GBASE-T EEE	1 = Advertise that the 5GBASE-T has EEE capability 0 = Do not advertise that the 5GBASE-T has EEE capability	126.4.2.5.10	R/W
7.62.0	2.5GBASE-T EEE	1 = Advertise that the 2.5GBASE-T has EEE capability 0 = Do not advertise that the 2.5GBASE-T has EEE capability	126.4.2.5.10	R/W

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

45.2.7.14aa.1 5GBASE-T EEE (7.62.1)

If the device supports EEE operation for 5GBASE-T as defined in 126.6.1, and EEE operation is desired, bit 7.62.1 shall be set to one.

45.2.7.14aa.2 2.5GBASE-T EEE (7.62.0)

If the device supports EEE operation for 2.5GBASE-T as defined in 126.6.1, and EEE operation is desired, bit 7.62.0 shall be set to one.

Insert new subclause (45.2.7.14ab including Table 45–211ab) as follows:

45.2.7.14ab EEE link partner ability 2 (Register 7.63)

All of the bits in the EEE LP ability 2 register are read-only. A write to the EEE LP ability 2 register shall have no effect. When the AN and training processes have been completed, this register shall reflect the contents of the link partner’s EEE advertisement 2 register. The assignment of bits in the EEE link partner ability 2 register and the correspondence with the bits in the Next Page messages (for PHYs that exchange EEE abilities during AutoNegotiation) are shown in Table 45–211ab.

The definitions for the contents of the EEE LP ability 2 register are given by the definitions for the contents on the link partner's EEE advertisement 2 register, 7.62 (see 45.2.7.14aa).

Table 45–211ab—EEE link partner ability 2 (Register 7.63) bit definitions

Bit(s)	Name	Description	Clause reference; Next Page bit number	R/W ^a
7.63.15:2	Reserved	Value always 0		RO
7.63.1	5GBASE-T EEE	1 = Link partner is advertising that the 5GBASE-T has EEE capability 0 = Link partner is not advertising that the 5GBASE-T has EEE capability	126.4.2.5.10	RO
7.63.0	2.5GBASE-T EEE	1 = Link partner is advertising that the 2.5GBASE-T has EEE capability 0 = Link partner is not advertising that the 2.5GBASE-T has EEE capability	126.4.2.5.10	RO

^aRO = Read only

45.2.7.14a MultiGBASE-T AN control 2 (Register 7.64)

Change the reserved row and insert two new rows in Table 45-211a, inserted by IEEE Std 802.3bq-2016, as follows (unchanged rows not shown):

Table 45–211a—MultiGBASE-T AN control 2 (Register 7.64) bit definitions

Bit(s)	Name	Description	R/W ^a
7.64.15:24	Reserved	Value always 0	RO
<u>7.64.3</u>	<u>2.5GBASE-T THP Bypass Request</u>	<u>1 = Local device requests link partner to initially reset THP during fast retrain</u> <u>0 = Local device requests link partner not to initially reset THP during fast retrain</u>	<u>R/W</u>
<u>7.64.2</u>	<u>5GBASE-T THP Bypass Request</u>	<u>1 = Local device requests link partner to initially reset THP during fast retrain</u> <u>0 = Local device requests link partner not to initially reset THP during fast retrain</u>	<u>R/W</u>

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

Insert 45.2.7.14a.a and 45.2.7.14a.b after 45.2.7.14a and before 45.2.7.14a.1 (both inserted by IEEE Std 802.3bq-2016) as follows:

45.2.7.14a.a 2.5GBASE-T THP Bypass Request

Bit 7.64.3 is valid only if 7.32.5 is set to one advertising fast retrain ability, and is used to request the link partner whether to initially reset the THP during fast retrain. THP Bypass Request is exchanged during link training, see 126.4.2.5.10. If bit 7.64.3 is set to zero the local device requests link partner not to reset THP during fast retrain. If bit 7.64.3 is set to one the local device requests link partner to initially reset THP during fast retrain.

45.2.7.14a.b 5GBASE-T THP Bypass Request

Bit 7.64.2 is valid only if 7.32.6 is set to one advertising fast retrain ability, and is used to request the link partner whether to initially reset the THP during fast retrain. THP Bypass Request is exchanged during link training, see 126.4.2.5.10. If bit 7.64.2 is set to zero the local device requests link partner not to reset THP during fast retrain. If bit 7.64.2 is set to one the local device requests link partner to initially reset THP during fast retrain.

45.2.7.14b MultiGBASE-T AN status 2 (Register 7.65)

Change the reserved row and insert two new rows in Table 45-211b (inserted by IEEE Std 802.3bq-2016) as follows:

Table 45–211b—MultiGBASE-T AN status 2 (Register 7.65) bit definitions

Bit(s)	Name	Description	R/W ^a
7.65.15:24	Reserved	Value always 0	RO
7.65.3	2.5GBASE-T Link Partner THP Bypass Request	1 = Link partner requests local device to initially reset THP during fast retrain 0 = Link partner requests local device not to initially reset THP during fast retrain	RO
7.65.2	5GBASE-T Link Partner THP Bypass Request	1 = Link partner requests local device to initially reset THP during fast retrain 0 = Link partner requests local device not to initially reset THP during fast retrain	RO

^aRO = Read only

Insert new subclauses 45.2.7.14b.a and 45.2.7.14b.b before 45.2.7.14b.1 (as inserted by IEEE Std 802.3bq-2016) as follows:

45.2.7.14b.a 2.5GBASE-T Link Partner THP Bypass Request

Bit 7.65.3 is valid only if 7.33.3 is set to one indicating that the link partner has fast retrain ability. THP Bypass Request is exchanged during link training, see 126.4.2.5.10. Bit 7.65.3 is updated after link is established. When read as a zero, the link partner requests local device not to reset THP during fast retrain. When read as a one, the link partner requests local device to initially reset THP during fast retrain.

45.2.7.14b.b 5GBASE-T Link Partner THP Bypass Request

Bit 7.65.2 is valid only if 7.33.0 is set to one indicating that the link partner has fast retrain ability. THP Bypass Request is exchanged during link training, see 126.4.2.5.10. Bit 7.65.2 is updated after link is established. When read as a zero, the link partner requests local device not to reset THP during fast retrain. When read as a one, the link partner requests local device to initially reset THP during fast retrain.

45.5 Protocol implementation conformance statement (PICS) proforma for Clause 45, Management Data Input/Output (MDIO) interface²

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

*Insert two rows for options *2.5T and *5T before the row for *10T, and two rows for options *2.5G and *5G before the row for *10G in 45.5.3.2 as follows (unchanged rows not shown):*

45.5.3.2 PMA/PMD MMD options

Item	Feature	Subclause	Value/Comment	Status	Support
*2.5T	Implementation of a 2.5GBASE-T PMA	45.2.1.14c		PMA:O	Yes [] No []
*5T	Implementation of a 5GBASE-T PMA	45.2.1.14c		PMA:O	Yes [] No []
*2.5G	Implementation of a 2.5 Gb/s PMA/PMD	45.2.1.4		PMA:O	Yes [] No []
*5G	Implementation of a 5 Gb/s PMA/PMD	45.2.1.4		PMA:O	Yes [] No []

45.5.3.3 PMA/PMD management functions

Change row for descriptions for MM111 and MM112 (as modified by IEEE Std 802.3bw-2015, IEEE Std 802.3by-2016, IEEE Std 802.3bq-2016, and IEEE Std 802.3bp-2016) as follows (unchanged rows not shown):

Item	Feature	Subclause	Value/Comment	Status	Support
MM111	Bit set to zero if PMA link_status=FAIL	45.2.1.62.1		<u>PMA*2.5T:M</u> <u>PMA*5T:M</u> PMA*10T:M PMA*25T:M PMA *40T:M	Yes [] No []
MM112	Skew delay register update rate	45.2.1.78	At least once per second	<u>PMA*2.5T:M</u> <u>PMA*5T:M</u> PMA*10T:M PMA*25T:M PMA*40T:M	Yes [] No []

²Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

45.5.3.9 Auto-Negotiation management functions

Insert rows for items AM65 through AM68 (as modified by IEEE Std 802.3bq-2016 and IEEE Std 802.3bp-2016), below AM64, inserted by IEEE Std 802.3bq-2016, as follows (unchanged rows not shown):

Item	Feature	Subclause	Value/Comment	Status	Support
AM65	Advertise 5GBASE-T PHY capability when bit is set to one	45.2.7.10.4ca		AN:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>
AM66	5GBASE-T PHY capability not advertised when bit is set to zero	45.2.7.10.4ca		AN:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>
AM67	Advertise 2.5GBASE-T PHY capability when bit is set to one	45.2.7.10.4ca		AN:M	Yes [<input type="checkbox"/> N/A [<input type="checkbox"/>
AM68	2.5GBASE-T PHY capability not advertised when bit is set to zero	45.2.7.10.4ca		AN:M	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.7:2017

46. Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)

46.1 Overview

Change the third paragraph of 46.1 as follows, to include 2.5 Gb/s and 5 Gb/s rates:

The RS adapts the bit serial protocols of the MAC to the parallel encodings of 2.5 Gb/s, 5 Gb/s, and 10 Gb/s PHYs. Though the XGMII is an optional interface, it is used extensively in this standard as a basis for specification. The 2.5 Gb/s, 5 Gb/s, and 10 Gb/s Physical Coding Sublayers (PCS) ~~is~~ are specified to the XGMII, so if not implemented, a conforming implementation shall behave functionally as if the RS and XGMII were implemented.

Change item a) in the list following the fourth paragraph of 46.1 to include 2.5 Gb/s and 5 Gb/s rates as follows:

The XGMII has the following characteristics:

- a) It is capable of supporting at least one of the following rates of operation: 2.5 Gb/s, 5 Gb/s, or 10 Gb/s-operation.

46.1.1 Summary of major concepts

Insert item i) after the last item in the list under 46.1.1 as follows:

- i) The XGMII is rate scalable and may support rates of 2.5 Gb/s, 5 Gb/s, and 10 Gb/s.

46.1.2 Application

Change the second paragraph of 46.1.2 as follows:

This interface is used to provide media independence so that an identical media access controller may be used with all 2.5GBASE, 5GBASE, and 10GBASE PHY types.

46.1.3 Rate of operation

Change 46.1.3 as follows:

The XGMII supports ~~only the 10 Gb/s~~ MAC data rates of 2.5 Gb/s, 5 Gb/s, and 10 Gb/s as defined within this clause. ~~A compliant device may implement any subset of these rates.~~ Operation at 10 Mb/s and 100 Mb/s is supported by the MII defined in [Clause 22](#) and operation at 1000 Mb/s by the GMII defined in [Clause 35](#).

~~PHYs that provide an XGMII shall support the 10 Gb/s MAC data rate.~~ 10GBASE-X and 10GBASE-R PHYs operate at a 10 Gb/s data rate. 10GBASE-W PHYs operate at the STS-192/VC-4-64c line rate of 9.95328 Gb/s, mapping the encoded data stream at a 9.58464 Gb/s payload rate. On transmit, this mapping is performed by discarding Idle control characters corresponding to the stretched interpacket gap created by the MAC in this mode of operation, and on receive, by adding interpacket gap Idle control characters as required to adapt to the XGMII RX_CLK rate.

Change the title and text of 46.3.1.1 as follows:

46.3.1.1 TX_CLK (10-Gb/s transmit clock)

TX_CLK is a continuous clock used for operation at 10 Gb/s. TX_CLK provides the timing reference for the transfer of the TXC<3:0> and TXD<31:0> signals from the RS to the PHY. The values of TXC<3:0> and TXD<31:0> shall be sampled by the PHY on both the rising edge and falling edge of TX_CLK. TX_CLK is sourced by the RS.

The TX_CLK frequency shall be $156.25 \text{ MHz} \pm 0.01\%$, one sixty-fourth of the MAC transmit data rate $1/64 \times f_{\text{MAC}} \pm 100 \text{ ppm}$, where f_{MAC} is the frequency (in Hz) corresponding to the MAC's nominal bit rate.

NOTE—For EEE capability, TX_CLK may be halted according to 46.3.1.5.

46.3.2.1 RX_CLK (receive clock)

Change text of second paragraph of 46.3.2.1 as follows:

The frequency of RX_CLK may be derived from the received data or it may be that of a nominal clock (e.g., TX_CLK). When the received data rate at the PHY is within tolerance, the RX_CLK frequency shall be $156.25 \text{ MHz} \pm 0.01\%$, one sixty-fourth of the MAC receive data rate $1/64 \times f_{\text{MAC}} \pm 100 \text{ ppm}$, where f_{MAC} is the frequency (in Hz) corresponding to the MAC's nominal bit rate.

46.5 XGMII electrical characteristics

Change first paragraph of 46.5 as follows:

The electrical characteristics of the XGMII are specified such that the XGMII can be applied within a variety of 2.5 Gb/s, 5 Gb/s, and 10 Gb/s equipment types. The electrical specifications are selected for an integrated circuit to integrated circuit application. The electrical characteristics specified in this clause apply to all XGMII signals.

46.6 Protocol implementation conformance statement (PICS) proforma for Clause 46, Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)³

46.6.3 PICS proforma Tables for Reconciliation Sublayer and 10 Gigabit Media Independent Interface

Change row G1 and insert rows G2 and G3 in 46.6.3.1 as follows:

46.6.3.1 General

Item	Feature	Subclause	Value/Comment	Status	Support
G1	PHY support of <u>10 Gb/s MAC data rate</u>	46.1.3	Support MAC data rate of 10 Gb/s	PHY:M-O.1	Yes [] N/A []
G2	PHY support of <u>5 Gb/s MAC data rate</u>	46.1.3	Support MAC data rate of 5 Gb/s	PHY:O.1	Yes [] N/A []
G3	PHY support of <u>2.5 Gb/s MAC data rate</u>	46.1.3	Support MAC data rate of 2.5 Gb/s	PHY:O.1	Yes [] N/A []

46.6.3.6 XGMII signal functional specifications

Change rows FS2 and FS9 as follows (unchanged rows not shown):

Item	Feature	Subclause	Value/Comment	Status	Support
FS2	TX_CLK frequency	46.3.1.1	$156.25 \text{ MHz} \pm 0.01\% \frac{1}{64} \times f_{\text{MAC}} \pm 100 \text{ ppm}$, where f_{MAC} is the frequency (in Hz) corresponding to the MAC's nominal bit rate	XGE:M	Yes [] N/A []
FS9	RX_CLK frequency	46.3.2.1	$156.25 \text{ MHz} \pm 0.01\% \frac{1}{64} \times f_{\text{MAC}} \pm 100 \text{ ppm}$, where f_{MAC} is the frequency (in Hz) corresponding to the MAC's nominal bit rate when received data rate is within tolerance	XGE:M	Yes [] N/A []

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.

78. Energy-Efficient Ethernet (EEE)

78.1 Overview

Insert two new rows into Table 78–1 (as modified by IEEE Std 802.3bp-2016, IEEE Std 802.3bq-2016, and IEEE Std 802.3by-2016) following the entry for 1000BASE-T1 (inserted by IEEE Std 802.3bp-2016) as follows:

78.1.4 PHY types optionally supporting EEE

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

PHY or interface type	Clause
2.5GBASE-T	126
5GBASE-T	126

78.2 LPI mode timing parameters description

Insert the following new rows into Table 78–2 for 2.5GBASE-T and 5GBASE-T EEE parameters after the row for 1000BASE-T1, inserted by IEEE Std 802.3bp-2016.

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
2.5GBASE-T	11.52	12.8	76.8	76.8	5.12	5.12
5GBASE-T	5.76	6.4	38.4	38.4	2.56	2.56

78.3 Capabilities Negotiation

Change the text in the first and second paragraphs of 78.3 (as modified by IEEE Std 802.3bq-2016) as follows:

The EEE capability shall be advertised during the Auto-Negotiation stage, except for PHYs that only support fast wake operation or PHYs that exchange EEE capability during link training. Auto-Negotiation 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. Auto-Negotiation is performed at power up, on command from management, due to link failure, or due to user intervention. Fast wake capability shall be advertised using L2 protocol frames as described in 78.4. The EEE capability for 2.5GBASE-T and 5GBASE-T shall be advertised during link training according to 126.4.2.5.10. The EEE capability for 25GBASE-T and 40GBASE-T shall be advertised during link training according to 113.4.2.5.10.

During Auto-Negotiation, both link partners indicate their EEE capabilities. EEE is supported only if during Auto-Negotiation both the local device and link partner advertise the EEE capability for the resolved PHY type. If EEE is not supported, all EEE functionality is disabled and the LPI client does not assert LPI. EEE deep sleep operation shall not be enabled unless both the local device and link partner advertise deep sleep capability during Auto-Negotiation for the resolved PHY type. If EEE is supported by both link partners for the negotiated PHY type, then the EEE function can be used independently in either direction. The same applies to 2.5GBASE-T, 5GBASE-T, 25GBASE-T, and 40GBASE-T except the EEE capabilities are exchanged and resolved during link training instead of during Auto-Negotiation.

78.5 Communication link access latency

Insert the following new rows into Table 78-4 with 2.5G/5GBASE-T LPI parameters after 1000BASE-KX (unchanged rows not shown):

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

PHY 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)
2.5GBASE-T	Case-1	29.44	29.44	17.92	0	11.52
	Case-2	17.92	17.92	6.4	0	11.52
5GBASE-T	Case-1	14.72	14.72	8.96	0	5.76
	Case-2	8.96	8.96	3.2	0	5.76

Insert a new Clause 125 as follows:

125. Introduction to 2.5 Gb/s and 5 Gb/s networks

125.1 Overview

125.1.1 Scope

This clause describes the general requirements for 2.5 Gigabit and 5 Gigabit Ethernet. 2.5 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer operating at a data rate of 2.5 Gb/s, coupled with any IEEE 802.3 2.5GBASE Physical Layer implementation. 5 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer operating at a data rate of 5 Gb/s, coupled with any IEEE 802.3 5GBASE Physical Layer implementation. 2.5 Gb/s and 5 Gb/s Physical Layer entities, such as those specified in Table , provide a bit error ratio (BER) better than or equal to 10^{-12} at the MAC/PLS service interface.

2.5 Gigabit and 5 Gigabit Ethernet are defined for full duplex operation only.

125.1.2 Relationship of 2.5 Gigabit and 5 Gigabit Ethernet to the ISO OSI reference model

2.5 Gigabit and 5 Gigabit Ethernet couple the IEEE 802.3 MAC to a family of 2.5 Gb/s and 5 Gb/s Physical Layers. The relationships among 2.5 Gigabit and 5 Gigabit Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in Figure 125–1. While this standard 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 XGMII, which, when implemented as a logical interconnection port between the MAC sublayer and the Physical Layer (PHY), uses a 32-bit-wide data path as specified in Clause 46.
- 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 MDI as specified in Clause 126 for 2.5GBASE-T and 5GBASE-T uses a 4 lane data path.

125.1.3 Nomenclature


The nomenclature employed by the 2.5 Gigabit and 5 Gigabit Physical Layers is explained in the following paragraphs.

The alpha-numeric prefix 2.5GBASE in the port type (e.g., 2.5GBASE-T) represents a family of Physical Layer devices operating at a speed of 2.5 Gb/s. The alpha-numeric prefix 5GBASE in the port type (e.g., 5GBASE-T) represents a family of Physical Layer devices operating at a speed of 5 Gb/s.

2.5GBASE-T represents Physical Layer devices using Clause 126 Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, and baseband medium, for transmitting 2.5 Gb/s Ethernet over a point-to-point 4-pair balanced twisted-pair medium. 2.5GBASE-T uses low density parity check (LDPC) FEC in its Physical Coding Sublayers mapped to a PAM16 constellation for transmission on 4-pair, twisted-pair copper cabling.

5GBASE-T represents Physical Layer devices using Clause 126 Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, and baseband medium, for transmitting 5 Gb/s Ethernet over a point-to-point 4-pair balanced twisted-pair medium. 5GBASE-T uses low density parity check (LDPC) FEC in its Physical Coding Sublayers mapped to a PAM16 constellation for transmission on 4-pair, twisted-pair copper cabling.

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation



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
 XGMII = 10 GIGABIT MEDIA INDEPENDENT INTERFACE
 * XGMII IS OPTIONAL

Figure 125-1—Architectural positioning of 2.5 Gigabit and 5 Gigabit Ethernet

125.1.4 Physical Layer signaling systems

Physical Layer devices listed in Table 125-1 are defined for operation at 2.5 Gb/s and 5 Gb/s.

Table 125-1—2.5 Gb/s and 5 Gb/s PHYs

Name	Description
2.5GBASE-T	2.5 Gb/s PHY using LDPC encoding and PAM16 modulation over balanced twisted-pair structured cabling systems (see Clause 126)
5GBASE-T	5 Gb/s PHY using LDPC encoding and PAM16 modulation over balanced twisted-pair structured cabling systems (see Clause 126)

This standard specifies a family of Physical Layer implementations. Table 125–2 specifies the correlation between nomenclature and clauses. Implementations conforming to one or more nomenclatures must meet the requirements of the corresponding clauses.

Table 125–2—Nomenclature and clause correlation

Nomenclature	Clause ^a				
	28	46		78	126
	Auto-Negotiation	RS	XGMII	EEE	2.5GBASE-T PCS/PMA 5GBASE-T PCS/PMA
2.5GBASE-T	M	M	O	O	M
5GBASE-T	M	M	O	O	M

^aO = Optional, M = Mandatory.

125.2 Summary of 2.5 Gigabit and 5 Gigabit Ethernet sublayers

125.2.1 Reconciliation Sublayer (RS) and Media Independent Interface

2.5 Gigabit and 5 Gigabit Ethernet use the 10 Gigabit Media Independent Interface defined in Clause 46, with the clock scaled to their respective data rates. The physical instantiation of the Media Independent Interface is optional, and can be used to logically connect layers within a device.

The XGMII supports 2.5 Gb/s and 5 Gb/s operation, in addition to 10 Gb/s operation, through its 32-bit wide transmit and receive data paths. The Reconciliation Sublayer (RS) provides a mapping between the signals provided at the Media Independent Interface (XGMII) and the MAC/PLS service definition.

While XGMII 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 Physical Coding Sublayers defined in Clause 126.

125.2.2 Physical Coding Sublayer (PCS)

2.5GBASE-T and 5GBASE-T PHYs contain a 65B-LDPC PCS that maps the data transferred across the XGMII interface to 64B/65B blocks encoded in a 2048-bit LDPC frame. This LDPC frame is then mapped to 512 Gray-coded PAM16 symbols for transfer to the 4-lane PMA.

125.2.3 Physical Medium Attachment sublayer (PMA)

The PMA provides a medium-independent means for the PCS to support the use of a range of physical media. The 2.5GBASE-T and 5GBASE-T PMAs perform the mapping of transmit and receive data streams between the PCS and PMA via the PMA service interface. The PMA provides a medium-independent means for the PCS to support the use of a range of physical media.

In the case of BASE-T, the 2.5GBASE-T and 5GBASE-T PMAs perform the mapping of transmit and receive data streams between the PCS and PMA via the PMA service interface, and the mapping and multiplexing of transmit and receive data streams between the PMA and MDI. In addition, the PMAs perform retiming of the received data stream when appropriate.

125.2.4 Auto-Negotiation, type BASE-T

Auto-Negotiation (Clause 28) is used by 2.5GBASE-T and 5GBASE-T devices to detect the abilities (modes of operation) supported by the device at the other end of a link segment, determine common abilities, and configure for joint operation. Auto-Negotiation is performed upon link startup through the use of a special sequence of fast link pulses.

125.2.5 Management interface (MDIO/MDC)

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

125.2.6 Management

Managed objects, attributes, and actions are defined for all 2.5 Gigabit and 5 Gigabit Ethernet components. These items are defined in Clause 30.

125.3 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 125-3 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.

See 31B.3.7 for PAUSE reaction timing constraints for stations at operating speeds of 2.5 Gb/s and 5 Gb/s.

Table 125-3—Sublayer delay constraints

Sublayer	Maximum (bit time) ^a	Maximum (pause_quanta) ^b	Maximum (ns)	Notes ^c
2.5GBASE-T PHY	12 800	25	5 120	Does not include delay of cable medium. See 126.11.
5GBASE-T PHY	14 336	28	2 867.2	Does not include delay of cable medium. See 126.11.

^a For 2.5GBASE-T, 1 bit time (BT) is equal to 400 ps and for 5GBASE-T, 1 bit time (BT) is equal to 200 ps. (See 1.4.117 for the definition of bit time.)

^b For 2.5GBASE-T, 1 pause_quantum is equal to 204.8 ns and for 5GBASE-T, 1 pause_quantum is equal to 102.4 ns. (See 31B.2 for the definition of pause_quanta.)

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

Insert a new Clause 126 as follows:

126. Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, and baseband medium, types 2.5GBASE-T and 5GBASE-T

126.1 Overview

The 2.5GBASE-T and 5GBASE-T PHYs are members of the 2.5 Gb/s and 5 Gb/s Ethernet family of high-speed network specifications respectively. The 2.5GBASE-T PCS, PMA, and baseband medium specifications are intended for operation over balanced twisted-pair structured cabling systems. The 5GBASE-T PCS, PMA, and baseband medium specifications are intended for operation over balanced twisted-pair structured cabling systems. 2.5GBASE-T and 5GBASE-T signaling both require four pairs of balanced cabling as specified in ISO/IEC 11801:2002 and ANSI/TIA-568-C.2.

This clause defines the types 2.5GBASE-T and 5GBASE-T PCS, PMA sublayers, and Medium Dependent Interfaces (MDI). Together, the PCS and PMA sublayers define a Physical Layer (PHY). Functional, electrical, and mechanical specifications for the type 2.5GBASE-T PMA, 5GBASE-T PMA, and MDI are provided in this clause. This clause also specifies the baseband media used with 2.5GBASE-T and 5GBASE-T. Management functions are optionally accessible through the management interface defined in Clause 45, or equivalent. Please refer to Table 125–2 for associated sublayers and options for assembling a 2.5 Gb/s or 5 Gb/s system with the 2.5GBASE-T or 5GBASE-T PHY, respectively.

This clause also specifies 2.5GBASE-T and 5GBASE-T Low Power Idle (LPI) as part of Energy-Efficient Ethernet (EEE). This allows the PHY to enter a low power mode of operation during periods of low link utilization as described in Clause 78.

2.5GBASE-T and 5GBASE-T PHYs may optionally support a fast retrain mechanism. Implementation of the fast retrain option is recommended. Configurations wishing to disable fast retrain on the link may do so by advertising lack of support during link startup, thus preventing the link partner from attempting fast retrain and potentially dropping the link, see 45.2.7.10.

126.1.1 Nomenclature

The 2.5GBASE-T and 5GBASE-T PHYs described in this clause represent two distinct PHY types that share the same PCS, PMA, and MDI specifications subject to frequency scaling. In order to efficiently describe the two PHYs, the nomenclature 2.5G/5GBASE-T is used to describe specifications that apply to both the 2.5GBASE-T and 5GBASE-T PHYs. Additionally, for parameters that scale with the PHYs data rate, the parameter S is used for scaling. For 2.5GBASE-T, $S = 0.5$ and for 5GBASE-T, $S = 1$.

126.1.2 Relationship of 2.5GBASE-T and 5GBASE-T to other standards

The relationships between the 2.5GBASE-T and 5GBASE-T PHYs, the ISO Open Systems Interconnection (OSI) reference model, and the IEEE 802.3 Ethernet model are shown in Figure 126–1. The PHY sublayers (shown shaded) in Figure 126–1 connect the IEEE 802.3 Ethernet MAC to the medium. The 2.5GBASE-T and 5GBASE-T PHY service interface is the XGMII, which is defined in Clause 46.




Figure 126-1—Types 2.5GBASE-T and 5GBASE-T PHYs relationship to the ISO Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

126.1.3 Operation of 2.5GBASE-T and 5GBASE-T

The 2.5GBASE-T PHY and 5GBASE-T PHY each employ full duplex baseband transmission over four pairs of balanced twisted-pair structured cabling. The aggregate data rates of 2.5 Gb/s or 5 Gb/s are achieved by transmitting one-quarter of the aggregate data rate in each direction simultaneously on each wire pair, as follows in Figure 126-2. Baseband 16-level PAM signaling with a modulation rate of 200 MBd for 2.5GBASE-T and 400 MBd for 5GBASE-T is used on each of the wire pairs. Ethernet data and control characters are encoded at a rate of 3.125 information bits per PAM16 symbol, along with auxiliary bits. Each transmitted PAM16 symbol is considered as a single one-dimensional (1D) symbol. After link startup, PHY frames consisting of 512 PAM16 symbols are continuously transmitted. The PAM16 symbols are determined by 4-bit labels, each comprising 4 LDPC-encoded bits. The 512 PAM16 symbols of one PHY frame are transmitted as 4×128 PAM16 symbols over the four wire pairs. Data and Control symbols are embedded in a framing scheme that runs continuously after startup of the link. For 2.5GBASE-T, the modulation symbol rate of 200 MBd results in a symbol period of 5 ns, and for 5GBASE-T, the modulation symbol rate of 400 MBd results in a symbol period of 2.5 ns.

A 2.5GBASE-T or 5GBASE-T PHY can be configured either as a MASTER PHY or as a SLAVE PHY. The MASTER-SLAVE relationship between two stations sharing a link segment is established during Auto-Negotiation (see Clause 28, 126.6, Annex 28B, Annex 28C, and Annex 28D). The MASTER PHY uses a local clock to determine the timing of transmitter operations. The MASTER-SLAVE relationship includes loop timing. The SLAVE PHY recovers the clock from the received signal and uses it to determine the timing of transmitter operations, i.e., it performs loop timing, as illustrated in Figure 126-3.

2.5GBASE-T and 5GBASE-T PHYs optionally provide support for LPI as part of EEE (see Clause 78). This extension allows PHYs to enter an LPI mode when either the local or link partner system requests low power operation. The transmit and receive functions may enter and leave the LPI mode independently so that both symmetric and asymmetric operation is supported. While the PHY is in the LPI mode, the PHY periodically transmits a refresh signal to allow the remote PHY to refresh its receiver state (e.g., timing recovery, adaptive filter coefficients) and thereby track long-term variation in the timing of the link or the underlying channel characteristics. An easily detectable alert signal is transmitted to signal an end to the LPI mode. The alert signal is followed by a wake signal to enable a rapid transition back to the normal operational mode.

2.5GBASE-T and 5GBASE-T PHYs may optionally support a fast retrain mechanism. This function allows PHYs to quickly recover from link degradation without a normal two-second retrain.

The PCS and PMA are summarized in 126.1.3.1 and 126.1.3.2. The EEE capability is summarized in 126.1.3.3. Figure 126-3 shows the functional block diagram.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd7:2017





Figure 126-2—2.5GBASE-T and 5GBASE-T topology

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation



NOTE 1—The recovered_clock arc is shown to indicate delivery of the received clock signal back to PMA TRANSMIT for loop timing.

NOTE 2—pcs_data_mode is required only for the EEE or fast retrain capabilities; alert_detect and rx_lpi_active are only required for the EEE capability; fr_active is only required for the fast retrain capability. Figures and capabilities only required for EEE are noted by dashed boxes.

Figure 126-3—Functional block diagram

126.1.3.1 Summary of Physical Coding Sublayer (PCS)

The 2.5GBASE-T or 5GBASE-T PCS couples a 10 Gigabit Media Independent Interface (XGMII), as described in Clause 46, to the 2.5GBASE-T or 5GBASE-T Physical Medium Attachment (PMA) sublayer.

In addition to the normal mode of operation, the PCS supports a training mode. Furthermore, the PCS contains a management interface.

In the transmit direction (see Figure 126–6), in normal mode, the PCS receives eight XGMII data octets provided by two consecutive transfers on the XGMII service interface on TXD<31:0> and groups them into 64-bit blocks with the 64-bit block boundaries aligned with the boundary of the two XGMII transfers. Each group of eight octets along with the data/control indications is transcoded into a 65-bit block. The resulting 65-bit blocks are scrambled and assembled in a group of 25 blocks, yielding an Ethernet payload of $25 \times 65 = 1625$ bits. Additionally, 97 zero-bits are appended, and a leading auxiliary bit is added to obtain a block of 1723 bits.

The 1723 bits are encoded by a systematic LDPC(1723,2048) encoder, which adds 325 LDPC check bits to form an LDPC codeword of 2048 coded bits. The 97 zero-bits are then replaced with vendor-defined random data. Implementers are cautioned that insufficient randomization can impact meeting PMA PSD mask requirements (see 126.5.3.4 for transmit PSD mask definition). The resulting 2048 bit LDPC frame is then divided into 512 4-bit labels, which are mapped into PAM16 modulation symbols.

The obtained PHY frame of 512 PAM16 symbols is passed on to the PMA as PMA_UNITDATA.request. The PMA transmits the PAM16 symbols over the four wire pairs in the form of 128 constituent PAM16 symbols per pair.

In the receive direction (see Figure 126–7), in normal mode, the PCS processes code-groups received from the remote PHY via the PMA in 128 four-dimensional (4D) symbol blocks and maps them to the XGMII service interface in the receive path. In this receive processing scheme, symbol clock synchronization is done by the PMA Receive function.

The signals provided by the PCS at the XGMII conform to the interface requirements of Clause 46.

Details of the PCS functions and state diagrams are covered in 126.3. The interface to the PMA is an abstract message-passing interface specified in 126.2.

126.1.3.2 Summary of Physical Medium Attachment (PMA) sublayer

The PMA couples messages from the PCS service interface onto the balanced cabling physical medium via the Medium Dependent Interface (MDI) and provides the link management and PHY Control functions. The PMA provides full duplex communications at $400 \times S$ Mb/s over four pairs of balanced cabling up to 100 m in length.

The PMA Transmit function comprises four transmitters to generate continuous time analog signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD, as described in 126.4.3.1. In normal mode, each 4D symbol received from the PCS Transmit function undergoes multiple stages of processing. First the symbol goes through a Tomlinson-Harashima precoder (THP), which maps the PAM16 input (as described in 126.3.2.2.18) in each dimension of the 4D symbol into a quasi-continuous discrete-time value in the range $-16 \leq x < 16$. This THP-processed 4D symbol stream may be further processed by a digital transmit filter and is then passed on to four digital-to-analog converters (DACs). The DAC outputs may be further processed with continuous time filters to roll off the high-frequency spectral response to limit high-frequency emissions and are then applied to each of the four balanced pairs via the MDI port.

The PMA Receive function comprises four independent receivers for pulse-amplitude modulated signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD, as described in 126.4.3.2. The receivers are responsible for acquiring symbol timing and, when operating in normal mode, for canceling echo, near-end crosstalk, far-end crosstalk, and equalizing the signal. The 4D symbols are provided to the PCS Receive function via the PMA_UNITDATA.indication message. The PMA also contains functions for Link Monitor.

The PMA PHY Control function generates signals that control the PCS and PMA sublayer operations. PHY Control begins following the completion of Auto-Negotiation and provides the startup functions required for successful 2.5GBASE-T and 5GBASE-T operation. PHY Control determines whether the PHY operates in a normal mode, enabling data transmission over the link segment, or whether the PHY sends special PAM3 code-groups that are used in the training mode.

PMA functions and state diagrams are specified in 126.4. PMA electrical specifications are given in 126.5.

The PMA sublayer may also support a fast retrain function. The fast retrain function is specified in 126.4.2.5.16.

126.1.3.3 Summary of EEE capability

A 2.5GBASE-T or 5GBASE-T PHY may optionally support the EEE capability, as described in 78.1.4. The EEE capability is a mechanism by which 2.5GBASE-T and 5GBASE-T PHYs are able to reduce power consumption during periods of low link utilization. PHYs can enter this mode of operation after reaching PCS data mode. Each direction of the full duplex link is able to enter and exit the LPI mode independently, supporting symmetric and asymmetric LPI operation. This allows power savings when only one side of the full duplex link is in a period of low utilization. No data frames are lost or corrupted during the transition to or from the LPI mode.

In the transmit direction, the transition to the LPI transmit mode begins when the PCS transmit function detects an LPI control character in all four lanes of two consecutive transfers of TXD[31:0] that is then mapped into a single 64B/65B block, according to the position in the 2.5GBASE-T or 5GBASE-T LDPC frame. Following this event a sleep signal is transmitted by the PMA. The sleep signal is composed of LDPC frames that contain only LP_IDLE 64B/65B blocks. The sleep signal indicates to the link partner that the transmit function of the PHY is entering the LPI transmit mode. Immediately after the transmission of the sleep frames, the transmit function of the local PHY enters the LPI transmit mode. While the transmit function is in the LPI mode the PHY may disable data path and control logic to save additional power. Periodically the transmit function of the local PHY transmits refresh frames that are used by the link partner to update adaptive filters and timing circuits in order to maintain link integrity. The LPI mode begins with quiet signaling or with a full refresh period. Partial refreshes (defined as a refresh signal shorter than eight LDPC frames) that immediately follow the transition to the LPI mode are replaced with quiet signaling. The quiet-refresh cycle continues until the PCS function detects IDLE characters on the XGMII. These characters signal to the PHY that the LPI transmit mode should end. The PMA Transmit function in the PHY then sends an alert message to the link partner. The alert signal begins on a LDPC 2-frame 256 4D-symbol boundary aligned to the inversion on pair A during PMA training, but has no fixed relationship to the quiet-refresh cycle. The alert signal wakes the link partner from sleep. The alert signal is followed by a wake signal, composed of LDPC frames containing only IDLE 64B/65B blocks. After a short recovery time the normal operational mode is resumed.

In the receive direction the transition to the LPI mode is triggered when the PCS Receive function detects LPI control characters within received LDPC frames. This indicates that the link partner is about to enter the LPI transmit mode. Following these frames the link partner ceases transmission and begins quiet-refresh signaling. During the quiet time it is highly recommended that the local receiver power off circuits to reduce power consumption. Periodically the link partner transmits refresh frames that are used by the receiver to update adaptive coefficients and timing circuits. This quiet-refresh cycle continues until the link partner transmits the alert signal, initiating a transition back to the normal operational mode. The alert signal is

detected in the PMA and signals that normal data frames will follow. The alert signal is followed by a wake signal that allows the local receiver time to prepare for the normal operational mode. The wake signal is composed of repeated IDLE 64B/65B blocks. After a short recovery time the normal operational mode is resumed.

Support for the EEE capability is advertised in the Infofield (Octet 12 bit 7) during link startup. Transitions to and from the LPI transmit mode are controlled via XGMII signaling. Transitions to and from the LPI receive mode are controlled by the link partner using sleep, alert, and wake signaling.

The PCS 64B/65B Transmit state diagram in Figure 126–14 and Figure 126–15 includes additional states for EEE. The PCS 64B/65B Receive state diagram in Figure 126–16 and Figure 126–17 includes additional states for EEE. The EEE Transmit state diagram is contained in the PCS Transmit function and is specified in Figure 126–18.

126.1.4 Signaling

2.5GBASE-T and 5GBASE-T signaling is performed by the PCS generating continuous code-group sequences that the PMA transmits over each wire pair. The signaling scheme achieves a number of objectives including:

- a) Forward error correction (FEC) coded symbol mapping for data.
- b) Algorithmic mapping from TXD<31:0> and TXC<3:0> to 4D symbols in the transmit path.
- c) Algorithmic mapping from the received 4D signals on the MDI port to RXD<31:0> and RXC<3:0> on the XGMII interface.
- d) Uncorrelated symbols in the transmitted symbol stream.
- e) No correlation between symbol streams traveling both directions on any pair combination.
- f) No correlation between symbol streams on pairs BI_DA, BI_DB, BI_DC, and BI_DD.
- g) Block framing and other control signals.
- h) Ability to signal the status of the local receiver to the remote PHY to indicate that the local receiver is not operating reliably and requires retraining.
- i) Ability to automatically detect and correct for pair swapping and crossover connections.
- j) Ability to automatically detect and correct for incorrect polarity in the connections.
- k) Ability to automatically correct for differential delay variations across the wire-pairs.
- l) Ability to support refresh, quiet and alert signaling during LPI operation.

The PHY operates in two modes—normal mode or training mode. In normal mode, PCS generates a continuous stream of 4D symbols that are transmitted via the PMA at one of eight power levels. In training mode, the PCS is directed to generate only PAM2 symbols for transmission by the PMA, which enable the receiver at the other end to train until it is ready to operate in normal mode. (See 126.3.2.2 for description of PCS transmit modes.)

PHYs may also support the EEE capability as described in 126.1.3.3. Transitions to the LPI mode are supported after reaching normal mode.

126.1.5 Interfaces

All 2.5GBASE-T and 5GBASE-T PHY implementations are compatible at the MDI and at the XGMII, if implemented. Implementation of the XGMII is optional. Designers are free to implement circuitry within the PCS and PMA in an application-dependent manner provided that the MDI and XGMII (if the XGMII is implemented) specifications are met. System operation from the perspective of signals at the MDI and management objects are identical whether the XGMII is implemented or not.

126.1.6 Conventions in this clause

The body of this clause contains state diagrams, including definitions of variables, constants, and functions. 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.

Default initializations, unless specified, are left to the implementer.

126.2 2.5GBASE-T and 5GBASE-T service primitives and interfaces

2.5GBASE-T and 5GBASE-T transfer data and control information across the following four service interfaces:

- a) 10 Gigabit Media Independent Interface (XGMII)
- b) Technology Dependent Interface
- c) PMA service interface
- d) Medium dependent interface (MDI)

The XGMII is specified in Clause 46; the Technology Dependent Interface is specified in Clause 28. The PMA service interface is defined in 126.2.2 and the MDI is defined in 126.8.

126.2.1 Technology Dependent Interface

2.5GBASE-T and 5GBASE-T use the following service primitives to exchange status indications and control signals across the Technology Dependent Interface as specified in Clause 28:

PMA_LINK.request (link_control)

PMA_LINK.indication (link_status)

126.2.1.1 PMA_LINK.request

This primitive allows the Auto-Negotiation algorithm to enable and disable operation of the PMA as specified in 28.2.6.2.

126.2.1.1.1 Semantics of the primitive

PMA_LINK.request (link_control)

The link_control parameter can take on one of three values: SCAN_FOR_CARRIER, DISABLE, or ENABLE.

- SCAN_FOR_CARRIER Used by the Auto-Negotiation algorithm prior to receiving any fast link pulses. During this mode the PMA reports link_status=FAIL. PHY processes are disabled.
- DISABLE Set by the Auto-Negotiation algorithm in the event fast link pulses are detected. PHY processes are disabled. This allows the Auto-Negotiation algorithm to determine how to configure the link.
- ENABLE Used by Auto-Negotiation to turn control over to the PHY for data processing functions.

126.2.1.1.2 When generated

Auto-Negotiation generates this primitive to indicate a change in link_control as described in Clause 28.

126.2.1.1.3 Effect of receipt

This primitive affects operation of the PMA Link Monitor function as defined in 126.4.2.6.

126.2.1.2 PMA_LINK.indication

This primitive is generated by the PMA to indicate the status of the underlying medium as specified in 28.2.6.1. This primitive informs the Auto-Negotiation algorithm about the status of the underlying link.

126.2.1.2.1 Semantics of the primitive

PMA_LINK.indication (link_status)

The link_status parameter can take on one of two values: FAIL or OK.

FAIL	No valid link established.
OK	The Link Monitor function indicates that a valid 2.5GBASE-T or 5GBASE-T link is established. Reliable reception of signals transmitted from the remote PHY is possible.

126.2.1.2.2 When generated

The PMA generates this primitive to indicate a change in link_status in compliance with the state diagram given in Figure 126–29.

126.2.1.2.3 Effect of receipt

Auto-Negotiation uses this primitive to detect a change in link_status as described in Clause 28.

126.2.2 PMA service interface

2.5GBASE-T and 5GBASE-T use the following service primitives to exchange symbol vectors, status indications, and control signals across the service interfaces:

- PMA_TXMODE.indication (tx_mode)
- PMA_CONFIG.indication (config)
- PMA_UNITDATA.request (tx_symb_vector)
- PMA_UNITDATA.indication (rx_symb_vector)
- PMA_SCRSTATUS.request (scr_status)
- PMA_PCSSTATUS.request (pcs_status)
- PMA_RXSTATUS.indication (loc_rcvr_status)
- PMA_REMRXSTATUS.request (rem_rcvr_status)

EEE-capable PHYs additionally support the following service primitives:

- PMA_ALERTDETECT.indication (alert_detect)

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation


PCS_RX_LPI_STATUS.request (rx_lpi_active)

PMA_PCSDATAMODE.indication (PCS_data_mode)

Fast retrain capable PHYs additionally support the following service primitive:

PMA_FR_ACTIVE.indication (fr_active)

The use of these primitives is illustrated in Figure 126–4. Connections from the management interface (signals MDC and MDIO) to the sublayers are pervasive and are not shown in Figure 126–4.



NOTE 1—PMA_PCSDATAMODE.indication is required only for the EEE or fast retrain capabilities.
 NOTE 2—PMA_ALERTDETECT.indication and PCS_RX_LPI_STATUS.request are only required for the EEE capability.
 NOTE 3—PMA_FR_ACTIVE.indication is only required for the fast retrain capability.

Figure 126–4—2.5GBASE-T and 5GBASE-T service interfaces

126.2.2.1 PMA_TXMODE.indication

The transmitter in a 2.5GBASE-T or 5GBASE-T link normally sends over the four pairs, 4D symbols that represent an XGMII data stream with framing, scrambling and encoding of data, control information, or idles.

126.2.2.1.1 Semantics of the primitive

PMA_TXMODE.indication (tx_mode)

PMA_TXMODE.indication specifies to PCS Transmit via the parameter tx_mode what sequence of code-groups the PCS should be transmitting. The parameter tx_mode can take on one of the following three values of the form:

SEND_N	This value is continuously asserted when transmission of sequences of 4D symbols representing an XGMII data stream in normal mode.
SEND_T	This value is continuously asserted in case transmission of sequences of code-groups representing the training mode is to take place.
SEND_Z	This value is continuously asserted in case transmission of zeros is required.

126.2.2.1.2 When generated

The PMA PHY Control function generates PMA_TXMODE.indication messages to indicate a change in tx_mode.

126.2.2.1.3 Effect of receipt

Upon receipt of this primitive, the PCS performs its transmit function as described in 126.3.2.2.

126.2.2.2 PMA_CONFIG.indication

Each PHY in a 2.5GBASE-T or 5GBASE-T link is capable of operating as a MASTER PHY and as a SLAVE PHY. MASTER-SLAVE configuration is determined during Auto-Negotiation (126.6.1). The result of this negotiation is provided to the PMA.

126.2.2.2.1 Semantics of the primitive

PMA_CONFIG.indication (config)

PMA_CONFIG.indication specifies to PCS and PMA Transmit via the parameter config whether the PHY operates as a MASTER PHY or as a SLAVE PHY. The parameter config can take on one of the following two values of the form:

MASTER	This value is continuously asserted when the PHY operates as a MASTER PHY.
SLAVE	This value is continuously asserted when the PHY operates as a SLAVE PHY.

126.2.2.2.2 When generated

PMA generates PMA_CONFIG.indication messages to indicate a change in config.

126.2.2.2.3 Effect of receipt

PCS and PMA Clock Recovery perform their functions in MASTER or SLAVE configuration according to the value assumed by the parameter config.

126.2.2.3 PMA_UNITDATA.request

This primitive defines the transfer of code-groups in the form of the tx_symb_vector parameter from the PCS to the PMA. The code-groups are obtained in the PCS Transmit function using the encoding rules defined in 126.3.2.2 to represent XGMII data and control streams or other sequences.

126.2.2.3.1 Semantics of the primitive

PMA_UNITDATA.request (tx_symb_vector)

During transmission, the PMA_UNITDATA.request simultaneously conveys to the PMA via the parameter tx_symb_vector the value of the symbols to be sent over each of the four transmit pairs BI_DA, BI_DB, BI_DC, and BI_DD. For IEEE-capable PHYs, the vector also requests the PMA to send the ALERT signal during LPI. The tx_symb_vector parameter takes on the following form:

SYMB_4D	A vector of four multi-level symbols, one for each of the four transmit pairs BI_DA, BI_DB, BI_DC, and BI_DD. In normal operation, each symbol may take on one of the values in the set $\{-15, -13, -11, -9, -7, -5, -3, -1, 1, 3, 5, 7, 9, 11, 13, 15\}$. The symbols may additionally take the value 0 when zeros are to be transmitted in the following two cases: 1) when PMA_TXMODE.indication is SEND_Z during PMA training, and 2) after data mode is reached, the transmit function is in the LPI transmit mode and lpi_tx_mode is QUIET
ALERT	A vector used to indicate that the PMA should transmit the alert sequence. ALERT is asserted for a time equal to 8 LDPC frames.

The symbols that are elements of tx_symb_vector are called, according to the pair on which each is transmitted, tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD].

126.2.2.3.2 When generated

The PCS generates PMA_UNITDATA.request synchronously with every transmit clock cycle.

126.2.2.3.3 Effect of receipt

Upon receipt of this primitive the PMA transmits on the MDI the signals corresponding to the indicated symbols after processing with the THP, the transmit filter and other specified PMA Transmit processing. The parameter tx_symb_vector is also used by the PMA Receive function to process the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD for canceling the echo and near-end crosstalk (NEXT).

126.2.2.4 PMA_UNITDATA.indication

This primitive defines the transfer of code-groups in the form of the rx_symb_vector parameter from the PMA to the PCS.

126.2.2.4.1 Semantics of the primitive

PMA_UNITDATA.indication (rx_symb_vector)

During reception the PMA_UNITDATA.indication simultaneously conveys to the PCS via the parameter rx_symb_vector the values of the symbols detected on each of the four receive pairs BI_DA, BI_DB, BI_DC, and BI_DD. The rx_symb_vector parameter takes on the following form:

SYMB_4D A vector of the four 1D symbols that is the receiver's best estimate of the symbols that were sent by the remote transmitter across the four pairs with reliability measures.

126.2.2.4.2 When generated

The PMA generates PMA_UNITDATA.indication (SYMB_4D) messages synchronously every four symbols received at the MDI. The nominal rate of the PMA_UNITDATA.indication primitive is $S \times 400$ MHz, as governed by the recovered clock.

126.2.2.4.3 Effect of receipt

The effect of receipt of this primitive is unspecified.

126.2.2.5 PMA_SCRSTATUS.request

This primitive is generated by PCS Receive to communicate the status of the descrambler for the local PHY. The parameter scr_status conveys to the PMA Receive function the information that the training mode descrambler has achieved synchronization.

126.2.2.5.1 Semantics of the primitive

PMA_SCRSTATUS.request (scr_status)

The scr_status parameter can take on one of the following two values of the form:

OK The training mode descrambler has achieved synchronization.
NOT_OK The training mode descrambler is not synchronized.

126.2.2.5.2 When generated

PCS Receive generates PMA_SCRSTATUS.request messages to indicate a change in scr_status.

126.2.2.5.3 Effect of receipt

The effect of receipt of this primitive is specified in 126.4.2.4, 126.4.2.5, and 126.4.6.1.

126.2.2.6 PMA_PCSSTATUS.request

This primitive is generated by PCS Receive to indicate the fully operational state of the PCS for the local PHY. The parameter pcs_status conveys to the PMA Receive function the information that the PCS is operating reliably in data mode.

126.2.2.6.1 Semantics of the primitive

PMA_PCSSTATUS.request (pcs_status)

The pcs_status parameter can take on one of the following two values of the form:

OK	The PCS is operating reliably in data mode.
NOT_OK	The PCS is not operating reliably in data mode.

126.2.2.6.2 When generated

PCS Receive generates PMA_PCSSTATUS.request messages to indicate a change in pcs_status.

126.2.2.6.3 Effect of receipt

The effect of receipt of this primitive is specified in 126.4.6.

126.2.2.7 PMA_RXSTATUS.indication

This primitive is generated by PMA Receive to indicate the status of the receive link at the local PHY. The parameter loc_rcvr_status conveys to the PCS Transmit, PCS Receive, PMA PHY Control function, and Link Monitor the information on whether the status of the overall receive link is satisfactory or not. Note that loc_rcvr_status is used by the PCS Receive decoding functions. The criterion for setting the parameter loc_rcvr_status is left to the implementer. It can be based, for example, on observing the mean-square error at the decision point of the receiver and detecting errors during reception of symbol streams.

126.2.2.7.1 Semantics of the primitive

PMA_RXSTATUS.indication (loc_rcvr_status)

The loc_rcvr_status parameter can take on one of the following two values of the form:

OK	This value is asserted and remains true during reliable operation of the receive link for the local PHY.
NOT_OK	This value is asserted whenever operation of the link for the local PHY is unreliable.

126.2.2.7.2 When generated

PMA Receive generates PMA_RXSTATUS.indication messages to indicate a change in loc_rcvr_status on the basis of signals received at the MDI.

126.2.2.7.3 Effect of receipt

The effect of receipt of this primitive is specified in Figure 126–26 and in 126.2 and 126.4.6.3.

126.2.2.8 PMA_REMRXSTATUS.request

This primitive is generated by PCS Receive to indicate the status of the receive link at the remote PHY as communicated by the remote PHY via its encoding of its loc_rcvr_status parameter. The parameter rem_rcvr_status conveys to the PMA PHY Control function the information on whether reliable operation of the remote PHY is detected or not. The criterion for setting the parameter rem_rcvr_status is left to the implementer. It can be based, for example, on asserting rem_rcvr_status is NOT_OK until loc_rcvr_status is OK and then asserting the detected value of rem_rcvr_status after proper PCS Receive decoding is achieved.

126.2.2.8.1 Semantics of the primitive

PMA_REMRXSTATUS.request (rem_rcvr_status)

The `rem_rcvr_status` parameter can take on one of the following two values of the form:

OK	The receive link for the remote PHY is operating reliably.
NOT_OK	Reliable operation of the receive link for the remote PHY is not detected.

126.2.2.8.2 When generated

The PCS generates `PMA_REMRXSTATUS.request` messages to indicate a change in `rem_rcvr_status` on the basis of signals received at the MDI.

126.2.2.8.3 Effect of receipt

The effect of receipt of this primitive is specified in Figure 126–26.

126.2.2.9 PMA_ALERTDETECT.indication

This primitive is generated by PMA Receive to indicate the status of the receive link at the local PHY when `rx_lpi_active` is TRUE. The parameter `alert_detect` conveys to the PCS receive function information regarding the detection of the LPI alert signal by the PMA receive function. The criterion for setting the parameter `alert_detect` is left to the implementer.

126.2.2.9.1 Semantics of the primitive

`PMA_ALERTDETECT.indication` (`alert_detect`)

The `alert_detect` parameter can take on one of the following two values of the form:

TRUE	The alert signal has been reliably detected at the local receiver.
FALSE	The alert signal at the local receiver has not been detected.

126.2.2.9.2 When generated

The PMA generates `PMA_ALERTDETECT.indication` messages to indicate a change in the `alert_detect` status.

126.2.2.9.3 Effect of receipt

The effect of receipt of this primitive is specified in 126.3.2.3, Figure 126–16, and Figure 126–17.

126.2.2.10 PCS_RX_LPI_STATUS.request

When the PHY supports the EEE capability this primitive is generated by the PCS receive function to indicate the status of the receive link at the local PHY. The parameter `PCS_RX_LPI_STATUS.request` conveys to the PCS transmit and PMA receive functions information regarding whether the receive function is in the LPI receive mode. The parameter is generated by the Receive 64B/65B state diagram in Figure 126–16.

126.2.2.10.1 Semantics of the primitive

`PCS_RX_LPI_STATUS.request` (`rx_lpi_active`)

The `rx_lpi_active` parameter can take on one of the following two values of the form:

TRUE	The receive function is in the LPI receive mode.
FALSE	The receive function is not in the LPI receive mode.

126.2.2.10.2 When generated

The PCS generates PCS_RX_LPI_STATUS.request messages to indicate a change in the rx_lpi_active variable as determined by the receive state diagram in Figure 126–16.

126.2.2.10.3 Effect of receipt

The effect of receipt of this primitive is specified in 126.3.2.3 and Figure 126–30.

126.2.2.11 PMA_PCSDATAMODE.indication

This primitive indicates whether or not the PCS state diagrams are able to transition from their initialization states. The pcs_data_mode variable is generated by the PMA PHY Control function. It is passed to the PCS Control function via the PMA_PCSDATAMODE.indication primitive.

126.2.2.11.1 Semantics of the primitive

PMA_PCSDATAMODE.indication (pcs_data_mode)

The pcs_data_mode parameter can take on one of the following two values of the form:

TRUE	PHY is in state PCS_Data (see Figure 126–26).
FALSE	PCS is not in state PCS_Data (see Figure 126–26).

126.2.2.11.2 When generated

The PMA PHY Control function generates PMA_PCSDATAMODE.indication messages continuously.

126.2.2.11.3 Effect of receipt

Upon receipt of this primitive, the PCS performs its transmit function as described in 126.3.2.2.

126.2.2.12 PMA_FR_ACTIVE.indication

This primitive indicates whether or not the PMA is currently performing a fast retrain. The fr_active variable is generated by the PMA PHY Control function. It is passed to the PCS Receive Control function via the PMA_FR_ACTIVE.indication primitive. This primitive is only supported by PHYs with the fast retrain capability.

126.2.2.12.1 Semantics of the primitive

PMA_FR_ACTIVE.indication (fr_active)

The fr_active parameter can take on one of the following two values of the form:

TRUE	PHY is currently performing a fast retrain.
FALSE	PCS is not currently performing a fast retrain.

126.2.2.12.2 When generated

The PMA PHY Control function generates PMA_FR_ACTIVE.indication messages continuously.

126.2.2.12.3 Effect of receipt

The effect of receipt of this primitive is specified in Figure 126–16.

126.3 Physical Coding Sublayer (PCS)

126.3.1 PCS service interface (XGMII)


The PCS service interface allows the 2.5GBASE-T or 5GBASE-T PCS to transfer information to and from a PCS client. The PCS Interface is precisely defined as the 10 Gigabit Media Independent Interface (XGMII) in Clause 46.

126.3.2 PCS functions

The PCS comprises one PCS Reset function and two simultaneous and asynchronous operating functions. The PCS operating functions are: PCS Transmit and PCS Receive. All operating functions start immediately after the successful completion of the PCS Reset function.

The PCS reference diagram, Figure 126–5, shows how the two operating functions relate to the messages of the PCS-PMA interface. Connections from the management interface (signals MDC and MDIO) to other layers are pervasive and are not shown in Figure 126–5.

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation



NOTE 1—pcs_data_mode is required only for the EEE or fast retrain capabilities.
 NOTE 2—alert_detect and rx_lpi_active are only required for the EEE capability.
 NOTE 3—fr_active is only required for the fast retrain capability.

Figure 126-5—PCS reference diagram

126.3.2.1 PCS Reset function

PCS Reset initializes all PCS functions. The PCS Reset function shall be executed whenever one of the following conditions occur:

- a) Power on (see 126.3.6.2.2).
- b) The receipt of a request for reset from the management entity.

PCS Reset sets pcs_reset = true while any of the above reset conditions hold true. All state diagrams take the open-ended pcs_reset branch upon execution of PCS Reset. The reference diagrams do not explicitly show the PCS Reset function.

126.3.2.2 PCS Transmit function

The PCS Transmit function shall conform to the PCS 64B/65B Transmit state diagram in Figure 126–14 and the PCS Transmit bit ordering in Figure 126–6.

Dashed rectangles in Figure 126–14 and Figure 126–15 are used to indicate states and state transitions in the PCS 64B/65B Transmit state diagram that shall be supported by PHYs with the EEE capability. PHYs without the EEE capability do not support these transitions.

When communicating with the XGMII, the PCS uses a 4-octet-wide, synchronous data path, with packet delimiting being provided by transmit control signals and receive control signals. Alignment to 64B/65B is performed in the PCS. The PMA sublayer operates independently of block and packet boundaries. The PCS provides the functions necessary to map packets between the XGMII format and the PMA service interface format.

When the transmit channel is in normal mode, the PCS Transmit process continuously generates 65B blocks based upon the TXD <31:0> and TXC <3:0> signals on the XGMII. The subsequent functions of the PCS Transmit process then scramble the bits of the 65B blocks, pack the resulting scrambled blocks, append 97 zeros, and attach a leading aux channel bit, all of which are then processed by a low density parity check (LDPC) encoder. The appended zeros are then replaced by vendor discretionary randomized bits. The resulting 2048-bit LDPC frame is then mapped into PAM16 symbols. Transmit data-units are sent to the PMA service interface via the PMA_UNITDATA.request primitive.

In each symbol period, when communicating with the PMA, the PCS Transmit generates a code-group (A_n , B_n , C_n , D_n) that is transferred to the PMA via the PMA_UNITDATA.request primitive. The PMA transmits symbols A_n , B_n , C_n , D_n over wire-pairs BI_DA, BI_DB, BI_DC, and BI_DD respectively. The integer, n , is a time index that is introduced to establish a temporal relationship between different symbol periods. A symbol period, T , is 2.5/ S ns.

If a PMA_TXMODE.indication message has the value SEND_Z, PCS Transmit passes a vector of zeros at each symbol period to the PMA via the PMA_UNITDATA.request primitive.

If a PMA_TXMODE.indication message has the value SEND_T, PCS Transmit generates sequences of code-groups (TA_n , TB_n , TC_n , TD_n) defined in 126.3.4.2 to the PMA via the PMA_UNITDATA.request primitive. These code-groups are used for training mode and only transmit the values $\{-9, 9\}$ to keep the transmit power in the training mode the same as the transmit power in normal mode.

During training mode an Infocfield is transmitted at regular intervals containing messages for startup operation. By this mechanism, a PHY indicates the status of its own receiver to the link partner and makes requests for remote transmitter settings. (See 126.4.2.5.)

In the normal-mode of operation, the PMA_TXMODE.indication message has the value SEND_N, and the PCS Transmit function uses a 65B-LDPC coding to generate at each symbol period code-groups that represent data or control. During transmission, the 65B encoded bits are scrambled by the PCS using a PCS scrambler, 97 zero bits and an auxiliary bit are added, then frames are encoded into a code-group of 4D symbols and transferred to the PMA. During data encoding, PCS Transmit utilizes an LDPC frame encoder.

After reaching the normal mode of operation, EEE-capable PHYs may enter the LPI transmit mode under the control of the MAC via the XGMII. The EEE Transmit state diagram is contained within the PCS Transmit function. The EEE capability is described in 126.3.2.2.19.

126.3.2.2.1 Use of blocks

The PCS translates between XGMII signals and 65-bit blocks inserted within an LDPC frame using a 65B-LDPC coding scheme. The PAM2 PMA training frame synchronization allow establishment of LDPC frame and 65B boundaries by the PCS Synchronization process. Blocks and frames are unobservable and have no meaning outside the PCS. During the LPI mode, LDPC frame boundaries delimit sleep, wake, refresh, quiet, and alert cycles. The PCS functions ENCODE and DECODE generate, manipulate, and interpret blocks and frames as provided by the rules in 126.3.2.2.2.

126.3.2.2.2 65B-LDPC transmission code

The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link and to support transmission of control and data characters. In addition, the code enables the receiver to achieve PCS synchronization alignment on the incoming PHY bit stream.

The relationship of block bit positions to XGMII, PMA, and other PCS constructs is illustrated in Figure 126-6 for transmit and Figure 126-7 for receive. These figures illustrate the processing of a multiplicity of blocks containing 8 data octets. See 126.3.2.2.5 for information on how blocks containing control characters are mapped.

126.3.2.2.3 Notation conventions

For values shown as binary, the leftmost bit is the first transmitted bit.

64B/65B encodes 8 data octets or control characters into a block. Blocks containing control characters also contain a block type field. Data octets are labeled D_0 to D_7 . Control characters other than /O/, /S/ and /T/ are labeled C_0 to C_7 . The control character for ordered set is labeled as O_0 or O_4 since it is only valid on the first octet of the XGMII. The control character for start is labeled as S_0 or S_4 for the same reason. The control character for terminate is labeled as T_0 to T_7 .

Two consecutive XGMII transfers provides eight characters that are encoded into one 65-bit transmission block. The subscript in the above labels indicates the position of the character in the eight characters from the XGMII transfers.

Contents of block type fields, data octets and control characters are shown in hexadecimal notation. The LSB of the equivalent binary value represents the first transmitted bit. For instance, the block type field 0x1E is sent from left to right as 01111000. The bits of a transmitted or received block are labeled $TxB\langle 64:0 \rangle$ and $RxB\langle 64:0 \rangle$ respectively where $TxB\langle 0 \rangle$ and $RxB\langle 0 \rangle$ represent the first transmitted bit. The value of the data/ctrl header is shown as a binary value. Binary values are shown with the first transmitted bit (the LSB) on the left.


126.3.2.2.4 Transmission order

The PCS Transmit bit ordering shall conform to Figure 126-6. Note that this figure shows the mapping from XGMII to 64B/65B block for a block containing eight data characters.

126.3.2.2.5 Block structure

Blocks consist of 65 bits. The first bit of a block is the data/ctrl header. Blocks are either data blocks or control blocks. The data/ctrl header is 0 for data blocks and 1 for control blocks. The remainder of the block contains the payload.


Data blocks contain eight data characters. Control blocks begin with an 8-bit block type field that indicates the format of the remainder of the block. For control blocks containing a Start or Terminate character, that



NOTE— This figure shows the mapping from the XGMII to a 64B/65B block for a block containing eight data characters.

Figure 126-6—PCS Transmit bit ordering

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation



NOTE 1—This figure shows the mapping from a 64B/65B block for a block containing eight data characters to the XGMII.

NOTE 2—Conversion from 4DPAM-16 symbols to bits occurs in the LDPC decoder.

Figure 126-7—PCS Receive bit ordering

character is implied by the block type field. Other control characters are encoded in a 7-bit control code or a 4-bit O Code. Each control block contains eight characters.

The format of the blocks is as follows in Figure 126–8. In the figure, the column labeled Input Data shows, in abbreviated form, the eight characters used to create the 65-bit block. These characters are either data characters or control characters and, when transferred across the XGMII interface, the corresponding TXC or RXC bit is set accordingly. Within the Input Data column, D_0 through D_7 are data octets and are transferred with the corresponding TXC or RXC bit set to zero. All other characters are control octets and are transferred with the corresponding TXC or RXC bit set to one. The single bit fields (thin rectangles with no label in the figure) are sent as zero and ignored upon receipt.

Bits and field positions are shown with the least significant bit on the left. Hexadecimal numbers are shown prepended with '0x', and with the least significant digit on the right. For example the block type field 0x1E is sent as 01111000 representing bits 1 through 8 of the 65-bit block. The least significant bit for each field is placed in the lowest numbered position of the field.

All unused values of block type field are reserved.⁴

126.3.2.2.6 Control codes

The same set of control characters are supported by the XGMII, 2.5GBASE-T and 5GBASE-T PCS. The representations of the control characters are the control codes. The XGMII encodes a control character into an octet (an 8-bit value). The 2.5GBASE-T and 5GBASE-T PCS encode the start and terminate control characters implicitly by the block type field. The 2.5GBASE-T and 5GBASE-T PCS encode the ordered set control codes using a combination of the block type field and a 4-bit O code for each ordered set. The 2.5GBASE-T and 5GBASE-T PCS encode each of the other control characters into a 7-bit C code.

The control characters and their mappings to 2.5GBASE-T and 5GBASE-T control codes and XGMII control codes are specified in Table 126–1. All XGMII, 2.5GBASE-T and 5GBASE-T control code values that do not appear in the table shall not be transmitted and shall be treated as an error if received.

126.3.2.2.7 Ordered sets

Ordered sets are used to extend the ability to send control and status information over the link such as remote fault and local fault status. Ordered sets consist of a control character followed by three data characters. Ordered sets always begin on the first octet of the XGMII. One kind of ordered set is used for 2.5 and 5 Gigabit Ethernet—the sequence ordered set (see 46.3.4). The sequence ordered set control character is denoted /Q/. An additional ordered set, the signal ordered set, has been reserved and it begins with another control code. The 4-bit O field encodes the control code. See Table 126–1 for the mappings.

126.3.2.2.8 Idle (/I)

Idle control characters (/I) are transmitted when idle control characters are received from the XGMII. Idle characters may be added or deleted by the PCS to adapt between clock rates. /I/ insertion and deletion shall occur in groups of 4. /I/s may be added following idle or ordered sets. They shall not be added while data is being received. When deleting /I/s, the first four characters after a /T/ shall not be deleted.

⁴The block type field values have been chosen to have a 4-bit Hamming distance between them. The only unused value that maintains the Hamming distance is 0x00.

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

Input Data	data ctrl header	Block Payload									
Bit Position:	0 1	64									
Data Block Format:											
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ D ₇	0	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆	D ₇		
Control Block Formats:		Block									
C ₀ C ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	1	0x1E	C ₀	C ₁	C ₂	C ₃	C ₄	C ₅	C ₆	C ₇	
C ₀ C ₁ C ₂ C ₃ /O ₄ D ₅ D ₆ D ₇	1	0x2D	C ₀	C ₁	C ₂	C ₃	O ₄	D ₅	D ₆	D ₇	
C ₀ C ₁ C ₂ C ₃ /S ₄ D ₅ D ₆ D ₇	1	0x33	C ₀	C ₁	C ₂	C ₃		D ₅	D ₆	D ₇	
O ₀ D ₁ D ₂ D ₃ /S ₄ D ₅ D ₆ D ₇	1	0x66	D ₁	D ₂	D ₃	O ₀		D ₅	D ₆	D ₇	
O ₀ D ₁ D ₂ D ₃ /O ₄ D ₅ D ₆ D ₇	1	0x55	D ₁	D ₂	D ₃	O ₀	O ₄	D ₅	D ₆	D ₇	
S ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ D ₇	1	0x78	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆	D ₇		
O ₀ D ₁ D ₂ D ₃ /C ₄ C ₅ C ₆ C ₇	1	0x4B	D ₁	D ₂	D ₃	O ₀	C ₄	C ₅	C ₆	C ₇	
T ₀ C ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	1	0x87		C ₁	C ₂	C ₃	C ₄	C ₅	C ₆	C ₇	
D ₀ T ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	1	0x99	D ₀		C ₂	C ₃	C ₄	C ₅	C ₆	C ₇	
D ₀ D ₁ T ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	1	0xAA	D ₀	D ₁		C ₃	C ₄	C ₅	C ₆	C ₇	
D ₀ D ₁ D ₂ T ₃ /C ₄ C ₅ C ₆ C ₇	1	0xB4	D ₀	D ₁	D ₂		C ₄	C ₅	C ₆	C ₇	
D ₀ D ₁ D ₂ D ₃ /T ₄ C ₅ C ₆ C ₇	1	0xCC	D ₀	D ₁	D ₂	D ₃		C ₅	C ₆	C ₇	
D ₀ D ₁ D ₂ D ₃ /D ₄ T ₅ C ₆ C ₇	1	0xD2	D ₀	D ₁	D ₂	D ₃	D ₄		C ₆	C ₇	
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ T ₆ C ₇	1	0xE1	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅		C ₇	
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ T ₇	1	0xFF	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆		

Figure 126-8—64B/65B block formats

Table 126-1—Control codes

Control character	Notation	XGMII control codes	2.5G/5GBASE-T control codes	2.5G/5GBASE-T O code
idle	/I/	0x07	0x00	
LPI	/LI/	0x06	0x06	
start	/S/	0xFB	Encoded by block type field	
terminate	/T/	0xFD	Encoded by block type field	
error	/E/	0xFE	0x1E	
Sequence ordered set	/Q/	0x9C	Encoded by block type field plus O code	0x0

Table 126–1—Control codes (continued)

Control character	Notation	XGMII control codes	2.5G/5GBASE-T control codes	2.5G/5GBASE-T O code
reserved0		0x1C	0x2D	
reserved1		0x3C	0x33	
reserved2		0x7C	0x4B	
reserved3		0xBC	0x55	
reserved4		0xDC	0x66	
reserved5		0xF7	0x78	
Signal ordered set ^a	/Fsig/	0x5C	Encoded by block type field plus O code	0xF

^aReserved for INCITS T11 Fibre Channel use.

126.3.2.2.9 LPI (/LI/)

Low power idle (LPI) control characters (/LI/) on the XGMII indicate that the LPI client is requesting operation in the LPI transmit mode. A continuous stream of LPI control characters (/LI/) is used to maintain a link in the LPI transmit mode. Idle control characters (/I/) are used to transition from the LPI transmit mode to the normal mode. PHYs that support EEE respond to the LPI XGMII control characters using the procedure outlined in 126.1.3.3. LPI characters may be added or deleted by the PCS to adapt between clock rates. /LI/ insertion and deletion shall occur in groups of four. /LI/s may be added following low power idle characters. They shall not be added while data is being received.

If EEE is not supported, then /LI/ is not a valid control character.

126.3.2.2.10 Start (/S/)

The start control character (/S/) indicates the start of a packet. This delimiter is only valid on the first octet of the XGMII (TXD<7:0> and RXD<7:0>). Receipt of an /S/ on any other octet of TXD indicates an error. Block type field values implicitly encode an /S/ as the fifth or first character of the block. These are the only characters of a block on which a start can occur.

126.3.2.2.11 Terminate (/T/)

The terminate control character (/T/) indicates the end of a packet. Since packets may be any length, the /T/ can occur on any octet of the XGMII interface and within any character of the block. The location of the /T/ in the block is implicitly encoded in the block type field. A valid end of packet occurs when a block containing a /T/ is followed by a control block that does not contain a /T/.

126.3.2.2.12 ordered set (/O/)

The ordered set control characters (/O/) indicate the start of an ordered set. There are two kinds of ordered sets: the sequence ordered set and the signal ordered set (which is reserved). When it is necessary to designate the control character for the sequence ordered set specifically, /Q/ is used. /O/ is only valid on the first octet of the XGMII. Receipt of an /O/ on any other octet of TXD indicates an error. Block type field

values implicitly encode an /O/ as the first or fifth character of the block. The 4-bit O code encodes the specific /O/ character for the ordered set.

Sequence ordered sets may be deleted by the PCS to adapt between clock rates. Such deletion shall only occur when two consecutive sequence ordered sets have been received and shall delete only one of the two. Only Idles may be inserted for clock compensation. Signal ordered sets are not deleted for clock compensation.

126.3.2.2.13 Error (/E/)

The /E/ is sent whenever an /E/ is received. The /E/ allows physical sublayers such as the PCS to propagate received errors. See R_BLOCK_TYPE and T_BLOCK_TYPE function definitions in 126.3.6.2.4 for further information.

126.3.2.2.14 Transmit process

The transmit process generates blocks based upon the TXD<31:0> and TXC<3:0> signals received from the XGMII. Two XGMII data transfers are encoded into each block. 50 XGMII data transfers are encoded into an LDPC frame. It takes 128 PMA_UNITDATA transfers to send an LDPC frame of data. Therefore, if the PCS is connected to an XGMII and PMA sublayer where the ratio of their transfer rates is exactly 25:64, then the transmit process does not need to perform rate adaptation. Where the XGMII and PMA sublayer data rates are not synchronized to that ratio, the transmit process needs to insert idles, delete idles, or delete sequence ordered sets to adapt between the rates.

The transmit process generates blocks as specified in the PCS 64B/65B Transmit state diagram (see Figure 126–16 and Figure 126–17). The contents of each block are contained in a vector tx_coded<64:0>, which is passed to the scrambler. tx_coded<0> contains the data/ctrl header and the remainder of the bits contain the block payload.

126.3.2.2.15 PCS Scrambler

The payload of the PCS PHY frame is scrambled with a self-synchronizing scrambler. The scrambler for the MASTER shall produce the same result as the implementation shown in Figure 126–9. This implements the scrambler polynomial:⁵

$$G(x) = 1 + x^{39} + x^{58} \quad (126-1)$$

The scrambler for the SLAVE shall produce the same result as the implementation shown in Figure 126–9. This implements the scrambler polynomial:

$$G(x) = 1 + x^{19} + x^{58} \quad (126-2)$$


The initial seed values for the MASTER and SLAVE are left to the implementer. The scrambler is run continuously on all payload bits.

126.3.2.2.16 LDPC framing and LDPC encoder

The resulting payload of scrambled 25 65B blocks, followed by the 97 zero bits and preceded by 1 auxiliary bit results in a total payload of $25 \times 65 + 97 + 1 = 1723$ bits. The use of the auxiliary bit is for vendor-specific

⁵The convention here, which considers the most recent bit into the scrambler to be the lowest order term, is consistent with most references and with other scramblers shown in this standard. Some references consider the most recent bit into the scrambler to be the highest order term and would therefore identify this as the inverse of the polynomial in Equation (126–1). In case of doubt, note that the conformance requirement is based on the representation of the scrambler in the figure rather than the polynomial equation.

PCS scrambler employed by the MASTER



PCS scrambler employed by the SLAVE




Figure 126-9—MASTER and SLAVE PCS scramblers

communication is outside the scope of this document. For the purposes of this standard it is ignored by the link partner. The 1723 bits shall be encoded by the LDPC(1723, 2048) generator matrix G. G is described in Annex 55A.

The LDPC encoding takes the 1723 bit input code vector $x = [x_0 \ x_1 \ x_2 \ \dots \ x_{1722}]$, and shall generate the 2048 bit codeword c represented by the matrix multiplication $c = x \times G$. For both x and c the leftmost element of the vector is the first bit into the LDPC encoder and the first transmitted bit.

126.3.2.2.17 Substitution for zero-bit fill

The 2048 LDPC-coded bits output from the LDPC encoder in Figure 126-6 are then divided into three groups: the first 1626 bits, representing the auxiliary bit and the 25 scrambled 65B blocks of Ethernet payload; the subsequent 97 bits, representing the zero-fill added prior to encoding, and the subsequent 325 LDPC check bits. The group of 97 zero-fill bits are then replaced by vendor discretionary, randomized bits. The randomized fill bits should approximate the autocorrelation properties of the PCS scrambler described in 126.3.2.2.15, so as not to generate tones violating the transmit spectral PSD masks in 126.5.3.4.

126.3.2.2.18 PAM16 bit mapping

The LDPC frame shall be mapped four bits at a time in bit order of transmission into Gray-coded PAM-16 as follows in Table 126-2.

Table 126–2—PAM16 to Gray coded PAM16 mapping

Bits (b ₀ b ₁ b ₂ b ₃)	Hex	Level
0100	0x4	+15
0101	0x5	+13
0111	0x7	+11
0110	0x6	+9
0010	0x2	+7
0011	0x3	+5
0001	0x1	+3
0000	0x0	+1
1000	0x8	-1
1001	0x9	-3
1011	0xB	-5
1010	0xA	-7
1110	0xE	-9
1111	0xF	-11
1101	0xD	-13
1100	0xC	-15

126.3.2.2.19 EEE capability

The optional 2.5GBASE-T or 5GBASE-T EEE capability allows compliant PHYs to transition to an LPI mode of operation when link utilization is low.

PHYs that support EEE shall conform to the EEE transmit state diagram, shown in Figure 126–18, within the PCS.

When PCS_Reset is asserted or pcs_data_mode is not asserted, the state diagram enters the TX_NORMAL state.

When a complete 64B/65B block of LPI characters is generated by the PCS transmit function, the PHY transmits the sleep signal to indicate to the link partner that it is transitioning to the LPI transmit mode. If the sleep signal begins on an even LDPC frame boundary aligned to the inversion on pair A during PMA training, then it contains 18 full LDPC frames each composed entirely of LDPC encoded LP_IDLE blocks. If the sleep signal does not begin on an even LDPC frame boundary, then it contains one to two LDPC frames partially composed of LP_IDLE blocks followed by 18 LDPC frames fully composed of LP_IDLE blocks.

Following the transmission of the sleep signal, quiet-refresh signaling begins, as described in 126.3.5.

After the sleep signal is transmitted LPI control characters shall be input to the PCS scrambler continuously until the PCS Transmit Function exits the LPI transmit mode.

While the PMA asserts SEND_N, the lpi_tx_mode variable shall control the transmit signal through the PMA_UNITDATA.request primitive described as follows:

When the PHY is not in the PCS_Data state, the lpi_tx_mode variable is ignored.

When the lpi_tx_mode variable takes the value NORMAL and the PMA asserts SEND_N, the PCS passes coded data to the PMA via the PMA_UNITDATA.request primitive as described in 126.3.2.2.

When the lpi_tx_mode variable takes the value QUIET and the PMA asserts SEND_N, the PCS passes zeros to the PMA through the PMA_UNITDATA.request primitive.

When the lpi_tx_mode variable takes the value REFRESH_A and the PMA asserts SEND_N, the PCS passes the PMA training signal to the PMA on pair A, to allow both the local and remote PHY to refresh adaptive filters and timing loops. The PCS passes zeros to all other pairs in this condition. REFRESH_B, REFRESH_C and REFRESH_D operate in an analogous manner for the other pairs.

When the lpi_tx_mode variable takes the value ALERT and the PMA asserts SEND_N, the PCS passes the ALERT vector to the PMA.

The quiet-refresh cycle is repeated until codewords other than LP_IDLE are detected at the XGMII. These codewords indicate that the local system is requesting a transition back to the normal operational mode. Following this event, the PMA_UNITDATA.request parameter tx_symb_vector is set to the value ALERT. The alert signal is not synchronized with respect to the quiet-refresh cycle but shall be synchronized so that the alert signal from the PMA begins on a LDPC 2-frame 256 4D-symbol boundary aligned to the inversion on pair A during PMA training.

The PHY also transitions back to the normal operation mode if an error condition occurs. This error condition is defined as the detection of any characters other than LPI or IDLE at the XGMII.

After the alert signal the PCS completes the transition from LPI mode to normal mode by sending a wake signal containing lpi_wake_time LDPC frames composed of IDLE 64B/65B blocks.

lpi_wake_time is a fixed parameter that is defined as 18 LDPC frames as follows in Table 126–3. The maximum PHY wake time when wake is requested before sleep has been completely transmitted is $14.72/S \mu\text{s}$ ($lpi_wake_timer = T_{w_phy}$ as defined by Clause 78). The maximum PHY wake time when wake is requested after sleep has been completely transmitted is $8.96/S \mu\text{s}$.

Table 126–3—LPI wake time

lpi_wake_time (frames)	lpi_wake_timer when wake starts before sleep signal is complete		lpi_wake_timer when wake starts after sleep signal is complete	
	(frames)	(μs)	(frames)	(μs)
18	46	$14.72 / S$	28	$8.96 / S$

126.3.2.3 PCS Receive function

The PCS Receive function shall conform to the PCS 64B/65B receive state diagram in Figure 126–16 and Figure 126–17 and the PCS Receive bit ordering in Figure 126–7 including compliance with the associated state variables as specified in 126.3.6.

The PCS Receive function accepts received code-groups provided by the PMA Receive function via the parameter `rx_symb_vector`. The PCS receiver uses knowledge of the encoding rules to correctly align the 65B-LDPC frames. The randomized bits are replaced with known zeros. The received 65B-LDPC frames are decoded with error correction; the auxiliary bit and the trailing zero-fill bits are then stripped; descrambling is then performed. This process generates the 64B/65B block vector `rx_coded<64:0>` which is then decoded to form the XGMII signals `RXD<31:0>` and `RXC<3:0>` as specified in the PCS 64B/65B Receive state diagram (see Figure 126–16 and 126–17). Two XGMII data transfers are decoded from each block. Where the XGMII and PMA sublayer data rates are not synchronized to a 25:64 ratio, the receive process inserts idles, deletes idles, or deletes sequence ordered sets to adapt between rates.

During PMA training mode, PCS Receive checks the received PAM2 framing and signals the reliable acquisition of the descrambler state by setting the `scr_status` parameter of the `PMA_SCRSTATUS` request primitive to OK.

When the PCS Synchronization process has obtained synchronization, the LDPC frame error ratio (LFER) monitor process monitors the signal quality asserting `hi_lfer` if excessive LDPC frame errors are detected (LDPC parity error). If 40 consecutive LDPC frame errors are detected, the `block_lock` flag is de-asserted. When `block_lock` is asserted and `hi_lfer` is de-asserted, the PCS Receive process continuously accepts blocks. The PCS Receive process monitors these blocks and generates `RXD <31:0>` and `RXC <3:0>` on the XGMII.

When the receive channel is in training mode, the PCS Synchronization process continuously monitors `PMA_RXSTATUS.indication (loc_rcvr_status)`. When `loc_rcvr_status` indicates OK, then the PCS Synchronization process accepts data-units via the `PMA_UNITDATA.indication` primitive. It attains frame and block synchronization based on the PMA training frames and conveys received blocks to the PCS Receive process. The PCS Synchronization process sets the `block_lock` flag to indicate whether the PCS has obtained synchronization. The PMA training sequence includes 1 bit pattern on pair A every 256 PAM2 symbols, which is aligned with the PCS boundary of two LDPC frames. When the PCS Synchronization process is synchronized to this pattern, `block_lock` is asserted.

PHYs with the EEE capability support transition to the LPI mode when the PHY has successfully completed training and `pcs_data_mode` is TRUE. Transitions to and from the LPI mode are allowed to occur independently in the transmit and receive functions. The PCS receive function is responsible for detecting transitions to and from the LPI receive mode and indicating these transitions using signals defined in 126.2.2.

The link partner signals a transition to the LPI mode of operation by transmitting 18 LDPC frames composed entirely of 64B/65B blocks of /LI/. When blocks of /LI/ are detected at the output of the 64B/65B decoder, `rx_lpi_active` is asserted by the PCS receive function and the /LI/ character is continuously asserted at the receive XGMII. These frames may be preceded by a frame composed partially of /LI/ characters. After these frames the link partner begins transmitting zeros, and it is recommended that the receiver power down receive circuits to reduce power consumption. The receive function uses LDPC frame counters to maintain synchronization with the remote PHY and receives periodic refresh signals that are used to update coefficients, so that the integrity of adaptive filters and timing loops in the PMA is maintained. LPI signaling is defined in 126.3.5. The quiet-refresh cycle continues until the PMA asserts `alert_detect` to indicate that the alert signal has been reliably detected. After the alert signal the link partner transmits repeated /I/ characters, representing a wake signal. The PHY receive function sends /I/ to the XGMII for 18 LDPC frame periods and then resumes normal operation.

126.3.2.3.1 Frame and block synchronization


When the receive channel is operating in normal mode, the frame and block synchronization function receives data via 4D-PAM16 `PMA_UNITDATA.indication` primitives. It shall form a 4D-PAM16 stream from the primitives by concatenating requests with the PAM16s of each primitive in order from

rx_symb_vector<0> to rx_symb_vector<127> (see Figure 126–7). It obtains block_lock to the LDPC frames during the PAM2 training pattern using synchronization bits provided on pair A. The 65-bit blocks are extracted based on their location in the LDPC frame.

126.3.2.3.2 PCS descrambler

The descrambler processes the payload to reverse the effect of the scrambler using the same polynomial. It shall produce the same result as the implementations shown in Figure 126–10 for the MASTER and the SLAVE.

PCS descrambler employed by the MASTER



PCS descrambler employed by the SLAVE




Figure 126–10—MASTER and SLAVE PCS descramblers

126.3.2.3.3 Invalid blocks

A block is invalid if any of the following conditions exists:

- a) The block type field contains a reserved value.
- b) Any control character contains a value not in Table 126–1.
- c) Any O code contains a value not in Table 126–1.
- d) The block contains information from the payload of an invalid received PHY frame or the first 64B/65B block following an invalid received PHY frame.

The PCS Receive function shall check the integrity of the LDPC parity bits defined in 126.3.2.2.16. If the check fails the PHY frame is invalid.


R_BLOCK_TYPE of an invalid block is set to E.

126.3.3 Test-pattern generators

The test-pattern generator mode is provided for enabling joint testing of the local transmitter, the channel and remote receiver. When the transmit PCS is operating in test-pattern mode it shall transmit continuously as illustrated in Figure 126–6, with the input to the scrambler set to zero and the initial condition of the scrambler set to any non-zero value. When the receiver PCS is operating in test-pattern mode it shall receive continuously as illustrated in Figure 126–7. After acquiring the self-synchronizing scrambler state, the output of the received scrambled values should ideally be zero. Any nonzero values correspond to receiver bit errors. This mode is further described as test mode 7 in 126.5.2.

126.3.4 PMA training side-stream scrambler polynomials

The PCS Transmit function employs side-stream scrambling for generating 2-level PAM PMA training sequences as follows in Figure 126–11. An implementation of MASTER and SLAVE PHY side-stream scramblers is shown in the “Main PN sequence” box. The bits stored in the shift register delay line at time n are denoted by $Scr_n[32:0]$. At each symbol period, the shift register is advanced by one bit, and one new bit represented by $Scr_n[0]$ is generated. The transmitter side-stream scrambler is reset upon execution of the PCS Reset function. If PCS Reset is executed, all bits of the 33-bit vector representing the side-stream scrambler state are arbitrarily set. The initialization of the scrambler state is left to the implementer. In no case shall the scrambler state be initialized to all zeros.



Main PN sequence:

$$\text{Scr}_n[32:1] = \text{Scr}_{n-1}[31:0]$$

$$\text{Scr}_n[0] = \begin{cases} \text{Scr}_{n-1}[12] + \text{Scr}_{n-1}[32] & \text{if PMA_CONFIG=MASTER} \\ \text{Scr}_{n-1}[19] + \text{Scr}_{n-1}[32] & \text{if PMA_CONFIG = SLAVE} \end{cases}$$

Derived sequences:

$$\text{Sa}_n = \begin{cases} \text{Scr}_n[0] \oplus 1 & \text{if } n \bmod 256 = 0 \\ \text{Scr}_n[0] & \text{otherwise} \end{cases}$$

$$\text{Sb}_n = \text{Scr}_n[3] \oplus \text{Scr}_n[8]$$

$$\text{Sc}_n = \text{Scr}_n[6] \oplus \text{Scr}_n[16]$$

$$\text{Sd}_n = \text{Scr}_n[9] \oplus \text{Scr}_n([14] \oplus \text{Scr}_n([19] \oplus \text{Scr}_n[24]))$$

Figure 126–11—A realization of PMA training PAM2 sequences

126.3.4.1 Generation of bits Sa_n , Sb_n , Sc_n , Sd_n

PMA training signal encoding rules are based on the generation, at time n , of the four bits Sa_n , Sb_n , Sc_n , Sd_n . These four bits are generated in a systematic fashion using the bits in $\text{Scr}_n[32:0]$, and an auxiliary generating polynomial. For both MASTER and SLAVE PHYs, they are obtained by the same linear combinations of bits stored in the transmit scrambler shift register delay line. These four bits are derived from elements of the same maximum-length shift register sequence of length $2^{33} - 1$ as $\text{Scr}_n[0]$, but shifted in time. The associated delays are all large and different so that there is no short-term correlation among the bits Sa_n , Sb_n , Sc_n , Sd_n . The four bits are generated using the bit $\text{Scr}_n[0]$ and the equations in Figure 126–11 in the “Derived sequences” box.

126.3.4.2 Generation of 4D symbols TA_n , TB_n , TC_n , TD_n

The four bits Sa_n , Sb_n , Sc_n , Sd_n are mapped to a 4D symbol (TA_n , TB_n , TC_n , TD_n) as follows in Figure 126–11.

The inversion on pair A at 256 intervals ($n = k \times 256, k = 0, 1, 2, \dots$) defines the LDPC boundary during data mode.

Notice that over the repeating time intervals of 16384 and of length 128, $m \times 16384 - 128 \leq n < m \times 16384, m = 1, 2, 3, \dots$, the PMA training pattern in pair A is XOR'ed with the Infocfield. Thus, pair A transmits the Infocfield, which communicates to the remote transceiver settings of THP and power backoff and other control information.

126.3.4.3 PMA training mode descrambler polynomials

The PHY shall acquire descrambler state synchronization to the PAM2 training sequence and report success through scr_status. For side-stream descrambling, the MASTER PHY shall employ the receiver descrambler generator polynomial $g'_M(x) = 1 + x^{20} + x^{33}$ and the SLAVE PHY shall employ the receiver descrambler generator polynomial $g'_S(x) = 1 + x^{13} + x^{33}$.

126.3.5 LPI signaling

PHYs with EEE capability have transmit and receive functions that can enter and leave the LPI mode independently. The PHY can transition to the LPI mode when the PHY has successfully completed training and pcs_data_mode is TRUE. The transmit function of the PHY initiates a transition to the LPI transmit mode when it generates 64B/65B blocks composed entirely of LPI control characters, as described in 126.3.2.2.19. The transmit function of the link partner signals the transition using the sleep signal. When the transmitter begins to send the sleep signal, it asserts tx_lpi_active and the transmit function enters the LPI transmit mode.

Within the LPI mode PHYs use a repeating quiet-refresh cycle (see Figure 126–12). The first part of this cycle is known as the quiet period and lasts for a time lpi_quiet_time equal to 120 LDPC frame periods. The quiet period is defined in 126.3.5.2. The second part of this cycle is known as the refresh period and lasts for a time lpi_refresh_time equal to 8 LDPC frame periods. The refresh period is defined in 126.3.5.3. A cycle composed of one quiet period and one refresh period is known as a single pair LPI cycle and lasts for a time lpi_qr_time equal to 128 LDPC frame periods. The time taken to complete a quiet-refresh cycle for all four pairs is known as a complete LPI cycle.

lpi_offset, lpi_quiet_time, lpi_refresh_time, lpi_qr_time, and lpi_allpairs_qr_time are timing parameters that are integer multiples of the LDPC frame period. lpi_offset is a fixed value equal to lpi_qr_time/2 that is used to ensure refresh signals are appropriately offset by the link partners.




Figure 126–12—Timing periods for LPI signals

PHYs begin the transition from the LPI receive mode when the alert signal is detected by the PMA as defined in 126.4.2.4.

126.3.5.1 LPI Synchronization

To maximize power savings, maintain link integrity, and ensure interoperability, EEE-capable PHYs must synchronize refresh intervals during the LPI mode. The transition to PCS_Test is used as a fixed timing reference for the link partners. Refresh signaling is derived by counting LDPC frames from the transition to PCS_Test.

In initial training, normal retraining, and fast retraining, with or without the EEE capability being supported, the master and slave signal when they transition to PCS_Test using the transition counter following the procedure described in 126.4.2.5.15.

A EEE-capable PHY in slave mode is responsible for synchronizing its PMA training frame to the master's PMA training frame during the transition to PMA_Training_Init_S. The slave shall ensure that its PMA training frames are synchronized to the master's PMA training frames within two LDPC frames, measured at the slave MDI on pair A. In addition, the slave shall initialize its transition counter so that it transitions to PCS_Test within two LDPC frames of the master PHY's transition to PCS_Test, measured at the slave PHY's MDI on pair A. This mechanism ensures that the refresh offset is bounded to a small value at both MDI interfaces, thus ensuring there is no overlap of master and slave signals when both transmit and receive are in the LPI mode.

Following the transition to PCS_Test, the PCS counts transmitted and received LDPC frames, and uses these counters to generate refresh and pair control signals for the transmit and receive functions. The transmitted LDPC frame count is named tx_ldpc_frame_cnt. The received LDPC frame count is named rx_ldpc_frame_cnt.

The master and slave shall derive the active pair and refresh_active signals from the LDPC frame counters as follows in Table 126-4 and Table 126-5.

Table 126-4—Synchronization logic derived from slave signal LDPC frame count

Slave-side variable	Master-side variable	for master u=rx_ldpc_frame_cnt for slave u=tx_ldpc_frame_cnt
tx_refresh_active=true	rx_refresh_active=true	$lpi_offset - lpi_refresh_time \leq \text{mod}(u, lpi_qr_time) < lpi_offset$
tx_lpi_full_refresh=true	N/A	$lpi_offset - lpi_refresh_time = \text{mod}(u, lpi_qr_time)$
tx_active_pair=PAIR_A	rx_active_pair=PAIR_A	$lpi_offset + lpi_qr_time \leq u < lpi_offset + 2 \times lpi_qr_time$
tx_active_pair=PAIR_B	rx_active_pair=PAIR_B	$lpi_offset + 2 \times lpi_qr_time \leq u < lpi_offset + 3 \times lpi_qr_time$
tx_active_pair=PAIR_C	rx_active_pair=PAIR_C	$lpi_offset + 3 \times lpi_qr_time \leq u < 4 \times lpi_qr_time$ OR $0 \leq u < lpi_offset$
tx_active_pair=PAIR_D	rx_active_pair=PAIR_D	$lpi_offset \leq u < lpi_offset + lpi_qr_time$

Table 126–5—Synchronization logic derived from master signal LDPC frame count

Slave-side variable	Master-side variable	for master $v=tx_ldpc_frame_cnt$ for slave $v=rx_ldpc_frame_cnt$
$rx_refresh_active=true$	$tx_refresh_active=true$	$lpi_quiet_time \leq \text{mod}(v, lpi_qr_time)$
N/A	$tx_lpi_full_refresh=true$	$lpi_quiet_time = \text{mod}(v, lpi_qr_time)$
$rx_active_pair=PAIR_A$	$tx_active_pair=PAIR_A$	$0 \leq v < lpi_qr_time$
$rx_active_pair=PAIR_B$	$tx_active_pair=PAIR_B$	$lpi_qr_time \leq v < 2 \times lpi_qr_time$
$rx_active_pair=PAIR_C$	$tx_active_pair=PAIR_C$	$2 \times lpi_qr_time \leq v < 3 \times lpi_qr_time$
$rx_active_pair=PAIR_D$	$tx_active_pair=PAIR_D$	$3 \times lpi_qr_time \leq v < 4 \times lpi_qr_time$

126.3.5.2 Quiet period signaling

During the quiet period the transmitters on all four pairs should be turned off. Average launch power (as measured from 56 LDPC frames after a refresh period to 56 LDPC frames before the next refresh period on the same lane) for each Transmitter shall be less than –41 dBm. This requirement does not apply to the periods when the alert signal is transmitted as defined in 126.4.2.2.1.

126.3.5.3 Refresh period signaling

During the LPI mode 2.5GBASE-T and 5GBASE-T PHYs use staggered, out-of-phase refresh signaling to maximize power savings. Two-level PAM refresh symbols are generated using the PMA side-stream scrambler polynomials described in 126.3.4 and exactly as is shown in Figure 126–11 with the exception that the Infocfield consists of a sequence of 128 zeros. The training sequence described in 126.3.4 shall be used during the LPI mode, with the scramblers free-running from PCS Reset.

Refresh signals shall be sent using the THP filter as described in 126.4.3.1. At the start of each refresh signal the THP feedback delay line shall be initialized with zeros.

While a transmit function is in the LPI transmit mode only one of the transmit pairs is active during a refresh period. tx_symb_vector for all transmit pairs that are not active shall be set to zero.

When tx_symb_vector has the value ALERT and the PHY is master, the transmitter on pair A shall be active and all other pairs shall be quiet. When tx_symb_vector has the value ALERT and the PHY is slave, the transmitter on pair C shall be active and all other pairs shall be quiet. If $lpi_tx_mode=REFRESH_A$ on a MASTER PHY or $lpi_tx_mode=REFRESH_C$ on a SLAVE PHY, and tx_symb_vector has the value ALERT, then the alert signaling shall be transmitted in place of the refresh signaling where the signals overlap.

126.3.6 Detailed functions and state diagrams

126.3.6.1 State diagram conventions

The body of this subclause is composed of state diagrams, including the associated definitions of variables, constants, and functions. 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.

126.3.6.2 State diagram parameters

126.3.6.2.1 Constants

EBLOCK_R<71:0>

72 bit vector to be sent to the XGMII interface containing /E/ in all the eight character locations.

EBLOCK_T<64:0>

65 bit vector to be sent to the LDPC encoder containing /E/ in all the eight character locations.

LBLOCK_R<71:0>

72 bit vector to be sent to the XGMII interface containing two Local Fault ordered sets. The Local Fault ordered set is defined in 46.3.4.

LBLOCK_T<64:0>

65 bit vector to be sent to the LDPC encoder containing two Local Fault ordered sets.

LPBLOCK_R<71:0>

72 bit vector to be sent to the XGMII containing /LI/ in all the eight character locations.

LPBLOCK_T<64:0>

65 bit vector to be sent to the LDPC encoder containing /LI/ in all the eight character locations.

IBLOCK_R<71:0>

72 bit vector to be sent to the XGMII containing /I/ in all the eight character locations.

IBLOCK_T<64:0>

65 bit vector to be sent to the LDPC encoder containing /I/ in all the eight character locations.

UBLOCK_R<71:0>

72 bit vector to be sent to the XGMII containing two Link Interruption ordered sets. The Link Interruption ordered set is defined in 46.3.4.

126.3.6.2.2 Variables

lfer_test_lf

Boolean variable that is set true when a new LDPC frame is available for testing and false when LFER_TEST_LF state is entered. A new LDPC frame is available for testing when the Block Sync process has accumulated enough symbols from the PMA to evaluate the next LDPC frame.

block_lock

Boolean variable that is set true when receiver acquires block delineation.

hi_lfer

Boolean variable that is asserted true when the lfer_cnt reaches 16 errors in one lfer_timer interval.

pcs_reset

Boolean variable that controls the resetting of the PCS. It is true whenever a reset is necessary including when reset is initiated from the MDIO, during power on, and when the MDIO has put the PCS into low-power mode.

rx_coded<64:0>

Vector containing the input to the 64B/65B decoder. The format for this vector is shown in Figure 126-8. The leftmost bit in the figure is rx_coded<0> and the rightmost bit is rx_coded<64>.

rx_raw<71:0>

Vector containing two successive XGMII output transfers. RXC<3:0> for the first transfer are taken from rx_raw<3:0>. RXC<3:0> for the second transfer are taken from rx_raw<7:4>. RXD<31:0> for the first transfer are taken from rx_raw<39:8>. RXD<31:0> for the second transfer are taken from rx_raw<71:40>.

lf_valid

Boolean indication that is set true if received LDPC frame is valid. LDPC frame is valid if and only if all parity checks of the LDPC code are satisfied.

tx_coded<64:0>

Vector containing the output from the 64B/65B encoder. The format for this vector is shown in Figure 126–8. The leftmost bit in the figure is tx_coded<0> and the rightmost bit is tx_coded<64>.

tx_raw<71:0>

Vector containing two successive XGMII transfers. TXC<3:0> for the first transfer are placed in tx_raw<3:0>. TXC<3:0> for the second transfer are placed in tx_raw<7:4>. TXD<31:0> for the first transfer are placed in tx_raw<39:8>. TXD<31:0> for the second transfer are placed in tx_raw<71:40>.

The following variables are required for PHYs that support the EEE capability:

tx_lpi_active

A Boolean variable that is set true when the PHY transmit function is operating in the LPI transmit mode and during transitions to and from the LPI transmit mode (i.e., at any time when the PHY is transmitting sleep, alert, wake, or quiet-refresh signaling). It is set false otherwise.

tx_lpi_qr_active

A Boolean variable that is set true during the LPI transmit mode, when the PHY is transmitting quiet-refresh signaling. Set false otherwise.

rx_lpi_active

A Boolean variable that is set true when the PHY receive function is operating in the LPI receive mode and set false otherwise. The LPI receive mode begins when the sleep signal is detected and lasts until the alert signal is detected. When the EEE capability is not supported, rx_lpi_active is set false.

tx_lpi_req

A Boolean variable that is set true when the LPI client indicates that it is requesting operation in the LPI transmit mode via the XGMII and set false otherwise.

alert_detect

Indicates that an alert signal from the link partner has been received at the MDI as indicated by PMA_ALERTDETECT.indication(alert_detect).

tx_lpi_alert_active

A Boolean variable that is set true when the PHY is transmitting ALERT signaling. Set false otherwise.

rx_lpi_wake

A Boolean variable that is set true when the PHY receiver is in the WAKE state and sending IDLE to the XGMII. Set false otherwise. When the EEE capability is not supported, rx_lpi_wake is set false.

tx_active_pair

A variable indicating the transmit active pair during the LPI transmit mode. The variable may take the values PAIR_A, PAIR_B, PAIR_C, PAIR_D. This variable is defined in 126.3.5.1.

lpi_tx_mode

A variable indicating the signaling to be used from the PCS to the PMA across the PMA_UNITDATA.request (tx_symb_vector) interface.

lpi_tx_mode controls tx_symb_vector only when tx_mode is set to SEND_N.

The variable is set to NORMAL when (!tx_lpi_qr_active * !tx_lpi_alert_active), indicating that the PCS is in the normal mode of operation and encodes code-groups as described in Figure 126–14 and Figure 126–15.

The variable is set to REFRESH_A when (tx_lpi_qr_active * (tx_active_pair=PAIR_A) * tx_refresh active).

The variable is set to REFRESH_B when (tx_lpi_qr_active * (tx_active_pair=PAIR_B) * tx_refresh active).

The variable is set to REFRESH_C when (tx_lpi_qr_active * (tx_active_pair=PAIR_C) * tx_refresh active).

tx_refresh active).

The variable is set to REFRESH_D when $(tx_lpi_qr_active * (tx_active_pair=PAIR_D) * tx_refresh_active)$.

The variable is set to QUIET when $(tx_lpi_qr_active * (!tx_refresh_active + tx_lpi_initial_quiet))$.

The variable is set to ALERT when $(tx_lpi_alert_active)$.

tx_refresh_active

A Boolean value. This variable is set true following the logic described in 126.3.5.1.

tx_lpi_full_refresh

A Boolean value. This variable is set true following the logic described in 126.3.5.1.

tx_lpi_initial_quiet

A Boolean value. This variable is set true when the transmit function enters the LPI transmit mode and a partial refresh is replaced by quiet signaling.

ldpc_two_frame_done

A Boolean value. This variable is set true when the final symbol of each even LDPC frame aligned to the inversion on pair A during PMA training is transmitted and is set false otherwise.

The following variable is only required for PHYs that support the fast retrain capability:

fr_sigtype

If fast retrain is supported, this variable controls the block type the PMA sends on the receive path during fast retrain. If MDIO is supported, this variable is set based on the value in 1.147.2:1 as follows:

- 00 IBLOCK_R
- 01 LBLOCK_R
- 10 UBLOCK_R
- 11 Reserved

If MDIO is not supported, an equivalent method of controlling fast retrain functionality should be provided.

126.3.6.2.3 Timers

State diagram timers follow the conventions described in 14.2.3.2.

lfer_timer

Timer that is triggered every $125/S \mu s +1\%, -25\%$. When the timer reaches its terminal count it sets lfer_timer_done = TRUE.

The following timers are required for PHYs that support the EEE capability:

lpi_tx_sleep_timer

This timer defines the time the local transmitter sends the sleep signal to the link partner.

Values: The condition lpi_tx_sleep_timer_done becomes true upon timer expiration.

Duration: This timer shall have a period equal to 18 LDPC frame periods.

lpi_tx_alert_timer

This timer defines the time the local transmitter transmits the alert signal.

Values: The condition lpi_tx_alert_timer_done becomes true upon timer expiration.

Duration: This timer shall have a period equal to 8 LDPC frame periods.

lpi_tx_wake_timer

This timer defines the time the local transmitter transmits the wake signal.

Values: The condition lpi_tx_wake_timer_done becomes true upon timer expiration.

Duration: This timer shall have a period equal to lpi_wake_time LDPC frame periods.

lpi_rx_wake_timer

This timer defines the time the receiver sends IDLE blocks to the XGMII after the alert signal is detected.

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

Values: The condition `lpi_rx_wake_timer_done` becomes true upon timer expiration.
 Duration: This timer shall have a period equal to `lpi_wake_time` LDPC frame periods.

126.3.6.2.4 Functions

DECODE(`rx_symb_vector<64:0>`)

In the PCS Receive process, this function takes as its argument 65-bit `rx_coded<64:0>` from the LDPC decoder and decodes the 65B-LDPC bit vector returning a vector `rx_raw<71:0>`, which is sent to the XGMII. The DECODE function shall decode the block based on code specified in 126.3.2.2.2.

ENCODE(`tx_raw<71:0>`)

Encodes the 72-bit vector received from the XGMII, returning 65-bit vector `tx_coded`. The ENCODE function shall encode the block as specified in 126.3.2.2.2.

R_BLOCK_TYPE = {C, S, T, D, E, I, LI, LII}

When the EEE capability is not supported, this function classifies each 65-bit `rx_coded` vector as belonging to one of the five types {C, S, T, D, E} depending on its contents.

When the EEE capability is supported, this function classifies each 65-bit `rx_coded` vector as belonging to the eight types depending on its contents. A vector may simultaneously belong to the C and I types when it contains eight valid control characters that are all /I/, but in every other case the vector belongs to only one type.

Values: C; The vector contains a data/ctrl header of 1 and one of the following:

- a) A block type field of 0x1E and eight valid control characters other than /E/ and /LI/;
- b) A block type field of 0x2D or 0x4B, a valid O code, and four valid control characters;
- c) A block type field of 0x55 and two valid O codes.

S; The vector contains a data/ctrl header of 1 and one of the following:

- a) A block type field of 0x33 and four valid control characters;
- b) A block type field of 0x66 and a valid O code;
- c) A block type field of 0x78.

T; The vector contains a data/ctrl header of 1, a block type field of 0x87, 0x99, 0xAA, 0xB4, 0xCC, 0xD2, 0xE1, or 0xFF and all control characters are valid.

D; The vector contains a data/ctrl header of 0.

I; If the optional EEE capability is supported, then the I type is a special case of the C type where the vector contains a data/ctrl header of 1, a block type field of 0x1e, and eight control characters of /I/.

LI: If the optional EEE capability is supported, then the LI type occurs when the vector contains a data/ctrl header of 1, a block type field of 0x1e, and eight control characters of /LI/.

LII: If the optional EEE capability is supported, then the LII type occurs when the vector contains a data/ctrl header of 1, a block type field of 0x1E, and one of the following:

- a) Four control characters of /LI/ followed by four control characters of /I/;
- b) Four control characters of /I/ followed by four control characters of /LI/.

E; The vector does not meet the criteria for any other value.

A valid control character is one containing a 2.5G/5GBASE-T control code specified in Table 126-1. A valid O code is one containing an O code specified in Table 126-1.

R_TYPE(`rx_coded<64:0>`)

Returns the R_BLOCK_TYPE of the `rx_coded<64:0>` bit vector.

R_TYPE_NEXT

Prescient end of packet check function. It returns the R_BLOCK_TYPE of the `rx_coded` vector immediately following the current `rx_coded` vector.

T_BLOCK_TYPE = {C, S, T, D, E, I, LI, LII}

When the EEE capability is not supported, this function classifies each 72-bit tx_raw vector as belonging to one of the five types {C, S, T, D, E} depending on its contents.

When the EEE capability is supported, this function classifies each 72-bit tx_raw vector as belonging to the eight types depending on its contents. A vector may simultaneously belong to the C and I types when it contains eight valid control characters that are all /I/, but in every other case the vector belongs to only one type.

Values: C; The vector contains one of the following:

- a) Eight valid control characters other than /O/, /S/, /T/, /E/, and /LI/;
- b) One valid ordered set and four valid control characters other than /O/, /S/, and /T/;
- c) Two valid ordered sets.

S; The vector contains an /S/ in its first or fifth character, any characters before the S character are valid control characters other than /O/, /S/, and /T/ or form a valid ordered set, and all characters following the /S/ are data characters.

T; The vector contains a /T/ in one of its characters, all characters before the /T/ are data characters, and all characters following the /T/ are valid control characters other than /O/, /S/ and /T/.

D; The vector contains eight data characters.

I; If the optional EEE capability is supported, then the I type is a special case of the C type where the vector contains eight control characters of /I/.

LI: If the optional EEE capability is supported, then the LI type occurs when the vector contains eight control characters of /LI/.

LII: If the optional EEE capability is supported, then the LII type occurs when the vector contains one of the following:

- a) Four control characters of /LI/ followed by four control characters of /I/;
- b) Four control characters of /I/ followed by four control characters of /LI/.

E; The vector does not meet the criteria for any other value.

A tx_raw character is a control character if its associated TXC bit is asserted. A valid control character is one containing an XGMII control code specified in Table 126–1. A valid ordered set consists of a valid /O/ character in the first or fifth characters and data characters in the three characters following the /O/. A valid /O/ is any character with a value for O code in Table 126–1.

T_TYPE(tx_raw<71:0>)

Returns the T_BLOCK_TYPE of the tx_raw<71:0> bit vector.

T_TYPE_NEXT

Prescient end of packet check function. It returns the FRAME_TYPE of the tx_raw vector immediately following the current tx_raw vector.

126.3.6.2.5 Counters

lfer_cnt

Count up to a maximum of 16 of the number of invalid LDPC frames within the current lfer_timer period.

The following counters are required for PHYs that support the EEE capability:

tx_ldpc_frame_cnt

An integer value that counts transmit LDPC frame periods. The counter is reset when the first symbol of the first LDPC frame crosses the MDI on pair A in the transmit direction after normal training or fast retraining. It is incremented after the last symbol of each transmitted LDPC frame.

tx_ldpc_frame_cnt is reset to 0 when $\text{tx_ldpc_frame_cnt} = \text{lpi_qr_time} \times 4$.

rx_ldpc_frame_cnt

An integer value that counts receive LDPC frame periods. The counter is reset when the first symbol of the first LDPC frame crosses the MDI on pair A in the receive direction after normal

IEEE Std 802.3bz-2016
Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
for 2.5 Gb/s and 5 Gb/s Operation

training or fast retraining. It is incremented after the last symbol of each received LDPC frame. rx_ldpc_frame_cnt is reset to 0 when rx_ldpc_frame_cnt = lpi_qr_time × 4.

lpi_rxw_err_cnt

An integer value that counts the number of receive wake on error conditions. lpi_rxw_err_cnt is reset to zero during PCS_Test. The counter is reflected in register 3.22 (see 45.2.3.10).

126.3.6.3 State diagrams

The LFER Monitor state diagram shown in Figure 126–13 monitors the received signal for high LDPC frame error ratio.

The 64B/65B Transmit state diagram shown in Figure 126–14 controls the encoding of 65B transmitted blocks. It makes exactly one transition for each 65B transmit block processed. Though the Transmit state diagram sends Local Fault ordered sets when reset is asserted, the scrambler and 65B-LDPC may not be operational during reset. Thus, the Local Fault ordered sets may not appear on the PMA service interface.

The 64B/65B Receive state diagram shown in Figure 126–16 controls the decoding of 65B received blocks. It makes exactly one transition for each receive block processed except for the transition from RX_WE to RX_E, which occurs immediately after the RX_WE processes are complete.

The PCS shall perform the functions of LFER Monitor, Transmit, and Receive as specified in these state diagrams. The PCS shall not perform the LFER Monitor function during LPI receive operation from the time that the PCS 64B/65B Receiver enters the state RX_L, until the state RX_W is exited.

Transitions surrounded by dashed rectangles indicate requirements for 2.5GBASE-T and 5GBASE-T EEE-capable implementations.

126.3.7 PCS management

The following objects apply to PCS management. If an MDIO Interface is provided (see Clause 45), they are accessed via that interface. If not, it is recommended that an equivalent access be provided.

126.3.7.1 Status

pcs_status

Indicates whether the PCS is in a fully operational state. It is only true if block_lock is true and hi_lfer is false. This status is reflected in MDIO register 3.32.12. A latch low view of this status is reflected in MDIO register 3.1.2 and a latch high of the inverse of this status, Receive fault, is reflected in MDIO register 3.8.10.

block_lock

Indicates the state of the block_lock variable. This status is reflected in MDIO register 3.32.0. A latch low view of this status is reflected in MDIO register 3.33.15.

hi_lfer

Indicates the state of the hi_lfer variable. This status is reflected in MDIO register 3.32.1. A latch high view of this status is reflected in MDIO register 3.33.14.

Rx LPI indication

For EEE capability, this variable indicates the current state of the receive LPI function. This flag is set to TRUE (register bit set to one) when the PCS 64B/65B Receive state diagram (Figure 126–17) is in the RX_L or RX_W states. This status is reflected in MDIO register 3.1.8. A latch high view of this status is reflected in MDIO register 3.1.10 (Rx LPI received).

Tx LPI indication

For EEE capability, this variable indicates the current state of the transmit LPI function. This flag is set to TRUE (register bit set to one) when the PCS 64B/65B Transmit state diagram

(Figure 126–15) is in the TX_L or TX_W states. This status is reflected in MDIO register 3.1.9. A latch high view of this status is reflected in MDIO register 3.1.11 (Tx LPI received).

126.3.7.2 Counters

The following counters are reset to zero upon read and upon reset of the PCS. When they reach all ones, they stop counting. Their purpose is to help monitor the quality of the link.

lfer_count

6-bit counter that counts each time LFER_BAD_LF state is entered. This counter is reflected in MDIO register bits 3.33.13:8. The counter is reset when register 3.33 is read by management. Note that this counter counts a maximum of 16 counts per lfer_timer period since the LFER_BAD_LF can be entered a maximum of 16 times per lfer_timer window.

errored_block_count

8-bit counter. When the receiver is in normal mode, errored_block_count counts once for each time RX_E state is entered. This counter is reflected in MDIO register bits 3.33.7:0.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd7:2017

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation




Figure 126–13—LFER monitor state diagram






Figure 126–14—PCS 64B/65B Transmit state diagram, part a

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation



NOTE—This figure is mandatory for PHYs with the EEE capability.


Figure 126–15—PCS 64B/65B Transmit state diagram, part b



NOTE—Signals and functions shown with dashed lines are only required for the EEE capability.


Figure 126–16—PCS 64B/65B Receive state diagram, part a

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation



NOTE—This figure is mandatory for PHYs with the EEE capability.

Figure 126–17—PCS 64B/65B Receive state diagram, part b



NOTE—This figure is mandatory for PHYs with the EEE capability.

Figure 126–18—EEE transmit state diagram

126.3.7.3 Loopback


The PCS shall be placed in loopback mode when the loopback bit in MDIO register 3.0.14 is set to a one. In this mode, the PCS shall accept data on the transmit path from the XGMII and return it on the receive path to the XGMII. In addition, the PCS shall transmit a continuous stream of 65B-LDPC encoded 4D-PAM16 symbols to the PMA sublayer, and shall ignore all data presented to it by the PMA sublayer.

126.4 Physical Medium Attachment (PMA) sublayer

126.4.1 PMA functional specifications

The PMA couples messages from a PMA service interface specified in 126.2.2 to the 2.5GBASE-T and 5GBASE-T baseband media, specified in 126.7.

The interface between PMA and the baseband medium is the Medium Dependent Interface (MDI), which is specified in 126.8.



NOTE 1—The recovered_clock arc is shown to indicate delivery of the recovered clock signal back to PMA TRANSMIT for loop timing.

NOTE 2—pcs_data_mode is required only for the EEE or fast retrain capabilities alert_detect and rx_lpi_active are only required for the EEE capability fr_active is only required for the fast retrain capability.

Figure 126–19— PMA reference diagram

126.4.2 PMA functions

The PMA sublayer comprises one PMA Reset function and five simultaneous and asynchronous operating functions. The PMA operating functions are PHY Control, PMA Transmit, PMA Receive, Link Monitor, and Clock Recovery. All operating functions are started immediately after the successful completion of the PMA Reset function.

The PMA reference diagram, Figure 126–19, shows how the operating functions relate to the messages of the PMA Service interface and the signals of the MDI. Connections from the management interface, comprising the signals MDC and MDIO, to other layers are pervasive and are not shown in Figure 126–19.

126.4.2.1 PMA Reset function

The PMA Reset function shall be executed whenever one of the two following conditions occur:

- a) Power on (see 126.3.6.2.2)
- b) The receipt of a request for reset from the management entity

All state diagrams take the open-ended pma_reset branch upon execution of PMA Reset. The reference diagrams do not explicitly show the PMA Reset function.

126.4.2.2 PMA Transmit function

The PMA Transmit function comprises four synchronous transmitters to generate four pulse-amplitude modulated signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD. While send_fail is FALSE and ALERT is not indicated by tx_symb_vector, PMA Transmit shall continuously transmit onto the MDI pulses modulated by the symbols given by tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD], respectively after processing with the THP, optional transmit filtering, digital to analog conversion (DAC) and subsequent analog filtering. When ALERT is indicated by tx_symb_vector, the alert signal is transmitted as specified in 126.4.2.2.1. When send_fail is TRUE, the link failure signal is transmitted as specified in 126.4.2.2.2. The four transmitters shall be driven by the same transmit clock, TX_TCLK. The signals generated by PMA Transmit shall follow the mathematical description given in 126.4.3.1, and shall comply with the electrical specifications given in 126.5.

When the PMA_CONFIG.indication parameter config is MASTER, for both normal and LPI operation, the PMA Transmit function shall source TX_TCLK from a local clock source while meeting the transmit jitter requirements of 126.5.3.3. The MASTER/SLAVE relationship includes loop timing. If the PMA_CONFIG.indication parameter config is SLAVE, the PMA Transmit function shall source TX_TCLK from the recovered clock of 126.4.2.8 while meeting the jitter requirements of 126.5.3.3.

The PMA Transmit fault function is optional. The faults detected by this function are implementation specific. If the MDIO interface is implemented, then this function shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

EEE-capable PHYs shall generate the alert signal as defined in 126.4.2.2.1. PHYs that support the fast retrain capability shall generate the link failure signal as defined in 126.4.2.2.2. If ALERT is indicated by tx_symb_vector at the same time as send_fail is TRUE, then link failure signaling is transmitted.

126.4.2.2.1 Alert signal

PHYs that support the optional EEE capability transmit the following PAM2 sequence when the PMA_UNITDATA.request parameter is set to ALERT. The alert signal is sent for a total of 8 LDPC frame periods and begins on a LDPC 2-frame 256 4D-symbol boundary aligned to the inversion on pair A during

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

PMA training. The alert signal is transmitted without THP filtering. The alert signal is transmitted on pair A when the PHY operates as a MASTER. The alert signal is transmitted on pair C when the PHY operates as a SLAVE. All other pairs transmit quiet as described in 126.3.5.

When the PMA_CONFIG.indication(config) is MASTER the alert signal is composed of 7 repetitions of the following 128 symbol PAM2 sequence, followed by 128 zero symbols.

xpr_master =

9	9	-9	-9	-9	-9	-9	-9	9	9	-9	-9	9	9	9	9
9	9	9	9	-9	-9	9	9	9	9	-9	-9	9	9	-9	-9
-9	-9	-9	-9	-9	-9	9	9	-9	-9	-9	-9	-9	-9	9	9
-9	-9	-9	-9	-9	-9	-9	-9	9	9	-9	-9	9	9	-9	-9
-9	-9	9	9	9	9	9	9	9	9	9	9	-9	-9	-9	-9
9	9	-9	-9	-9	-9	9	9	9	9	-9	-9	9	9	-9	-9
-9	-9	-9	-9	-9	-9	-9	-9	9	9	9	9	-9	-9	9	9
9	9	-9	-9	9	9	-9	-9	9	9	9	9	-9	-9	-9	-9

When the PMA_CONFIG.indication(config) is SLAVE the alert signal is composed of 7 repetitions of the following 128 symbol PAM2 sequence, followed by 128 zero symbols.

xpr_slave =

-9	-9	-9	-9	9	9	9	9	-9	-9	9	9	-9	-9	9	9
9	9	-9	-9	9	9	9	9	-9	-9	-9	-9	-9	-9	-9	-9
-9	-9	9	9	-9	-9	9	9	9	9	-9	-9	-9	-9	9	9
-9	-9	-9	-9	9	9	9	9	9	9	9	9	9	9	-9	-9
-9	-9	9	9	-9	-9	9	9	-9	-9	-9	-9	-9	-9	-9	-9
9	9	-9	-9	-9	-9	-9	-9	9	9	-9	-9	-9	-9	-9	-9
-9	-9	9	9	-9	-9	9	9	9	9	-9	-9	9	9	9	9
9	9	9	9	-9	-9	9	9	-9	-9	-9	-9	-9	-9	9	9

The alert signal is followed by a wake signal composed of repeated IDLE characters encoded using the 64B/65B encoding technique. At the start of the wake signal all THP feedback delay lines are initialized with zeros.

126.4.2.2.2 Link failure signal

PHYs that support the fast retrain capability transmit the link failure signal under the control of the Fast Retrain state diagram. The link failure signal indicates to the link partner that a link failure has been detected and that the link partners should begin the fast retrain procedure.

The link failure signal is sent for 8 LDPC frames and begins on an even LDPC frame boundary aligned to the inversion on pair A during PMA training. The link failure signal is transmitted without THP filtering. The link failure signal is transmitted on pair A when the PHY operates as a MASTER. The link failure signal is transmitted on pair C when the PHY operates as a SLAVE. All other pairs transmit quiet as described in 126.3.5.

When the PMA_CONFIG.indication(config) is MASTER, the link failure signal is composed of 7 repetitions of the following 128 symbol PAM2 sequence, followed by 128 zero symbols.

$$xfr_master = xpr_master \times (-1)$$

When the PMA_CONFIG.indication(config) is SLAVE the link failure signal is composed of 7 repetitions of the following 128 symbol PAM2 sequence, followed by 128 zero symbols.

$$xfr_slave = xpr_slave \times (-1)$$

126.4.2.3 PMA transmit disable function**126.4.2.3.1 Global PMA transmit disable function**

The Global_PMA_transmit_disable function allows all of the transmitters to be disabled. It is used in either of the following cases:

- a) When a Global_PMA_transmit_disable variable is set to TRUE, this function shall turn off all of the transmitters so that the each transmitter Average Launch Power of the OFF Transmitter is less than -53 dBm.
- b) If a PMA_transmit_fault is detected, then the PMA may set the Global_PMA_transmit_disable to TRUE, turning off the transmitter on each pair.

126.4.2.3.2 PMA pair by pair transmit disable function

The PMA_transmit_disable function allows the transmitters on each pair to be selectively disabled.

When a PMA_transmit_disable_N variable is set to TRUE, this function shall turn off the transmitter associated with that variable so that the transmitter Average Launch Power of the OFF Transmitter is less than -53 dBm.

126.4.2.3.3 PMA MDIO function mapping

The MDIO capability described in Clause 45 defines several variables that provide control and status information for and about the PMA. Mapping of MDIO control variables to PMA control variables is shown in Table 126–6. Mapping of MDIO status variables to PMA status variables is shown in Table 126–7.

Table 126-6—MDIO/PMA control variable mapping

MDIO control variable	PMA register name	Register/bit number	PMA control variable
Reset	Control register 1	1.0.15	PMA_reset
Global transmit disable	Transmit disable register	1.9.0	Global_PMA_transmit_disable
Transmit disable pair D	Transmit disable register	1.9.4	PMA_transmit_disable_D
Transmit disable pair C	Transmit disable register	1.9.3	PMA_transmit_disable_C
Transmit disable pair B	Transmit disable register	1.9.2	PMA_transmit_disable_B
Transmit disable pair A	Transmit disable register	1.9.1	PMA_transmit_disable_A

Table 126-7—MDIO/PMA status variable mapping

MDIO status variable	PMA register name	Register/bit number	PMA status variable
Fault	Status register 1	1.1.7	PMA_fault
Transmit fault	Status register 2	1.8.11	PMA_transmit_fault
Receive fault	Status register 2	1.8.10	PMA_receive_fault

126.4.2.4 PMA Receive function

The PMA Receive function comprises four independent receivers for pulse-amplitude modulated signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD. The PMA Receive contains the circuits necessary to both detect symbol sequences from the signals received at the MDI over receive pairs BI_DA, BI_DB, BI_DC, and BI_DD and to present these sequences to the PCS Receive function. The signals received at the MDI are described mathematically in 126.4.3.2. The PMA translates the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD into the PMA_UNITDATA.indication parameter rx_symb_vector. The quality of these symbols shall allow an LFER of less than 3.2×10^{-9} after LDPC decoding, over a channel meeting the requirements of 126.7. The receiver shall correct for differential delay variations of up to 50 ns across the wire-pairs. The delay skew is removed by computing the relative received delay of the four known transmit patterns described in 126.3.4.

To achieve the indicated performance, it is highly recommended that PMA Receive include the functions of signal equalization, echo and crosstalk cancellation. The sequence of code-groups assigned to tx_symb_vector is needed to perform echo and self near-end crosstalk cancellation.

The PMA Receive function uses the scr_status parameter and the state of the equalization, cancellation, estimation, and LPI functions to determine the quality of the receiver performance, and generates the loc_rcvr_status variable accordingly. The precise algorithm for generation of loc_rcvr_status is implementation dependent.

The receiver uses the sequence of symbols during the training sequence to detect and correct for pair swaps and crossovers. The receiver pairs BI_DA, BI_DB, BI_DC, and BI_DD may be connected in any manner described in 126.4.4 to the corresponding transmit pairs. The receiver also detects and corrects for polarity mismatches on any pairs and corrects for differential delay variations across the wire-pairs.

The PMA Receive fault function is optional. The PMA Receive fault function is the logical OR of the link_status = FAIL and any implementation specific fault. If the MDIO interface is implemented, then this function shall contribute to the receive fault bit specified in 45.2.1.7.5.

PMA receive functions that support the optional EEE capability shall generate alert_detect when the alert signal is detected at the receiver. The PMA receive function asserts alert_detect after the entire alert signal (7 LDPC frame periods of the xpr_master or xpr_slave sequence and 1 frame of silence) has been detected. The alert signal is specified in 126.4.2.2.1. The criterion used to generate alert_detect is left to the implementer.

PHYs that support the fast retrain capability shall set link_fail_detect to TRUE when the link failure signal is reliably detected at the receiver. The PMA receive function asserts link_fail_detect after the entire link failure signal (7 LDPC frame periods of the xfr_master or xfr_slave sequence and 1 frame of silence) has been detected. The link failure signal is specified in 126.4.2.2.2. The criterion used to generate link_fail_detect is left to the implementer. It is highly recommended that the generation of link_fail_detect is qualified with repeated errored frames at the LDPC decoder output.

126.4.2.5 PHY Control function

PHY Control generates the control actions that are needed to bring the PHY into a mode of operation during which frames can be exchanged with the link partner. PHY Control shall comply with the state diagram description given in Figure 126–26.

During PMA training (includes PMA_Training_Init_M, PMA_Training_Init_S, PMA_PBO_Exch, PMA_Coeff_Exch, and PMA_Fine_Adjust states in Figure 126–26), PHY Control information is exchanged between link partners with a 16 octet Infocfield, which is XOR’ed with the last 128 bits of the PMA 16384 PAM2 frame on pair A (see Figure 126–11). The link partner is not required to decode every Infocfield transmitted but is required to decode Infocfields at a rate that enables the correct actions to timer expiration times, transition counter values, etc. described in Figure 126–26, Figure 126–27, and Figure 126–28.

The 16 octet Infocfield shall include the fields in 126.4.2.5.2 through 126.4.2.5.14, also shown in the overview Figure 126–20, and the more detailed Figure 126–21, Figure 126–22, and Figure 126–23.

Start of Frame Delimiter 0xBBA70000	3 Transmitter Settings	Message Field	SNR Margin	Message Field Dependent	Message Field Dependent	CRC16
4 octets	3 octets	1 octet	4 bits	1.5 octets	4 octets	2 octets

Figure 126–20—Infocfield format

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

Start of Frame Delimiter 0xBBA70000	3 Transmitter Settings	Message Field	SNR Margin	Reser- ved	Transition Counter	Reser- ved/ Ability	Vendor Specific	CRC16
4 octets	3 octets	1 octet	4 bits	2 bits	10 bits	2 octets	2 octets	2 octets

Figure 126–21—Infocfield transition counter format

Start of Frame Delimiter 0xBBA70000	3 Transmitter Settings	Message Field	SNR Margin	Coefficient Exchange	Coefficient Field	CRC16
4 octets	3 octets	1 octet	4 bits	1.5 octets	4 octets	2 octets

Figure 126–22—Infocfield coefficient exchange format

Start of Frame Delimiter 0xBBA70000	3 Transmitter Settings	Message Field	SNR Margin	Reserved	Reserved /Ability	Vendor Specific	CRC16
4 octets	3 octets	1 octet	4 bits	1.5 octets	2 octets	2 octets	2 octets

Figure 126–23—Infocfield not transition counter and not coefficient exchange format

126.4.2.5.1 Infocfield notation

For all the Infocfield notation below, Reserved<bit location> represents any unused values and shall be set to zero and ignored by the link partner. For all PBO Infocfield values below, the PBO<6:4> are unsigned 3-bit values 000, 001, 010, 011, 100, 101, 110, and 111 shall indicate power backoffs of 0 dB, 2 dB, 4 dB, 6 dB, 8 dB, 10 dB, 12 dB, and 14 dB respectively. The Infocfield is transmitted following the notation described in 126.3.2.2.3 where the LSB of each octet is sent first and the octets are sent in increasing number order (that is, the LSB of Octet 1 is sent first).

126.4.2.5.2 Start of Frame Delimiter

The start of Frame Delimiter consist of 4 octets [Octet 1<7:0>, Octet 2<7:0>, Octet 3<7:0>, Octet 4<7:0>] and shall use the hexadecimal value 0xBBA70000. 0xBB corresponds to Octet 1<7:0> and so forth.

126.4.2.5.3 Current transmitter settings

Current transmitter setting (1 octet). Represented by Octet 5{Valid<7>, PBO<6:4>, Reserved<3:0>} and shown in Figure 126–24. Used to announce the current fixed PBO setting during PMA_Training_Init_M, PMA_Training_Init_S and PMA_PBO_Exch, and the current programmable PBO setting during PMA_Coeff_Exch. For every other state this octet is set to zero and ignored by the link partner. The bit Valid shall be set to one if the corresponding octet information is valid and shall be set to zero if it the octet information is not valid. If Valid is set to zero, the octet is ignored by the link partner.

Single transmitter setting detail (one for current, next, or requested)



Figure 126–24—Infofield transmitter setting format

126.4.2.5.4 Next transmitter settings

Next transmitter setting (1 octet). Represented by Octet 6{Valid<7>, PBO<6:4>, Reserved<3:0>} and shown in Figure 126–24. Used to announce the next programmable PBO setting during PMA_PBO_Exch that takes effect upon entering PMA_Coeff_Exch state. For every other state, this octet is set to zero and ignored by the link partner. The bit Valid shall be set to one if the corresponding octet information is valid and shall be set to zero if it the octet information is not valid. If Valid is set to zero, the octet is ignored by the link partner.

126.4.2.5.5 Requested transmitter settings

Requested remote transmitter setting (1 octet). Represented by Octet 7{Valid<7>, PBO<6:4>, Reserved<3:0>} and shown in Figure 126–24. Used to request the remote transmitter programmable PBO setting during PMA_PBO_Exch that takes effect upon entering PMA_Coeff_Exch state. For every other state, this octet is set to zero and ignored by the link partner. The bit Valid shall be set to one if the corresponding octet information is valid and shall be set to zero if it the octet information is not valid. If Valid is set to zero, the octet is ignored by the link partner.

126.4.2.5.6 Message field

Message field (1 octet). For the MASTER, this field is represented by Octet 8{PMA_state<7:6>, loc_rcvr_status<5>, en_slave_tx<4>, trans_to_Coeff_Exch<3>, Coeff_exchange<2>, trans_to_Fine_Adjust<1>, trans_to_PCS_Test<0>}. For the SLAVE, this field is represented by Octet 8{PMA_state<7:6>, loc_rcvr_status<5>, timing_lock_OK<4>, trans_to_Coeff_Exch<3>, Coeff_exchange<2>, trans_to_Fine_Adjust<1>, trans_to_PCS_Test<0>}.

The two state-indicator bits PMA_state<7:6> shall indicate the state of the transmitting transceiver to the link partner as follows: PMA_state<7:6>=00 indicates PMA_Training_Init_M or PMA_Training_Init_S, PMA_state<7:6>=01 indicates PMA_PBO_Exch, PMA_state<7:6>=10 indicates PMA_Coeff_Exch, and PMA_state<7:6>=11 indicates PMA_Fine_Adjust.

All possible Message field settings are listed in Table 126–8 for the MASTER and Table 126–9 for the SLAVE. No other value shall be transmitted, and all other values shall be ignored at the receiver. The Message field setting for the first transmitted PMA frame shall be the first row of Table 126–8 for the MASTER and the first row of Table 126–9 for the SLAVE. Moreover, for a given Message field setting, the following Message field setting shall be the same Message field setting or the Message field setting corresponding to a row below the current setting. When loc_rcvr_status=OK the Infofield variable is set to loc_rcvr_status<5>=1 and set to 0 otherwise.

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

Table 126–8—Infofield message field valid MASTER settings

PMA_state<7:6>	loc_rcvr_status	en_slave_tx	trans_to_Coeff_Exch	Coeff_exchange	trans_to_Fine_Adjust	trans_to_PCS_Test
00	0	0	0	0	0	0
00	0	1	0	0	0	0
01	0	1	0	0	0	0
01	0	1	1	0	0	0
10	0	1	0	0	0	0
10	0	1	0	1	0	0
10	0	1	0	0	0	0
10	0	1	0	0	1	0
11	0/1	1	0	0	0	0
11	1	1	0	0	0	1

Table 126–9—Infofield message field valid SLAVE settings

PMA_state<7:6>	loc_rcvr_status	timing_lock_OK	trans_to_Coeff_Exch	Coeff_exchange	trans_to_Fine_Adjust	trans_to_PCS_Test
00	0	0	0	0	0	0
00	0	0/1	0	0	0	0
01	0	1	0	0	0	0
01	0	1	1	0	0	0
10	0	0/1	0	0	0	0
10	0	1	0	1	0	0
10	0	1	0	0	0	0
10	0	1	0	0	1	0
11	0	0/1	0	0	0	0
11	0/1	1	0	0	0	0
11	1	1	0	0	0	1

126.4.2.5.7 SNR_margin

SNR_margin (4 bits). Represented by Octet 9<7:4>, which reports received decision point SNR margin in 1/2 dB steps. SNR_margin is relative to the SNR required for reception of LDPC-coded PAM16 at an LDPC frame error ratio of less than 3.2×10^{-9} . The SNR_margin<7:4> 4-bit values, 0010, 0011, 0100, 0101, 0110, 0111, 1000, 1001, 1010, 1011, 1100, 1101, 1110 shall indicate the decision point SNR margin values of –1.5,

−1, −0.5, 0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5 dB respectively. The value 0001 shall indicate a margin of −2 dB or less, and the value 1111 shall indicate 5 dB or more. Finally, the value 0000 shall indicate that the SNR margin value is unknown.

126.4.2.5.8 Transition counter

Transition counter (10 bits). Represented by the 1.25 octets [Octet 9<1:0>, Octet 10<7:0>]. When configured as Transition counter (Coeff_exchange<2>=0 and a transition is announced to PMA_Coeff_Exch, PMA_Fine_Adjust or PCS_Test) this field is used as a 10-bit counter that counts the number of remaining frames until the next transition (PMA_Coeff_Exch, PMA_Fine_Adjust, PCS_Test).

126.4.2.5.9 Coefficient exchange handshake

Coefficient exchange handshake (12 bits). Represented by the 1.5 octets [Octet 9<3:0>, Octet 10<7:0>]. If Coeff_exchange<2>=1, this field is configured as a Coefficient exchange handshake and is used as a handshake control channel during programmable THP coefficient exchange. The details of the coefficient exchange are described in 126.4.2.5.15.

126.4.2.5.10 Ability fields

Ability field (1 octet). Represented by Octet 12{EEE Ability<7>, THP Bypass Request<6>, Fast Retrain<5>, Reserved<4:0>}. Used to advertise the abilities of the PHY during the PMA_PBO_Exch state when Message<7:6> = 01.

For every other state, this octet is set to zero and ignored by the link partner. The Ability bits are defined as follows:

Octet 12<4:0> = Reserved

Octet 12<5> = Fast Retrain

0 = Fast Retrain not supported

1 = Fast Retrain supported

Octet 12<6> = THP Bypass Request in PMA_Coeff_Exch state

0 = Local device requests link partner not to bypass THP during fast retrain

1 = Local device requests link partner to bypass THP during fast retrain

Octet 12<7> = EEE Ability

0 = EEE not supported

1 = EEE supported.

126.4.2.5.11 Reserved fields

All Infield fields denoted Reserved in Figure 126–21, Figure 126–22, and Figure 126–23 are reserved for future use. This includes Octet 11 and Octet 12 when Coeff_exchange<2>=0 and Message<7:6> is not equal to 01; Octet 9<3:2> when transition counter is announced; and [Octet 9<3:0>, Octet 10<7:0>] when no transition is announced and no coefficients are exchanged.

126.4.2.5.12 Vendor-specific field

If Coeff_exchange<2>=0, Octet 13 and Octet 14 are vendor-specific fields. If during Auto-Negotiation both transceivers agree on the use of the two vendor-specific octets, they may be used as a PHY communication

channel; otherwise they are set to zero and ignored by the link partner. They are represented by Octet 13<7:0> and Octet 14<7:0>.

126.4.2.5.13 Coefficient field

The Coefficient field (4 octets) is represented by Octet 11<7:0>, Octet 12<7:0>, Octet 13<7:0>, and Octet 14<7:0>. When Coeff_exchange<2>=1, this field is used to exchange programmable THP coefficients. It transmits four 8-bit THP coefficients out of the total of 64 (16 coefficients over each of the 4 pairs). The order is pair A, coefficients 0:3, followed by coefficients 4:7, followed by 8:11 and 12:15. For all cases the first coefficient (indices 0, 4, 8, and 12) is mapped to Octet 11, the second coefficient (indices 1, 5, 9, 13) is mapped to Octet 12 and so on. The same coefficient order is followed to transmit the coefficients for pair B, followed by pair C, and finally pair D. The details of the coefficient exchange are described in 126.4.2.5.15.

126.4.2.5.14 CRC16

CRC16 (2 octets). This field shall contain the CRC16 value calculated using the polynomial $(x+1)(x^{15}+x+1)$ of the previous 10 octets, Octet 5<7:0>, Octet 6<7:0>, Octet 7<7:0>, Octet 8<7:0>, Octet 9<7:0>, Octet 10<7:0>, Octet 11<7:0>, Octet 12<7:0>, Octet 13<7:0>, and Octet 14<7:0>. The CRC16 shall produce the same result as the implementation shown in Figure 126–25. In Figure 126–25 the 16 delay elements S0,..., S15, shall be initialized to zero. Afterwards Octet 5 through Octet 14 are used to compute the CRC16 with the switch set to CRCgen in Figure 126–25. After all the 10 octets have been processed, the switch is set to CRCout and the 16 values stored in the delay elements are transmitted in the order illustrated, first S15, followed by S14, and so on, until the final value S0.




Figure 126–25—CRC16

126.4.2.5.15 Startup sequence

The startup sequence shall comply with the state diagram description given in Figure 126–26 and the transition counter state diagrams Figure 126–27 and Figure 126–28.

During Auto-Negotiation, PHY Control is in the DISABLE_2.5G/5GBASE-T_TRANSMITTER state and the transmitters are disabled. During normal training, prior to enabling the transmitter, the THP coefficients are set to zero.

When the Auto-Negotiation process asserts link_control=ENABLE, PHY Control enters the INIT_MAXWAIT_TIMER state. Upon entering this state, the maxwait_timer is started and PHY Control enters the SILENT state, which starts the minwait_timer and forces transmission of zeros by setting tx_mode=SEND_Z.

In MASTER mode, after expiration of the `minwait_timer`, PHY Control transitions to the `PMA_Training_Init_M` state.

Upon entering the `PMA_Training_Init_M` and `PMA_Training_Init_S` states, the PHY Control forces transmission into the training mode by asserting `tx_mode=SEND_T`, which includes the transmission of Infobfields.

Upon entering state `PMA_Training_Init_M`, the MASTER starts transmission with a fixed transmit power level, `PBO=4` (corresponding to a power backoff of 8 dB). The PBO variable is communicated to the link partner via the current transmitter octet of the Infobfield.

Initially the MASTER is not ready for the SLAVE to respond and sets `en_slave_tx=0`, which is communicated to the link partner via the Infobfield. After the MASTER has sufficiently converged the necessary circuitry, the MASTER must set `en_slave_tx=1` to allow the SLAVE to transition to `PMA_Training_Init_S`.

In SLAVE mode, PHY Control transitions to the `PMA_Training_Init_S` state only after the SLAVE PHY acquires timing, converges its equalizers, acquires its descrambler state and sets `loc_SNR_margin=OK`. The SLAVE shall respond using the fixed PBO transmit power level, `PBO=4` (corresponding to a power backoff of 8 dB). For PHYs with the EEE capability, further requirements for this transition are described in 126.3.5.1.

While in states `PMA_Training_Init_S`, `PMA_PBO_Exch`, or `PMA_Coeff_Exch`, whenever a SLAVE operating in loop timing mode loses the MASTER timing reference (for example, after transmit power level transitions) it sets `timing_lock_OK=0`, which is communicated to the link partner via the Infobfield. Otherwise, `timing_lock_OK` is set to one.

In MASTER mode, PHY Control enters the `PMA_PBO_Exch` state after `loc_SNR_margin=OK` and in SLAVE mode PHY Control enters the `PMA_PBO_Exch` state after the `loc_SNR_margin=OK` and `minwait_timer` expires. In the `PMA_PBO_Exch` state while Infobfield Message<7:6> = 01, the PHY advertises EEE and Fast Retrain capability in octet 12 of the Infobfield. When both the local device and remote device advertise EEE capability then EEE is supported. When both the local device and remote device advertise Fast Retrain capability then Fast Retrain is supported. In the `PMA_PBO_Exch` state, after the MASTER has computed the final desired programmable PBO level, it shall request a PBO change using the requested transmitter setting in the Infobfield (octet 7). In SLAVE mode, after the MASTER has requested the desired PBO level, the SLAVE shall request a desired PBO level that is within two levels (within 4 dB) of the requested MASTER PBO level.

Following PBO exchange for both transceivers, each PHY shall announce the next PBO setting using the next transmitter setting (octet 6). Afterwards, each PHY announces a transition to the `PMA_Coeff_Exch` state using the `trans_to_Coeff_Exch=1` and `transition_count` as described in 126.4.5.1. MASTER initiates the transition to `PMA_Coeff_Exch` count with the `trans_to_Coeff_Exch=1` flag and a transition counter value of $S \times 2^8$. The SLAVE responds prior to the MASTER transition counter reaching $S \times 2^5$ by setting `trans_to_Coeff_Exch=1` flag and a transition counter value matching the MASTER. The PMA frame after each transceiver `transition_count` reaches zero, the PHYs shall enter the `PMA_Coeff_Exch` state and enable the requested PBO. Therefore, both PHYs enter the `PMA_Coeff_Exch` state within one PMA frame.

While both MASTER and SLAVE are in state `PMA_Coeff_Exch`, when either end has computed the programmable THP settings, the programmable THP coefficient exchange process can begin, using the 1.5-octet Coefficient exchange handshake and the 4-octet Coefficient field as follows:

- a) During `PMA_Coeff_Exch` each PHY begins a coefficient exchange by setting the `Coeff_Exchange` flag to 1 in the Message field.

- b) During coefficient exchange, the transition counter bits are used as the Coefficient Exchange Handshake
 - 1) Octet 9{Reserved<3:0>}: unused
 - 2) Coefficient Pair Received, Octet 10<7:6>: 01 for local transmitter pair A, 10 for B, 11 for C and 00 for D (default). This is the handshake to tell the remote unit the last coefficients received.
 - 3) Coefficient Group Received, Octet 10<5:4>: 01 for coefficients 0:3, 10 for 4:7, 11 for 8:11 and 00 for 12:15 (default). This is the handshake to tell the remote unit the last coefficients received.
 - 4) Coefficient Pair Sent, Octet 10<3:2>: 01 for remote transmitter pair A, 10 for B, 11 for C and 00 for D (default). This is the handshake to tell the remote unit the current coefficients being sent.
 - 5) Coefficient Group Sent, Octet 10<1:0>: 01 for 0:3, 10 for 4:7, 11 for 8:11 and 00 for 12:15 (default). This is the handshake to tell the remote unit the current coefficients being sent.
- c) The Coefficient field is used to send four 8-bit coefficients in each frame designated by the Coefficient Pair Sent and Coefficient Group Sent bits. The coefficient format is as follows:
 - 1) 8 bits per coefficient. Use one octet per coefficient in two's complement notation
 - 2) Coefficient range is -2.0 to 1.984375 in steps of 0.015625
 - 3) The sign of the coefficients shall be consistent with Equation (126-3)
- d) Each PHY begins the exchange by sending pair A coefficients 0:3 with Coefficient Pair Sent=01 and Coefficient Group Sent=01.
- e) The remote unit acknowledges by setting Coefficient Pair Received=01 and Coefficient Group Received=01.
- f) Following each acknowledgment, the PHY increments through the Coefficient Group and then Coefficient Pair settings until Coefficient Pair Sent=00 and Coefficient Group Sent=00 and Coefficient Pair Received=00 and Coefficient Group Received=00. At this time, coefficient exchange is done and both PHYs set Coeff_Exchange=0.

Following coefficient exchange for both transceivers, each PHY announces a transition to the PMA_Fine_Adjust state ($\text{trans_to_Fine_Adjust}=1$) and starts the transition_count as described in 126.4.5.1. During the first PMA frame after the transition_count reaches zero, the PHYs enter the PMA_Fine_Adjust state and enable the THP precoders with the requested coefficients. At the closure of the THP feedback loop, the initial state of the THP feedback filters shall be the last 16 symbols from the state PMA_Coeff_Exch.

The THP coefficients and PBO setting may not be changed during PMA_Fine_Adjust. The final convergence of the adaptive filter parameters is completed in the PMA_Fine_Adjust state.

After the PHY completes successful training and establishes proper receiver operations, PCS Transmit conveys this information to the link partner via transmission of the parameter Infocfield value loc_rcvr_status. The link partner's value for loc_rcvr_status is stored in the local device parameter rem_rcvr_status. When the condition $\text{loc_rcvr_status}=\text{OK}$ and $\text{rem_rcvr_status}=\text{OK}$ is satisfied, each PHY announces a transition to the PCS_Test state ($\text{trans_to_PCS_Test}=1$) and start the transition counter as described in 126.4.5.1. For PHYs with the EEE capability, further requirements for this transition are described in 126.3.5.1.

The normal mode of operation corresponds to the PCS_Data state, where PHY Control asserts $\text{tx_mode}=\text{SEND_N}$ and transmission of data over the link can take place.

PHY Control may force the transmit scrambler state to be initialized to an arbitrary value by requesting the execution of the PCS Reset function defined in 126.3.2.1.

The operation of the maxwait_timer requires that the PHY complete the startup sequence from state SILENT to PMA_Fine_Adjust in the PHY Control state diagram (Figure 126-26) in less than 2000 ms to

avoid link_status being changed to FAIL by the Link Monitor state diagram (Figure 126–29). To ensure interoperability the timing in Table 126–10 should be observed.

After reaching the PCS_Data state, PHYs with the EEE capability can transition to the LPI receive mode under the control of the link partner and to the LPI transmit mode under control of the local LPI client.

Table 126–10—Recommended startup sequence timing

Master	Recommended maximum time (ms)	Recommended average time (ms)	Slave
SILENT plus (PMA_Training_Init_M state AND en_slave_tx = 0)	350	315	SILENT
(PMA_Training_Init_M state AND en_slave_tx = 1) plus PMA_PBO_Exch state	480	432	PMA_Training_Init_S state plus PMA_PBO_Exch state
PMA_Coeff_Exch state	100	90	PMA_Coeff_Exch state with timing_lock_OK=0
	520	468	Total for PMA Coeff Exch state
PMA_Fine_Adjust state	650	585	PMA_Fine_Adjust state
Total	2000	1800	

126.4.2.5.16 Fast retrain function

PHYs that support the fast retrain capability shall conform to the fast retrain state diagram shown in Figure 126–31. PHYs may request a fast retrain by setting the variable loc_fr_req to TRUE. This causes the transmission of an easily-detected link failure signal specified in 126.4.2.2.2. After completing the link failure signal the PHY shall transition to the PMA_INIT_FR state followed immediately by the PMA_Coeff_Exch state. If the link partner requested THP bypass for fast retrain the PHY bypasses the THP (or set THP coefficients to zero). Otherwise the PHY shall keep its THP turned on with its previously exchanged coefficients, and send PAM2 signaling within a time period equivalent to 18 LDPC frame periods.

After the detection of the link failure signal, a PHY shall transition to the PMA_Coeff_Exch state and respond with PAM2 signaling within a time period equivalent to 18 LDPC frame periods after receiving the link failure signal.

The PAM2 symbols are generated using the PMA sidestream scrambler polynomials shown in Figure 126–11. The training sequence described in 126.3.4 shall be used during fast retraining, with the scramblers free-running from PCS Reset.

Note that reliable traffic on the transmitter may be interrupted when the local receiver requests a fast retrain.

Following the link failure signal, the two link partners transition back to the PMA_Coeff_Exch state and follow the training procedure described in 126.4.2.5.15, with the exception that the initial Infocfield countdown values are reduced as indicated in Figure 126–27 and Figure 126–28.

To ensure interoperability the training times in Table 126–11 should be observed during the fast retrain.

Table 126–11—Recommended fast retrain sequence timing

State	Recommended maximum time (ms)
PMA_Coeff_Exch state	20
PMA_Fine_Adjust state	10

126.4.2.6 Link Monitor function

Link Monitor determines the status of the underlying receive channel and communicates it via the variable link_status. Failure of the underlying receive channel typically causes the PMA’s clients to suspend normal operation.

The Link Monitor function shall comply with the state diagram of Figure 126–29.

Upon power on, reset, or release from power down, the Auto-Negotiation algorithm sets link_control=SCAN_FOR_CARRIER and, during this period, sends fast link pulses to signal its presence to a remote station. If the presence of a remote station is sensed through reception of fast link pulses, the Auto-Negotiation algorithm sets link_control=DISABLE and exchanges Auto-Negotiation information with the remote station. During this period, link_status=FAIL is asserted. If the presence of a remote 2.5GBASE-T or 5GBASE-T station is established, the Auto-Negotiation algorithm permits full operation by setting link_control=ENABLE. As soon as reliable transmission is achieved, the variable link_status=OK is asserted, upon which further PHY operations can take place.

126.4.2.7 Refresh Monitor function

The Refresh monitor is required for PHYs that support the EEE capability. The Refresh monitor operates when the PHY is in the LPI receive mode. The Refresh monitor shall comply with the state diagram of Figure 126–17. The function forces a link retrain if a refresh signal is not reliably detected within a moving time window equivalent to 50 complete quiet-refresh cycles (nominally equal to 16.384/S ms), when the PHY is in the lower power receive mode.

126.4.2.8 Clock Recovery function

The Clock Recovery function couples to all four receive pairs. It may provide independent clock phases for sampling the signals on each of the four pairs.

The Clock Recovery function shall provide clocks suitable for signal sampling on each line so that the LDPC FER indicated in 126.4.2.4 is achieved. The received clock signal should be stable and ready for use when training has been completed (loc_rcvr_status=OK). The received clock signal is supplied to the PMA Transmit function by received_clock.

126.4.3 MDI

Communication through the MDI is summarized in 126.4.3.1 and 126.4.3.2.

126.4.3.1 MDI signals transmitted by the PHY

The symbols to be transmitted by the PMA on the four pairs BI_DA, BI_DB, BI_DC, and BI_DD are denoted by tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and

tx_symb_vector[BI_DD], respectively. The modulation scheme used over each pair is PAM16. PMA Transmit generates a pulse-amplitude modulated signal on each pair in the following form:

$$x_n = M\left(a_n - \sum_{k=1}^{16} x_{n-k} c_k\right) = a_n + 32m_n - \sum_{k=1}^{16} x_{n-k} c_k \quad (126-3)$$

$$s(t) = \sum_{n=0}^{\infty} x_n h_T(t-nT) \quad (126-4)$$

In Equation (126-3), a_n is the PAM16 modulation symbol from the set $\{-15, -13, -11, -9, -7, -5, -3, -1, 1, 3, 5, 7, 9, 11, 13, 15\}$ to be transmitted at time nT . Each of the 16 THP coefficients c_1, c_2, \dots, c_{16} per wire pair is represented in two's complement form by 8 bits described in 126.4.2.5. The nonlinear THP operation given by $M(\alpha) = (\alpha + 16) \bmod 32 - 16$ corresponds to changing the modulation symbol a_n to an augmented modulation symbol $\tilde{a}_n = a_n + 32m_n$ with the integer m_n chosen such that the THP output lies in the interval $-16 \leq x_n < 16$. Equation (126-4) describes the convolution of the THP output signals with the transmitter symbol response $h_T(t)$ to obtain the transmit signal $s(t)$ at the MDI. The values of the programmable THP coefficients are exchanged in the Infocfield during PMA_Coeff_Exch. The THP filter coefficients shall be fixed after startup.

The nominal power (denoted P_{tx}) and the symbol response of the PMA transmitted signal $s(t)$, shall comply with the electrical specifications given in 126.5. When the link segment does not experience the maximum insertion loss (IL), each transceiver indicates to the link partner that the link partner PMA Transmit signal shall be reduced in increments of 2 dB. The minimum power backoff level requested shall comply with the power backoff schedule in Table 126-12. If a given receiver has sufficient decision point SNR margin, it may choose to request from the link partner larger power backoff (up to 14 dB) than shown in Table 126-12. Additionally, the Slave shall select a PBO level as described in the PMA_PBO_Exch state of 126.4.2.5.15. The PMA Transmit shall be capable of eight power backoff settings in approximately 2 dB steps. The difference between each consecutive power setting shall be 2 ± 0.25 dB, and each step shall be centered at $2 \times n$ dB ($n = 0$ to 7) reduction from nominal, with a maximum error of ± 1 dB.

The received signal power at the MDI, P_r (dBm), in Table 126-12, should be the estimate of the average received power across all four pairs from the remote transmitter when the link partner PMA Transmit is at nominal power (after accounting for local transmitter power). If the remote transmitter is not at nominal power during the measurement, the estimate of the received power should be incremented by the amount of power backoff of the link partner transmitter during the measurement. Nominal power refers to the transmit power without any power backoff and is specified in 126.5.3.4. The estimate of the received signal power is stored in registers 1.141 to 1.144 as described in 45.2.1. The values in the length, L (m), column in Table 126-12 are for reference only (not required for power backoff evaluation).

Table 126–12—Power backoff schedule

5GBASE-T		
Received signal power at MDI, <i>P</i> (dBm)	Length <i>L</i> (m) (reference)	Minimum power backoff (dB)
$-5.8 < P$	$0 \leq L < 35$	8
$-7.0 < P \leq -5.8$	$35 \leq L < 45$	6
$-9.2 < P \leq -7.0$	$45 \leq L < 65$	4
$-11 < P \leq -9.2$	$65 \leq L < 85$	2
$P \leq -11$	$85 \leq L$	0
2.5GBASE-T		
Received signal power at MDI, <i>P</i> (dBm)	Length <i>L</i> (m) (reference)	Minimum power backoff (dB)
$-4.3 < P$	$0 \leq L < 45$	2
$P \leq -4.3$	$45 \leq L$	0

126.4.3.2 Signals received at the MDI

Signals received at the MDI can be expressed for each pair as pulse-amplitude modulated signals that are corrupted by noise as follows:

$$r(t) = \sum_{n=0}^{\infty} \tilde{a}_n h_R(t-nT) + w(t) \tag{126-5}$$

In Equation (126–5), \tilde{a}_n are the augmented PAM16 modulation symbols described in 126.4.3.1, $h_R(t)$ denotes the symbol response of the overall channel from the THP precoder to the MDI at the receiver, and $w(t)$ represents the contribution of various noise sources including uncanceled crosstalk. The four signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD are processed within the PMA Receive function to yield the received symbols *rx_symb_vector*.

126.4.4 Automatic MDI/MDI-X configuration

Automatic MDI/MDI-X configuration is intended to eliminate the need for crossover cables between similar devices. Automatic MDI/MDI-X configuration is required for 2.5GBASE-T and 5GBASE-T devices and shall comply with 40.4.4.1 and 40.4.4.2.

Having established MDI/MDI-X configuration, the receiver shall detect and correct for several configurations of pair swaps and crossovers and arbitrary polarity swaps. The receiver pairs BI_DA, BI_DB, BI_DC, and BI_DD might be connected to the corresponding transmit pairs in any of the following ways with arbitrary polarity:

- a) No crossover
- b) A/B crossover only
- c) C/D crossover only
- d) A/B crossover and C/D crossover

For EEE-capable PHYs, the MDI/MDIX function configuration shall apply to refresh and alert signaling. For PHYs with the fast retrain capability, the MDI/MDIX function configuration shall apply to link failure signaling.

126.4.5 State variables

126.4.5.1 State diagram variables

coeff_exchange_done

This variable reports that both transceivers have received the corresponding coefficients from the link partner.

Values: TRUE: The coefficient exchange has completed.

FALSE: The coefficient exchange has not completed.

config

The PMA shall generate this variable continuously and pass it to the PCS via the PMA_CONFIG.indication primitive.

Values: MASTER or SLAVE.

link_control

The link_control parameter generated by Auto-Negotiation and passed to the PMA via the PMA_LINK.request primitive (see 126.2.1.1).

link_status

The link_status parameter set by PMA Link Monitor state diagram and communicated through the PMA_LINK.indication primitive.

Values: OK or FAIL.

loc_rcvr_status

Variable set by the PMA Receive function to indicate correct or incorrect operation of the receive link for the local PHY.

Values: OK: The receive link for the local PHY is operating reliably.

NOT_OK: Operation of the receive link for the local PHY is unreliable.

loc_SNR_margin

This variable reports whether the local device has sufficient SNR margin to continue to the next state. The criterion for setting the parameter loc_SNR_margin is left to the implementer.

Values: OK: The local device has sufficient SNR margin.

NOT_OK: The local device does not have sufficient SNR margin.

master_transition_counter

This variable reports the current value of the MASTER's transition counter reported in the Infofield defined in 126.4.2.5.

Values: 0 to 2⁹.

MessageField_IF

This variable reports that a receiver has successfully received and decoded the Infofield from the remote device. This variable takes on the value contained in the Message field. If the Message field cannot be decoded or no explicit action is outstanding the value Null is returned.

Values: trans_to_Coeff_Exch, trans_to_Fine_Adjust, trans_to_PCS_Test or Null.

PBO

PBO is a variable that can take any integer value from 0 to 7 and indicates the power backoff level. Denoting P_{tx} as the maximum nominal power, the PBO values are as follows:

Values: 0, 1, 2, 3, 4, 5, 6, 7, which correspond to transmit power levels of

P_{tx} , $P_{tx}-2$ dB, $P_{tx}-4$ dB, $P_{tx}-6$ dB, $P_{tx}-8$ dB, $P_{tx}-10$ dB, $P_{tx}-12$ dB, $P_{tx}-14$ dB respectively.

PBO_next

PBO_next is a variable that can take any integer value from 0 to 7 and indicates the next power backoff level to be used at the local transmitter. The value is taken from the fixed set of values during PMA_Training_Init_M and PMA_Training_Init_S as described in 126.4.2.5. The value is

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation

taken from the decoded value of the link partner Infocfield during PMA_PBO_Exch
 Values:0, 1, 2, 3, 4, 5, 6, 7, which correspond to transmit power levels of
 P_{tx} , $P_{tx}-2$ dB, $P_{tx}-4$ dB, $P_{tx}-6$ dB, $P_{tx}-8$ dB, $P_{tx}-10$ dB, $P_{tx}-12$ dB, $P_{tx}-14$ dB
 respectively.

PBO_tx

PBO_tx is a variable that can take any integer value from 0 to 7 and indicates the power backoff
 level currently used at the local transmitter.

Values:0, 1, 2, 3, 4, 5, 6, 7, which correspond to transmit power levels of
 P_{tx} , $P_{tx}-2$ dB, $P_{tx}-4$ dB, $P_{tx}-6$ dB, $P_{tx}-8$ dB, $P_{tx}-10$ dB, $P_{tx}-12$ dB, $P_{tx}-14$ dB
 respectively.

PBO_exchange_done

This variable reports that both transceivers have received the corresponding PBO levels from the
 link partner.

Values:TRUE: The PBO exchange has completed.
 FALSE: The PBO exchange has not completed.

pcs_status

The pcs_status parameter generated by the PCS and passed to the PMA via the PMA_PCSSTA-
 TUS.request primitive (see 126.2.2.6).

pma_reset

Allows reset of the PHY Control and Link Monitor state diagrams.
 Values:ON or OFF.

rem_rcvr_status

Variable set by the PCS Receive function to indicate whether correct operation of the receive link
 for the remote PHY is detected or not.

Values:OK: The receive link for the remote PHY is operating reliably.
 NOT_OK: Reliable operation of the receive link for the remote PHY is not detected.

THP_next

THP_next is a variable that contains sixteen 8-bit values and describes the next transmitter setting
 of the THP coefficients. It refers to the programmable THP coefficients selected during coefficient
 exchange described in 126.4.2.5.

Values:16 coefficients of 8-bit values each. Range is -2.0 to 1.984375 in steps of 0.015625 .

THP_tx

THP_tx is a variable that contains sixteen 8-bit values and describes the current transmitter setting
 of the THP coefficients. It refers to the programmable THP coefficients selected during the coeffi-
 cient exchange described in 126.4.2.5.

Values:16 coefficients of 8-bit values each. Range is -2.0 to 1.984375 in steps of 0.015625 .

trans_to_Coeff_Exch

Message field variable defined in 126.4.2.5 that flags a transition by the local device to the
 PMA_Coeff_Exch state.

Values:1: The local device transitions to the PMA_Coeff_Exch state on expiration of the transition
 counter.
 0: The local device does not transition to the PMA_Coeff_Exch state.

trans_to_Fine_Adjust

Message field variable defined in 126.4.2.5 that flags a transition by the local device to the
 PMA_Fine_Adjust state.

Values:1: The local device transitions to the PMA_Fine_Adjust state on expiration of the transition
 counter.
 0: The local device does not transition to the PMA_Fine_Adjust state.

trans_to_PCS_Test

Message field variable defined in 126.4.2.5 that flags a transition by the local device to the
 PCS_Test state.

Values:1: The local device transitions to the PCS_Test state on expiration of the transition counter.
 0: The local device does not transition to the PCS_Test state.

transition_count

This variable reports the value of the transition counter contained in the Infocfield sent to the remote device. Transition_count must comply with the state diagram description given in 126.4.6.2. When the Message field contains a flag for a state transition, the transition counter denotes the remaining number of Infocfield until the next state transition. MASTER initiates the transition to PMA_Coeff_Exch count with the trans_to_Coeff_Exch=1 flag and a counter value of $S \times 2^8$. The SLAVE responds prior to the counter reaching $S \times 2^5$ with the same flag and a count value matching the MASTER. Then both PHY's transition to PMA_Coeff_Exch within one PMA frame. The same sequence is performed in the transition to PMA_Fine_Adjust state and PCS_Test state using the trans_to_Fine_Adjust=1 and trans_to_PCS_Test=1 flags respectively. In EEE-capable PHYs, synchronization of the PMA frames is tightly controlled as described in 126.3.5.1. When the Message field does not contain a flag for a state transition, the transition counter is set to zero and ignored by the receiver.

Values: 0 to 2^9 .

tx_mode

PCS Transmit sends code-groups according to the value assumed by this variable.

Values: SEND_N: This value is continuously asserted when transmission of sequences of code-groups representing a XGMII data stream take place.

SEND_T: This value is continuously asserted when transmission of sequences of code-groups representing the sequences of code-groups (TA_n , TB_n , TC_n , TD_n) defined in 126.3.4.2 is to take place.

SEND_Z: This value is asserted when transmission of zero code-groups is to take place.

The following variables are required only for PHYs that support the EEE capability:

lpi_refresh_detect

Set TRUE when the receiver has reliably detected refresh signaling and FALSE otherwise. The exact criteria left to the implementer.

pcs_data_mode

Generated by the PMA PHY Control function and indicates whether or not the local PHY may transition its PCS state diagrams out of their initialization states. The current value of the pcs_data_mode is passed to the PCS via the PMA_PCSDATAMODE.indicate primitive. In the absence of the optional EEE and fast retrain capabilities, the PHY operates as if the value of this variable is TRUE.

mtc

mtc is the transition count for a MASTER PHY during normal training and fast retraining. mtc shall be equal to $S \times 2^8$ for normal training and $S \times 2^5$ for fast retrain.

stc

stc is the transition count for a SLAVE PHY during normal training and fast retraining. stc shall be equal to $S \times 2^5$ for normal training and $S \times 2^4$ for fast retrain.

The following six variables are required only for PHYs that support the fast retrain capability:

fr_enable

This variable is set to TRUE if fast retrain is supported. The variable is set to FALSE otherwise. If MDIO is supported, this variable is based on the value of 1.147.0 with the value of TRUE corresponding to 1.147.0 set to 1. If MDIO is not supported, an equivalent method of controlling fast retrain functionality should be provided.

loc_fr_req

Set TRUE when the receiver has detected a link failure condition and is requesting a fast retrain; set FALSE otherwise.

loc_fr_detect

Set TRUE when the receiver has reliably detected the link failure signal. It is highly recommended that loc_fr_detect is qualified with the reception of errored blocks at the LDPC decoder output. Set FALSE when the link failure signal is not detected.

send_link_fail

When TRUE indicates that the PMA should send the link failure signal. When FALSE the variable has no effect.

fr_active

Set TRUE when the PHY is performing a fast retrain and set FALSE otherwise.

fast_retrain_flag

Set TRUE after the PHY generates or detects a link failure signal and set FALSE otherwise.

126.4.5.2 Timers

All timers operate in the manner described in 14.2.3.2.

maxwait_timer

A timer used to limit the amount of time during which a receiver dwells in the SILENT and TRAINING states. The timer shall expire 2000 ms \pm 10 ms after being started. This timer is used jointly in the PHY Control and Link Monitor state diagrams. The maxwait_timer is tested by the Link Monitor to force link_status to be set to FAIL if the timer expires and loc_rcvr_status is NOT_OK. See Figure 126–26 and Figure 126–29.

minwait_timer

A timer used to determine the minimum amount of time the PHY Control stays in the SILENT, PMA_Training_Init_S, PCS_Test and PCS_Data states. The timer shall expire 1 ms \pm 0.1 ms after being started.

The following timer is required only for PHYs that support the EEE capability:

lpi_refresh_rx_timer

This timer is used to monitor link quality during the LPI receive mode. If the PHY does not reliably detect reliable refresh signaling before this timer expires then a full retrain is performed.

Values: The condition lpi_refresh_rx_timer_done becomes true upon timer expiration.

Duration: This timer shall have a period equal to 50 complete quiet-refresh signal periods, equivalent to 8.192/S ms.

The following two timers are required only for PHYs that support the fast retrain capability:

link_fail_sig_timer

Determines the period of time the PHY sends the link failure signal.

Values: The condition link_fail_sig_timer_done becomes true upon timer expiration.

Duration: This timer shall have a period equal to 8 LDPC frame periods.

fr_maxwait_timer

Determines the period of time the PHY has to transition its PCS Control State to PCS_Test following a fast retrain before the fast retrain is aborted and a full retrain performed.

Values: The condition fr_maxwait_timer_done becomes true upon timer expiration.

Duration: This timer shall have a period equal to 30 ms.

126.4.5.3 Functions

Exchange_Final_PBO

This function transmits and receives the final PBO settings using the Infofield as described in 126.4.2.5.

Exchange_THP_Coefficients

This function compiles and sends to the link partner and receives from the link partner the desired programmable THP coefficients using the Infofield as described in 126.4.2.5.

126.4.5.4 Counters

The following two counters are required only for PHYs that support the fast retrain capability:

fr_tx_counter

Counts the number of times the PHY initiates a fast link retrain by transmitting the link failure signal. This counter is reflected in MDIO register 1.147.10:6 specified in 45.2.1.79.2.


fr_rx_counter

Counts the number of times the PHY begins a fast link retrain in response to the detection of link failure signaling from the link partner. This counter is reflected in MDIO register 1.147.15:11 specified in 45.2.1.79.1.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd7:2017

126.4.6 State diagrams


126.4.6.1 PHY Control state diagram



NOTE—For PHYs that do not support the fast retrain capability, the variable fast_retrain_flag is set to FALSE.

Figure 126–26—PHY Control state diagram

126.4.6.2 Transition counter state diagrams




NOTE—For PHYs that do not support the fast retrain capability, the variable fast_retrain_flag is set to FALSE.

Figure 126–27—MASTER transition counter state diagram

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation


MessageField_IF = trans_to_Coeff_Exch *
 master_transition_counter > $S \times 2^5$



NOTE—For PHYs that do not support the fast retrain capability, the variable fast_retrain_flag is set to FALSE.

Figure 126–28—SLAVE transition counter state diagram


126.4.6.3 Link Monitor state diagram



NOTE 1—maxwait_timer is started in PHY Control state diagram (see Figure 126-26).
 NOTE 2—The variables link_control and link_status are designated as link_control_2p5GigT and link_status_2p5GigT, respectively for 2.5GBASE-T, and link_control_5GigT and link_status_5GigT, for 5GBASE-T (by the Auto-Negotiation Arbitration state diagram (Figure 28-16)).

Figure 126-29—Link Monitor state diagram

126.4.6.4 EEE Refresh monitor state diagram




NOTE—This state diagram is only required when the PHY supports the EEE capability.

Figure 126–30—EEE Refresh monitor state diagram

STANDARDSISO.COM :: Click to view the full PDF of ISO/IEC/IEEE 8802-3:2017/Amd7:2017

126.4.6.5 Fast retrain state diagram



NOTE—This state diagram is only required when the PHY supports the fast retrain capability.

Figure 126–31—Fast retrain control state diagram

126.5 PMA electrical specifications

This subclause defines the electrical characteristics of the PMA and specifies PMA-to-MDI interface tests.

126.5.1 Isolation requirement

The PHY shall provide electrical isolation between the port device circuits, including frame ground (if any) and all MDI leads. This electrical isolation shall withstand at least one of the following electrical strength tests:

- a) 1500 V rms at 50 Hz to 60 Hz for 60 s, applied as specified in Section 5.2.2 of IEC 60950-1:2001.
- b) 2250 V dc for 60 s, applied as specified in Section 5.2.2 of IEC 60950-1:2001.
- c) A sequence of ten 2400 V impulses of alternating polarity, applied at intervals of not less than 1 s. The shape of the impulses is 1.2/50 μs (1.2 μs virtual front time, 50 μs virtual time or half value), as defined in Annex N of IEC 60950-1:2001.

There shall be no insulation breakdown, as defined in Section 5.2.2 of IEC 60950-1:2001, during the test. The resistance after the test shall be at least 2 MΩ, measured at 500 V dc.

126.5.2 Test modes

The test modes described next shall be provided to allow for testing of the transmitter waveform, transmitter distortion, transmitted jitter, transmitter droop, and BER testing.

For a PHY with an MDIO management interface, these modes shall be enabled by setting bits 1.132.15:13 (MultiGBASE-T test mode register) of the MDIO Management register set as follows in Table 126–13. These test modes shall only change the data symbols provided to the transmitter circuitry and shall not alter the electrical and jitter characteristics of the transmitter and receiver from those of normal (non-test mode) operation. PHYs without a MDIO shall provide a means to enable these modes for conformance testing.

Table 126–13—MDIO management register settings for test modes

1.132.15	1.132.14	1.132.13	Mode
0	0	0	Normal operation.
0	0	1	Test mode 1—Setting of MASTER transmitter required by SLAVE for transmit jitter test in SLAVE mode.
0	1	0	Test mode 2—Transmit jitter test in MASTER mode.
0	1	1	Test mode 3—Transmit jitter test in SLAVE mode.
1	0	0	Test mode 4—Transmit distortion test.
1	0	1	Test mode 5—Normal operation with no power backoff. This is for the PSD mask and power level test.
1	1	0	Test mode 6—Transmitter droop test mode.
1	1	1	Test mode 7—Pseudo-random test mode for BER Monitor.

Test mode 1 is a mode provided for enabling testing of timing jitter on a SLAVE transmitter. When test mode 1 is enabled, the PHY shall transmit the PMA training pattern (PRBS 33) continually from all four transmitters with the THP turned off, with no power backoff and with the transmitted symbols timed from its local clock source.

Test mode 2 is for transmitter jitter testing when transmitter is in MASTER timing mode. When test mode 2 is enabled, the PHY shall transmit {two +16 symbols followed by two –16 symbols} continually from all four transmitters with the THP turned off, with no power backoff and with the transmitted symbols timed from its local clock source. The transmitter output is a $S \times 100$ MHz signal.

When test mode 3 is enabled on a PHY, the PHY shall transmit, with THP turned off, the data symbol sequence {two +16 symbols followed by two –16 symbols} repeatedly on pair D with the symbols timed from its recovered receive data clock in SLAVE timing mode. A PHY operates in test mode 3 when there is no input signal on pair D. The transmitter output is a $S \times 100$ MHz signal on pair D and shall be silence on pairs A, B, and C.

Test mode 4 is for transmitter nonlinear distortion testing. When test mode 4 is enabled, the PHY shall transmit, with the THP turned off, transmitted symbols, timed from a transmit clock (as specified in 126.3.3.5) in the MASTER timing mode, defined by the bits 1.132.12:10 and Table 126–14.

Table 126–14—MDIO management register settings for transmit frequencies in Test mode 4

1.132.12	1.132.11	1.132.10	Output waveform frequencies in MHz
			Two tone frequency pairs
0	0	0	Reserved
0	1	1	Reserved
1	1	1	Reserved
0	0	1	$S \times (400/1024) \times 47, S \times (400/1024) \times 53$
0	1	0	$S \times (400/1024) \times 101, S \times (400/1024) \times 103$
1	0	0	$S \times (400/1024) \times 179, S \times (400/1024) \times 181$
1	0	1	$S \times (400/1024) \times 277, S \times (400/1024) \times 281$
1	1	0	$S \times (400/1024) \times 397, S \times (400/1024) \times 401$

The peak-to-peak levels used in this test shall correspond to the ± 16 symbol levels and the relative amplitudes of the tones in a two-tone pair shall be within 0.5 dB of each other.

Test mode 5 is for checking whether the transmitter is compliant with the transmit PSD mask and the transmit power level. When test mode 5 is enabled, the PHY shall transmit as in normal operation but with the power backoff disabled.

Test mode 6 is for testing transmitter droop. When test mode 6 is enabled, the PHY shall transmit the following sequence of data symbols A_n, B_n, C_n, D_n , of 126.4.3.1 continually from all four transmitters, with the THP turned off:

{One hundred twenty eight +16 followed by one hundred twenty eight –16 symbols}.

Test mode 7 is for enabling measurement of the bit error ratio of the link including the LDPC encoder/decoder, the transmit and receive analog front ends of the PHY and a cable connecting two PHYs. This mode shall use the 2.5GBASE-T and 5GBASE-T scrambler and is defined in detail in 126.3.3.

126.5.2.1 Test fixtures

The following fixtures (illustrated by Figure 126–32, Figure 126–33, and Figure 126–34), or their functional equivalents, can be used for measuring the transmitter specifications described in 126.5.3.




Figure 126–32—Transmitter test fixture 1 for transmitter droop measurement




Figure 126–33—Transmitter test fixture 2 for linearity measurement, power spectral density measurement and transmit power level measurement

The high impedance probe shown in Figure 126–32 in transmitter test fixture 1 has resistance $> 10\text{ k}\Omega$ and capacitance $< 1\text{ pF}$ over the frequency range of 1 MHz to $S \times 200\text{ MHz}$. Figure 126–33 includes a power summer or balun device to couple the $100\ \Omega$ differential output of the transmitter to the $50\ \Omega$ single-ended input typically found in a spectrum analyzer input. The center frequency (f_c) of the band pass filter shown in Figure 126–34 is $S \times 100\text{ MHz} \pm 200\text{ kHz}$ and the band pass filter noise bandwidth (B_n) is $2\text{ MHz} \pm 200\text{ kHz}$. The center frequency (f_c) of the band pass filter shown in Figure 126–35 is $45\text{ MHz} \pm 200\text{ kHz}$ and the band pass filter noise bandwidth (B_n) is $2\text{ MHz} \pm 200\text{ kHz}$.




Figure 126–34—Transmitter test fixture 3 for transmitter jitter measurement




Figure 126–35—Transmitter test fixture 4 for linearity measurement of 2.5GBASE-T with sine wave injected

126.5.3 Transmitter electrical specifications

The PMA provides the Transmit function specified in 126.4.2.2 in accordance with the electrical specifications of this clause. The PMA shall operate with AC-coupling to the MDI.

Where a load is not specified, the transmitter shall meet the requirements of this clause with a 100 Ω resistive differential load connected to each transmitter output.

126.5.3.1 Maximum output droop

With the transmitter in test mode 6 and using the transmitter test fixture 1, the magnitude of both the positive and negative droop shall be less than $(7.5 + 5/S)\%$, measured with respect to an initial value at 10 ns after the zero crossing and a final value at $(10 + 160/S)$ ns after the zero crossing.

126.5.3.2 Transmitter nonlinear distortion

When in test mode 4 and observing the spectrum of the differential signal output at the MDI using transmitter test fixture 2, for each pair, with no intervening cable, the transmitter nonlinear distortion mask is defined as follows:

The SFDR of the transmitter, with dual tone inputs as specified in test mode 4, shall meet the requirement shown in Equation (126-6).

$$\text{SFDR} \geq 2.5 + \min\{52, 58 - 20 \times \log_{10}(f/25)\} \quad (126-6)$$

where

f is the maximum frequency of the two test tones in MHz
SFDR is the ratio in dB of the minimum RMS value of either input tone to the RMS value of the worst intermodulation product in the frequency range of 1 MHz to $S \times 200$ MHz

Additionally, for 2.5GBASE-T, when in test mode 4, at 0 dB PBO, and observing the spectrum of the differential signal output at the MDI using transmitter test fixture 4, for each pair, while injecting a 45 MHz sine wave from the signal generator so that it has an amplitude 7 dB below the peak of the transmitter at the MDI, with no intervening cable, the transmitter nonlinear distortion mask is defined as follows: The SFDR of the transmitter, with dual tone inputs as specified in test mode 4, shall meet the requirement shown in Equation (126-7).

$$\text{SFDR} \geq -5.5 + \min\{52, 58 - 20 \times \log_{10}(f/25)\} \quad (126-7)$$

where

f is the maximum frequency of the two test tones in MHz
SFDR is the ratio in dB of the minimum RMS value of either input tone to the RMS value of the worst intermodulation product in the frequency range of 1 MHz to 100 MHz

This specification on transmit linearity is derived from the requirement for interoperability with the far-end device.

126.5.3.3 Transmitter timing jitter

When in test mode 2, the PHY transmits {two +16 symbols followed by two -16 symbols} continually with the THP turned off and with no power backoff. In this mode, the transmitter output should be a $S \times 100$ MHz signal and the RMS period jitter measured at the PHY MDI output shall be less than 7.2 ps for 5GBASE-T and 10.0 ps for 2.5GBASE-T. The RMS period jitter is measured as per the test configuration shown in Figure 126-34 over an integration time interval of $2/S$ ms \pm 10%.

The SLAVE mode RMS period jitter test is measured using the test configuration shown in Figure 126-34. For this test, the MASTER PHY is in test mode 1 and the SLAVE PHY is in test mode 3. The MASTER is transmitting the PMA training pattern (PRBS 33) to the SLAVE PHY on pairs A, B, and C. The SLAVE PHY is in loop timing mode, synchronizing its transmit clock to the signals received from the MASTER PHY on pairs A, B, and C. In this configuration, the transmitter output on pair D should be a $S \times 100$ MHz signal and the RMS period jitter measured at the SLAVE PHY MDI output shall be less than 7.2 ps for 5GBASE-T and 10.0 ps for 2.5GBASE-T. The RMS period jitter is measured over an integration time interval of $2/S$ ms \pm 10%.

RMS period jitter over an integration time interval of $2/S$ ms \pm 10% is defined as the root mean square period difference from the average period ($T - T_{avg}$), accumulated over a sample size of $200\,000 \pm 20\,000$, as shown in Equation (126-8).

$$\text{RMS period jitter} = \sqrt{\frac{\sum [(T - T_{avg})^2]}{\text{Sample size}}} \quad (126-8)$$

126.5.3.4 Transmitter power spectral density (PSD) and power level

In test mode 5 (normal operation with no power backoff), the transmit power shall be in the range 1.0 dBm to 3.0 dBm and the power spectral density of the transmitter, measured into a $100\ \Omega$, load using the test fixture shown in Figure 126-33 shall be between the upper and lower masks specified in Equation (126-9) and Equation (126-10). In the highest frequency segment, the PSD mask is the maximum of the PSD specified for 2.5G/5GBASE-T, or 6 dB less than that specified in Clause 55 by Equation 55-9. The masks are shown in Figure 126-36.

$$\text{PSD1}(f) \leq \begin{cases} -77.7 - 10 \times \log_{10}(S) & \text{dBm/Hz} & 0 < 2\frac{f}{S} \leq 70 \\ -77.7 - 10 \times \log_{10} S - \frac{(2\frac{f}{S} - 70)}{80} & \text{dBm/Hz} & 70 < 2\frac{f}{S} \leq 150 \\ -78.7 - 10 \times \log_{10} S - \frac{(2\frac{f}{S} - 150)}{58} & \text{dBm/Hz} & 150 < 2\frac{f}{S} \leq 730 \\ -78.7 - 10 \times \log_{10} S - \frac{(2\frac{f}{S} - 330)}{40} & \text{dBm/Hz} & 730 < 2\frac{f}{S} \leq 1822 - 400 \times \log_{10}(S) \\ -116 & \text{dBm/Hz} & S \times (911 - 200 \times \log_{10}(S)) < f \leq 3000 \end{cases}$$

$$\text{UpperPSD}(f) \leq \max(\text{PSD1}(f), \text{Equation 55-9}) - 6 \text{ dB}$$

(126-9)

and

$$\text{Lower PSD}(f) \geq \begin{cases} -82.2 - 10 \times \log_{10}(S) & \text{dBm/Hz} & 5 < 2\frac{f}{S} \leq 50 \\ -82.2 - 10 \times \log_{10} S - \frac{(2\frac{f}{S} - 50)}{50} & \text{dBm/Hz} & 50 < 2\frac{f}{S} \leq 200 \\ -85.2 - 10 \times \log_{10} S - \frac{(2\frac{f}{S} - 200)}{25} & \text{dBm/Hz} & 200 < 2\frac{f}{S} \leq 400 \end{cases} \quad (126-10)$$

where

f is in MHz

IEEE Std 802.3bz-2016
 Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation




Figure 126-36—Transmitter power spectral density mask

126.5.3.5 Transmit clock frequency

The symbol transmission rate on each pair of the MASTER PHY shall be within the range $S \times 400$ MHz \pm 50 ppm.

For a MASTER PHY, when the transmitter is in the LPI transmit mode or when the receiver is in the LPI receive mode the transmitter clock short-term rate of frequency variation shall be less than 0.1 ppm/second. The short-term frequency variation limit shall also apply when switching to and from the LPI mode.

126.5.4 Receiver electrical specifications

The PMA provides the Receive function specified in 126.4.2.4 in accordance with the electrical specifications of this clause using cabling that is within the limits specified in 126.7.

126.5.4.1 Receiver differential input signals

Differential signals received at the MDI that were transmitted from a remote transmitter within the specifications of 126.5.3 and have passed through a link specified in 126.7 shall be received with a BER less than 10^{-12} after LDPC decoding, and sent to the XGMII after link reset completion. This specification can be verified by a frame error ratio less than 7.8×10^{-9} for 800 octet frames with minimum IPG or greater than 220 octet IPG.

126.5.4.2 Receiver frequency tolerance

The receive feature shall properly receive incoming data, per the requirements of 126.5.4.1, with a symbol rate within the range $S \times 400$ MHz \pm 50 ppm.

126.5.4.3 Rejection of External EM Fields

When the cabling system is subjected to electromagnetic fields, currents are generated that may be converted to interference. This specification is provided to limit the sensitivity of the PMA receiver to external EM fields picked up by the cabling and interconnect system. It provides an assessment method of the electromagnetic performance of the link segment and the PHY, including the MDI.

An 80 MHz to 1000 MHz test can be made based on the cable clamp test described in Annex 113A, a 30 meter plug-terminated link segment that meets the requirements of 126.7, and suitable broadband ferrites. All components that are exposed to the induced fields should remain over the ground reference plane. A sine wave with the amplitude held constant over the whole frequency range from 80 MHz to 1000 MHz, with the amplitude calibrated so that the signal power measured at the output of the clamp does not exceed 6 dBm, is used to generate the external electromagnetic field and corresponding currents.

A system integrating a 2.5GBASE-T or 5GBASE-T PHY may perform this test to evaluate anticipated performance in regulatory test environments. Operational requirements of the transceiver during the test are determined by the manufacturer.

NOTE—The 6 dBm limit includes the 10% frequency-dependent variation mentioned in Annex 113A.3.

126.5.4.4 Alien crosstalk noise rejection

While receiving data from a transmitter compliant with specifications in 126.5.3, through a 100 m link segment compliant with the specifications in 126.7, a receiver shall operate with an Ethernet frame error ratio less than 7.8×10^{-9} for 800 octet frames with either a minimum IPG or greater than 220 octet IPG with four noise sources at the specified levels representing alien crosstalk, one connected to each of the four pairs. Independent noise sources should be injected into each MDI input using couplers that do not significantly alter the link segment characteristics. The injected noise shall have a flat spectrum within the following limits: a 3 dB bandwidth extending over at least 10 MHz to $200 \times S$ MHz and a power spectral density such that at the MDI port of the device under test the power spectral density of the injected noise is -137 dBm/Hz and -125 dBm/Hz for 5GBASE-T and 2.5GBASE-T, respectively. A flat noise source is chosen to model the sum of all alien noise sources. See Figure 126–37.




Figure 126–37—Alien crosstalk noise rejection test

The structure shown for injecting the noise in Figure 126–37 is illustrative and alternative approaches are possible. The loss of the coupling structure shown in Figure 126–37, which consists of two baluns and a

coupler, is approximately 2.5 dB. The overall insertion loss of the link segment together with the insertion loss of the coupling structure should be adjusted to match the insertion loss specified in 126.7.2.1 to within ± 0.5 dB. The balun-coupler-balun structure shown in Figure 126–37 can be replaced by resistively coupling a balanced noise source to the twisted pair using 500 Ω resistors. In either case, calibration of the test setup is required to confirm the overall insertion loss and the injected noise power at the MDI of the receiver under test.

126.6 Management interfaces

2.5GBASE-T and 5GBASE-T make extensive use of the management functions that may be provided by the MDIO (Clause 45), and the communication and self-configuration functions provided by Auto-Negotiation (Clause 28). Additional Auto-Negotiation requirements are set forth within this subclause.

126.6.1 Support for Auto-Negotiation

All 2.5GBASE-T and 5GBASE-T PHYs shall provide support for Auto-Negotiation (Clause 28) and shall be capable of operating as MASTER or SLAVE. All 2.5GBASE-T and 5GBASE-T PHYs shall provide support for Extended Next Pages as defined in 28.2.3.4.2 and shall support and use optimized FLP Burst to FLP burst timing as defined in 28.2.1.1.1, and `nlp_link_test_min_timer` and `link_fail_inhibit_timer` as defined in 28.3.2.

Auto-Negotiation is performed as part of the initial set-up of the link, and allows the PHYs at each end to advertise their capabilities (speed, PHY type, half or full duplex) and to automatically select the operating mode for communication on the link. Auto-Negotiation signaling is used for the following primary purposes for 2.5GBASE-T and 5GBASE-T:

- a) To negotiate that the PHY is capable of supporting 2.5GBASE-T or 5GBASE-T transmission.
- b) To determine the MASTER-SLAVE relationship between the PHYs at each end of the link.

126.6.1.1 2.5GBASE-T and 5GBASE-T use of registers during Auto-Negotiation

When Clause 45 registers are implemented, a 2.5GBASE-T and 5GBASE-T PHYs shall use the management register definitions and values specified in Table 126–15.

Table 126–15—2.5GBASE-T and 5GBASE-T registers

Register	Bit	Name	Description	Type ^a
7.0	7.0.15:0	AN control register	Defined in 45.2.7.1	R/W
7.1	7.1.15:0	AN status register	Defined in 45.2.7.2	RO
7.2, 7.3	7.2.15:0, 7.3.15:0	AN device identifier registers	Defined in 45.2.7.3	R/W
7.5, 7.6	7.5.15:0, 7.6.15:0	AN devices in package registers	Defined in 45.2.7.4	R/W
7.14, 7.15	7.14.15:0, 7.15.15:0	AN package identifier registers	Defined in 45.2.7.5	R/W
7.16	7.16.15:0	AN advertisement register	Defined in 45.2.7.6	R/W

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

Table 126–15—2.5GBASE-T and 5GBASE-T registers (continued)

Register	Bit	Name	Description	Type ^a
7.19	7.19.15:0	AN LP Base Page ability register	Defined in 45.2.7.7	RO
7.22, 7.23, 7.24	7.22.15:0, 7.23.15:0, 7.24.15:0	AN XNP transmit register	Defined in 45.2.7.8	R/W
7.25, 7.26, 7.27	7.25.15:0, 7.26.15:0, 7.27.15:0	AN LP XNP ability register	Defined in 45.2.7.9	RO
7.32	7.32.15:0	MultiGBASE-T AN control register	Defined in 45.2.7.10	R/W
7.33	7.33.15:0	MultiGBASE-T AN status register	Defined in 45.2.7.11	RO

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

126.6.1.2 2.5GBASE-T and 5GBASE-T Auto-Negotiation page use

2.5GBASE-T and 5GBASE-T PHYs shall exchange a MultiGBASE-T and 1000BASE-T formatted Extended Next Page, as specified in Table 126–16, immediately following the exchange of the Base Page.

Note that the Acknowledge 2 bit is not utilized and has no meaning when used for the 2.5GBASE-T and 5GBASE-T message page exchange.

Table 126–16—2.5GBASE-T and 5GBASE-T Base and Next Pages bit assignments

Bit	Name	Description
Base Page		
D15	Next Page	Defined in 28.2.1.6
D14	Acknowledge	Defined in 28.2.1.5
D13	Remote Fault	Defined in 28.2.1.4
D12	Extended Next Page	Defined in 28.2.1.3
D11:D5	Technology Ability Field	Defined in 28.2.1.2
D4:D0	Selector Field	Defined in 28.2.1.1
Extended Next Page (Message Code Field and Flags Field)		
M10:M0	Next Page message code	Defined in Annex 28C
T	Toggle	Defined in 28.2.3.4.7
Ack2	Acknowledge 2	Defined in 28.2.3.4.6
MP	Message Page	Defined in 28.2.3.4.5
Ack	Acknowledge	Defined in 28.2.3.4.4
NP	Next Page	Defined in 28.2.3.4
Extended Next Page (Unformatted Message Code Field)		
U31:U29	Reserved, transmit as 0	